A Formal Model of Infants' Acquisition of Practical Knowledge from Observation

by Alexandra Varga

Submitted to Central European University

In partial fulfillment of the requirements for the degree of Doctor of Philosophy in Philosophy

Internal supervisors: Prof. Gergely Csibra
and Prof. György Gergely
External supervisor: Prof. Michiel van Lambalgen,
Institute for Logic, Language and Computation,
University of Amsterdam

Budapest, Hungary
2012

CEU eTD Collection

Undersigned declares that this dissertation contains no material accepted for any other degrees in any other institution.

October 27, 2012

Undersigned declares that this dissertation contains no material previously written and/ or published by another person, except where appropriate acknowledgement has been made in the form of bibliographic reference.

October 27, 2012

CEU eTD Collection

CEU eTD Collection

Acknowledgements

First and foremost I would like to express my deepest gratitude to my supervisors, Professors Csibra, Gergely, and van Lambalgen, whose valuable feedback and constructive criticisms have helped me to overcome times and again the conundrums I found myself in along the way. I wish to thank Prof. van Lambalgen, without whom I would not have embarked on this octopus-like project in the first place, and who has constantly helped me to frame my closed-world reasoning, Prof. Gergely without whom I would not have thought of asking some of the right questions with the right words, and Prof. Csibra without whom I would have not found the right answers, perhaps any answers.

I would also like to thank Prof. ben Yami for insightful discussions and useful skeptical overtones. Many thanks to Dr. Király for supporting my attempts to try out the logical-philosophical tool in the EEG lab, for her moral support, and for her contagious positive look at things. I thank Prof. Kowalski for commenting on an earlier version of the model, and for providing much inspiration for further developments.

Last but not least, I wish to thank my friends and colleagues who have put up with my never-ending 'why'-list meant to verify what it is that people expect from explanations, and who patiently bore with my overly-enthusiastic faith in the omnipresence of rational thought.

CEU eTD Collection

CEU eTD Collection

Abstract

Social observational learning is a cognitive phenomenon that manifests itself throughout a lifetime, starting from early in development. From the first year of life human infants acquire practical knowledge from observation of other agents. Imitative behavior is one way by which learning is displayed. Infants imitate more or less accurately the actions that they observe other agents performing, and thereby develop their practical knowledge about actions, artifacts and their effects.

Imitative learning is selective, that is, human infants do not indiscriminately blind-copy; rather, different kinds of factors such as the understanding of the causal structure of the situation, the understanding of what counts as goals in the context, or the social attitude of the agents involved, modulate what is stored from the current observation, and used further in a different context. The empirical studies of Gergely et al. (2002) and Király et al. (2013) instantiate this claim in the case of 14-month-old infants. The results of such studies have been explained from several theoretical vantage points in the developmental psychology literature. Simplifying a little, task performance is accounted either via bottom-up mechanisms (motor activation and simulation processes), or via top-down ones (teleological processing of actions). The latter theories suggest that infants interpret the observed instrumental actions with respect to their goals. The inference-driven teleological interpretation underlies infantsÕ selectivity relative to the aspects of an action context that are worth learning. In the thesis I pursue by and large this rationalist approach to imitative learning.

The thesis is concerned with modelling observational learning as exhibited by infants in the above-mentioned studies. The primary aim is to construct a process model that accounts for the empirical results from within the rationalist explanatory framework. I assume that young agents' instrumental behavior is supported by practical reasoning processes; I justify this assumption by a review of empirical data from developmental psychology. The full-fledged model corroborates developmentalists' claim regarding the inferential support for imitative learning. It strengthens the argument that infants' learning of practical skills and artifact functions from observation of adult human agents is an instance of teleological instrumental rationality. However the roles of perceptual, motor, or memory (encoding and retrieval) capacities of the human mind are also considered; such capacities partake in the formalization. They constitute the infrastructure for the execution of practical knowledge acquisition via reasoning processes.

CEU eTD Collection

The model is built in the thesis in two steps. First I spell out the proposed conceptualization of the processes that support infants' performance in the head-touch task. Subsequently I provide a technical depiction of those processes in the logic of choice.

I put forward a view of practical reasoning as not necessarily explicit, i.e. consciously available to the agent. It is rather a normative method of context-relative information processing that deals with uncertainty. Moreover, reasoning does not hinge on the linguistic ability, as it has been traditionally assumed in the rationalist traditions in philosophy, linguistics, or cognitive sciences.

Goals and efficiency with respect to goals are hallmarks of practical reasoning. I set forth a wide-scope understanding of goals; goals may range from concrete physical state changes, to abstract volitional states. I focus in particular on epistemic intentions, such as the adults' pedagogical intention, or infants' drive to learn. The notion of efficiency is correspondingly broad; it is relative to human actions which are goal-directed in the above-mentioned broad understanding of goals. To express this, I use a conceptual distinction imported from AI, between achievement goals (here-and-now, episodic, concrete changes of state), and maintenance goals (abstract, generalizable, universally quantified over time intervals). I show the utility of the distinction for an explanation of selective imitative learning underpinned by teleological practical reasoning.

The two modes of practical reasoning presumed to generate infants' imitative behavior are teleological action explanation, and action planning. They are motivated by the epistemic drives to learn and to understand. Explanation is undertaken during action observation in the experimental phase. It amounts to interpretation of actions. Planning is modulated by interpretation and prescribes performance in the test phase. Both modes are realized by a planning algorithm, i.e. backwards reasoning from end-states assigned the teleological status of goals, to the actions that should be performed. I propose to view the human capacity for planning as fundamental. It is recruited for teleological action explanation.

All in all I account for imitative learning in the head-touch task via a single mechanism of practical reasoning, which is applied to an extended class of goals (physical, and also higher order, abstract goals). The model suggests a single-mechanism multi-level approach to teleological processing, which explains infants' learning of new skills and artifact functions. Social aspects of a situation, e.g., communication, ostension, are integrated in teleological explanation when physical details do not afford a satisfactory interpretation. The acquisition of practical knowledge from observation is pictured as an epitome of rational action that goes beyond practical efficiency.

Planning and teleological explanation are instances of reasoning with uncertainty; they are inverse problems. They are manageable by real agents with finite cognitive and time resources if construed as closed-world reasoning (CWR), i.e. reasoning with closed-world assumptions. The assumptions frame the inferential space and allow real-time computations despite informational complexity. The formal system that embeds CWR is constraint logic programming. It

CEU eTD Collection

works as an inferential engine over the event calculus language. These are the means used to formalize infants' reasoning processes.

The model speaks from the infants' perspective. Their observations and bits of prior knowledge are represented as event calculus clauses; the CLP inferential rules based on backwards chaining of actions connect observations and prior knowledge to behavior in the test phase.

A formal model has salutary effects in cognitive science. Although built on a limited dataset, it is productive, i.e., it has explanatory potential for related empirical results. Also it may provide better-structured and more detailed theories for the investigated cognitive phenomena. Consequently, fine-grained theoretical generalizations and empirical predictions are expected from the model. Some of these are discussed in the concluding chapter of the thesis.

Apart from its intrinsic interest, the formalization also has methodological implications regarding the role of logical systems in the study of the human mind.

CEU eTD Collection

CEU eTD Collection

Contents

Acknowledgements iii

Abstract v

List of Figures xiii

List of Tables xv

Introduction and road map 1

Chapter 1. Practical knowledge acquisition
– a perspective from developmental psychology 5
1.1. Manifestations of infant rationality. 7
1.1.1. The teleological stance. 7
1.1.2. Related domains of rationality. 15
1.1.3. The pedagogical stance. 19
1.2. The task to be modelled. 24
1.2.1. What is to be explained in the task. 27
1.2.2. A direction for the explanation to be pursued in modelling. 29

Chapter 2. Practical reasoning and the acquisition of practical knowledge 33
2.1. Focal points of human agents' practical reasoning. 36
2.2. A stepping stone for representing goals. 41
2.3. Planning. 44
2.4. Explanation. 49
2.5. Practical reasoning as closed-world reasoning. 55
2.5.1. Closed-world reasoning for goal attribution. 59
2.5.2. Closed-world reasoning for goal validation. 61
2.5.3. Closed-world reasoning for action-oriented planning. 63
2.6. Observational learning through practical reasoning. 64

Chapter 3. The Model I: The Reasoning Steps 71
3.1. Motivational factors for the teleological processing
in action interpretation and action planning. 75
3.2. Teleological interpretation of the context. 82
3.2.1. Action-to-goal reasoning: goal attribution. 83

CEU eTD Collection

3.2.2. Goal-to-action reasoning: offline planning. 87
3.2.3. Goal-to-action reasoning: plan recognition. 91
3.3. Online planning. 96

Chapter 4. Some logic for a refined account of the behavioral phenomena 101
4.1. Constraint logic programming and the event calculus
for a theory of observational learning in preverbal infants. 103
4.1.1. The representational environment of the event calculus. 105
4.1.2. The axioms. 109
4.1.3. A logical format for knowledge databases. 111
4.1.4. Some relevant aspects of the system's semantics and notion of validity. 113
4.1.5. Computational strategies in CLP. 120
4.1.6. The role of integrity constraints in computations. 126
4.1.7. A concrete example of CLP derivations. 129
4.2. Cognitive relevance of CLP. 132
4.2.1. Logic as a formal expression of learning in the head-touch task. 133

Chapter 5. The Model II: The Reasoning Formalized 139
5.1. The ontology of the head touch task. 139
5.2. Database for teleological interpretation in a Communicative context. 142
5.2.1. Integrity constraints on teleological interpretation. 143
5.2.2. The scenarios. 147
5.3. Goal derivation. 150
5.4. Offline planning. 153
5.5. Plan recognition. 157
5.5.1. Plan recognition in the HandsOccupied condition. 158
5.5.2. Plan recognition in the HandsFree condition. 164
5.6. Reasoning in the test phase: online planning. 168
5.7. 'What if she doesn't talk to, nor look at me, nor try to show me anything fun?' 171
5.8. Generalizability of the formalization. 172

Chapter 6. Concluding remarks 179
6.0.1. Sketching an answer to 'why model?'. 183

Bibliography 189

Appendix A. Electrophysiological signatures of teleological action processing and
memory encoding of novel means actions 201
A.1. Method. 201
Stimuli. 201
Experimental design. 202

CEU eTD Collection

A.2. Objectives. 204
Results in the pilot experiment. 204
Hypotheses. 204

CEU eTD Collection

CEU eTD Collection

List of Figures

1 Three types of inference that infants can draw based on a teleological representation of actions. (Reprinted from *Trends in Cognitive Sciences* 7(7) July 2003 with authors' permission.) 12

2 The objects used in Experiment 3 of Futó et al. (2010). B is the double-function object used during familiarization, and in the one-object presentation in the test phase. C are the two objects in the two-object presentation in the test phase. (Reprinted from *Cognition* 117 (2010) with authors' permission.) 16

1 Three-valued connectives 116

2 Reduction of maintenance goal represented as an integrity constraint. 130

3 Reduction of achievement goal represented as goal clause. 131

1 Achievement goal derivation from the maintenance goal $change_d$. 151

2 Goal derivation for the agent's contact action. 152

3 Offline planning: Derivation from the goal *light-box is on now*. 155

4 Offline planning: Resolution of the contact event query. 156

5 Plan recognition: Reduction of the background achievement goal $scarf_on$ in HandsOccupied. 159

6 Plan recognition: Reduction of the foreground goal *light-on* in HandsOccupied. 160

7 Plan recognition: Resolution of the contact query in HandsOccupied. 161

8 Branching derivations from the goal *light-box is on at time 9*. 162

9 Failing default branch in reduction of goal *light-box is on at time 9* in HandsOccupied. 163

10 Reduction of background goal $scarf_on$ in HandsFree. 165

11 Resolution of the knot query in HandsFree condition. 165

12 Reduction of goal *instantiate learning in own instrumental behavior*. 170

13 Self-Discovery: Goal derivation for the agent's contact action. 175

14 Self-Discovery: Branching derivations from the goal. 176

15 Self-Discovery: Failing default branch in plan recognition. 176

CEU eTD Collection

CEU eTD Collection

List of Tables

1 The experimental results under focus
Top row: Gergely et al. (2002), Bottom row: Király et al. (2013, 2004). 26

1 The reasoning processes involved in task performance. 74

CEU eTD Collection

Introduction and road map

The current work is an interdisciplinary approach to the epistemological problem of knowledge acquisition. I use logic as an explanatory tool; I apply it to model formally psychological results regarding infants' observational learning about actions and artifact affordances. The narrower object of interest is the acquisition of practical knowledge in infancy. In even more concrete terms, I model and thereby explain infants' acquisition of practical knowledge about a novel artifact (i.e. its functions, and action affordances) from the demonstrative behavior of an adult agent.

The acquisition of knowledge from observation of other human agents' actions is known as social observational learning. It is a well-documented phenomenon in the empirical literature in developmental psychology, whether it is concerned with language acquisition (Tomasello 2000), artifact categorization and object individuation (Futó et al. 2010), or practical knowledge acquisition (Gergely et al. 2002; Brugger et al. 2007). I take the latter to denote learning what to do, when, what for. Among various behavioral manifestations of observational learning, I focus on imitative behavior and empirical studies thereof. Imitation provides a robust learning ability from the first year of life – according to Barr and Hayne (2003), infants learn from observing people around them and doing what they do one to two behaviors on a daily basis. In connection with the intrinsic link between imitation and learning, it is interesting to note that the Chinese and Japanese ideographic characters for the verb 'to imitate' also mean 'to learn' (according to Nehaniv and Dautenhahn (2002a)).

Imitative behavior provides a diagnosis of infants' learning[1]. Imitative behavior is a measure of the addition of novel generalizable information to an observer's knowledge database. The experimental findings of Gergely et al. (2002) with respect to infants' imitative behavior ground the current modelling endeavor, i.e. they constitute the facts to be accounted for. The authors explain imitative performance as a result of infants' teleological interpretation of observations, i.e. of the artifact-directed actions performed by a human agent on a novel object. But imitative learning appears even more complex than that. In further related experiments, social factors of the context of presentation of the novel artifact, e.g., the communicative attitude of the observed actor, are shown to have a decisive influence on what is learnt from observations (Király et al. 2004, 2013, to be submitted). As a result of explaining observed actions in a goal-centered fashion, infants get to acquire knowledge about a novel artifact – what can be done with it, by

[1]I wish to thank Gyorgy Gergely – personal communication – for the suggestion of conceptualizing the explanatory role of learning for the experimental results in imitation – emulation studies in terms of diagnosis. An argument for this take on the issue is presented in Section 2.6.

CEU eTD Collection

what means, under what circumstances; precisely what infants learn is a matter of how the actor communicatively frames her instrumental artifact directed (i.e. transitive) behavior.

My main goal in the thesis is to develop a computational model that provides an abstract representation of the teleological reasoning processes involved in imitative learning. This representation is meant to provide an integrated explanation of empirical findings. The model does not attempt to give an exhaustive account of imitative behavior. For instance, the correspondence problem of imitation, i.e.

> how can the imitator identify, generate, and evaluate appropriate mappings ... between its own behavior and the behavior of the model? (Dautenhahn and Nehaniv 2002),

although an important issue of ongoing debate (Brass and Heyes 2005; Nehaniv and Dautenhahn 2002b,a) is beyond my scope. What the model addresses head on is infants' capacity to

> interpret and reproduce an agent's selection of some novel action as a means to achieving her goal (Gergely and Jacob 2012),

under the assumption that observations are mapped onto observers' corresponding motor programs, without spelling out the mechanism by which this is done[2]. Furthermore, the model attempts to pinpoint the effect of interpretation on infants' action (re)enactment by contending that it is prescribed by planning computations. To this end it uses a non-classical formalism, i.e. constraint logic programming, as an inferential engine over the language of the event calculus (van Lambalgen and Hamm 2005; van Lambalgen and Stenning 2008).

The logical tool formalizes goal centered explanation of observed instrumental actions, and its effect on subsequent planning of own actions performed on a novel object. As McCarthy and Hayes (1969) put it, modelling provides

> a method for better understanding of the intellectual mechanisms and the identification of new ones (p.465).

It supports a clearer understanding of empirical findings, which is generalizable, and which fosters the development of new experimental predictions. The model is meant to connect in a principled way the input variables, i.e. infants' observations of another actor's instrumental behavior using a novel object, to infants' behavioral output, i.e. their actions onto the object, by spelling out formally the reasoning processes that occur given the input. By so doing it corroborates the developmentalists' contention that imitative learning is indeed supported by inferential processes. Imitative learning as evidenced by the results of Gergely et al. (2002) is best understood in a framework according to which infants act upon the novel object *as if* they understood the observed actions by explaining them teleologically, thereby storing a certain teleological structure of the context, and planned their own actions according to its retrieval. The methodological upshot of the modelling endeavor is that the use of logic is salutary for the

[2]The interested reader is directed to Nehaniv and Dautenhahn (2002b) for an insightful mathematical formalization that provides a framework for the mechanism by which different systems match their behaviors.

CEU eTD Collection

empirical sciences, contrary to the popular line of research of Wason (1968) or Byrne (1989) according to which human agents do not reason by logical rational standards, and hence logic is irrelevant.

I lead off the discussion in Chapter 1 with a presentation of relevant issues in developmental psychology. I introduce a conception of imitative learning in infancy that manifests itself behaviorally as action imitation and goal emulation. Since the model formalizes reasoning processes, I devote a section to presenting several studies and the corresponding theoretical frameworks that speak for the rationality of the youngest members of our species. It is meant to justify the psychological plausibility of an inference based model of imitative learning. The primary locus of concern is infants' use of teleological reasoning. I contrast the rationalist account of action interpretation with respect to goals with lower level accounts, and argue that the former should be pursued because of its explanatory productivity. Nevertheless, I suggest that the virtues of the latter approaches should be taken into account when modelling, as long as the executive role of inferences is at all times maintained. I detail the preconditions of the teleological capacity and refer to studies that show their availability at early ages (in the first year of life). I also present some relevant outcomes of goal based action processing. I then introduce the task of Gergely et al. (2002) and the relevant subsequent manipulations of Király et al. (2013) that I model. I end the Chapter with outstanding questions that the model aims to address, and with a certain direction for the modelling explanatory strategy.

Chapter 2 outlines the logical philosophical framework of the strategy. The theory is backed up by real world examples. I elaborate on the topic of practical reasoning, and state what I mean by high-level, inferential processes. I explicate their most relevant features for the model. Goals, the main element of concern of practical reasoning, and their representation by human agents are devoted a separate section. I then zero in on two modes of manifestation of practical reasoning, namely, action planning and action explanation. Planning is an equivocal term for both a cognitive function, and a cognitive mechanism that subserves human agents' teleological explanation practices. The context of usage should clarify what is meant each time I use the term. I present an argument for the utility of planning as an explanatory algorithm. I discuss the computational concerns of planning and explanation, and suggest a solution – reasoning with closed-world assumptions. I show how practical reasoning becomes computationally tractable, and thus formalizable, if understood as closed-world reasoning. In the end, I discuss the acquisition of practical knowledge as a cognitive outcome of teleological practical reasoning.

Chapter 3 instantiates the considerations laid out in the previous Chapter for the concrete case of infants' action processing in the task of concern. I first spell out the two dimensions along which the model attempts to refine the developmentalist explanatory strategy of the empirical findings. One has to do with the use of goal representations at different levels of abstraction, the other one – with the inclusion of planning computations in the account of the empirical data. I then describe in detail the reasoning processes proposed to underpin infants' behavior

CEU eTD Collection

in the experiments under focus. In the end I argue that a model that posits infants' reliance on these processes is justified given the data presented in Chapter 1.

In Chapter 4 I introduce the logical system chosen as a formalization method for closed-world practical reasoning: constraint logic programming derivations over knowledge databases represented in the language of the event calculus. It is a non-monotonic formal system, particularly suitable to express reasoning with uncertainty such as teleological practical reasoning has been explained to be. Throughout the Chapter, I emphasize the manifestations of closed-world reasoning, which recommend it as an appropriate formalism for planning and explanation. I describe the three parameters of a logical system: syntactic (the inferential rules), semantic (the notion of truth), and the notion of validity. A particularly important notion to be introduced is that of integrity constraints. It serves the formalization of the goals represented at different levels of abstraction. In order to put some flesh on the dry logical notions, I provide a concrete example of a derivation in the system, using the same working example introduced in Chapter 2. I end with an argument why the system is cognitively relevant and thus apt to be used in modelling human data. I emphasize the suitability of this particular logical tool to formalize learning.

Chapter 5 presents the full-fledged logical model. It follows the sequence of the reasoning steps put forth in Chapter 3. The input variables are expressed as components of a knowledge database in the language of the event calculus. There are four databases, denoting observations in the four experimental conditions. I elaborate the two most complex ones. Derivation trees along the lines of constraint logic programming are presented for each reasoning step, for the two more complex conditions. They show how, starting from each condition-dependent observational database, an agent is prescribed to behave precisely as the majority of the experimental subjects do behave. I explain what changes upon observation of the input in the other two conditions. I end by presenting a concrete case to which the theory generalizes; I apply the model to explain a recent experimental manipulation to the head-touch task in Gergely et al. (2002), i.e. the study of Pinkham and Jaswal (2011).

I wrap up in Chapter 6. I first zoom in succinctly on the answers proposed by the model for the outstanding questions raised by the empirical findings. I show the practical utility of the formalization by reference to the design of a currently undergoing event-related brain potential (ERP) study based on refined predictions with respect to action understanding (the research proposal and some preliminary confirming results are presented in the Appendix). The predictions arose from the current model of teleological action interpretation. Relatedly, I emphasize new issues that are worth investigating, to which I was led by the research involved in assembling the full fledged formal model.

CEU eTD Collection

CHAPTER 1

Practical knowledge acquisition – a perspective from developmental psychology

I start with some conceptual and terminological clarifications with respect to imitative learning. Imitative learning overarches two related behavioral phenomena, imitation and emulation. The terms have been used in the empirical literature in developmental and comparative psychology in various ways, emphasizing different nuances; I briefly present the understanding that I adopt for the current discussion. Emulation as a learning mechanism is a term coined by Tomasello (1990); it denotes copying the results of actions, while disregarding the instrumental actions themselves. Learning by imitation, on the other hand, means that the observer performs a recognizable copy of the witnessed behavior (although not necessarily an accurate copy), which essentially includes the observed agent's instrumental actions (Whiten and Ham 1992). It is crucial for my current purposes to acknowledge that the distinction between imitation and emulation upon observation of goal directed actions resides in the teleological status of the information that the observer reproduces from the observed agent (Csibra 2007). While emulation refers to adopting the goal, and bringing it about via the observer's own behavioral strategies, imitation refers to performance of the observed means actions. Csibra (2007) makes a helpful suggestion along these lines, to view emulation as action reconstruction, and imitation as action duplication. Because of the hierarchical nature of teleological (means – ends) relations[1], the contrast between the two phenomena is not absolute. I endorse the suggestion that they be understood as the two ends of a continuum of decreasing fidelity of action reproduction (Whiten and Custance 1996). Emulation is a less faithful copy of an action, because of its primary focus on the action's goal. Imitation is close reproduction of the witnessed action. Imagine, for instance, an agent who takes a coffee mug from the cupboard in the morning, pours coffee in it, and drinks. An alien observer who takes a wine glass and drinks coffee from it, is said to emulate the agent's goal to drink coffee, without thereby imitating precisely her method for so doing. The alien has thus learnt that in the morning agents on Earth drink coffee. Another alien who takes a coffee mug and pours coffee in it imitates the mug-taking action, while at the same time emulating the goal to drink coffee. She has learnt not only that agents on Earth drink coffee in the morning, but also that they drink coffee from coffee mugs. If the mug she takes is of a different kind than the one taken by the human agent, we would still say that she imitated the mug-taking action (less accurately), but we would think that she is more concerned about the finality of the action than about the action itself. Both aliens have shown evidence

[1] A more in depth discussion of the hierarchy is in Chapter 2, in particular in Section 2.2.

CEU eTD Collection

for learning, the former via emulation, the latter via imitation. The bottom line is that the behavioral manifestations of imitative learning of practically relevant information are twofold, as goal emulation and means imitation.

Meltzoff (1988) ran a seminal study whereby imitative behavior is suggested to be not only a nonverbal measure of memory and representation (as previously believed), but also to provide a potential avenue for social learning. More specifically, because the experiment looks at infants' ability to perform imitation of artifact-directed (i.e., transitive) instrumental actions demonstrated by an adult agent, it gives a good example of practical knowledge acquisition from observation. I describe the experiment in what follows. 14-month-olds were shown six novel toy objects. In the experimental condition an adult model demonstrated a critical action upon each of them that produced a salient effect by activating the objects. There were also two control conditions – the baseline control, where infants were handed the toys with no demonstration, and the adult-manipulation control, where the adult handled the objects without performing the critical means action. The rationale of the study is that if infants produce more actions that correspond to the demonstrated one in the experimental than in the control conditions, the conclusion of imitative learning is warranted. The criteria for novelty given in this article have been adopted in subsequent studies assessing imitative learning from exposure to new instrumental action contexts. I list below the elements of novelty.

(1) the triad object – action – effect has not been seen by the infant before;
(2) the action has not been performed before;
(3) the action has not been imitated before;
(4) although possibly performed in passing, it has not become a familiar, well practiced game for the infant;
(5) it has not been previously put in relation with a particular class of objects before;
(6) it occurs with near 0 probability in spontaneous play during a baseline period.

Among the six triads object – action – effect, one was particularly engaging for the infants, namely a light panel (a high wooden box with a translucent orange plastic panel for a surface), which is illuminated by a light bulb inside the box being pressed by the experimenter with the forehead. In the test phase, which took place one week after the demonstration session, 67% of the infants in the experimental group imitated the novel head action, whereas no infant from the control groups did so. Imitation was scored when infants did the following:

(1) lean toward the panel,
(2) strain to touch it,
(3) miss contact by no more than 10 cm.

The results led Meltzoff to conclude that observed instrumental actions guide infants' subsequent behavior, and hence that

> imitation is well enough articulated to play an important role in early learning and development (Meltzoff 1988, p.475).

CEU eTD Collection

This Chapter is concerned with empirical data from developmental psychology, and with developmentalists' proposed explanatory theories. I acknowledge from the beginning an empirically warranted preference for cognitivist – rationalist explanatory strategies for infants' learning behavior. In Section 1.1 I review experimental findings that witness manifestations of rational behavior at very early ages, and the corresponding theories. In particular, I discuss teleological reasoning, and I go on with related abilities that either support it or apply it to certain domains: counterfactual reasoning, causal reasoning, artifact categorization, and planning. I end with a discussion of infants' proneness to learn from observation of communicating agents, labeled in the literature the pedagogical stance. Section 1.2 introduces the experiments that the model addresses. The two experiments in question are based on manipulations of the design of

1.1. Manifestations of infant rationality.

A view of infants' cognitive development as guided in a top-down manner by rational capacities is in line with the theory theory (Gopnik and Wellman 1995) approach to the human psychological architecture. One of the hypothesis put forth by this approach is that the

> theory formation system may have evolved specifically to allow human children to learn (Gopnik 1998, p.102).

In other words, one of the main functions of the rationality skills evidenced even in the youngest members of our species is epistemic in nature, and refers to the acquisition of knowledge. Human agents are naturally born learners, and this capacity exhibits much of their rational endowment.

1.1.1. The teleological stance. The basic aspect of action processing that may result, under certain conditions, in imitative behavior was hypothesized by Gergely et al. (2002) to be an inferential use of the teleological stance to make sense of the observed actions. I elaborate first the theory, and then introduce supporting empirical results.

The teleological stance with respect to action understanding is an interpretive scheme; it is the label assigned to the teleological obsession (Csibra and Gergely 2007) shown by humans in their action-related interpretive practices. It refers to infants' expectation that actions have a goal, i.e. that they are performed in order to bring about a particular change of state in a given context (Csibra and Gergely 1998, 2007; Gergely and Csibra 1997, 2003). This expectation manifests itself in teleological interpretation of action contexts: goals are assigned to actions, and actions are understood as means to fulfill goals in a particular context. The input for processing is either an observed action (action-to-goal direction of processing), or a (previously computed) goal state (goal-to-action direction of processing) (Csibra and Gergely 2007). Goal-centered action processing aggregates a triad composed of an observed action, relevant aspects of the context of observation, and a particular action effect (Csibra 2003; Gergely and Csibra 2003). The outcome of bidirectional processing is an explanatory structure whereby an

CEU eTD Collection

observed action is categorized as a means for the end-state conjectured as its goal in the context. Teleological interpretation helps infants to make sense of an observed context by linking means actions, goals, and contextual features. If two elements of the triad are available, the third can be extracted. It is a computational strategy for explanation, and online prediction of goals (when presented with actions) and means actions (when a goal is known), as well as for social learning about actions and their affordances. It is this latter, epistemic function of teleological interpretation that Gergely et al. (2002) refer to when suggesting that the teleological stance is involved in infants' imitative learning.

Human agents' teleological hierarchies are to a great extent context dependent – almost any state of affairs can be a goal in certain circumstances, and there are usually more than one means for its pursuit. Intelligent instrumental action then has to deal with multifinality, and equifinality. An action may lead to achieving several outcomes in different circumstances (multifinality), and a desired end-state can be achieved via several action courses depending on the circumstances (equifinality). These are difficulties that make teleological processing an inverse problem (Csibra and Gergely 2007). Inverse problems are such that the information available does not constrain the sought solution sufficiently. For instance, according to Marr (1982), since one 2D image can be produced by an unlimited number of 3D arrangements, the visual system must be somehow equipped to solve the inverse problem of singling out a tridimensional interpretation for the 2D projection in the retina. Similarly, an action can be performed to achieve a large number of goals among its multitude of effects, so in order to assign it a goal an observer must select wisely among those. The problem holds as well in the converse direction of teleological processing, goal-to-action: because several actions can lead to one and the same effect, choosing the appropriate thing to be done in a given context in order to achieve a desired end-state is not trivial. Additional constraints on the processing of teleological queries are needed to adaptively deal with the inverse problem inherent in the complex human teleology.

Teleological interpretation of actions as an inverse problem is unsolvable, on the one hand, by basic perceptual and storage mechanisms in a bottom-up manner, and on the other hand, by inferential deductive methods. Teleological reasoning (Kelemen 1999a; Gergely and Csibra 1997, 2003; Csibra and Gergely 2007) fares better as an interpretive tool, particularly in the case of new action contexts where bottom-up mechanisms stumble. Deductive strategies are not suitable for teleological processing as an under-constrained problem of reasoning with uncertainty; this issue is elaborated in Chapter 2.

Regarding the non-inferential strategies, Csibra and Gergely (2007) argue that the two such proposals in the recent psychological literature – action-effect associations (Hommel et al. 2001; Elsner 2007; Elsner and Aschersleben 2003), and simulation or direct matching procedures (Sommerville and Woodward 2005; Meltzoff 2002, 2005, 2007) – cannot account for goal-based action processing especially when the context of observation involves novelty. The theory of action-effect associations has it that actions are represented in association with the

CEU eTD Collection

effect that is most frequently observed to follow their performance. Highly frequent effects are categorized as the actions' goals precisely because of their high frequency of co-occurrence. Once an association between an action and its effect is stored, representation of a desired effect directly triggers activation of the corresponding motor program, and, conversely, perceiving the action automatically anticipates the effect. Thereby action-effect associations fulfill two functions: they help agents to decide online which are the goals of observed actions, and to form expectations about the actions to be performed by another agent once a goal has been assigned. Perception of frequently observed behaviors and resultant states, and associative storage, would thus provide an adequate mechanism for observational learning in terms of goals or functions, and means actions. However, the mechanism lacks flexibility, and is thus not likely to support teleological understanding of novelty where experience does not help much. Moreover, it does not address head-on the issue of multifinality because it does not provide a mechanism of goal selection among different effects of actions.

Simulation procedures propose that means – ends interpretation of action contexts is subserved by the observers' own experience with the world, projected over the observed agents. It is supported by empirical evidence according to which infants are better able to assign goals to actions that they themselves can perform, than to those which are not yet among their abilities (Sommerville and Woodward 2005). Own experience is cast over observed agents and facilitates teleological understanding. A "like me" assumption (Meltzoff 2002, 2005, 2007) warrants this transfer: if the observed agent is assumed to be like me in relevant aspects (e.g., motor constraints, preferences), then it suffices to simulate what I would have done or what I would have wanted in a similar situation in order to come up with an acceptable teleological interpretation. Simulation however does not fare well as a mechanism for teleological interpretation when no first-person experience with a similar situation is available to the observer. That is, like the action-effect associations theories but for different reasons, it has difficulties to account for successful processing of novel situations. Direct simulation is by and large a non-inferential procedure, based on automatic perception of similarity between the observer and the observed agent. In fact such behavioral theories are connected with neural accounts that speak for the bottom-up role of mirror neurons in mind-reading, empathy, and imitation (Meltzoff 2007; Rizzolatti 2005). It is possible that by inference-driven application of "like me" assumptions to novel situations, the scope of the theory is extended, and thereby its explanatory power covers novelty as well.

A non-deductive, yet inferential mechanism provides a broader account of interpretation via the teleological stance. The generalizability potential of reasoning recommends it as a more powerful explanation for behavioral phenomena than the procedures introduced above. In particular, it gives a satisfactory account of teleological interpretation of new contexts, or situations where prior experience is absent. For this reason it is recommended as a useful tool for explaining the acquisition of knowledge from observation. It sets forth solutions for the difficulties of

CEU eTD Collection

goal-based action processing through more productive normative principles. Teleological processing is an instance of reasoning with uncertainty. It deals with the epistemic opacity (Gergely and Jacob 2012; Király et al. 2013) of novel action contexts by adapting flexibly some general principles to relevant perceptual features of the context.

Csibra, Gergely, and colleagues hold that teleological reasoning about actions deals with the inverse problems by appeal to assumptions as a constraining factor. The assumptions are used to select and single out goal and means hypotheses, and thereby generate a satisfactory teleological understanding of action contexts[2]. The basic assumption is the principle of rationality: that actions are performed by rational agents, so they comply with rationality standards. It offers a heuristic to deal with the inverse problem of teleological processing, by selecting means among instrumental actions, and goals among outcomes in the circumstances observed. Rationality is operationalized very early on by the principle of efficiency: a rational agent pursues the most efficient action to achieve her goal in the context where she acts (Gergely and Csibra 2003; Csibra and Gergely 1998) (see also discussion in Section 2.5). An outcome qualifies as a goal if it is the state most efficiently brought about by an action in the observed context. An instrumental action is a means insofar as it is the most efficient one available for the desired outcome in the given context, i.e. insofar as it approximates the ideal of efficiency. In this theoretical framework, (physical) efficiency and rationality are often used interchangeably.

> The rationality assumption serves as ... an inferential principle guiding and constraining the construction of action interpretations ... The principle asserts that a teleological explanation is well formed (and therefore acceptable) if and only if the action realizes the goal state in a *rational (or efficient)* manner within the particular situational constraints (Gergely 2003, p.195) [added emphasis].

Teleological interpretation with the rationality assumption constitutes a normative appraisal of observed actions with respect to their outcomes in given circumstances[3]. Goals are expected to contextually justify goal-directed actions. The justificatory role of a hypothesized goal is verified by evaluation of causal relations in terms of efficiency within that context. The principle of efficiency thus promotes a teleological explanation of the context. Upon availability of an explanatory goal hypothesis, it guides the inferences that generate the means as the most efficient course of action towards the goal state within the constraints of the given situation. This is the action that *should* be performed for that goal, in that context. The normative relation between means actions and goals obtains if the action is inferentially rendered the most efficient one in the context of interest, i.e. if no more efficient alternative is available within the constraints of

[2]It is noteworthy that a similar strategy, based on constraining assumptions, is involved on the simulationist side, where the "like me" assumption narrows down the potential goals to those that I would have wanted, and the potential means actions to those that I would have done. The difference resides in that the inferential theory is more flexible because the efficiency principle doesn't tie action interpretation to one's own motor repertoire, so it makes novel goal-directed actions interpretable irrespective of what one already knows.

[3]A theoretical take on normativity is presented throughout Chapter 2.

CEU eTD Collection

the situation. More generally, teleological interpretation derives from two given elements of the triad means – goal – contextual features, what the third one should be. Each of the elements is justified, i.e. deemed necessary, by the other two. The context relative normativity of means – ends relations in this rationalist view of action understanding constitutes the main difference from contingent action – effect associations.

An ample body of empirical evidence shows that from the first year of life infants exhibit a tendency to interpret actions as being *for* a purpose (Csibra et al. 1999; Gergely et al. 1995; Woodward and Sommerville 2000; Meltzoff 1995). Teleological reasoning provides a good account of the computational mechanism at work in action interpretation, both theoretically (for reasons mentioned above), and empirically. The justified choice of what is worth learning from observations (Gergely et al. 2002) is but one example of a cognitive outcome of teleological reasoning used proactively. I will come back to that in Section 1.2.2. I instantiate the three kinds of goal-centered inferences that account for empirical findings. Many of these findings are obtained in a classical paradigm in developmental psychology – violation-of-expectation.

Let me first introduce briefly the rationale of the paradigm. Infants are presented with a stimulus while their looking times are measured. When looking times drop below a certain level, they are deemed habituated with that stimulus. This interpretation relies on the assumption that looking time indexes interest in a given situation; the longer one looks at something, the more interesting that something must be. Upon repeated presentation of a stimulus, infants loose interest and thus tend to look away after less time. Shorter looking times upon habituation are taken to demonstrate that the infants have built expectations to observe 'the same story' over again – they look away because they expect to see something that they already know, so there is no point to attend to the stimulus anymore. When their looking times reach a certain threshold of habituation, the stimulus changes with either a congruent/ compatible or an incongruent/ incompatible one. For instance, upon habituation with a certain object in a certain spatial location, a congruent stimulus would be the same object in the same place, while an incongruent one would be an empty space. The latter is incompatible with certain expectations that the infants are assumed to have built during habituation, in this case, the expectation of object permanence. An incongruent or incompatible stimulus violates their expectations. Hence they are surprised, and surprise heightens their interest. Consequently they look longer at the incongruent than at the congruent stimulus. Longer looking times are also referred to as dishabituation – infants dishabituate to incompatible stimuli, and, correspondingly, they remain habituated to compatible ones. This paradigm provides a way to test if infants do indeed build certain expectations: dishabituation to incompatible stimuli in contrast to compatible ones shows that they do.

In all of the studies to be referred to, infants were habituated to an event that included two differently sized circles moving on a computer screen. The test events presented two situations, of which only one was congruent with teleological inferences as described above. The question was whether infants' expectations during habituation are guided by use of the teleological stance to interpret the event. An affirmative answer is supported if infants look significantly longer at

CEU eTD Collection

a test event which is incompatible with the idea that actions are the most efficient ones for their goals within contextual constraints.

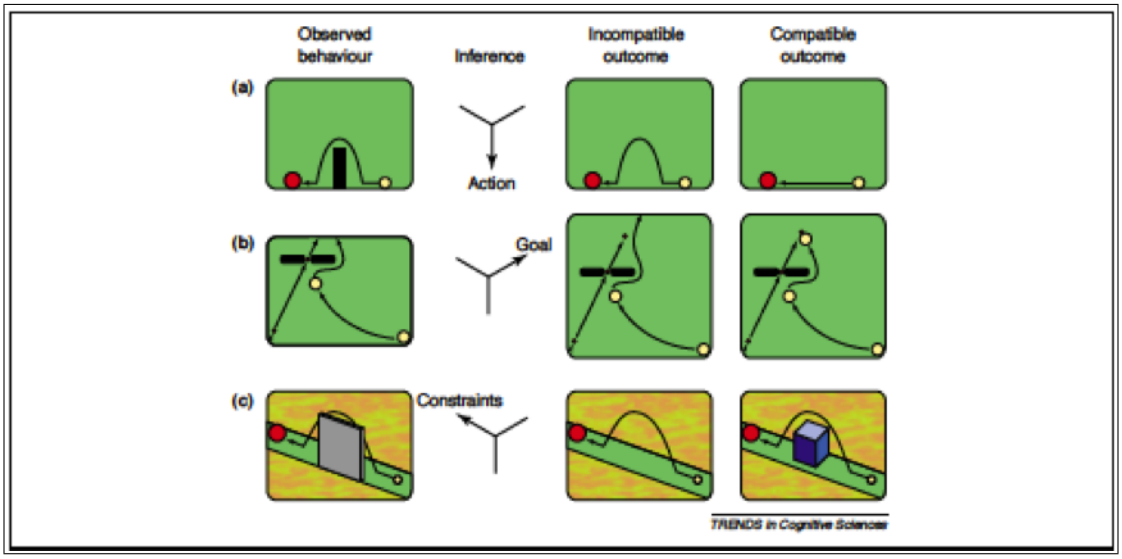

Figure 1. Three types of inference that infants can draw based on a teleological representation of actions. (Reprinted from *Trends in Cognitive Sciences* 7(7) July 2003 with authors' permission.)

I start with inferring goal states from observation of actions and contextual features. In experiment 1A of Csibra et al. (2003) 12-month-olds were habituated to an animated sequence of a large circle following a small one in a heat-seeking fashion. Adults had catalogued this as a chasing event. In the upper part of the screen there two black bars, with a small gap in between where the small circle (the chasee) but not the large one can go through. The large circle approaches the small one, and upon getting closer to the gap it changes direction and makes a detour of the right bar. The two circles disappeared from the screen before the babies could see any outcome of the chasing event. In the test phase of the experiment two different outcomes were presented. In the congruent, the small circle stops and the large circle 'catches' it. In the incongruent one, after the small one stops, the large circle continues its movement as if it 'didn't care' about the small one. 12-month-olds look significantly longer at the latter test event. Presumably they decided that the purpose of the event is a catch end-state, and if shown a different outcome they are surprised. The results suggest that the 2D event is interpreted as a 3D scene with solid physical object; such 'naive physics' kind of interpretation is a prerequisite of applying the efficiency principle to teleological computations. The goal can only be inferred if the situational constraints (i.e. circle sizes, gap dimension) are taken into account upon observation of the habituation event. That is, only if observers realize that the detour of the large circle occurs because the gap between the two bars is too small for the large circle to pass through, can they infer that its goal is to catch the small circle, and thus be surprised when it does not do so. The goal is the efficient outcome of a 'chasing' event in the current context, whereby the dimensions of different entities are taken into account relative to

CEU eTD Collection

each other. This latter issue had been checked in a control experimental group (in experiment 1 of the same study), which observed during habituation that the between-bars gap is large enough to allow the passage of the chaser. In this case, the detour is not justified by the inferred goal and the contextual constraints. The looking times of infants in the control group evidence that this is indeed how they interpreted the action: they do not look longer at the test events than they had in the habituation phase, and also they do not differentiate between the congruent and incongruent outcomes. It is as if observations during habituation do not support any teleological interpretation, hence the test events are not perceived as bringing in relevant information about the outcome of the situation. Such results confirm the hypothesis that infants assign goals to actions based on the expectation that it is the end-state most efficiently brought about by that action in the given context. Relatedly, they rule out lower level accounts of the looking time differences between the congruent and the incongruent outcome. Had it been perceptual complexity, similarity, familiarity, or symmetry preference that determined looking times, than infants in the control group should have looked longer at the test events (simply because they are less perceptually familiar and more complex than observations during habituation), and also at the incongruent than at the congruent test event (because the former is more complex and less similar with the habituation events).

The next type of teleological inference is when, based on knowledge of a goal state and of situational constraint, a means action is derived. An illustrative experiment whose results support the theoretical commitment to a reason-based use of the teleological stance for interpretation is the one in Gergely et al. (1995). Infants, 13-month old, observe a small circle following the larger one from the right side of the screen towards the left, where the large circle stands still. There is an obstacle in the middle of the screen, in between the two circles. The small circle starts rolling horizontally towards left, and is stopped midway by the obstacle. It turns back to the right-side starting position, then re-initiates moving, and jumps over the obstacle. Afterwards it continues moving in a straight line towards the large circle, until it contacts it and then it stops. Infants are likely to interpret the visible end-state which remains on screen for several moments (the small circle contacting the large one) as the goal of the small circle's behavior. For the test events the obstacle is removed from the screen. In the congruent one, the small circle follows a rectilinear trajectory towards the large one. The incongruent event shows an event more perceptually similar to the habituation one, namely, the small circle performing a jump in the middle of the screen, and continuing on a straight horizontal path towards the large circle. All infants observe both test events. Confirming the rationalist hypothesis, they dishabituate (i.e. look longer) when observing the incongruent test event, but not the congruent one. They are not surprised by a perceptually dissimilar action (going in a straight line) as long as the previously inferred goal is achieved. They are, however, surprised by a non-efficient action of the small circle in order to get near to the big circle – nothing calls for the jump in the middle of the screen. It could have done something more efficient for its goal (i.e. rolling on a straight horizontal path) so the small circle's behavior is not justified in the context of the test

CEU eTD Collection

event by the previously assigned goal. Again, 13-month-olds' longer looking times confirm the predictions of action interpretation via teleological reasoning.

Finally, when both the goal and the means actions are made available, the appropriate contextual constraints that justify the means given the goal can be inferred. A relevant test of the hypothesis is experiment 2 in Csibra et al. (2003). The experimental setup is very similar with that of Gergely et al. (1995). The relevant difference is that the obstacle in the middle of the screen was occluded. Infants thus saw the jumping action of the small circle, which led to the outcome of contact with the large circle. However, during the habituation stage, the crucial contextual constraint was hidden. In the congruent test event, the occluder was removed so that the obstacle could be seen. The jumping action was justified as an efficient action for the goal in a context where its path was hindered by an obstacle. In the incompatible test, there was nothing behind the occluder, which made the jumping action of the small circle redundant from the perspective of efficient goal-directed behavior. 12-month-olds' looking times were significantly shorter for the congruent event, indicating that they were not surprised by the obstacle because of previously having formed the corresponding expectation. This shows that infants in the first year of life are able to use teleological reasoning to infer contextual features relevant for the rationality of the goal pursuit.

All in all, infants looked longer across the line at the events that were incongruent with a normative evaluation of actions in terms of the goals that they are meant to bring about under particular contextual constraints. Their surprise evidences that upon observing instrumental actions they build teleological expectations according with principles of rational or efficient action. The under-constrained teleological inferences are supported by assumptions about rational agents' efficient goal-directed behavior. Infants manifest themselves as skilled, flexible reasoners by tailoring these general principles to the peculiarities of the context of observation; their expectations take into account physical contextual features. Another important aspect of these studies using animated geometric figures as stimuli shows that teleological interpretation of actions under rationality assumptions is not limited to human actions. Human action constitutes just one of the several domains of applicability of the teleological stance. It is precisely the domain of concern in observational learning, and thereby in the present modelling work.

I end with some remarks about the role of low-level cognition in the interpretation of human action with respect to goals. I presented an argument that the inferential use of the teleological stance is more productive an account of goal-based interpretation than the low-level proposals. Nevertheless I acknowledge that memory and perceptual processes facilitate teleological reasoning about human action. Stored associations between frequent effects of actions, perceptual salience, or own experience with action production, may highlight confidence in goal attributions to human agents. The inferential strategy works as an executive instance that ensures the productive usage of perceptual, motor, and memory capacities, e.g., their applicability to understanding novel behaviors. Interpretation involves the exploitation of an intricate notion of agency, whereby humans act in order to fulfill their intentions. The conception of human

CEU eTD Collection

agency seems to come in a pack with reference to mental states, such as intentions and desires (although see fervent resistance to such a model in Wittgenstein's *Philosophical Investigations*, and in anti-psychologist theory of action, e.g., Dancy (2000)), in virtue of which humans are thought to act instrumentally. Simulationist theories, for example, argue for a mentalistic understanding of human action, with which infants are pre-equipped (Meltzoff 2002). In contrast, the mechanism described here need not overload the core of the theory with mentalistic commitments. The actions of geometric figures are understood in terms of their goals, and this most likely does not require the possibility to put oneself in the observed agent's shoes, via some shared motor mechanisms and perceived body similarity. Consequently simulation of others' mental states is not a necessary precondition for understanding goal-directed actions. However, simulation, backed up by the use of teleological reasoning, may foster understanding human agents' instrumental behavior. Something like inferential simulation of mental states may be a further development of teleological reasoning. This idea aligns with the suggestion of Gergely and Csibra (2003), that teleological reasoning is a precursor of reasoning about intentional states of human agents. The physical goals of actions correspond to the mentalistic notion of agents' intended physical end-states; they are the content of intentional states. The teleological mechanism manifests itself initially, i.e. in the first year of life, in applications of the teleological stance to causal evaluations of the physical world. With time, it may gradually extend over infants' increasingly sophisticated knowledge about other agents' intentional states, and co-opt those in the explanatory endeavor of human actions. As a basic adaptation to intentional action understanding, the teleological stance model offers a non-mentalistic version of a generative action interpretation system that may have evolved in non-human species a well, and that is also manifest in the case of autistic children (whose difficulty with making sense of intentional actions in terms of corresponding mental states is a diagnosis criterion). Nevertheless currently converging evidence on early attribution of both motivational and epistemic states to other agents (Onishi and Baillargeon (2005) – false belief ascription based on attribution of motivation, Luo and Baillargeon (2007) – preferences ascription) suggests that normally developing human infants may in fact attribute the corresponding mental representations as well.

1.1.2. Related domains of rationality.

1.1.2.1. *Artifact categorization and normativity of means–goals relations.* Understanding the functions of objects or artifacts[4] involves teleological reasoning, since functions are goals brought about by action upon the object. The objects are part of the means, or the method for achieving the goal that is their function. When an object is generally used by agents in order to achieve a particular end-state, then most likely the end-state is the function of that object. Enduring and consistent usage of an object for a particular goal fosters understanding the goal as the function of that artifact. Such functional understanding of artifact kinds is evidenced in

[4]I use the two terms interchangeably throughout.

CEU eTD Collection

adults (Kelemen and Carey 2007); teleo-functional thought seems to be a 'crucial feature of intuitive theory building' (Kelemen 1999a, p.461). The skill seems to develop from early infancy. The issue is relevant for infant rationality because functional information is, like goal-related information, not available in perception – it can only be inferred based on a causal evaluation of the physical features of objects, and its recurrent or enduring occurrence in instrumental actions. Furthermore, the use of functions for kind-based categorization of objects shows refined integrative inferential skills.

In experiment 3 of Futó et al. (2010) 10-month-olds in the critical experimental condition were presented with one object on which two functions were successively demonstrated. At first infants heard a greeting in infant-directed speech. They were then familiarized with the object: a bell-shaped box with a handle, a dial, and three lights on top (B in Fig. 1.1.2.1).

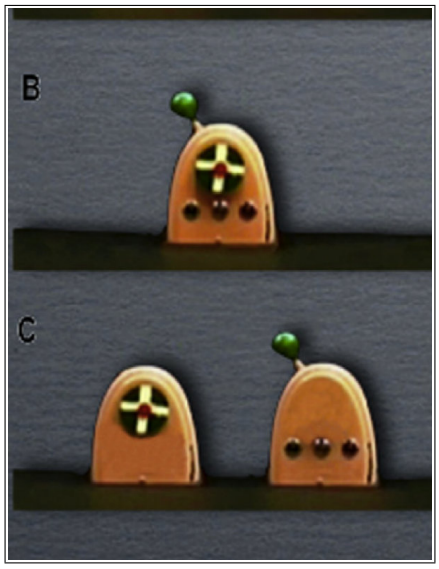

Figure 2. The objects used in Experiment 3 of Futó et al. (2010). B is the double-function object used during familiarization, and in the one-object presentation in the test phase. C are the two objects in the two-object presentation in the test phase. (Reprinted from *Cognition* 117 (2010) with authors' permission.)

When the object was introduced by a human hand from the left side of the screen, the green dial on it was turned, producing a sound effect. When the same object was introduced from the right, the hand pulled the handle, which resulted in the lights flashing. Each function was demonstrated three times in a row. Afterwards, an occluder covered the scene. During occlusion infants heard again the 'motherese' greeting they had heard at the beginning. In the test phase, the hand lifted the occluder and infants saw either the object they had been familiarized with (B in Fig. 1.1.2.1), or two objects (C in Fig. 1.1.2.1): a bell-shaped box without the dial, or a bell-shaped box without the handle and the top lights. The looking times of infants in the experimental condition were significantly longer when presented with one object than with two. Surprise by one object indicates that observation of two functions, although performed on the very same object, made infants build illusory expectations of two artifact kinds. The

CEU eTD Collection

results rule out reliance only on perceptual features in object categorization, since the two effects were presented on the very same object. They support the idea that functional information is what drives infants' categorization of artifacts (Tráuble and Pauen 2007; Kelemen 1999a). The teleo-functional information is inferred based on taking into account perceptually salient effects brought about by the observed artifacts. Interestingly, 10-month-olds' inference of two object kinds points to the use of an assumption according to which there is a one-to-one mapping between function and artifact kind, the

> assumption that basic level artifact kinds are defined by a single essential function (Futó et al. 2010, p.7).

This is an application of the expected normativity of means – ends relations to the domain of functional artifact understanding. Because an object and the action performed on it can be seen as the means to achieve the goal – the object's function, this assumption may point to infants' expectation of a one-to-one mapping between goals and corresponding means. The function of an artifact is the one end-state most efficiently brought about by the means. The experimental results referred to point to the use of such an assumption, which generates the expectation of two bell-shaped boxes, one that is for making sounds, and another one that is for making the lights light-up.

Evidence from even younger infants attests that, upon habituation with an action as a means for a particular purpose, they are surprised if, without any other contextual modifications, the same action is performed and a different outcome ensues. That is, 5- and 6-month-olds show the conservative expectation that an action is done in order to achieve one and the same goal, irrespective of whether the location of the object of the goal changes (cf. Leslie and Biró (2007); Luo and Baillargeon (2005)). Such results also point to a guise of the rationality principle as the expectation of flexible adaptability of goal-directed behavior to changes in the context where the goal object is (e.g., its change of spatial location).

The expectation of one-to-one mappings between goals achieved by artifact usage, and instrumental actions on the artifact appears to be a bijective function in the first year of life: two goals trigger expectation of two methods of achievement, and one means is expected to be performed for one goal only. This issue has not been directly investigated in empirical research so far, but it may be a promising avenue of research. The formal model proposed here recommends its pursuit.

1.1.2.2. *Causal reasoning.* A goal is a particular effect of an action. Hence teleological reasoning about actions is underpinned by action representations in terms of causes and effects. That is, grasping physical causal relations is a precondition of attributing goals to actions, and understanding some actions as means for some states of the world. Teleological categorization of a novel instrumental action as a means, and of the end brought about as a goal

> requires a causal analysis of the affordance structure of the observed actions and artifacts they involve (Csibra and Gergely 2007, p.72).

CEU eTD Collection

Teleological interpretation supervenes on perceiving actions and states of the world as causally related. Empirical evidence demonstrates that infants reason in terms of causes and effects about motion events and qualitative state changes before 12 months of age (Saxe and Carey 2006). Moreover, by 2 years, causal reasoning is co-opted for learning, e.g., about artifact functions (Gopnik et al. 2004).

Two principles that appear to guide causal reasoning in infancy are of special interest to the current modelling enterprise: contact causality, and inertia. The former refers to the fact that observation of spatio-temporal contiguity between events signals a relation of instantaneous causation: the effect ensues immediately after a momentary cause event. The commonsense principle of inertia states that a property persists unless acted on by an outside force. I briefly introduce the empirical evidence that shows infants' reliance on the two causal principles. I begin with contact causality. Because of its early development (around 6 months of age (Leslie 1987)) and its generative power, it has been argued that it serves as the foundation of human causal understanding (but see also Saxe and Carey (2006) for a critical discussion). When 7.5-month-olds observe repeatedly that contact between a moving object and a stationary one, is followed by the latter's being set in motion, they take the resulting motion state to be an effect of the contact event. This is evidenced by surprise (i.e. longer looking times) if contact is not followed by motion (i.e. if the effect does not co-occur with the cause), and if the second object starts to move although contact between the two is rendered impossible by the presence of a barrier (i.e. if the effect occurs without the cause) (Kotovski and Baillargeon 2000). The latter condition of surprise may also be taken as evidence for some rudimentary use of inertia in causal inferences.

Several months later, 10-month-olds' causal inferences show clearer signs of using the commonsense principle of inertia; the evidence becomes even more reliable in 1-year-olds (Spelke et al. 1994). For example, infants are surprised by objects moving rectilinearly, and changing their rectilinear path (e.g., by making a sudden U-turn) 'out of the blue', without any interference.

1.1.2.3. *Counterfactual reasoning.* Explicit counterfactual reasoning is evidenced later in development (after 3 years of age) and is intrinsically linked with a quite high level of linguistic proficiency. However empirical results attest that the capacity to use such reasoning implicitly in order to understand other agents' pretense play, is in place in the first one and a half years of life, even before infants engage themselves in such fictional games (Onishi et al. 2007; Leslie 1994). For example, in a study of Onishi et al. (2007), 15-month-olds were surprised by an agent's inconsistent instrumental behavior, even if it involved no real objects, but merely imagined ones; that is, they were surprised by an agent's incongruent pretense actions. After familiarization with an agent pretending to drink from a shoe, the infants observed her pretending to pour from a jug into one of two shoes. They looked reliably longer if the test event was inconsistent, i.e. if the agent pretended to drink from the other shoe than the one in which she pretended to pour. Given that other potential explanations based on perceptual highlighting were ruled out

CEU eTD Collection

in one of the 3 experiments in the study, infants' surprise stands as evidence for their capacity to apply general principles, e.g., of consistent instrumental behavior, in interpreting pretense events, which require counterfactual imagination of their object.

> To do something rational that is nevertheless not efficient in actual reality, one needs to act in a fictional or counterfactual world (Gergely and Csibra 2003, p.290).

It appears that before actually acting in pretense counterfactual worlds, infants are at least able to make sense of grown ups' pretense behavior via assumptions of rational behavior.

1.1.2.4. *Planning.* Planning is a notable kind of goal-based processing, which leads to proactive behavior. It is the mechanism underlying intelligent agents' action production. It is, in fact, the epitome of practical reasoning. Upon availability of a goal, a sequence of actions is derived taking into account specific features of the context. In an action-effect association framework, it can be argued that this derivation is non-inferential, based on direct triggering of actions by the stored association with a particular end-state (Hommel et al. 2001). However, a fine-grained adaptation of an agent's means to the context in which she acts is a manifestation of flexible processing of such stored knowledge, which is accountable for its planning is understood as a reasoning-based process. In the longitudinal study of von Hofsten and Fazel-Zandy (1984), infants' capacity to adapt their hands' position to the size and orientation of an object to be grasped was found to develop gradually between 4 and 9 months of age. After habituation with a goal to be achieved – grasping a toy, even at the youngest age the position of infants' hands tended to be adapted to the orientation of the toy so that they could attain the goal. More recently, Claxton and her colleagues (2003) have proven that 10-month-olds show an impressive level of sophistication in adapting the kinematics of their reaching actions to the particular goal they set in a given context (i.e. throwing a ball into a tub, or fitting it into a tube). That is, infants picked up a ball differently depending on the purpose of using the ball: either throwing it into a tub, or fitting it into a tube. The ability to plan motor responses suitably adapted for a desired outcome is available in early infancy, and shows a rapid development in the first year of life.

1.1.3. The pedagogical stance. Aside from own experience with the world, the prime source of knowledge is constituted by other human agents. If the former represents the hard way to learn (i.e. by trial and error – a slow and, at times, potentially harmful process), the latter is the easy way for learning, which saves learners from the inconveniences of the trial-and-error method. Humans seem to be particularly skilled at teaching the young, who, in turn, show a propensity to learn implicitly from the older, more experienced ones. This empirically documented fact is the prime justification for the developmentalist thesis of natural pedagogy (Csibra and Gergely 2009; Gergely et al. 2007; Csibra and Gergely 2006). According to it, human pedagogy is a species-specific system of learning from conspecifics, which allows for fast

CEU eTD Collection

and efficient transfer of knowledge, and thereby provides an answer to the puzzling query regarding the transmission of incommensurable amount of information in a very limited time that characterizes human culture. Human infants adopt the pedagogical stance by being particularly sensitive and receptive to pedagogical cues encountered in other agents' behavior. Pedagogy uses communication (verbal, as well as non-verbal) as the leading method. Communication fulfills

> . . . the primarily *epistemic* function of actively seeking out and cooperatively providing reliable, new and relevant information by knowledgeable adults to ignorant infants about the generalizable properties of referent objects and their kinds (Gergely et al. 2007, p.140).

Furthermore, I suggest that infants' pedagogical stance upon communicative demonstrations may be subserved, at least in the more complex cases, by reasoning mechanisms. More concretely, infants' acquisition of practical knowledge from observations of communicating agents is made possible by exercising the capacities for rational thought depicted in the previous Sections of this Chapter. Imitative learning instantiates this claim.

Behavior of other human agents in a communicative setting is assumed to be informative, and its content triggers credulous interpretations from the part of the recipients. Sperber and Wilson (1995) refer to ostension as the cues that point to an agent's communicative intention to convey new and relevant information. Ostensive signals direct infants' attention to the fact that it's them who are addressed so they should attend to the message. They point to an informative referential intention that the communicator has, and trigger the expectation to encounter relevant knowledge by decoding the message – "the presumption of relevance". If grounded in an episodic goal jointly pursued by the communicator and the addressee, the presumed relevance can be local. However, if relevance is not episodically grounded in a joint goal and novel referents are involved, then ostension triggers the presumption of enduring relevance, beyond here-and-now (the genericity bias in Csibra and Gergely (2009)). The expectation of enduring relevance naturally calls for willingness to put more effort in decoding the message to find out where precisely relevance may lie, than if it had been presented in a non-ostensive, aloof fashion. It also activates infants' proneness to knowledge acquisition, since the rational thing to do with information that promises to be relevant beyond the current circumstances is to store it for future usage.

The theory of natural pedagogy initiated a new line of research in developmental psychology. That is, social-communicative features of the context of stimuli presentation started to be manipulated in experiments, in an attempt to dissociate their effect on infants' behavior from the effect of other dependent variables. The manipulations corroborated the thesis that infants in the first two years of life, that is, preverbal infants, show robust signs of being adapted to play the receptive role in pedagogical relations (Senju and Csibra 2008; Futó et al. 2010; Southgate et al.

CEU eTD Collection

2009, 2010; Nielsen 2006; Brugger et al. 2007; Király 2009). They exhibit particular sensitivity to being addressed in communication by ostensive signals, such as infant-directed gaze or speech (motherese), and infant-contingent distal responsiveness. Sensitivity manifests itself in the preference to attend to, and to learn from such signals. They may even produce themselves communicative gestures asking for more information about items presented communicatively (Kovács et al. in preparation; Southgate et al. 2007; Begus and Southgate 2012). The sensitivity to such social cues manifests itself in building referential expectations (as tested in violation-of-expectation or looking time paradigms), in selecting relevant behaviors that are worth learning from observation (as tested in imitation studies), in performing proto-interrogative non-verbal behaviors (as tested in action production studies), etc.

Understanding adults' pedagogical intention may be looked at as a form of implicit mental state attribution to other agents. Infants learn *from* others' intention to convey expectedly relevant knowledge about something, without thereby necessarily learning *about* other agents as possessors of mental states, i.e. without explicitly assigning them intentions (cf. Gergely et al. (2007)). If natural pedagogy holds, knowledge acquisition is the natural response of an entity who takes the pedagogical stance when faced with another entity who behaves as if intending to teach novel relevant information. A primary manifestation of the pedagogical stance is that infants expect an agent's communicative intention to have an object, to refer to something. Evidence for this comes from studies showing that, for example, 6-month-olds follow an experimenter's gaze shift between two objects if preceded by ostensive communicative signals (infant-directed gaze and speech), but they do not do so if demonstration is aloof (Senju and Csibra 2008). Communication carries a referential import, and thereby activates infants learning mode about the referent.

Moreover, infants expect to be able to extract relevant, useful knowledge that is worth storing about a communicatively signaled novel referent. In relation with learning, 'relevance' means potentially generalizable knowledge, and not just transient, episodically important information, such as the spatial location of a referent. Infants expect that communicatively presented information, coded by ostension, bares enduring properties, which are worth retaining for use in future situations.

An interesting example of the effect of communicative behavior along these lines is the study of Topál et al. (2008). The authors put to the test of communicative stimulus presentation the classical results of Piaget (1954) about infants' perseverative error in the A-not-B task. Infants 8- to 12-months of age are engaged in a 'hide'-and-seek kind of game by an experimenter who repeatedly places a toy under one of two opaque containers A and B under infants' eyes[5], say, under A. After each hiding event, the infant is given the chance to retrieve it. In the test phase, she changes the hiding place to the other container. As if not noticing the change of location, infants persevere in mistakenly searching the toy under A. The task generated much debate and many explanations have been put forth throughout time. Topál et al. (2008) had

[5]This is why I use the term 'hide' in hide-and-seek between inverted commas.

CEU eTD Collection

10-month-olds play the game in three conditions, which differed along the social dimension of the infant-directed attitude of the experimenter: she was either communicating ostensively, or acting aloof (non-communicative condition), or hidden behind an occluder so that only her hands could be observed manipulating the toy and the containers (non-social condition). The Piagetian findings were replicated only in the communicative condition, where 86% of the infants kept on searching under the initial hiding place. The rate of the error dropped significantly in the other two conditions (43% in the non-communicative, and 59% in the non-social). The authors explain the error as an outcome of infants' expectation that communicative demonstrations introduce relevant information, which ought to be generalized over time. The context must have been interpreted as a pedagogical situation. Even a typically transient feature of spatial location, is interpreted in an ostensive context as a generic property of the referent object – the toy, and learnt as such. Infants' interpretation reads as a normative rule – *The toy ought to be under A*. Repeated observations of the mistaken character of the rule do not falsify it; communication induces some sort of immunity of the acquired information to falsification.

Action demonstration in a communicative manner, accompanied by ostensive cues from the performing agent, has been shown to shape imitative learning in the second year of life (Nielsen 2006; Brugger et al. 2007; Király 2009; Király et al. 2013, to be submitted). Imitation studies also provide a clearer example of the pedagogical stance because infants' imitative behavior provides a direct diagnosis of learning. In experiment 2 of Brugger et al. (2007) for instance, 14- to 16-month-olds were presented with two actions A and B that resulted in a salient outcome by using a toy tool. For instance, a rake could be used to obtain an out-of-reach toy. There were two experimental manipulations on action A. It could be either necessary (push a plastic tray to close a pit trap on the rake) or not necessary (push a plastic rake although the rake is intact) for B to yield a particular effect (e.g., obtain toy), and it could be either ostensively cued or not. Infants' performance of actions was measured in the four resulting action contexts. Results showed that most infants (53%) learnt only B as a means for producing the effect when A was neither ostensively cued, nor necessary. Even when it was necessary but not accompanied by communicative ostensive signals, only 13% of the infants did A prior to B. This result is significantly different from the one in experiment 1 of the same study, where 29% of the infants performed A when necessary to reach the goal. Even when it was unnecessary but communicatively presented, 47% of the infants' re-enacted A, as opposed to only 11% when it was not socially cued. Overall, infants were more likely to learn socially cued actions, even when they were causally unnecessary for goal attainment. In these cases, infants' credulous interpretation that the action is worth doing must have been triggered by the ostensive cues, given the fact that they do discriminate between causal necessity and causal redundancy (as shown in the results of experiment 1). The authors conclude that causal knowledge and sensitivity to communicative signals conveyed by the agent are two contributors to the imitative learning tendency, neither of which is dominant. Infants are selective with respect to the observed actions that they re-enact,

CEU eTD Collection

and selectivity is shaped by a causal analysis of actions in terms of necessity for desired effects, and by social-communicative cues from the experimenter.

Imitation studies offer a good example of how the pedagogical stance is subserved by teleological reasoning mechanisms. Upon demonstrations of instrumental actions in communicative settings, infants expect to encounter relevant information. This in turn leads to processing actions in their context, seeking for the relevant content of the agent's informative, pedagogical intention. Selectivity of imitative behavior that evidences learning shows that relevance is not automatically assigned to whatever the experimenter is observed to be doing. Had this been so, then in the above example infants should have performed A whenever ostensively signaled, irrespective of whether it was or not causally necessary for the goal. This shows that the pedagogical stance manifests itself by weighing observations with respect to goals. In this experiment, causal knowledge is involved in action processing that results in learning some, but not other actions. The effects of actions, with respect to which actions may be necessary or unnecessary, are salient end-states that are presumably categorized by infants as goals. The causal analysis is a form of teleological evaluation. But according to the arguments in Section 1.1.1 teleological processing is subserved by inferential mechanisms. The acquisition of practical knowledge from observations of agents performing instrumental actions ostensively attests that human infants' positioning on the receivers' place in pedagogical interactions is fostered by their capacity for rational thought.

Further support for the robust role of infants' pedagogical stance in inter-human interactions comes from recent studies where it is shown that not only are infants sensitive to communicative cues in an epistemic drive to learn, but they also produce communicative signals themselves (e.g., pointing) asking for relevant information about novel objects. In a recent study, Kovács et al. (in preparation) showed that 12-month-olds use pointing as a proto-interrogative means, which reflects the motivation to acquire useful practical knowledge. To this end, the number of infants' pointing gestures was counted after they received different kinds of feedback from the experiment with which they had been engaged in a triadic communicative interaction with a toy object: sharing attention, or teaching by providing new valence information about the object. The two kinds of feedback were divided between two experimental conditions (Sharing and Teaching). There were eight trials in each condition. Initially there was no significant difference between the rate of infants' pointing; however, in the last two trials they pointed significantly more in the Teaching condition. The results support the hypothesis that infants' produce pointing gestures as calls for knowledge worth acquiring from the part of the observed agent.

In conclusion, ostensive cues transmitted by adults in triadic interactions with infants (adult – infant – referent) trigger infants' natural predisposition to learn from the demonstration information about the referent that is expected to be relevant for future use. This predisposition is shaped in intricate manners, yet to be accounted for (see Section 1.2.1), by manifestations of infants' rational skills, such as teleological interpretations of behavior.

CEU eTD Collection

1.2. The task to be modelled.

The experiment of Gergely et al. (2002) is a follow up of Meltzoff's (1988), using the toy object that was most engaging for infants – the light panel. The authors' aim is to shed light on to the mechanisms that underpin infants' imitative behavior. The results of Meltzoff (1988) are consistent with the framework of natural pedagogy (i.e. infants are prone to do what is shown to them in a communicative fashion), but on a first glance at least they make one doubt the claims about teleological action processing. Why would a rational interpreter learn by reproducing an instrumental action that is not the most efficient for the physical end-state that it seems to be meant for?

I proceed with some terminological specifications meant to explain why it is that the odd artifact-directed action with the head demonstrated in the experiment of Meltzoff (1988) is not the most efficient for the resultant end-state. A *prepotent* action is a default response, most strongly associated with a certain action effect. It has priority over other response tendencies especially by virtue of developmental primacy. Commonsensically, human agents' default way to act upon artifacts is with their hands. Empirical results backup this claim: hands are prepotent agentive relata[6] for instrumental action from very early ages. Infants tend to act on objects with their hands (Király 2009). Moreover, infants assign causal power to human hands when they observe them to be involved in goal-directed actions (Saxe et al. 2005, 2007; Saxe and Carey 2006). Therefore I henceforth refer to hand actions as prepotent responses. A prepotent action is, until further evidence, optimal with respect to its goal: it is readily available in an agent's motor repertoire, hence it requires the least amount of processing effort prior to execution. In this sense it provides the most efficient route to a desired physical end state. Accordingly, any alternative to a prepotent action is suboptimal.

In Meltzoff (1988) infants' imitation of a suboptimal means for lighting up the light-box although nothing hindered a prepotent hand action does not seem, prima facie at least, the response of an entity who evaluates actions' efficiency for goals in the context of occurrence, and whose subsequent behavior is guided by this teleological interpretation (Gergely et al. 1995; Csibra et al. 1999). Gergely et al. (2002) modified the design of Meltzoff's experiment in an attempt to settle the issue. The unusual behavior was presented in two conditions: one in which it was called for by contextual factors and thus it was the rational or most efficient thing to do under the circumstances, and one in which it was not. Since the action performed with the forehead is unusual relative to a prepotent hand touch, the contextual factor manipulated was the availability of the agent's hands for action performance. The expectation of a rationalist take on imitative learning was that infants' interactions with the novel object would be significantly different when the hands are objectively constrained, relative to when they are not. This would

[6] An agentive relatum is a fine-grained, localized representation of agency. It is the body part with which agents act. More on agency in the preamble to Chapter 2.

CEU eTD Collection

show infants' selectivity with respect to the kind of knowledge they acquire from observation of manifold novelty depending on evaluations of the means action's efficiency for its purpose.

I describe the task of Gergely et al. (2002) in more detail below; throughout the thesis I refer to it as the head-touch task. I also depict an experimental manipulation by Király et al. (2004) and Király et al. (to be submitted) that challenges a purely rationalist account of the findings. The formalization in Chapter 5 models these versions of the task.

In the experimental phase, 14-month-olds are presented with the novel artifact – the funny looking table light-box from Meltzoff (1988) – in two conditions. The experimenter, an adult, enters the room, sits down at the table in front of the infant, and pretends to be cold. She rubs her upper arms, shivers, and shows corresponding facial expressions. She takes a thick scarf from the back of her chair and wraps it around her shoulders. In the Hands Occupied condition, she keeps the scarf around her with hands until the end of the demonstration session. In the HandsFree condition, she ties the scarf and markedly places her free hands on the table next to the light-box to ensure that their availability for action is noted. The adult then demonstrates the critical novel behavior – a forehead touch performed instrumentally as a means to turn on the light-box. She does this three times in a row. Half the infants see that the demonstrator's hands are occupied while executing the head action (Hands Occupied), the other half observe her acting with hands visibly free (HandsFree). After a one-week delay subjects come again to the lab for the test phase; they are given the chance to act upon the light-box themselves. The dependent measure is performance of a head action (scored according to the criteria used by Meltzoff (1988), listed at p. 6). Since the novel action is justified by achieving the goal light-on when the agent's hands are occupied, but not when they are free, the hypothesis supported by the theory of teleological reasoning is that infants, whose prepotent agentive relata are not constrained, would not perform the head action if having previously observed the Hands Occupied presentation.

The experimental results are plotted in Table 1. The findings of Meltzoff (1988) have been replicated only in the HandsFree condition. That is, imitation of the observed novel head action is selective depending on contextual factors: 69% of the infants in the HandsFree context, and only 21% in the Hands Occupied. In an attempt to light-up the light-box, all infants perform a hand action onto the light-box, most of them before even performing a head touch. Goal emulation with prepotent, optimal means actions is always the case. It is also worth mentioning once again that infants' unconstrained exploratory behavior upon the novel artifact, i.e. without demonstration by another agent, hardly ever includes similar head touches (see item (6) among the list of Meltzoff's criteria for novelty, at p. 1). This strengthens the interpretation of the results that imitation evidences the acquisition of means actions from the demonstration. Another relevant aspect regarding the imitation tendency is that infants' head actions appear in various forms, many times different from the demonstrated behavior (e.g., lifting up the light-box, touches with mouth, nose, or cheek), i.e. action reproduction shows low fidelity. Although the rationalist negative hypothesis regarding lack of imitation in the Hands Occupied context

CEU eTD Collection

Table 1. The experimental results under focus
Top row: Gergely et al. (2002),
Bottom row: Király et al. (2013, 2004).

Action performance	Communicative		Incidental	
	HandsFree	Hands Occupied	HandsFree	Hands Occupied
Head action	69%	21%	–	–
	60%	11%	11%	29%
No head action	31%	79%	–	–
	40%	89%	89%	71%

was confirmed, one still has to explain the high rate of reenactment of the forehead action in the HandsFree condition.

Király et al. (2004) and Király et al. (2013) performed more experimental manipulations from the direction of natural pedagogy, which were meant to clarify the whys and the wherefores of imitative learning. In experiment 1 of their study, Király et al. (2013) showed that the selectivity of imitative learning exhibited in the study of Gergely et al. (2002) is contingent on a ostensive action demonstration by a communicative agent. This involves that throughout the demonstration session the experimenter communicates with the infant seated in front of her using non-verbal ostensive-referential cues (eye contact, followed by gaze shift to the target object) as well as verbal ones ('Look, Baby, I'll show you something!'). In the original study of Gergely et al. (2002) the stimuli had also been presented ostensively. Király and her colleagues (2004, in press) manipulated the social context in which the novel object was acted upon. In these experiments, the original HandsFree – Hands Occupied conditions were subsumed to either a Communicative, or an Incidental presentation context. When the action was presented in a Communicative context, the previous results were replicated. In the Incidental conditions communicative cues are absent. The experimenter is more spatially distant from the infant – while she sits down at the table, the infant plays with its mother in a different corner of the room. Its attention is grabbed via a sound stimulus. During the agent's action performance, whenever the infant looks aways, the sound is heard again. Performance is aloof, in the sense that the experimenter does not speak to or look at the infant. The effect of manipulating the social context of presentation is that the rate of re-enactement of the novel head action is much lower, and there is no significant difference between the Hands Occupied and HandsFree conditions in any of the two social contexts. That is, no significant rate of imitative behavior is elicited unless the experimenter performs the instrumental actions in a communicative manner (contrast 11% in the Incidental – HandsFree condition with 60% in the Communicative – HandsFree).

> It is only when the head-touch actions are demonstrated by a communicative agent addressing the infants from a second-person perspective using ostensive referential gestures, did infants selectively imitate the head-action more in the hands-free than in the hands-occupied context. (Gergely and Jacob 2012).

CEU eTD Collection

1.2.1. What is to be explained in the task. The list below synthesizes the empirical data of the head-touch task. An explanation is required for all the items taken together.

(1) The experimenter performs a forehead touch as a new means to turn on a new object (a light-box). The new means is less efficient for the goal than the default alternative – a prepotent hand action;
(2) If the experimenter's hands are occupied, infants perform a prepotent hand action;
(3) If the experimenter's hands are free, infants re-enact observed head action;
(4) Unless the demonstration is communicative imitation rate of the head touch is low;
(5) Unless the demonstration is communicative imitation rate of the head touch is not selective between the HandsFree and Hands Occupied conditions;
(6) Across conditions, all infants use their hands at least once, always before head touch;
(7) Imitation of novel head action exhibits low fidelity.

I take the crucial explanatory quest to be an integrated assessment of factors that lead to practical knowledge acquisition manifested in imitative behavior (causal, and social), and the interaction between these factors and the particular behavioral manifestations of learning (emulation, and imitation). I elaborate on this below.

I first briefly expand upon the notion of learning as it surfaces in the experiments in question (a more extended discussion is in Section 2.6). Because both the light-box and the action performed on it are new to the babies, at least two different kinds of things can be learnt upon the experimental demonstration: what the light-box can do, and how to make it work. The former kind is evidenced by goal emulation (item (2)). It is learning about the functional properties, or the affordances of the new object. Functional knowledge is a kind of knowing *that* the object may produce such-and-such outcomes. It constitutes a junction point between theoretical and practical cognition: it is practical because it constrains the actions to be performed on the novel object depending on the desired outcome, and theoretical to the extent that it is shown to play an essential role in the categorization of artifact kinds (Kelemen 1999a; Tráuble and Pauen 2007). The latter form of learning is evidenced by imitation of the agent's means action ((3) in the list). It is the acquisition of a new skill. Knowledge of means actions is an epitome of practical knowledge, or of know *how*.

The emulative behavioral response occurs in the test phase across the board – in all four experimental conditions infants act instrumentally upon the novel object, as if expecting to produce a change of state. Using the prepotent agentive relatum to act upon the artifact (item (6)) provides evidence of the instrumentality of behavior. There are no preconditions for emulative learning of the end-state that can be achieved using the novel object – mere demonstration of the functionality of a novel object, irrespective of the social and physical context, triggers attempts to activate the object. It seems that the enlargement of practical knowledge with novel action goals, or artifact functions, is the default; I elaborate on this in Section 3.1 and put forth a formal treatment of the idea in Chapter 5. In contrast, infants imitate selectively. They re-enact the

CEU eTD Collection

observed new means action, and thereby show that they learnt it, only if observed in a context where the agent's hands are not involved in another goal-directed activity, and when the agent displays communicative signals. Items (3) and (4) spell out preconditions of selective imitative learning of skills. Item (7) corroborates the above mentioned interpretation of (6); it supports the claim that participants' behavior is goal-driven. Even when infants do imitate, they they imitate approximately; infants appear to be more interested about the effect of their action, than about reproducing it in the precise form observed. But if this is so, and infants' prepotent hand actions onto the light-box do cause it to light up, then it remains an open question why 69% of them went on to imitate the odd head-touch in the Communicative – HandsFree condition after successful goal attainment. This fact is even more puzzling if one acknowledges infants' assumption of a one-to-one mapping between goals and means actions.

Observational learning may be recommended solely by a causal evaluation of events. If an action is observed to be necessary for a desired effect, then it may be learnt by a wise observer as a means to be enacted whenever a similar end-state is desired in similar circumstances (e.g., someone might learn that in order to reach an object placed on too high a shelf, she should rise onto her toes). This would be relatively easy to explain within a rationalist framework. Such learning would occur when action interpretation removest the epistemic opacity of observation, and recommends it as necessary in certain circumstances. However, if the interpretive endeavor based on explanatory strategies applied to pre-existing knowledge does not render the observed novelty intelligible, then learning appears to be a more intricate phenomenon. It requires additional credulity in the rationality, epistemic authority, and bona fide of the observed agent. Experiment 1 of Király et al. (2013) showed that in the HandsFree condition of the head-touch task credulity was elicited by the agent's benevolent communicative attitude. That is, only agents who provided communicative ostensive cues were granted a credulous interpretation of their suboptimal means actions, and the unusual means that they demonstrated were imitatively learnt. It still remains to be accounted for the cognitive mechanisms that make possible this complex multi-factored learning of new skills.

The experimental findings of the two studies, plotted in Table 1, raise the following concrete questions. They constitute the empirical explanandum that the formal model engages with.

Question 1:

Why is learning of novel instrumental actions by imitation manifested in the Communicative – HandsFree condition, given that the adult could have used a more efficient means for the goal? Why learn what is not fully comprehensible?

Question 2:

Why is the prepotent hand action always used to act upon the light-box? Why is it used first, i.e. why does imitation follow goal emulation? In particular, if the hand action does achieve the desired end-state, why do infants produce the novel action?

Question 3:

What shapes the interaction between the two kinds of preconditions of imitation –

CEU eTD Collection

> one social (i.e. communicative ostensive presentation), and one physical-causal (i.e. availability of hands)? How are the two factors weighed in infants' decision what to do, when, what for?

Chapter 5 presents a formal computational model of practical knowledge acquisition as exhibited in the described versions of the head-touch task. Its aim is to offer a wide-scope explanation for the behavioral findings as manifestations of learning about what to do, when, what for. The principled answers to the questions above are thus generalizable over a larger array of empirical data (a concrete example is given in the Section 5.8). In the next Section I introduce the general outline of the explanatory direction pursued in modelling.

1.2.2. A direction for the explanation to be pursued in modelling. The idea that goals guide imitative learning about actions and their resultant states reaches consensus beyond the differences between concrete explanatory strategies (Call et al. 2005; Buttelmann et al. 2008; Gergely et al. 2002; Paulus et al. 2011b; Bekkering et al. 2000). Since infants' imitative behavior is no longer viewed as blind copying, it is generally acknowledged that it is shaped by how observations are understood. Understanding instrumental actions is inherently tied to understanding their goals (Sommerville and Woodward 2005; Meltzoff 1988, 1995; Zmyj et al. 2009; Csibra and Gergely 2007; Csibra 2008; Schwan and Garsoffky 2008; Kelemen 1999a).

Two main explanatory suggestions have been proposed in the developmental psychology literature for questions such as the ones listed in the previous Section. The two are based on different kinds of cognitive mechanisms suggested to underpin infants' imitative learning as a goal-driven behavior. One favors bottom-up processes; it links the acquisition of practical knowledge from observation to direct perception of actions, access to prior experience, and motor resonance. The other one is a top-down approach; it focuses on teleological rational interpretation of actions in the context of their occurrence. I argued in Section 1.1.1 that the former has difficulties as an account of understanding observed new actions in terms of goals. This is particularly the case for observations of novelty. Making sense of actions as goal-directed, though very early evidenced in development, is tremendously difficult because of the equifinality and multifinality that characterize the means – ends hierarchies of human actions. Non-inferential theories do not fare well with goal-based processing of novel actions, hence their explanatory power for imitative learning is reduced compared with the top-down ones.

I will however say a few words about one of them. The theory of motor resonance (Paulus et al. 2011a,b) is the most recent bottom-up approach to imitative behavior. It draws on non-inferential strategies based on action-effect associations (Hommel et al. 2001; Elsner 2007; Elsner and Aschersleben 2003) and simulation (Sommerville and Woodward 2005) to explain imitative behavior:

> infants' imitation relies on the interplay of motor resonance and salient action effects (Paulus et al. 2011b, p.908).

CEU eTD Collection

Motor resonance is a simulationist procedure by which observing another agent's action directly (automatically) activates a corresponding motor program in the observer's motor repertoire, and leads to motor resonance. Imitation is thus not blind copying, but it is selective because of the low-level motor constraints: only those actions are imitated with which infants can resonate. As a second stage after resonance is established (if supported by infants' pre-existing database of actions in the motor repertoire), then it is stored associatively with the observed effect, and retrieved as a compound action – effect when circumstances call for producing the effect. So the original findings of Gergely et al. (2002) are explained as follows: infants resonate with the head action in the HandsFree but not in the HandsOccupied condition (because 14-month-olds' motor system does not allow performance of a head touch while having the arms bent), so the effect light-on is stored in association with the head action only in the Hand Free condition, and thereby performed upon encountering the novel light-box in the test phase. This account has been empirically challenged by Király et al. (to be submitted), where further manipulations prove that it fails to account for a more refined manner in which contextual constraints shape the tendency to learn actions by imitation. Moreover, it does not distinguish between goal emulation and action imitation, and hence it does not recommend a mechanism that explains learning about novel artifact functions via emulation. Finally, it does not address the question regarding the influence of the social context of action demonstration on infants' learning.

Better sense may be made of the empirical explanandum if the infants' observational learning via imitation and emulation, supported by goal-centered action processing, is explained in a top-down approach. This explanatory strategy, pursued by, e.g., (Gergely et al. 2002; Schwier et al. 2006; Brugger et al. 2007), has it that the acquisition of practical knowledge instantiated by imitation and emulation is a form of rational behavior, supported by teleological interpretation of the observed actions in the context of their occurrence.

> The rationality principle is a well-formedness constraint that can apply to different domains of knowledge (Csibra et al. 2003, p.131);

one of these domains is practical knowledge. Its application in this domain serves epistemic functions, such as understanding and learning. The model to be developed here adopts this explanatory line. The role of perceptual motor processes is nevertheless acknowledged. The teleological inferential procedures lead to selective acquisition of knowledge by integrating automatic sensori-motor and storage procedures. Both simulationist and action – effect associations play their role in action interpretation and action production. However, I argue in Chapter 2 and instantiate formally in Chapter 5, they are guided in a top-down manner by an inference-driven mechanism. For instance, instead of saying that a certain action is automatically activated once an agent is set on a goal, the model spells out in detail the steps by which the action is selected, thereby providing a principled and generalizable procedure.

The three questions to be answered with respect to the behavioral findings shed doubts on a purely rationalist view. Gergely and Jacob (2012) distance themselves from the initial theory

CEU eTD Collection

building on infant rationality, because it is insufficient to account for the role of relevance principles in selective acquisition of practical knowledge from observation (cf. discussion in Section 1.1.3). The new approach builds on a separate treatment of instrumental and communicative agency of Gergely (2010). On this view, imitative learning as exhibited in the head-touch tasks is supported on the one hand by third-person teleological interpretation of instrumental actions, and on the other hand, by second-person understanding of relevant information conveyed by ostensive communicative agents. The two constitute separate cognitive mechanisms that contribute to the acquisition of generalizable practical knowledge in infancy.

The suggested model accounts for the behavioral findings described above as grounded in the pedagogical and the teleological stance with respect to action understanding. Human infants' readiness to learn actions imitatively from agents whose communicative ostensive behavior provides cues of knowledgeability and epistemic trustworthiness is a manifestation of the pedagogical stance. Reasoning about goals with abstract assumptions such as the rationality principle constrains the teleological explanation of actions in the context of occurrence, and guides decision as to the information bits that are worth acquiring, or relevant. In contrast with the above mentioned approach of Gergely and Jacob (2012), relevance and efficiency considerations are integrated via a single mechanism of teleological practical reasoning, where goals include desirable changes of state here and now, as well as states of affairs that need to be carried out or maintained throughout time. The latter can be understood as second order goals, i.e. as generic motivations that are instantiated by episodic achievement of concrete physical end-states (see Section 2.2). Two epistemic drives are crucial: understanding observations, and learning from them. I label this the single-mechanism multi-level approach to teleological processing; the model to be developed attempts to substantiate the explanatory proposal. This mechanism hinges on the rational capacities introduced in Section 1.1, whose availability at ages around 14 months has been empirically evidenced. Therefore, the proposed explanation is not like 'putting the cog in infant cognition' (Haith 1998).

CEU eTD Collection

CEU eTD Collection

CHAPTER 2

Practical reasoning and the acquisition of practical knowledge

Human beings are agents who act on the world, which also includes other agents. Actions are goal-directed, purposeful movements. Agents' actions take effect, i.e. produce changes, on the physical and mental world. The physical level of agency is developmentally primary (Gergely 2002). The basic understanding of agency that I make use of is at the personal level. I shall also appeal in modelling to a more fine-grained view of agency – an agent's body parts that actually cause changes, e.g., her hands, legs, head, shall be referred to as local agentive relata. These have agential status in virtue of belonging to a global human agent, and of being the actual source of physical causation. For instance, if I slice a loaf of bread, I am the global agent of the slicing action, whereas my hands are the executive local agentive relata. I use the term 'agent' to refer to global agency, and make the required specifications when talking about agency at the sub-personal level.

Agents' rational instrumental actions are meant to make their goals true properties of the world. The actions performed by other agents also become intelligible if assumed to bring about goals (see discussion of the teleological stance with respect to action understanding in Section 1.1.1). Instrumental actions contextually fit for goal achievement are underpinned by practical reasoning; that is, they are permeated with rationality (McDowell 2007). Under this assumption the terms 'agent' (standing for global agent) and 'reasoner' are used interchangeably throughout the Chapter. By practical reasoning I thus mean the commonsense reasoning by which agents process actions with respect to desired end-states or goals. I focus on practical reasoning in the form of action explanation and action planning. Practical reasoning links observations of other agents, their goals and reasoners' own goals, as well as appropriate actions to be performed by each. This is reasoning about what one (the reasoner herself, and other agents) should do, when, what for. I focus on the processing of transitive instrumental actions, i.e. on actions meant to achieve goals by means of artifact usage.

Practical reasoning is teleological reasoning whose practical utility is a crucial component. In relation with its practical utility, reasoning is viewed from the perspective of cognitive economy: it is worth performing if its epistemic or practical benefit surpasses the processing cost. Practical reasoning organizes actions and states of affairs along a relative means – ends hierarchy. An action can be a means with respect to higher levels in the hierarchy (e.g., a precondition, such as boiling water is for making tea), and also an end with respect to lower levels (e.g., boiling water relative to pouring water in the kettle). The picture becomes even more complex when tools or artifacts are considered; they are means for some further goals (the functions that

CEU eTD Collection

they are made for), but they are also part and parcel of the goal of an agent's action upon them. Generally goals are states of affairs (end-*states*), and means are actions[1].

The conclusions of practical reasoning are action rules[2]. They connect a type of goal to a sequence of actions. They have the form 'In order to bring about state s, agent X should perform actions $a_1, \ldots, a_n$ at times $t_1, \ldots, t_n$', where s is a goal and $a_1 \ldots a_n$ at times $t_1, \ldots, t_n$ is an ordered sequence of actions – the plan for s. The action sequence is, until further evidence, the most efficient or optimal one for the goal. As the expression action *rule* and the modal formulation (*should*) suggest, the relation between state s and the sequence $a_1 \ldots a_n$ at times $t_1, \ldots, t_n$ is normative: given a goal, a particular action plan is necessary, and the optimal such plan is also sufficient (i.e. the action sequence is exhaustive). Such conclusions prescribe X's performance of a_i at t_i in the case of planning, or another agent's Y explanation for X's doing a_i at t_i, given that X and Y are assumed to be rational agents.

Practical reasoning is bidirectional; it can proceed forward – from actions to their effects, or backward – from effects to cause actions[3]. Agents may reason forward about the effects of actions, or about future developments of current states of affairs. For instance, if an agent knows that 'If the button is pushed, then the light is on' and she observes someone pushing the button, then she expects that the light be on afterwards. This is called reactive reasoning (Kowalski 2006, 2011). It is the kind of processing that allows agents to respond to happenings and changes in the surrounding environment. Forward processing is involved in reasoning towards an interpretation (van Lambalgen and Stenning 2008) of an observed context. One example is goal assignment to an action, given particular features of the context where it is observed. Conversely, reasoning from goals to means actions or planning is based essentially on backward processing. A future desired end-state, or goal initiates agents' computations, and it is reduced stepwise in a backwards procedure to executable actions to be performed at a moment in time closer to, or equal with the present time. Let's take the action rule 'In order to turn light on, push button' and suppose that the agent had the goal to turn on the light. Given the rule, backward reasoning leads her to push the button. Such reasoning is proactive (Kowalski 2006, 2011). The backward and forward directions of action processing are intricately interconnected and usually co-occur in human agents' practical reasoning. As Bratman (1987, p.47) put it,

[1]However, things are not necessarily so. Because of the complex human teleological hierarchy, supervenient on two modes of causation – instantaneous and continuous, actions may be ends – e.g., crossing the finish line, and states of affairs may be means – e.g., water running from a tap is a means for filling a bucket. More is said on this issue in Sections 4.1.1 and 4.1.2.

[2]This view is almost universally attributed to Aristotle. According to Streumer (2010), in contemporary philosophy it is set forth by Dancy (2004) or Tenenbaum (2007).

[3]Note similarity with the psychological problematization of the directions of teleological action processing in Gergely and Csibra (1997, 2003) introduced in Section 1.1.1. It should be mentioned, however, that what counts as 'forward' and 'backward', respectively, is a matter of theoretical standpoint on action regarding the role of psychological notions such as intentions in action explanation. For instance, those authors who emphasize in a mentalistic, Davidsonian fashion the causal role of intentions in action generation, and equate goals with intentions, would hold that inferring goals from actions proceeds backwards, because it amounts to re-tracing the motor intention Jacob and Jeannerod (2005) or the intention in action (Searle 1983) that caused an agent to act.

CEU eTD Collection

> we reason forward from our existing beliefs, and backward from what we are interested in establishing, and try to bring the two chains of reasoning together,

where 'what we are interested in establishing' are our goals. The inferential path is not linear, rather, practical reasoning involved in deciding what to do, when, what for, usually involves both.

I focus on two manifestations of practical utility of reasoning: to change the world as intended via appropriate action courses (i.e. practical reasoning moves agents to action), as well as to understand other agents behaviors (i.e. practical reasoning ensures that agents understand each others *qua* agents). In certain kinds of social contexts, both may contribute to the development of agents' practical knowledge, i.e. to learning what to do, when, what for; such is the case in the task to be modelled. *Planning* is backwards teleological reasoning from goals to basic actions that presumably attain the desired end-state in a particular context; it amounts to goal reduction to basic actions. In the philosophy of action, basic actions, as introduced by Danto and Morgenbesser (1963), are those that are not performed by performing some other action. For instance, the non-basic action of hitting the ball is performed by the basic action of hitting the ball with the leg. For the current purposes, I take actions described in terms of local agentive relata as basic. Basic actions are executable in the sense that they result in direct activation of an agent's motor system. *Teleological explanation* is an application of teleological reasoning to understand the actions performed by other agents in the surrounding environment, in terms of the reasons that agents must have had for performing them. In a multi-agent world based on interactions between agents, the two forms of practical reasoning are tightly connected. Reasoning is meant to find the appropriate means actions to be performed in order to bring about a certain goal, in a certain context – the reasoner's own goal in the case of planning, or the observed agent's goal in the case of explanation. Planning and explanation pose similar computational problems, and consequently call upon similar computational solutions. This is why I suggest that agents' capacity for planning is recruited for explaining other agents' actions. The kind of justificatory reasoning that is of interest for the current modelling endeavor is subserved by planning computations. Along Marr's (1982) distinction between levels of analysis of a cognitive process, the reduction of goals to actions constitutes an algorithm that fulfills two functions: planning an agent's own actions, and teleological explanation of another agent's actions.

The current Chapter provides a theoretical framework for the reasoning processes that constitute action planning, their recruitment for action explanation, and their cognitive outcome in the form of (practical) knowledge acquisition. By 'theory' I mean philosophical, as well as logical takes on planning and explanation. From the vast array of available theories, those were picked that provide a notion of competence that is closest to the performance of the young practical reasoners, and is thus expected to provide a best-fit computational process model. I begin in Section 2.1 by introducing briefly the main features of practical reasoning, which are drawn

CEU eTD Collection

upon in the formalization enterprise. I elaborate on the notion of goal by introducing a useful distinction in Section 2.2. Section 2.3 and 2.4 describe the two main reasoning processes of interest – planning and explanation, respectively. I begin with the basic planning computations, and then show how action explanation capitalizes on these. I emphasize the computational problems involved in planning and explanatory goal derivation. I then propose a solution via a system called closed-world reasoning in Section 2.5. I end in Section 2.6 with a discussion of the crucial cognitive outcome of practical reasoning – the acquisition of practical knowledge.

2.1. Focal points of human agents' practical reasoning.

The features listed below are equally applicable to planning and explanation processes. They are developments of the ideas hinted at in the preamble. Except for (1) in the list, the other characteristics are constraining factors on reasoners' inferential strategies.

(1) Practical reasoning is not necessarily conscious reasoning.

More often that not, from Aristotle onwards the expression 'practical reasoning' has been used interchangeably with 'deliberation'. However, I do not use 'deliberation' because of its connotation of conscious processing. Practical reasoning, on the other hand, as is understood throughout this thesis, need not be consciously performed by agents with linguistic means at their disposal (and indeed is not). Reasoning does not require awareness. One's mind may well be engaged in cognitive processing without one being explicitly aware that this is the case. This perspective justifies one to attribute reasoning capacities even to pre-verbal entities.

Practical reasoning is rather a form of intuitive (Kowalski 2006) or automatic thinking. Both planning per se, and the use of planning strategies for action explanation, may be automatic and efortless, something like an agent's rational 'best guess' (van Lambalgen and Stenning 2008, p.115) as to the action to be performed, or the reasons why another agent does what she does. The choice to act in a particular way, or to assign a certain reason for another agent's actions is 'non-consciously rational with quantitative precision' (Hempel 1961–1962, p.22). Relatedly, the cognitive outcome of practical reasoning is implicit observational learning (see below Section 2.6).

The kind of reasoning that is mostly relevant for a model of infants' observational learning is the 'humble' motor planning (e.g., hitting the ball with the leg in order to score), and not the kind of problem solving (e.g., playing chess in order to defeat one's opponent) that usually receives more attention in discussions of human reasoning and rationality. Although motor planning is not explicit in human agents' awareness, that is, although it is an automatic process, it is reasoning nonetheless – a form of practical reasoning. According to Dennett (1987, p.50),

> even the trivial problems of planning and bodily guidance that are beneath our notice (though in some sense we 'face' them) are solved by similar processes

CEU eTD Collection

as the complex planning involved in problem solving. Moreover, motor planning is a kind of processing that, besides resulting in overt action performance, also subserves the epistemic functions of explaining the actions performed by other agents, and learning from them. Arguments for this rationalist view of human behavior as inference-driven have been set forth in the psychological literature (Gergely and Csibra 2003; Csibra and Gergely 1998, 2007) reviewed in Section 1.1, as well as in the broader field of cognitive science (Thagard 2005; van Lambalgen and Stenning 2008; Miller 2005; Chomsky 2006). The arguments are backed up by the logical philosophical framework exposed throughout the current Chapter. The core of the argument is that

> we may call an automatic process such as planning an instance of [. . .] reasoning if it can be viewed as the construction of a (representation of a) model from which conclusions are read off. (van Lambalgen and Stenning 2008, p.126).

The model of implicit processing may be extended to account for conscious, deliberative planning at larger temporal scales.

(2) Practical reasoning is *normative*.

Practical reasoning is reasoning about what an agent (either the reasoner himself, or an observed agent whose actions call for explanation) *should* do in order to bring about a certain desired end-state. Inferential normativity is supported empirically. The data reviewed in Section 1.1.1 and the corresponding theory of teleological reasoning used for action interpretation, functional understanding of objects, or online action anticipation, supports the view that means – ends relations are taken by the young reasoners to have normative status. That is, if a particular end-state is desired in a particular context then a particular action should be performed. The 'should' is brought about inferentially, since such modal statements are not derivable in a bottom-up manner, exclusively from perceptual and motor processes. I make this assumption along Kantian lines –

> experience itself is thus possible only in so far as we subject the succession of appearances to the law of causality (Second Analogy, *Critique of Pure Reason*).

Paraphrasing, normative practical reasoning is made possible by the application of generic, law-like principles to the integration of sensorimotor inputs. These principles are spelled out below, in Section 2.5.

In the case of explanation, normativity obtains between reasons for action and actions. An action rule subserves an explanatory function when a particular end-state and the action that causes it are linked not just by the descriptive 'is', but allow for the normative 'should'. A goal provides sufficient reasons to perform a certain action in a particular context if practical reasoning recommends that it should be performed in that context. In the case of planning, normativity between ends and means is a

CEU eTD Collection

manifestation of the idea that means actions are rationally prescribed by a plan started from a given end-state (Pollock 1995), whose outcome is the action rule including that particular action. Means are preconditions for goal attainment. They are sub-goals that must be made to obtain for the top goal to be achieved. A means action should be performed for a goal, and it is justified by that goal, if it is the unique action candidate singled out by reasoning processes grounded in the goal.

> If an agent has knowledge that *one* of its actions will lead to *one* of its goals, then the agent will select *that* action(Newell 1982) [added emphasis].

In infancy, this take on teleological relations can be seen as a precursor of the functional fixedness (Kelemen 2004) exhibited by agents older than 5 years of age, who seem to expect that actions as well as artifacts have one particular function.

(3) Practical reasoning is *context-relative*.

Planning and explanation are intrinsically linked with the features of the situation where they are applied. To begin with, the notion of goal is itself context-relative. In fact goals are contextual features; I label them dynamic features, to distinguish them from others such as color, shape, spatial location, etc., which are static. An effect of an action might be an agent's goal in one context, or merely a side effect in another. Take, for instance, someone running every day from A to B. This behavior has (among others) two causal effects: strengthening the agent's limb muscles, and reaching point B. If in the context of training for a marathon, then the former is likely to be the agent's goal. If in the context of working in location A and living in location B, the latter is more likely to be assigned as her goal (although increased muscular strength would also come about). A rational observer's interpretation of the very same action as to the agent's reason for running would be different in the two contexts. The runners might also plan their route from A to B differently – the one who runs for training might choose a hilly path, whereas the one who wants to go home might choose the shortest, neatest one. Both plans are rational in their context because the aspects that shape reasoning are different.

The input for an agent's practical reasoning is not abstracted away from concrete contextual details. This view is along the lines of McDowell's reading of the Aristotelian *phronesis* as 'concretely situation-specific discernment' (McDowell 2007, p.340). The influence of universal, or generally desirable state of affairs on practical reasoning that results in action performance or explained actions, is always mediated by an assessment of the context where they must apply. Because planning is meant to change something relative to the current state of affairs, it is concerned with the properties that obtain in a given context. An instance of the general good must be found, which fits that context. Bratman (1992) also sets forth that a model of practical reasoning should incorporate, along with agents' context-independent beliefs, informational

CEU eTD Collection

bits that are context-relative and that come about under 'practical context-relative pressures'; he calls this 'context-relative acceptance'. For instance, if I know that I am going to spend the summer in the Caribbean, my packing the suitcase can be viewed as behavior that exhibits planning. When so doing I am constrained by the assumption that the weather will be hot, hence I do not need to consider the possibility of packing gloves and scarf.

Practical reasoning is constrained by the goal that an agent acts upon or assigns another agent in a particular context, as well as by other contextual features.

(4) Practical reasoning deals with incomplete knowledge states; it is a form of *reasoning with uncertainty*.

To my knowledge, this feature receives full consensus in the philosophical, logical, AI, or cognitive science literature. Human agents' reasoning about what to do, when, what for, must deal with a presumably infinite number of things unknown, and things that might change while reasoning. More, agents do not have the certainty that they have isolated all the relevant features of a situation, and only those. The hierarchical relative structure of the relations between means and ends also adds to the uncertainty of practical reasoning. I elaborate on these issues in Sections 2.3 and 2.4 when introducing the computational concerns of practical reasoning.

As John L. Pollock (2009, p.173) puts it,

> One of the most striking characteristics of human beings is their ability to function successfully in complex environments about which they know very little.

Uncertainty is tied to the idea of bounded rationality of human agents (Pollock 1995). It must be framed or circumscribed (McCarthy 1980b; McCarthy and Hayes 1969) in order to facilitate reasoning, so that it leads to appropriate action performance and explanation in a reasonable amount of time. Human agents act as if doing so automatically; a logical formalism suitable to represent such reasoning must make use of strategies that allow coping with uncertainties.

(5) Practical reasoning is *non-monotonic*.

It is rational for an agent to keep in mind that their practical inferences are uncertain,

> while at the same time not allowing this uncertainty to inhibit one's actions. (van Lambalgen and Stenning 2008, p.229).

Such allowance is reflected in the flexibility that characterizes human agents' practical reasoning. The conclusions are open to revision; reasoning is defeasible as new information becomes available (via observation or additional reasoning). In fact,

> all interesting reasoning outside mathematics involves defeasible steps (Pollock 1995, p.41),

and it is beyond doubt that reasoning about what to do, when, what for, is interesting.

CEU eTD Collection

The logical reflection of flexibility is non-monotonicity (van Lambalgen and Hamm 2005; McCarthy 1980b; Mueller 2006). It means that the validly derived conclusion of a certain argument may become invalid when new premises are added to the database on which reasoning is performed. The conclusions of practical inferences are (action) rules that allow exceptions; they are formally represented as non-monotonic conditionals. Exceptions to action rules arise from unusual, or abnormal contexts of action performance (see below Section 2.5). They can lead to deriving exceptional means, or to revising expectations of certain goals, without thereby canceling the validity of the action rule that links a particular means – ends pair in a normal context. In the former case, non-monotonic reasoning singles out a non-default action and recommends it as an alternative novel means for a particular goal in an exceptional context. The latter manifestation of non-monotonicity leads to canceling an expected goal when a means action does not achieve it, presumably due to abnormal contextual features.

Action rules (either action- or explanation-oriented) arrived at via non-monotonic processing involve increased computational effort. Means and goals revision involves cancelation of the original content of the means or goal placeholder in the default action rule, followed by replacement with novel content selected from available evidence. Baggio et al. (2008) show in a neurological study that teleological re-computations involve more processing effort than monotonic goal-centered reasoning, evidenced by the temporally sustained character of EEG signals. One must have this in mind when thinking in terms of cognitive economy at the reasoning processes posited by a model.

Flexibility in teleological action interpretation is crucial for a cognitive outcome such as the inference-driven acquisition of practical information about instrumental actions and artifact functions discussed below in Section 2.6.

(6) Practical reasoning is concerned with *efficiency*.

Rational agents' instrumental actions are expected to be tailored for the goals that they are meant to bring about in the particular context. The standard of apropriateness of actions as means for goals is efficiency.

A crucial feature of planning is that actions are selected as means for goals based on their contextual efficiency in bringing about the goal. Likewise, a goal is justificatory for another agent's action depending on the efficiency of the action for that goal. I showed in Section 1.1.1 empirical evidence showing that efficiency is intrinsic to the rationality assumption (Gergely and Csibra 1997, 2003) used by infants in explanatory evaluations of observed actions.

Efficiency is a graded concept; that is, different actions can be more or less efficient for a particular goal. The most efficient one is optimal[4]. A goal is a good (i.e. sufficient) reason for an action in case the action is optimal for goal attainment. Tokens of default action types stored in agents' motor programs, i.e. in their databases of motor

[4]The two notions should be read as synonymous; they are used interchangeably in the thesis.

CEU eTD Collection

knowledge, are, until further evidence, optimal with respect to the types of physical goals with which they are connected.

It is important to note that efficiency considerations involved in practical reasoning need not be restricted to physical efficiency, although this is indeed the primary understanding that it bears. Efficiency of actions is goal relative, but goals can be represented at different levels of abstraction (see Section 2.2 below). Also the more abstract, second-order goals, i.e. the intentions to bring about physical goals, afford descriptions in terms of efficiency. For example, one's utterance of truth and nothing but the truth may be more or less efficient for one's intention to stay friends with some other person. Moreover, efficiency of actions for physical goals need not be on a par with efficiency of actions for more abstract desires to change another agent's mental state. That is, an agent's opening a suitcase with a beer opener may be suboptimal relative to the agent's physical goal to have the suitcase open in the given circumstances, but optimal relative to the same agent's higher-order intention to teach another agent a funny, creative, or exciting way to open the suitcase. Rational action is not reduced to physically efficient instrumental action.

2.2. A stepping stone for representing goals.

I proceed to specify the primary focus of practical reasoning, the notion of goal, which links reasoning to action (cf. Duff et al. (2006)) and action understanding. The discussion puts flesh on the notion of efficiency introduced in the preceding section as more complex than efficiency for physical goals.

Goals 'focus planning' (Pollock 1995, p.24), in the sense that they are set when planning own actions, and they are assigned to another agent in order to explain her actions. Some goals are states of affairs desirable in virtue of the biological structure of living organisms – e.g., food, water. Agents are by default set on those. In other cases, the choice of goals itself involves reasoning – e.g., when they are derived as sub-goals of the reasoner's higher-order intentions to bring about goals, or when they are inferred in an explanatory enterprise. Goals constitute the input for backward practical reasoning that leads to action rules and subsequent action performance, and that grounds the justificatory analysis in action explanation.

The background idea is that desired states of affairs, or goals, provide agents with internal reasons for action. This teleological understanding of the notion of 'reason for action' is applicable in the case of both planning and action explanation. Goals motivate agents to plan their actions, and to eventually execute them. They provide a rationale for actions, and thus motivate one's own actions, and explain others'. In fact it is because reasons for action are grounded in agents' prior motivations (Williams 1981), that goals fulfill an explanatory role. Explaining actions by assigning reasons to observed agents amounts to assuming that they are motivated by the goals that push them to action. The teleological stance as presented in Section 1.1.1 elicits the expectation that the observed agents' actions are performed in order to attain efficiently

CEU eTD Collection

physical goals, in conformity with the particular features of the situation. More generally, states of affairs assigned to other agents as goals, be those physical or more abstract, justify their actions from an observer's point of view[5].

Concrete physical goals, based on causal changes of state in the physical world, are but one kind of motivational and explanatory factors for human action. I bring to the fore a view of practical reasoning where the inferential processes connect agential motivations represented at different levels of abstraction, and thereby enable action explanation and action performance. This view is based on taking seriously the hierarchical structure of teleological relations. The teleological hierarchy presents itself as a continuum between concrete physical goals and abstract ones; the former are the content of the latter. It can be viewed as a network of causal relations linking basic actions to generic abstract objectives. I propose the expression *teleological thread* to denote a vertical structure (though not linear) whose nodes are sub-goals of more abstract states (located at a higher position), and goals of more concrete states (located at a lower position). All nodes along a thread are causally related in a certain context. For instance, typing capital A, submitting a PhD thesis, and financial well-being may constitute a single teleological thread as far as a PhD student is concerned, unlike, for instance, typing capital A, mopping the floor, and emotional well-being.

Physical goals are desired changes of state in the physical world, e.g., turning on a light at time t, drinking water now. Agents reduce them through backwards chaining to executable basic actions. They motivate action in the sense that they directly call for action upon them, in order to achieve the targeted state; they generate behaviors. They are episodic, i.e. anchored by a particular moment in time. They will henceforth be referred to as *achievement* goals.

At a more abstract, higher-order level, agents also have *maintenance* goals, which ensure some desired, stable, balanced relation with the ever changing state of the world. They refer to states of affairs that must be made to obtain generally, across time. For instance, staying alive is such a goal, which ultimately grounds all the other goals that human rational agents act upon (when they act rationally). Such 'ultimately' abstract goals are not the issue in the current discussion of practical reasoning, nevertheless they do illustrate well maintenance goals[6]. Their primary function is to motivate agents to set achievement goals upon which to act, so that the higher-order goal remains true throughout time, or so that it becomes true again if previously falsified. That is,

CEU eTD Collection

[5]This understanding of goals assumes Bernard William's (1981) internalism with respect to reasons for the current discussion of practical reasoning.

[6]One could remark that such a 'goal' is not at all characteristic to rational human agents, rather it is an objective of any living organism, and hence even using the term 'goal' to describe it can be misleading. An in-depth treatment of this issue is beyond my scope here. I reply only that although the aim to stay alive is shared by all living creatures, just like any other evolutionary trait, the way in which it grounds human practical reasoning, the way it establishes connections with lower-level goals, and ultimately guides action performance, as well as its explanatory force for observed actions of human agents, justifies using it as an example of maintenance goal. The interested reader is advised to consult Chapter 8 in Kowalski (2011).

> maintenance goals either trigger corrective actions to be taken, or constrain the choice of actions so that the maintenance goals will not be violated (Duff et al. 2006, p.1034).

They can be conceptualized as second order goals, or intentions to bring about corresponding concrete goals in a particular context. The teleological structure that grounds practical reasoning is represented as *achievement goal FOR maintenance goal*.

Maintenance goals set generic constraints on agent's actions. Corresponding achievement goals are inferred from them, as concrete objectives that can be satisfied in a number of steps through action performance in a given context. Maintenance goals can have a physical (e.g., stay warm), as well as a psychological (e.g., expand one's knowledge) component. I suggest understanding them as second-order goals, or intentions to bring about corresponding concrete achievement goals in a certain context. The latter constitute the concrete content of the former. I wish to stress that such talk about intentions does not commit me to a causal – psychologist theory of action, along the lines of Davidson (1963), Searle (1983), or Searle (2001). Reasoning about intentions resorts to the same mechanism employed for physical goals, i.e. backwards processing: they are reduced to a concrete physical content (the achievement subgoal), and further – to actions. The categorization of end-states as 'mental' or 'physical' makes a difference only with respect to the processing strategy of the unique algorithm. A good illustration is an agent's generic objective to acquire knowledge about the surrounding world, when faced with a reliable source of acquisition. This abstract intention is fulfilled at a particular moment via transformation into a concrete goal in the current circumstances, e.g., 'do as the teacher does now', which is a sub-goal of the maintenance goal. Physical maintenance goals (e.g., stay warm) are also more remote from action than achievement goals. They must be operationalized into a concrete state of affairs that can be acted upon. Suppose that an agent feels cold, so she must act in order to revert to the generally desirable state of warmth. To this end, she could use the physical knowledge that thick textiles warm one up, and hence reduce the negative formulation of the maintenance goal – 'do not stay cold' – to a concrete objective to put on some thick material as soon as possible.

The distinction between achievement and maintenance goals originates in the AI literature with respect to autonomous intelligent agents and multi-agent systems – e.g., Duff et al. (2006); Hindriks and van Riemsdijk (2007). It provides a way of improving agents' rationality by grounding action performance in a wider class of goals (Duff et al. 2006). It has also been used in applications of computational logic to human practical reasoning in Kowalski (2011) and Kowalski (2006). I suggest that the same distinction fosters the understanding of the multi-level character of teleological practical reasoning: it applies equally well to intentional motivational states represented as maintenance goals, as to concrete changes of physical states represented as achievement goals.

Neither concrete, nor higher-level goals are activated *ex nihilo*. All goals, as desirable states of affairs that call for action, have some conditions of activation. These conditions specify the

CEU eTD Collection

representation of the desired state, by tying it to particular states that obtain in the world. In the above example of a maintenance goal, 'when faced with a reliable source of acquisition' is the condition that triggers the activation of the goal to learn. The fact that goals are activated under some circumstances is captured by a conditional representation of the kind *IF certain conditions obtain THEN a certain state of affairs should be pursued*. The state of affairs in the conclusion is the maintenance goal. Once the antecedent is made true in the context of interest, the consequent is activated quasi-automatically via reactive forward reasoning. It is made true by backwards reasoning towards an achievement sub-goal, and then towards action (see below Section 2.3). There are also unconditional goals, whose 'conditions of activation' are a tautology $\top$, e.g., the goal to sustain life is active throughout life span.

Practical reasoning, and consequently its behavioral and cognitive outcomes, is shaped by maintenance goals. The scope of maintenance goals is practically unlimited, given the complex human teleology in which pretty much any state of affairs can be a goal in a particular context. However, the model set forth in the thesis focuses on epistemic maintenance goals. More specifically, the model is developed on the empirically grounded assumption that infants' maintenance goals in the tasks under scrutiny are understanding observations, and acquiring practical knowledge (Section 3.1). When these goals are triggered by observations, they guide the explanatory enterprise, and subsequently the motor planning that results in overt behavior. The fact that reasoning about goals represented at different levels of abstraction is subserved by the same inferential mechanisms allows an integrated account of instrumental and communicative agency (Gergely 2010) in the current model. The distinction between achievement and maintenance goals instantiates the multi-level teleological processing underpinned by the reasoning strategies to be introduced in what follows.

2.3. Planning.

Adopting a goal initiates planning for its achievement (Pollock 1995, p.25). Planning refers to the inferential process involved in devising a sequence of actions that presumably, to the best of the agent's knowledge, attain her goal optimally. In computational terms, planning means goal reduction to sub-goals, via backward chaining of effects to cause actions, all the way down the teleological hierarchy to basic actions (or motor programs). This definition of planning is used throughout the thesis. Because of the broad range of goal representations considered (see Section 2.2 above), the practical reasoning mechanism applies to basic motor planning (e.g., motor activation in order to raise one's arm), and to larger scale plans (e.g., going to the university today, obtaining a place on the podium at the marathon next month, maintaining a healthy and aesthetically pleasing weight). However it is only the former kind of planning that is the current locus of interest; in the subsequent discussion I focus on planning for immediate action. An agent's plan is expressed as an action rule of the form 'In order to bring about state s at time $t + \delta$ do a at time t', where s is a goal, a is a sequence of temporally ordered actions, t stands for their corresponding moments of occurrence, and δ is a temporal duration.

CEU eTD Collection

According to Bratman's observation in (Bratman 1987), individual actions represent the limiting case of plans. An action sequence obtained via planning derivations from a goal is prescribed for performance, as a means for the goal.

As stated in the preamble of the Chapter, functionality is a key aspect of practical reasoning. In the case of planning, it manifests itself in the form of derivation of action sequences which are most fit for concrete goal achievement, and abstract goal maintenance. The optimal such plan, the one that should be performed, does not allow disjunction between more alternative courses of actions within one action rule. An action rule like 'In order to bring about state s do action a or action b' derived by backward reasoning from goal s does not furnish an optimal plan[7]. I refer to a plan composed of a disjunctive action rule as a *branching* plan. The suboptimality of branching plans fares well with the intuitive axiomatic idea that *the* optimal instrumental action to be performed for attaining a desired state is unique.

> A noteworthy fact about human beings is that in many (perhaps most) cases of routine planning, no more than a single plan for achieving a particular goal is ever produced. We do not generate a large number of different candidate plans and then choose among them. We produce a single plan, and ... we are willing to act upon it without engaging in further search (Pollock 1995, p.272).

Disjunctive action rules require further processing; that is, they require decision as to which one of the several alternative actions is optimal for goal attainment. Hence branching plans are more demanding in terms of both storage (i.e. both actions *a* and *b* must be maintained in working memory during planning), and computational resources. It is only natural that a resource-saving rational agent, with a cognitive economic interest, prefers the simplest and most direct route from a desired end-state to action performance. This route is via a plan composed of a unique action sequence, presumably the most appropriate in the given context.

An action sequence that works as a well-formed plan is thus to be understood as a single, non-branching sequence, grounded in the desirable end-state that initiates derivations. 'No disjunction in conclusion' is a rule derived from the functionality of planning for rational agents who wish to proceed to action, and not remain blocked in decision loops. The one-to-one mapping of actions to goals operationalizes the normativity of practical reasoning in its guise as planning: an action should be performed for a goal (with no more ado in the form of decision making) if and only if it is prescribed as a unique action candidate by backward reasoning processes grounded in the goal. This means that optimal means under the 'no disjunction' constraint are both necessary and sufficient for a goal. A rational agent need do nothing more and nothing less than the one ordered action sequence obtained from successful planning derivations.

[7]The idea has important consequences when agents' capacity for planning is recruited for action explanation purposes. See Section 2.4 below.

CEU eTD Collection

I now elaborate on the reasoning processes that result in action rules, or plans. I first explain how maintenance goals fit in the picture of goal-to-action planning derivations. Once a maintenance goal is triggered in a given context, an achievement goal is derived from it. The higher-order goal is instantiated via reasoning processes in a more concrete, physical end-state to be achieved in the current context. For example, suppose there is an agent who feels unpleasantly cold; the formulation *put an extra sweater now FOR change unpleasant cold state FOR pleasantly warm state* expresses the two transitions from the abstract maintenance goal 'pleasantly warm state', to a need for change with respect to an unpleasant state of coldness, and further to a concrete subgoal to be achieved at a particular moment in time, perhaps as soon as possible – putting on an extra sweater.

The first step of planning is the derivation of a concrete instance of a maintenance goal, which agents set as the achievement goal to be acted upon. Inferring an achievement goal fit for action in the current context amounts to specializing the more abstract goal to a concrete state of affairs to be achieved, taking into account contextual features. Suppose, for example, that an agent who is a PhD student aims for financial well-being; this generic desire is her maintenance goal[8]. It is too abstract to be acted upon. The reasoner, knowing that fulfilling the PhD requirements is likely to lead to financial well-being, reduces it to the sub-goal of fulfilling the PhD requirements. Given that all other requirements (e.g., credit acquisition, course auditing) have been satisfied, the sub-goal amounts to submitting the PhD dissertation before the deadline. This is a concrete target to be pursued, which becomes the agent's achievement goal. Backward chaining then reduces it to a temporally ordered sequence of actions, which the reasoner executes: initiating the writing at a particular time (given the agent's knowledge of her general daily schedule, of her writing routine and speed, etc.) and continuing writing until the thesis is done. The initiation of writing is further reduced to basic acts, e.g., directing the right hand towards the laptop keyboard, more particularly towards SHIFT and *A* keys for instance.

From a computational point of view planning raises a number of concerns, which need be addressed in a modelling endeavor. In AI these concerns surfaced in the late 1960s, when researchers were devoted to constructing intelligent non-organic agents in the form of robots able to plan their paths. They constitute the notorious frame problem (Shanahan 2009). Originally, the term was used to denote a technical problem in logic-based artificial intelligence related to the use of logical methods to compute action sequences (McCarthy and Hayes 1969). Later on, philosophers took these concerns to point towards deeper, epistemological problems, e.g., the problem of induction (more specifically, Goodman's 'new riddle of induction'), or 'Hamlet's problem' of relevance and when to stop thinking (Fodor 1987). In the 1980s the different understandings of the frame problem, the narrow technical one in AI, and the extended one in philosophy, generated heated debates between representatives of different fields of inquiry. In what follows, I introduce the concerns that give rise to the logical problem. Irrespective of

[8]The reasoning steps are presented in Section 4.1.7 and formalized in Section 4.1.7.

CEU eTD Collection

controversial extensions of the frame problem in epistemology and philosophy of mind, I suggest that a formal solution has salutary effects in cognitive science; that is, it provides better structured and more detailed theories for cognitive phenomena attested in the empirical sciences. The logical model of practical reasoning set forth in Chapter 5 is meant to evidence such positive effects.

The capacity for planning rests on some underlying causal knowledge, because it involves computations of the causal development of actions and states of affairs over time. In order to form ordered sequences of actions meant to achieve particular states, planning presupposes reasoning about change. Given causal explosion – the fact that one event may have, and usually does have, more than just one (intended) effect, it is difficult to know in advance (while planning, and before acting) what changes and what remains unchanged upon action performance. From a logical point of view 'difficult' means computationally intractable, or at least so it did before the advent of non-classical formalisms, when classical logic was the monolith of formalization attempts.

In the process of planning, for sound and functional goal reduction via causal reasoning, agents must update their beliefs about the effects of actions. They must stay away from those undesired effects that might form obstacles in reaching the goal. But the frame problem points to the fact that it is impossible to enumerate a priori all the unforseen obstacles (indeed, more often than not, unforseeable). The processing difficulty increases under the time constraints facing real agents' action performance in order to achieve their goals.

> No planning formalism is likely to give 'provably correct plans' in the sense of classical logic (van Lambalgen and Hamm 2005, p.42),

whereby all possible models of the planning context, including the unforseeable effects, are considered.

Dennett's misfortunate robots (Dennett 1987) point to a threefold understanding of the frame problem that faces planning agents:

(1) computing unintended effects of actions (and taking them into account in the process of goal reduction),
(2) computing effects and distinguishing them from non-effects of actions (and gracefully avoiding thought about what remains unchanged in the world as an effect of particular actions), and
(3) sorting relevant from irrelevant effects (and again, avoiding thought about irrelevant ones).

The third understanding is an issue of debate between AI researchers who hold a narrow view of the frame problem (e.g., Hayes and McCarthy) and philosophers who favor a broad view (e.g., Fodor). The first two are essential and generally accepted as constitutive of the frame problem. The fact of the matter is that intelligent agents are somehow able to ignore most conceivable updating questions, to confront only realistic possibilities (Haugeland 1987, p.82),

CEU eTD Collection

and to proceed to action based on those. Actions are indeed rationally prescribed by plans (Pollock 1995), but only until further evidence (e.g., about an unintended side-effect of an action, about previously unconsidered contextual details). Relatedly, action plans bring about goals *presumably*, to the best of a selected part of an agent's knowledge. It is noteworthy that the issue is not of deciding between relevant and irrelevant details to be taken into account when planning, rather of avoiding such decision and focusing straightforwardly on the relevant ones. Retraction of plans is possible as a result of additional reasoning or additional information, so an intelligent planner must reason flexibly and stay open to revision. This points to the fact that planning requires non-monotonic reasoning skills.

The suggested example of planning to submit the thesis before the deadline should make these considerations easier to understand. The agent need not (and hopefully, for the sake of action, does not) consider the possibility that the lightning strikes while her hand is approaching key *A* on the computer keyboard, causing the window to break, and the shattered glass to cut her index finger, thus preventing performance of the initiating event and, further, the success of the plan for financial well-being altogether. Also, once she starts writing, she should not digress into considering the logical (but unfortunately unrealistic) possibility that if she closed her eyes for two seconds and removed hands from the keyboard, the good fairy will make such that the writing activity goes on without her will and motor control of her hands. The plan to start writing and to continue writing until submission should be somehow 'exonerated' from dealing with unintended and unexpected lightnings, or actions from the part of fairies or gremlins. Otherwise, most likely that actual writing would never start. However, should there be breaking news on the CNN about good fairies walking around and making people's wishes come true for nothing, she might retract the plan to start writing immediately, and allow a couple more minutes for the fairy to show up. More realistically, should there be a problem with the electrical system causing her computer to shut down, she should be able to adapt and re-compute her plan in the new, abnormal, context, e.g., revert to pen and paper for a while.

The bottom line is that a formal model of the planning processes must capture both the pragmatically wise ignorance and its counterpart – flexibility. Planning can be modelled in a formal system that addresses the above mentioned concerns.

> What is needed is a system that genuinely ignores most of what it knows, and operates with a well-chosen portion of its knowledge at any moment (Dennett 1987, p.55).

Closed-world reasoning is such a system, hence it provides a suitable expression of intelligent planning. It is informally introduced below, in Section 2.5. Its implementation in a full-fledged logical system is described in Chapter 4.

The overt, behavioral outcome of planning derivations is performance of the action sequence prescribed by the action rule. Plans can also be used for further practical reasoning, e.g., as a constraining element on other planning derivations as in Bratman (1987). Moreover, the

CEU eTD Collection

reasoning structure, i.e. backwards reasoning, used for plan derivation provides an explanatory tool. An explanation-oriented plan computed by an observer and assigned to an observed agent draws on the same reasoning mechanism.

2.4. Explanation.

Action explanation is the use of reasoning capacities in order to understand other agents' behaviors. I will talk about two kinds of agents: observers, or the ones who explain, and observed agents, or those whose actions and achieved results call for justification from the part of the observers. Throughout the Section, in order to facilitate linguistic expression, the former is referred to with "he", while for the latter I use "she". The explanatory queries that I am interested in are observed actions performed by human agents.

Explanation paves the way towards satisfactory interpretations of human actions.

> Explanation is a goal-directed human activity. It depends on what is relevant or important to the explainer, it satisfies a special kind of explanation-seeking curiosity, it answers 'why?' questions (Gopnik 1998, p.108).

The function of practical reasoning as it manifests itself in action explanation is to make sense of other agents' actions in the context of observation[9]. Furthermore, a long-term cognitive outcome of the human explanatory practice is knowledge acquisition. All in all the explanatory practice is pragmatically justified (see also Leake (1995)). Practical reasoning answers 'why' questions by finding reasons for the actions performed by other agents.

Explanation is based on teleological practical reasoning about actions and their effects. For the current modelling purposes, I assume that, inasmuch as goals focus agents' planning, they also focus the explanatory processing of other agents' actions. The assumption is supported by empirical evidence about humans' justificatory practices; Chapter 1 gives several examples of infants' goal centered interpretation of action contexts. The cited evidence maps well on teleological approaches to explanation (Sehon 2010; Hacker 2009) in the philosophical literature[10]. According to these, agents' goals are the crucial explanatory factors. They justify action performance in a given context, and thereby contribute to other agents' understanding of it. Goals provide agents with reasons for action, and, correspondingly, reasoning in terms of goals promotes other agents' explanatory attempts. Other contextual features also contribute to teleological explanation; actions may be justified with respect to a certain goal because of specific features of the context. For instance, drinking water from one's joined palms is not justified with respect to one's goal to drink water if in a fully equipped kitchen; the same goal however explains the action if performed next to a forrest spring with no containers around. In this case

[9]Because I discuss practical reasoning in an epistemic context (i.e. understanding), and not in a moral one (e.g., excusing, exonerating), I use the terms 'justification' and 'explanation' interchangeably.

[10]In the current endeavor I do not take sides in the debate between causalist and irreducibly teleological views of explanation for human agents' actions. I model a kind of action explanation, which is best understood in teleological terms; it is focused by goals. Moreover it fares well with empirical evidence of human infants' teleological justificatory tendencies.

CEU eTD Collection

the goal is a forward-looking reason, whereas the contextual features are backward-looking reasons. If concrete achievement goals are forward-looking reasons for actions, the more abstract maintenance goals with their intentional component are backward-looking reasons[11]. The teleological approach to explanation grounds explanatory processes on some 'independent rationality principles', which, together with principles of physical science, underly goal-centered explanations (Sehon 2010, p.127). The assumptions to be described in Section 2.5 are principles of the former kind.

The explanatory process for actions performed by observed agents is a form of model[12] construction via reasoning *towards* an interpretation van Lambalgen and Stenning (2008). It involves two inferential steps: the generation of an explanatory hypothesis, followed by the validation of the hypothesis. The two steps are a form of inference to the best explanation (Harman 1965). Generation of an explanatory hypothesis is the tip of the iceberg. It is followed by a validation procedure, attempting to confirm that the hypothesis is 'the best'. The inferential outcome is an action rule: 'In order to bring about s at time $t + n$ the agent X should do a at time t', or, more simple, 'action a is performed at t FOR s at $t + n$'; the goal state s justifies the observed action a at time t. The goal s gives reasons for a if the action must have been performed by X in order to attain s at some later time $t + n$, given the features of the context.

The generation of hypotheses to explain observations is called abduction (Kowalski 2006; Kakas et al. 1998). Abductive reasoning is the essential inferential tool for diagnosis (Kakas et al. 1998; Douven 2011). It is also useful for explanatory interpretation of the actions that agents observe around them; in this sense, it contributes to living in an intelligible social world. I use *abduction* to refer to the first stage of explanatory computations: the weighing of evidence (Charniak and McDermott 1985), i.e. of observations, to derive the purpose of actions. This use is more restricted compared with the one coined by Charles Sanders Peirce, who meant abduction as a general inferential pattern in explanation, which included both the formation and the acceptance of explanatory hypothesis (Douven 2011; Frankfurt 1958).

Abduction, or forming a goal conjecture for the observed actions, initiates the explanatory interpretation. An observer hypothesizes the observed agent's reasons for action, which, in a teleological framework, is the agent's goal in performing an action. By forward reasoning from observations an effect of the action is conjectured as the goal to be achieved. As discussed in Section 1.1.1, given the multifinality of human action, goal assignment faces reasoners with an inverse problem. The core of the problem is that actions have a number of causal effects, of which only one is the goal[13] for which the action is performed, and which justifies the action in the particular context. A flexible reasoning mechanism is needed to produce a goal hypothesis given contextual evaluation from the perspective of a relevant set of background knowledge.

[11]The direction of reasons (backward or forward) takes as a temporal point of reference the moment of action performance by the agent. See also footnote 2 in Chapter 1.

[12]This is a semantic notion of model; see Section 4.1.4 for more formal details.

[13]Abstracting away from end-states higher in the teleological hierarchy, for which the goal in question is a sub-goal.

CEU eTD Collection

Evaluation is constrained by normative principles or assumptions. In Section 2.5 I present for consideration a kind of such principles.

Returning to the example used in the previous Section, if I saw my friend restlessly typing on her computer, I could not be sure whether she is writing a love letter to her long-time not seen boyfriend, or attempting to submit her dissertation before the deadline. That is, given the truth of the observational premises, I cannot deduce (i.e. infer with necessity) the goal. However, given that I know that the submission deadline approaches, I might form the hypothesis that she is writing in order to finish the thesis before the deadline. I abduce the tentative explanation that her purpose in typing restlessly is thesis submission. Depending on how interested I am in the issue, I may then engage in validating the hypothesis.

> The existence of multiple explanations is a general characteristic of abductive reasoning, and the selection of 'preferred' explanations is an important problem (Kakas et al. 1998, p.237).

If abduction is understood as generation of a goal conjecture, the 'selection of "preferred" explanations' amounts to validation of the goal as the best explanation. Reasoning proceeds in an attempt to confirm (until further evidence) that it is precisely the hypothesized goal that justifies the agent's action in the context of performance. The goal grounds the inferential path. The validation procedure is an attempt to decide whether the goal is a good reason for the agent to perform the action that called for justification. Like in the case of planning, a 'good' reason is a sufficient reason. The best explanation is one that sets forth sufficient reasons for an agent to act the way she does, in the context she does. A goal is sufficient reason for an action if no other action would have achieved it more efficiently in that context.

The capacity for planning is recruited at this point in the goal validation process. Based on the functionality of goals as reasons, and on the computational isomorphism between planning and the validation of goals as reasons for action (i.e. backwards computations from goals to actions), I suggest that action explanation is subserved by human agents' capacity for planning. A planning algorithm is useful for human action explanation because it helps observers, as Donald Davidson (1963) put it, to see "what the agent saw in her action". It provides a counterfactual simulationist[14] procedure for the validation of a goal as an explanatory conjecture. Explanation through planning takes into account what I, the observer, would have done if I were to pursue the goal hypothesized for the observed agent, and compares this self-plan with actual observations. The validation of the explanatory goal hypothesis may involve two planning steps. At each step, a plan for the conjectured goal is derived, using the backward reasoning mechanism described in the previous Section over a different knowledge database.

Before introducing the validation strategy, I articulate first its expected computational outcomes. The well-formedness condition for planning introduced in the previous Section, i.e. 'no disjunction in conclusion', also applies to the planning computations that subserve action

[14]This is my own suggestion for construing simulation. It diverges from the way is generally treated in the recent cognitive science literature.

CEU eTD Collection

explanation. Validation of the conjectured goal amounts to the expectation that precisely the observed means is prescribed by backwards planning from the goal in the given circumstances. The abduced goal provides an acceptable or good enough (Lipton 1993) explanation insofar as its corresponding action rule, or the plan assigned to the agent, is composed of a unique sequence of actions that bring about the goal in the context of observation. The unique action rule amounts to concluding that the observed agent could not have done better in order to achieve the hypothesized goal in her context. Inference to the best explanation is thus conceptualized as assignment of a goal, which justifies the observed behavior because it represents the one means that the agent could have enacted for its optimal achievement, given contextual constraints. Well-formed explanations based on sufficient reasons are related to the normativity of means – ends relations. The criterion 'no disjunction in conclusion' means that teleological explanation is satisfactory if the plan assigned to the agent is non-branching. If no other action sequence but the one that has been observed constitutes a plan for the goal abduced at the beginning of the explanatory process, from the perspective of a rational agent acting in the context of the observed one, then the goal counts as a sufficient reason for action.

Planning-for-explanation proceeds as follows. Upon availability of a goal conjecture, the observer first reasons backwards towards what he himself would have done if he had the same goal as the agent whose actions he interprets. This is an offline use of the self planning capacity, abstracting away from specific contextual details. It is offline because it is not action-, but explanation-oriented.

> The term 'offline' refers to situations in which action-related processes are temporally separated from the relevant perceptual processing (Schütz-Bosbach and Prinz 2007, p.351).

The observed agent's goal is simulated abstracting away from the specific features of the context where she acts. In this sense offline planning is coarsely simulationist. The output is what the observer would have done normally, in habitual circumstances, to pursue the goal abduced for the agent's actions. It is a default action rule. It reads as 'In order to make s be the case, I would have done a', where s is the goal abduced in the first explanatory stage, and a is a default action. Now, there are two possibilities to continue the validation procedure.

On the one hand, a, the resultant action from offline planning, may be precisely the a that the other agent performs in her context, i.e. the action to-be-explained. This is because the agent performs a habitual means action. In this case, the observer's offline planning suffices to validate the goal conjecture. The goal assigned to the observed agent justifies her action performance because the explainer would have normally done the same. Presumably, the context where the observed agent pursues goal state s is normal (or at least not significantly abnormal), just like the observer's.

On the other hand, she may have been observed doing b, in which case s is the goal abduced as an explanation for b. In this case the result of the explainer's offline planning computations

CEU eTD Collection

do not match observations – he would have done a for s, whereas she did b. The observer's counterfactual plan conflicts with the actual conduct to be explained. So far, s does not fulfill the justificatory role for b. The two actions a and b are represented as a disjunction in the action rule grounded by the conjectured goal s – 'In order to make s be the case, do a or do b'. They are alternative means to achieve s – one that would have resulted in s in normal conditions, and one that did result in s in the observed context. The well-formedness condition is not satisfied. Hence teleological interpretation is not explanatorily satisfactory at this point. Hacker (2009, p.83) seems to describe the explanatory instability arising from the detection of such a mismatch when he says that

> [. . .] sometimes it is not obvious what another person is doing, or it is opaque what his reasons are for doing what he is doing.

More processing is needed for an acceptable explanation. Two potential strategies are available:

(1) cancel the goal conjecture, and engage in goal attribution anew (i.e. re-compute the teleological structure of the situation from scratch);
(2) keep the goal conjecture, and perform a more extended justificatory contextual evaluation of efficiency.

Regarding the former, I wish to emphasize that adopting a new goal hypothesis requires additional computations in the form of planning *de novo* what the explainer would have done if he wanted to achieve that different outcome.

The latter option is a continuation of the pursuit of sufficient reasons. It amounts to searching in the context an additional constraint that singles out the action observed as *the* means for the initial conjecture. It does not require canceling the explanatory goal hypothesis. In a Quinian fashion, (2) is an attempt to construct a coherent web of teleological belief targeted at updating the more peripheral zones, not at the core containing the goal. The objective is to supplement the initial goal conjecture with contextual constraints, which, taken together, constitute sufficient reasons for the agent's observed behavior.

The choice among (1) and (2) is determined by cognitive economy considerations: the least cognitively demanding, and most cognitively rewarding, is pursued. Baggio et al. (2008) provide neurological evidence that teleological re-computation is cognitively demanding in the domain of discourse interpretation when a goal conjecture is called upon by the semantics of the English progressive. However this does not show that goal revision is necessarily more costly and hence (2) is the default initial strategy. The computational costs may be evaluated in a case-by-case fashion, depending on the degree of confidence in the initial goal conjecture: low confidence recommends goal revision, whereas high confidence prompts enhanced contextual evaluation. I list some conditions of high confidence (without claiming exhaustivity):

a) the action effect categorized as a goal is highly perceptually salient,

CEU eTD Collection

b) prior experience with similar actions is available, and it is associated with similar goals (i.e. strength of the action – effect association),
c) the goal also explains other behaviors.

When at least one condition obtains, contextual re-evaluation is preferred. A rational explainer thus attempts to validate the teleological explanation of b via s by enhanced inspection of the context, looking for aspects that may justify b under the conjecture that b *is for* s. The other disjunct, the observer's default action a, may be cancelled by previously disregarded contextual features, so that a single action, b, is ultimately prescribed by planning from s. In the second step of the planning-for-explanation procedure, the observer derives a plan by taking the other agent's perspective. Currently observed contextual features constrain the stepwise reduction of the hypothesized goal to means actions. All in all the observer's backwards reasoning is constrained by the conjectured goal (which grounds computations), plus the features of the context where the other agent acts (which shape computations). The resultant action rule is a plan to be assigned to the agent in pursuit of the hypothesized goal in that context. It is explanatorily satisfactory if the explainer's default is ruled out in the context of observation.

I instantiate these considerations in the case of the working example. Suppose that my offline planning computations ended in 'In order to submit the the thesis in two weeks, write with small breaks for walks in the forrest'. The writing with small breaks differs from the observation of restless writing. That is, the offline planning computations of a rational agent – myself – yield a distinct cause action for the desired outcome, or an alternative means. I am left with the disjunctive rule 'In order to submit the the thesis in two weeks, write restlessly *or* write with small breaks for walks in the forrest'. The default, automatic planning computations run by me show that a rational agent could have normally done otherwise than the agent. Therefore the goal 'submit the the thesis before the deadline' does not explain my friend's restless writing. If I were in a frenetically explanatory mood, I would have to engage in further computations. Suppose that I am quite confident that the agent's goal is thesis submission in two weeks, because this also explains a number of other behaviors (or lack of) of hers, e.g., not joining her group of friends to the performance of their favorite rock band. Therefore I try out the less effortful alternative first – I attempt to find some contextual constraints, neglected before, that would favor one of the two disjunct causes over the other. For instance, suppose there is no forrest around my friend's writing place, a contextual aspect that constrains her action. Because of it, the second disjunct is no longer a rational choice, so planning from her perspective reduces the goal to a single action, namely restless writing. Her restless writing is thereby explained by her desire to submit within two weeks.

If no contextual feature singles out the observed action as necessary for the abduced goal, then re-computation of the initial explanatory stage remains the only way to go about explaining. The observer runs again the goal attribution abductive process, and assigns a different goal to the observed agent. The action context is teleologically re-interpreted. A satisfactory

CEU eTD Collection

justification is sought by generating a different goal conjecture. I wish to stress that the distinction between degrees of abstraction of goal representations does not furnish a distinct goal hypothesis. By abduction of a different goal I mean initiation of a causally unrelated teleological thread, and not merely moving up or down along the hierarchy of the initial hypothesis. For the case of the observation of restless writing, the intention to submit the dissertation before the deadline (i.e. a higher-order maintenance goal) to replace the initial concrete goal to submit the dissertation before the deadline is not based on re-computation of the teleological structure. Such re-computation occurs when a hierarchically unrelated goal, e.g., finishing a love letter before the departure of the post office ship towards the residence of the beloved, is conjectured as a potential justification of the restless writing. Upon a novel goal conjecture, the validation process via planning computations starts anew.

Throughout this Section I often use expressions that signal the presence of uncertainty in the explanatory reasoning processes – e.g., tentative explanation, I *might* conclude that state s is a goal, the observed agent's context is *presumably* normal. This use of language shows that explanatory processes, just like planning, are an instance of reasoning with uncertainty. A more appropriate name for 'the best' explanation of an action a would be 'the *presumably* best' explanation, where *presumably* should be understood just like in the case of planning: to the best of a restricted set of the explainer's knowledge, agent X's goal s is the best explanation of agent X's performance of action a in her context. The practical reasoning processes meant to render observed actions intelligible by inferring a best teleological explanation, are not deductive; nonetheless they can be represented in logical form. To this end, in modelling human explanatory practices, I adopt an abductive approach to action explanation (Shanahan 1993) in its guise as inference to the best explanation. Hypothesis formation regarding an agent's goals for action has been linked with a logical treatment of abduction, and thereby it has been widely used in logic-based approaches to intention recognition (Sadri 2011).

The inferential processes face reasoners with several computational problems. For one, abduction understood as goal hypothesis generation must deal with multifinality. Second, validation via planning shares all the computational difficulties presented in the previous Section.

In the next Section I spell out a kind of reasoning that suggests a meaning of *presumably* and, as a consequence, of how *the best* explanation can be inferred – closed-world reasoning.

CEU eTD Collection

2.5. Practical reasoning as closed-world reasoning.

*Motto:*To know how to reason is in part to know when to reason (or not to reason) about what. (Bratman 1987, p.66)

The concerns of goal-centered reasoning presented in Sections 2.3 and 2.4 may hinder attempts to explicate formally the practical reasoning processes involved in action planning and action explanation. How can one engage in practical reasoning for a finite amount of time, i.e. plan ahead one's own actions for one's goals, attribute goals to actions performed by other agents, or interpret actions as means for goal attainment, when one is faced with multifinality,

equifinality, and quasi infinite possibilities of unexpected events? My proposed answer is: by closing one's 'world of reasoning', that is, by using constraining assumptions to minimize the model of teleological interpretation, and to go on reasoning with a minimal amount of data, minimally extending it when necessary. Closed-world assumptions (henceforth CWA) are feasible solutions for the concerns of planning and explanation.

> A wide range of everyday tasks involve closed-world reasoning: planning, and adapting to failures of plans during their execution, diagnosis of causes, causal reasoning itself, reasoning about mental behavior and states, interpreting speaker's intentions underlying discourse (van Lambalgen and Stenning 2008, p.45).

The computational concerns recommend the utilization of closed-world reasoning (henceforth CWR) as a formal framework for practical reasoning. CWR is a principled, smart way of jumping to conclusions (McCarthy 1980a), and thereby it fosters the actual manifestation of reasoning in action performance. It is computationally less complex than classical logic (van Lambalgen and Stenning 2008; van Lambalgen and Hamm 2005), and therefore it is more accessible a strategy for real agents' reasoning and acting in the real world. It provides the simplest formalization tool for default, commonsense reasoning. CWR amounts to the use of assumptions to frame the inferential space, e.g, by disregarding unexpected and implausible effects of actions. It is non-monotonic, because it allows revisions of conclusions as new information is added to the agent's reasoning database. In this Section I describe the kinds of mechanisms by which CWR frames the inferential space, and thereby facilitates the human practice of practical reasoning.

Despite the intricacy that manifests itself as computational intractability, even the youngest members of the human species show in their behavior the ability to bypass gracefully the problem, according to the evidence presented in Section 1.1.1. The developmental psychologists' suggested solution for the cognitive mechanisms that support teleological reasoning in infancy has also been detailed there. The core of the idea is that infants' proneness and ease to reason about goals is explained by the use of constraining assumptions on inferences (Csibra 2008; Csibra and Gergely 2007; Gergely and Csibra 2003). The function of these assumptions is to reduce the inferential scope to manageable dimensions; they *frame*, as it were, the computational complexity. Consequently they allow human infants to come up with explanatorily useful teleological interpretations, to predict other agents' goal-directed behaviors, or to set up their own goals and action plans in response to these interpretations. The main example of such a strategy is spotting causal relations of the kind 'action x causes y_1, y_2, . . ., y_n', and then picking y_k to be the goal of performing x based on evaluation of the causal efficiency of x in bringing about y_k. The assumption that validates goal selection among several effects is that the goal y_k is most efficiently, or optimally brought about in the given physical context by the cause action

CEU eTD Collection

x. An assumption of this kind was labeled either the principle of rational action, or the principle of efficiency (the two being used interchangeably).

> The outcome (the effect) of an action may, or may not, be seen as the goal, depending on whether the outcome is judged to justify the action in the given situation. Such normative evaluation of actions is based on the principle of rational action (Csibra and Gergely, 1998; Gergely and Csibra, 2003), which allows for the assessment of the relative efficiency of the action performed to achieve the goal within the situational constraints given (Csibra and Gergely 2007, p.7).

With certain reserves regarding a blunt equation of rationality with processing in terms of physical efficiency (see discussion of efficiency at p. 40, and Section 2.2 below), I adopt the proposed strategy for goal-centered reasoning, and formalize it in the model. More specifically, along the lines of van Lambalgen and Stenning (2008) and van Lambalgen and Hamm (2005), I propose that the inferential processes at work in teleological practical reasoning are framed and guided by closed-world assumptions (CWA's), and that practical inferences are a form of CWR. The principle of efficiency is one example of a CWA. Another example are the hidden assumptions that agents need to make when they plan (Dennett 1987), and thus also when they use backwards planning for explanatory reasons. Such are the assumptions about information that does not change as a result of actions, or about the non-effects of actions – e.g., that my skin color remains the same after pouring beer from a bottle into a glass. These facilitate planning because, should I want to have a sip of beer from the glass, I need not spend an infinite amount of time considering the infinite chain of logically possible causal relations. I should merely incline the bottle and pour, and then bring the glass towards my mouth. Only in case I notice that something abnormal or unusual is the case, e.g., that the glass remains empty although I executed the pouring action, I should consider the realistic possibility that, e.g., I poured from an empty bottle, and re-compute a plan for the target of sipping beer. Similar framing principles have been dubbed in the psychological literature 'real-world assumptions that are maintained unless otherwise specified' (Onishi et al. 2007, p.125), or 'normality assumptions' (Gergely and Jacob 2012).

CWR has appeared in the logic-based AI literature precisely due to the need to deal with the frame problem as the problem of an infinite amount of negative causal information (i.e. what remains unchanged), which leads to both storage and processing insurmountable problems in the course of planning and/ or explanation. The answer suggested by CWR is, at least apparently, bluntly simple: assume that all information that is not explicitly mentioned as true is false, and proceed reasoning from these grounds (Reiter 1978, 1992). It is important to note that 'assuming' means not considering, not engaging in a decisional process. CWA's do not provide a heuristics for deciding what should be considered, they simply stipulate disregarding everything that is not explicitly mentioned, or derivable from the information mentioned. In fact

CEU eTD Collection

precisely because of this CWA's facilitate reasoning in reasonable amounts of time, and allow agents to stop thinking and proceed to action. CWA's are most often implicit; only in special cases (e.g., dealing with one's known paranoid tendencies) they are consciously used to frame inferential processes. The CWR solution to the frame problem belongs to the framework of the sleeping dog strategy, which is based on the motto 'let the sleeping dogs lie' (McDermott 1987). The 'sleeping dogs' are the information that is neither explicit, nor inferred from it. It is not considered in the process of reasoning. For instance, if upon checking today's attendance list a teacher does not see the name X, she calls to answer a question Y whose name is on the list, upon the implicit assumption that X is not there because his name is not on the list. X's name is a 'sleeping dog', and the teacher's behavior illustrates action upon CWR. Initially the logical means for CWR was predicate and domain circumscription (McCarthy 1980b). Many others have followed. For the current modelling purposes, I use logic programming (van Lambalgen and Hamm 2005) for reasons that shall become explicit as the story continues.

The frame problem surfaced in the discussion of goal-based reasoning, in particular of planning, as the concern with not wasting time considering the non-effects of actions (recall the second version of Dennett's robots mentioned in Section 2.3 at p. 47) when computing a plan. The basic way by which CWR makes this possible is the CWA for reasoning about abnormalities – CWA_{ab}, which prescribes that, unless positive information is available, one may assume that no abnormality with respect to the habitual, default running of a process occurs. A more accurate expression of action rules is then 'In order to bring about state s, agent X should perform actions $a_1, \ldots, a_n$ at times $t_1, \ldots, t_n$, unless there is positive information of abnormality'. The conditional in the action rule is strengthened with an 'unless' proviso, which points to the flexible use of the CWA's.

'Positive information of abnormality' may refer to current observation, or consequences derived from it based on background knowledge. Attached to the CWA_{ab} used for concrete instances of reasoning is a disclaimer, i.e. a set of conditions defining the kind of positive evidence that may constitute abnormality. They specify what the indefinite pronoun 'something' in the 'unless' proviso may plausibly refer to. In relation to the name of the strategy of the 'sleeping dogs', these conditions make explicit when some of the dogs may, in fact, awake. The set of conditions is open, because so is the set of logical possibilities – in principle, there could be gremlins, fairies, the laws of gravity may stop holding, etc. Some of these abnormal cases however are more realistic than gremlins, and it is those which constitute the abnormality conditions. They are context dependent bits of commonsense knowledge (Lifschitz 1985), hence the application of the CWA_{ab} to practical reasoning depends on the world knowledge of the agent (van Lambalgen and Stenning 2008). Abnormality conditions specify exceptions (e.g., In order to have a pet that flies buy a bird unless *the bird is a penguin*), and thereby spell out when the assumption is overridden (e.g., when a little penguin is the only bird available in the pet-bird shop window). In the example, the 'unless' proviso contains a concrete case of abnormality; however, abnormality conditions may encompass a broader range of exception via

CEU eTD Collection

rule-like formulations over variables. It goes without saying that abnormality conditions are checked by reasoners if and only if planning or explanation computations fail (i.e. do not fulfill well-formedness criteria) under the assumption that nothing abnormal is the case. Otherwise the CWA_{ab} would be self-defeating.

The selection of relevant factors for the current reasoning process is done via interaction between working and long term memory, under the closed-world kind of constraint that the sphere of relevance is always maintained minimal. This interaction allows a context-dependent delineation of abnormality conditions. CWR deals explicitly with the wise ignorance needed for action performance and action explanation in real time, and gives over the problem of relevance to the constrained interaction and resultant information flow between different storage systems.

CWA's are a reasoner's auto-noetic presumptions of complete knowledge (Kowalski 2011): disregarding abnormalities (unless explicit evidence to the contrary) is based on epistemic trust in the completeness of one's own knowledge database (background knowledge, as well as episodic observational representations), and in the rationality of the observed agent. This implicit epistemic trust justifies the label 'credulous reasoning' (van Lambalgen and Stenning 2008) for CWR. The CWA's for action explanation have a constraining effect on interpretation, which is analogous with the effect of Gricean pragmatic maxims for discourse understanding (Grice 1975). A closed-world formulation of Grice's maxim of quantity, for instance, reads as 'The rational agent whom I observe must have done everything he should, and nothing more than he should in order to attain her goal, unless I'm missing something'.

I now take turns to elaborate on how different facets of CWR enable two-staged teleological explanation (conjecturing a goal, and checking the validity of the goal), and planning one's own actions. In the next Section I address the use of CWA's for assigning goals. Goal validation by means of backwards derivation mechanism using CWA's is presented in Section 2.5.2. In Section 2.5.3 I discuss action-oriented planning with CWA's.

2.5.1. Closed-world reasoning for goal attribution. The CWA's require the reasoner-explainer to construct a minimal teleological interpretation of the situation, based on positive evidence from the context of observation, to which a minimal set of relevant prior knowledge is applied. Action explanation with CWA's proceeds by construction of a minimal model (van Lambalgen and Stenning 2008) for the actions that must be made sense of[15].

Goal assignment initiates reasoning towards a teleological interpretation. As explained in Section 2.2, goals can be represented at different levels of abstraction. So it is a legitimate question as to the kind of goal that is sought. Because change in the world is perceived based on evaluation of physical processes and states of affairs, a physical change of state fills the goal placeholder to begin with. That is, human agents assign observed agents physical ends, be those achievement goals (e.g., turn on the light), or more general maintenance goals (e.g., stay warm). At this first step, a resultant state is selected as a plausible goal via causal individuation

[15]The semantic notion of minimal model is discussed in Section 4.1.4.

CEU eTD Collection

of events, perceptual salience, and prior knowledge. The reasoning mechanism is according to the procedure from developmental psychology literature described above at p. 57. The proposal of such a mechanism builds on a view of cognitive architecture where low-level processes (i.e. the sensory-motor input, memory searches) and high-level ones (i.e. the use of CWA in causal processing) are joined in a bidirectional flow of information (cf. Shanahan (2005)).

The causal individuation of events is a prerequisite of generating goal hypotheses. The empirical evidence depicted in Section 1.1.2 implies that actions are perceived causally from very early on in development. The inferential processes that I am about to describe are highly automatized, and justify the use of the expression 'causal perception'[16]. Causal individuation is grounded by application of the common-sense law of physical inertia (McCarthy 1986) to observations. The law of inertia states that things and properties remain as they are, unless there is explicit information to the contrary. This means that an ongoing event or state is assumed to continue unless a force that ceases it is exerted. The termination of the event or state is a change that is perceived causally. The law has the format of a CWA – positive evidence for causal link between events is given by perception of a force that changes the inertial state. By application of the assumption of inertia to reasoning about actions, agents expect that observed actions go on until something happens that stops the motion, and a change of state occurs. When such evidence occurs, i.e. when an action ends, the inertial CWA is overridden. This licenses the conclusion that a state occurring in its proximity[17] is to be thought as its effect. For instance, if an agent observes a hand in motion, its stopping on a lamp switch violates inertia. As a consequence, the successive change of state, e.g., the state of the light in the room, is perceived as its causal effect.

Once causes and effects have been spotted in a certain context of interest, reasoners abduce one resultant state as a plausible goal conjecture. The main computational problem of abduction was described above as arising from multifinality, because a goal hypothesis must be selected among several effects of an action. Reliance on low-level processes such as memory, or perceptual feature individuation helps distinguish between mere effects, and goals, i.e. the effects *for* which actions are performed. If available, prior experience with the cause action, either of the resultant states, or the context of observation, fosters goal attribution[18]. Perceptual processes support goal abduction too[19]. The perceptually most salient (e.g., brightest, biggest, loudest) effect of an action is recommended as a potential goal, and increases confidence in the goal hypothesis (see discussion in Section 2.4). That is, the termination of an action triggers

[16]This is not meant as an argument in support of a nativist view of causal cognition. I remain agnostic in this respect. For an insightful discussion, see Chapter 6 in Carey (2009).

[17]The constraint of proximity is derived from some of the earliest developmental evidence for causal perception, according to which spatial contact and temporal order are the primary ingredients of the notion of causality in human cognition (Saxe and Carey 2006).

[18]Action – effect associations, and own experience with action production as a facilitating background for teleological interpretation are discussed in Section 1.1.1.

[19] The proposal is inspired by that of Shanahan (2005), whereby perception itself is formally described as abduction, or hypothesis formation about the raw sensory input.

CEU eTD Collection

assignment of the teleological status of goal to its most salient resultant effect, individuated by means of inertia. Consider again the case of the motion of the hand through the air. Suppose that upon its stop on the switch on the wall, two changes of state occur: for one, the light turns on, and second, a yellowish spot appears on the agent's palm. Suppose also (for the sake of the argument) that the context of observation is limited to just these details (of the cause action, and its effects), and that the observer has seen before hands touching wall switches. In this case, I propose that perceptual salience, together with similarity with previously observed effects of switch presses recommend the light-on state of the room, and not the spot on the palm, as a goal conjecture. In the case of novel actions, reliance on exact prior knowledge is minimally helpful; however, it can be compensated by remotely similar experience, and by additional credence to the salience of the present effect.

All in all, inertial reasoning about ongoing actions – a form of CWR – promotes causal individuation of actions and their effects. Perceptual processes that emphasize salience, and stored action – effect associations accessed in relation with the current context, support goal abduction.

2.5.2. Closed-world reasoning for goal validation. Explanatory interpretation of the context continues with the process of goal validation. It is done under the constraint that the assigned goal conjecture is the (sufficient) reason why the agent performs the observed actions. The goal is assumed to be a state of the physical world that is optimally brought about by the actions observed. Hypothesis validation verifies whether the actual choice of goal obeys this criterion, and thereby fulfills the justificatory role.

A useful formulation of the CWA for action interpretation is *All actions are forced to occur by the data, unless positive evidence to the contrary*. In the context of teleological evaluation 'the data' refers to the constraining elements on inferences, namely, the goal of actions, and other contextual features. This formulation of the assumption captures the core of teleological explanation, i.e. that goals justify actions in the context of occurrence, and are expected to provide reasons for action. An action is forced to occur by its goal if it is optimal for the goal, given contextual constraints. In observers' initial reading, 'maximal efficiency' stands for ' maximal physical efficiency'. Only physical features of the context are relevant for considerations of actions' efficiency in bringing about the physical achievement goals conjectured by abduction. The CWA expresses the normativity of the relation between 'the data' – the abduced goal and other contextual features, and the actions observed. An action is *forced to occur* by its goal when it *should* be performed in order to bring about the goal in the specific context. The goal validation process amounts to twofold backwards planning from the goal conjecture. A well-formed (non-branching) plan grounded in that goal and assigned to the agent provides a presumably best explanation.

Offline planning is planning under the observer's CWA_{ab} that nothing abnormal is the case; it yields a prepotent, default action rule. Defaults constitute, until further evidence, presumably

CEU eTD Collection

best action plans. It is noteworthy that CWR for planning is similar with using the principle of efficiency (Csibra et al. 1999; Gergely and Csibra 2003) to select means actions from one's practical database. The default instrumental action for a physical goal is, to the best of one's knowledge and until further evidence, the most efficient one. Moreover, planning over a restricted database is the most efficient in terms of processing effort. Recruitment of habitual actions from one's motor repertoire requires the least amount of processing effort, and enacting the resultant action rule is the most accessible. Motor activation is quickest when default motor programs are called for. Default motor responses provide, *caeteris paribus*, the most efficient route to a desired physical end state.

The corresponding CWA for interpretation of other rational agents' behavior is the expectation that their behavior is standard, i.e. as previously known, unless there is positive evidence for abnormality in the context. If the explainer's default, standard plan does not overlap with the agent's observed actions, then a branching plan is assigned to the agent: she could have done the default, but she did something else. The branching plan is equivocal, and it overrides the CWA that all actions are forced to occur by the data in the minimal, default model. So far the conjectured goal does not provide a satisfactory explanation. The agent's action is exceptional with respect to the default action rule. It remains unclear to the observer why she did the particular novel action she did, and not the one activated in the observer's offline plan. Observation of non-standard goal-directed behavior activates the expectation that there is an abnormality in the context, which overrode the observed agent's CWA when planning her actions; it must have been this abnormality that made her deviate from the standard action rule.

The action may after all be contextually justified by its goal insofar as an exception or abnormality is encountered in the circumstances of performance. A difference between a habitual context where the default plan holds, and the context in which the observed agent acts must be found. Observations are made sense of if a concrete 'something abnormal' is found, i.e. if the abnormality conditions attached to the reasoner's default action rule obtain. If this is the case, the default means is retracted and replaced by the novel means action performed by the agent. The novel means is justified by the goal conjecture, given abnormal features of the agent's context of action. The agent is justified in doing something unusual in her abnormal context, different from the default that she shares with the explainer. Abnormality does not cancel the reasoner's action rule altogether, it merely shows why it should not apply in an exceptional context. That is, satisfactory explanatory processes for exceptional means are non-monotonic. Should the explainer himself pursue that goal in normal circumstances, he would rationally return to the default and disregard the modified action rule that explained the observed behavior.

If the abnormality conditions do not obtain, action explanation in a minimal model of the situation fails. The non-default observed action is under-explained with respect to physical situational constraints: the assigned achievement goal, and the physical features of the situation. It is not forced to occur by these constraints, hence it remains teleologically opaque.

CEU eTD Collection

2.5.3. Closed-world reasoning for action-oriented planning. CWR for action-oriented planning instantiates reasoning *from* a minimal interpretation (van Lambalgen and Stenning 2008) computed during action explanation. The CWA_{ab} for planning prescribes that, *Nothing abnormal is the case, unless there is positive information to the contrary.*. It allows that the backwards chaining of actions from a goal be blind to fairy tales, gremlins, and indeed everything that is not explicitly mentioned in, or derivable from the knowledge database of the current situation (that is, contextual features) and a relevant subset of world knowledge about actions and their effects.

The application of the CWA_{ab} to action-oriented planning licenses agents to proceed to action according to previously entertained (default) action rules, selected according to the particular features of the context in which planning takes place. Paraphrasing McCarthy (1980a), the known means action can be used to attain a goal, unless something (abnormal) prevents it. The CWA_{ab} is used at each step backwards in planning derivations, whenever a sub-goal is sought for an end higher in the teleological hierarchy all the way down to basic actions. CWR licenses at each step of goal reduction processing only the direct implications of actions, and ignoring abnormalities, unless an abnormality manifests itself explicitly and requires that it be considered (McDermott 1987). In fact, precisely because no abnormality is expected to interfere, goals can be reduced to sub-goals. Otherwise, there could always be a gremlin around the corner not allowing for expected effects to ensue. Cause actions or end-states are selected as sub-goals from the database of background knowledge of causal relations. Concluding action rules thus normally contain means from the planning agents' motor repertoire, and these are the ones that they should execute. Rational agents' optimal route towards a goal is along habitual or default means actions. CWR licenses execution of a default action plan, via activation of a motor response corresponding to the basic action obtained in the final step of a planning derivation. The default action is, until further evidence, optimal for goal attainment.

Plans derived by means of CWR are only presumably, or to the best of one's knowledge, correct for the attainment of the goal that grounds computations. Human knowledge is after all fallible. The use of CWA for practical reasoning is tied with acceptance of information as relevant within the limits of the particular context (cf. Bratman (1992)). The openness to revision of plans, as based on non-monotonic processing, justifies calling them 'partial solutions to the problem of achieving a goal' (Bratman 1987, p.28).

Agents' non-standard yet rational behavior may be a consequence of contextual abnormalities, which provide exceptions to action rules. Non-monotonic planning prescribes performance of exceptional behavior in abnormal contexts, where abnormality conditions obtain and provide exceptions to action rules.

As a concrete instance of CWR for planning, I return to the working example of thesis submission. I focus on the case of the agent's plan for the achievement goal to submit the thesis before the deadline. Suppose that writing restlessly is the habitual thing to do when one has very limited time to finalize a paper. Therefore the agent computed the plan 'In order to submit

CEU eTD Collection

the thesis before the deadline, type restlessly on the computer'. This plan is warranted as long as the CWA_{ab} holds. The action rule is more appropriately expressed as 'In order to submit the thesis before the deadline, write restlessly, unless something abnormal prevents writing restlessly'. Suppose now that the agent lives in communist Democrat Republic of China, and knows that there are frequent electricity shortcuts. In case there is one, the computer can no longer be used for writing. This information is part of the agent's commonsense knowledge relevant for the plan that she computes, so it provides an abnormality condition leading to a potential exception to the default plan. An electricity shortcut constitutes positive evidence of an abnormality interfering in the normative causal connection between the prescribed means 'type restlessly on the computer' for the goal 'submit the thesis before the deadline'. It might attach a plan-B to the main plan of typing restlessly, in case the electricity shortcut occurs and overrides the CWA_{ab}. In plan-B the exception provides additional constraint on derivations from the original goal. The plan is thus grounded in the initial goal 'submit the thesis before the deadline', and derivations are constrained by the exception 'there is an electricity shortcut'. The complete action plan could then be 'In order to submit the thesis before the deadline type restlessly, and if there is an electricity shortcut have a break, and resume immediately after you are able to switch on the computer'.

I now proceed to discussing learning about actions and their effects from observation of other agents as a crucial cognitive outcome of practical reasoning, capitalizing on the flexible use of CWA's.

2.6. Observational learning through practical reasoning.

Observation of other agents is a central way by which knowledge is transmitted in the human culture (Gergely and Csibra 2006; Tomasello et al. 2005). Observational learning manifested in imitative behavior is empirically well documented throughout development (Gergely et al. 2002; Gergely and Csibra 2006; Heyes 2001; Csibra and Gergely 2006; Tomasello et al. 1993; Zmyj et al. 2009; Gergely 2003; Bekkering et al. 2000). Knowledge acquisition from observation is implicit, because observation does not occur in a formally established student – teacher environment, where these roles are acknowledged and assumed by each party. Learning is not conscious, just like the underlying processes of practical reasoning and the corresponding assumptions that guide it are not (cf. discussion at p. 36 in Section 2.1).

Along the dimension of the number of observations (single vs. multiple) implicit learning from observation is of two kinds; the reasoning processes that underpin learning are correspondingly different. On the one hand, imitative learning from multiple observations of instrumental actions is likely subserved by inductive reasoning mechanisms, and formalizable in corresponding systems such as Inductive Logic Programming (Muggleton 1999; De Raedt 1997), or probabilistic Bayesian models (Baker et al. 2009; Rao et al. 2004)[20]. Learning selectively from one

[20]The methods that I associate with learning from multiple observation, e.g., Bayesian models, have been used to model also imitative learning upon single presentation with novelty Rao et al. (2004).

CEU eTD Collection

observation, on the other hand, requires enhanced processing of the context from which things are learnt. This is the kind of learning that I model, where observers are faced on just one occasion[21] with a novel instrumental action upon a novel object. Learning is based on teleological explanation of observations, which recommends storing relevant information about what to do, when, what for.

I propose that observational learning is supported by teleological interpretation of observed actions, processed along the lines of Section 2.5. The practical reasoning processes result in activation of the observers' natural tendency to learn what to do, when, what for. The behavioral manifestation of learning is the result of the observer's online planning computations, shaped by particular formats of the epistemic drive triggered in teleological interpretation. The performance of an instrumental action, or bringing about a goal by acting on a novel artifact in another context than the one observed qualifies as practical knowledge acquisition. When observers instantiate in their own behavior the novelty that they are presented with, after some time passes from the demonstration session, they provide a cue for having generalized it. The extent of the generalization is an empirical question worth pursuing – e.g., in what respects must two contexts be similar to foster between-context transgression of information, or learning[22]. Henceforth learning is understood as the generalization of instrumental actions and achieved end-states from a teleologically structured observed context. The learners' behavior in a subsequent similar context is prescribed by planning computations grounded in an epistemic drive activated according to the outcome of teleological interpretation. This view corroborates the basic claim that action imitation or goal emulation provide a diagnostics of observational learning.

Young human agents' approach to novelty is curiosity, and an intrinsic motivation to learn (according to the thesis of natural pedagogy discussed in Section 1.1.3). As stated in the preamble, practical reasoning is concerned with structuring action contexts along means – ends hierarchies. This is the role of teleological explanation discussed in Section 2.4. Abstracting away from the relative character of means – ends structures, reasoning supports roughly two kinds of practical knowledge acquisition: about goal states, and about instrumental actions. Suppose teleological explanation of an action a with respect to an end-state s in context c yields *In order for state s to be the case in context c, do a.* The observer instantiates goal learning if acting so as to bring about s in $c^{'}$, and skill learning if doing a in $c^{'}$. The behavioral manifestation of the former is goal emulation, and of the latter – means imitation[23]. Learning novel goals fosters the acquisition of functional information about artifacts, which is used productively for

[21]Although the novel action is presented three times in a row, as described in Section 1.2, there is no variation whatsoever between performances – they are three identical tokens of the same action, in the same context, during one presentation session, with an interval of seconds in-between. This justifies my calling it single-observation learning. An example of multiple-observation imitative learning is the study of Schulz et al. (2008), where the deterministic vs. probabilistic efficiency of actions for goals is manipulated.

[22] In fact the question is being addressed presently in the developmental psychology lab of the CEU Cognitive Development Center, Budapest.

[23]According to discussion of imitation and emulation in the preamble to Chapter 1.

CEU eTD Collection

the categorization of artifacts starting in the first year of life and constantly developing throughout childhood (Tráuble and Pauen 2007; Casler and Kelemen 2007). Moreover, presumably due to the scarce use of tools even among other primates, the learning of novel goals of transitive actions is extremely rare (Csibra and Gergely 2006). Hence learning of goal states that can be brought about via artifact usage is a particularly relevant instance of expanding one's practical knowledge with respect to novel goal states. In respect to the object of learning, or to the teleological status of the information stored during observation and manifested behaviorally in another context, the model developed here focuses on two symptoms of infants' epistemic drive:

(1) *learn artifact functions*, i.e. infants have a motivation for knowledge update with the effect brought about by using a novel artifact;
(2) *learn practical skills*, i.e. infants have a motivation for knowledge update with a transitive goal-directed action.

Classifying knowledge acquisition along the dimension of the object of learning can well be understood in relation with Ryle's classical distinction between *knowing how* and *knowing that* (Ryle 1945–1946). Function learning is learning *that* an artifact can be used for a particular purpose, or *that* an action has a particular outcome; action learning is the acquisition of know-how or of practical skills. The former, given reliance on functions for categorization of artifact kinds, may provide a convergence point between practical and theoretical knowledge.

I propose that learning about artifact functions is the primary manifestation of practical knowledge acquisition from observation. Acquiring knowledge about artifact functions has intrinsic relevance – presumably, knowing that (action upon) certain objects produces certain desirable results is generally better for a rational agent. The suggestion is backed by empirical evidence regarding the crucial role of functions in infants' knowledge of artifact kinds (discussion in Section 1.1.2). The motivation to learn functions is triggered by observation of a novel artifact, irrespective of the means – ends structure of the action context where it occurs, or the nature of the social environment. The only interpretive reasoning process that is a precondition of learning *that* artifacts are good for specific things is goal abduction, since a function is an enduring goal of an artifact. Because of the strong preference for teleo-functional aspects of the world is well evidenced across development (Kelemen and Carey 2007; Kelemen 1999a), and because knowing the function of an object facilitates agents' goal attainment via tool use (the pragmatic argument), information about object functions primarily calls for storage in the practical knowledge database of rational agents. When a new object is observed as used for a specific purpose, the goal information is generalized over a different context in which the same object is presented affording action performance from the observer. I thus suggest that the genericity bias (Csibra and Gergely 2009), although enhanced by ostension and communication, is not restricted solely to such demonstrations, at least not when instrumental behaviors

CEU eTD Collection

are performed on novel artifacts resulting in salient changes of state that afford teleological categorization as goals.

Novelty prompts infants' manifestation of the pedagogical stance (Gergely and Csibra 2006; Topál et al. 2008) in the form of readiness to learn from the demonstrator. However, the stance is manifested selectively. In Section 1.1.3 I referred to empirical evidence showing that human infants exhibit a propensity to learn in particular from communicating adult agents. The tendency is quite robust, and learning occurs even when it involves apparently irrational generalization in habitual contexts of teleologically opaque, under-explained (Gergely et al. 2002), or altogether unnecessary (Nielsen 2006) behaviors. It can be said that the motivation to learn outplays the explanatory success when the source of novelty is a communicative agent. In fact, at very early ages at least[24], communicative ostensive demonstration of novelty is a necessary condition for generalizing instrumental behaviors (i.e. imitative learning of practical skills) that do not seem to comply with physical efficiency (Király et al. 2013; Southgate et al. 2009). I ended Section 1.2.1 (p. 29) with some questions raised by empirical evidence that the formal model attempts to address; one of those concerned precisely the mechanism by which social factors, such as communication, interact with inference-driven teleological interpretation to shape infants' acquisition of practical knowledge manifested in imitative behavior.

The core of the answer set forth in modelling is the suggestion that different goal representations, one's own as well as another agent's, are handled by the same computational mechanism described in Section 2.5 – the multi-level single-mechanism view of teleological processing. The idea is along the following lines. Communication supports activation of the epistemic motivation to learn from novelty. When an agent demonstrates novel transitive instrumental behaviors upon novel artifacts while conveying communicative ostensive signals, young observers expect there to be something relevant that the agent communicates about the referred object (according to the principle of relevance in Sperber and Wilson (1995)). This goes together with the expectation of something to be learnt from observations – the agent's communicative intention is pedagogical. I suggest that infants implicitly assign such higher-order goals to the agent, and thereby become ready to learn the practical content of the conveyed message. As explained above, the object of learning is either a physical goal state, or an instrumental action. Independently of the motivation to learn from observations, infants show a robust drive to understand actions with respect to goals (according to the theory of the teleological stance). They assign physical goals, which they expect to justify actions contextually, as described in Section 2.4. Aside from the further closed-world computations involved in goal validation, conjectured physical goals are stored in relation to the object through which they are attained. Furthermore, infants' expectation of a relevant content of the agent's teaching intention manifests itself in the

[24]Later on, the tendency changes and older children and toddlers appear to be more prone to overimitation – e.g., Lyons et al. (2007).

CEU eTD Collection

form of epistemic trust, or a credulous lowering of the threshold for relevance. For instance, observed novelty may be deemed relevant or 'worth learning' despite failed teleological explanation with respect to physical goals of actions (as described in Section 2.5.2). In fact, explanatory failure with respect to physical goals might be precisely what is needed to make infants learn teleologically opaque behaviors from adults who manifest themselves as reliable pedagogues through communicative and ostensive behavior. In a credulous understanding, fostered by the lastly mentioned manifestations, opacity relative to physical goals may be taken to point to the adult's intention to change the observer's knowledge state. Consequently, an observed novel behavior is encoded as merely potentially useful for concrete purposes, but certainly useful for the abstract motivation to expand one's knowledge. The implicit assignment of a pedagogical intention to the agent, supported by her communicative demonstration, automatically activates the "natural born learner's" maintenance goal to acquire knowledge.

I end with a concrete example of how multi-level teleological reasoning may support knowledge updates in a pedagogical setting also in the case of adult observers. Imagine that I take a photography course, and I missed the previous theoretical class. The next class, the whole group goes on the field during the day, to practice the theory taught the time when I skipped. We fixate some windmills that the teacher has told us to photograph. I observe a colleague, whom I trust as diligent and who, for all I know, did not miss the class before, turning on the flash button. She does this markedly while gazing at me, as if making sure that I notice what she is doing. All the reasoning steps that I am now about to describe go on extremely fast, below awareness. I assign her the goal of photographing the wind airscrew x. Her turning on the flash before shooting is unusual, given that it is daylight. I run offline the backward chaining with CWA's from the goal to photograph x and I obtain my own default means action 'push the SHOOT button unless it is dark'. When I assign my colleague an action plan in the observed context, I obtain the action rule 'In order to photograph x, turn on the flash before shooting'. Current observations conflict with my default action rule, so I seek a justificatory abnormality in the observed context. The abnormality condition does not obtain in the context; because it is daytime, there is no constraint in the physical context that singles this action rule as the action that should be performed for the goal. No exception to the default rule is obvious in the context. My colleague's means sequence is not forced to occur by the data in my minimal, default model of the situation. The novel action is not justified by the physical goal conjecture formed in this model, namely, that she wants to photograph x now, in daylight. However, because I know that she was here the time when I was not so I trust that she might know more theory than I do, and because of the accentuated way in which she gazes at me when doing the new action, I assign her the intention to teach me something. That is, given my trust in her epistemic reliability, and the explicitly ostensive communicative cues, the unusual means action can be justified as my friend's pedagogical intention, despite being opaque from the point of view of efficiency for the physical goal. She acted differently than I would have because she wanted not just to photograph x, but also to share with me relevant knowledge about how to achieve this

CEU eTD Collection

goal. This makes me implicitly assign her the corresponding intention to teach me how to take the picture correctly. Consequently, after I shoot first without the flash, i.e. the default way, I turn on the flash and take another picture. Online planning from my maintenance goal 'learn' prescribed that I make use of my trusted friend's behavior. I thereby learn a new, more complex means, which I perform when I photograph my own wind airscrew y. All in all, the imitative performance of the novel action, is the execution of a plan grounded in my maintenance goal to learn, automatically triggered by my colleague's intention to teach me, which justifies the observed unusual behavior.

CEU eTD Collection

CEU eTD Collection

CHAPTER 3

The Model I: The Reasoning Steps

The theoretical account of observational learning exhibited through selective imitative behavior in the head-touch task set forth by Gergely et al. (2002) proposes that infants' teleological interpretation of the observed actions supports their artifact directed action performance, which evidences observational learning (as detailed in Section 1.2.2). The theory is centered on the top-down role of cognitive processes underlying behavior. The model that I set forth pursues this explanatory strategy. It acknowledges the role of inferential processes in shaping a behavior that demonstrates selective observational learning. The primary aim of the modelling enterprise is to refine the theory by spelling it out in a formal fashion. This Chapter provides a non-formal depiction of the reasoning processes claimed to support infants' behavior. They are presented as constraint logic programming derivations in Chapter 5.

The model elaborates the developmentalist account along two directions. First, it extends the notion of goals involved in teleological reasoning along the lines of the distinction introduced in Section 2.2 between achievement and maintenance goals. Second, it sets forth that the practical reasoning processes involve planning, or reasoning backwards from goals to actions that should be performed. I explain the role of these extensions below.

As anticipated in Section 2.6), the first elaboration is connected with the proposal that the teleology grounding infants' action interpretation, and shaping their imitative behavior, is not merely physical. Infants' teleological processing of instrumental artifact directed actions may rely on non-physical (i.e. mental, psychological) factors when actions remain opaque with respect to physical goals and situational features. The same teleological reasoning mechanism applies to higher order, more abstract goals, e.g., implicitly assigning an agent the intention to change one's knowledge state. Empirical evidence shows that the social communicational environment of the demonstration influences the imitative responses that index 14-month-olds' observational learning Király et al. (2004); Király (2009); Király et al. (2013); Southgate et al. (2009). Such results support the idea that learning from other agents' instrumental behaviors is not exhausted by action evaluation in terms of efficiency for physical goals. Correspondingly, young actors themselves may act upon physical goals because they are motivated by generic abstract objectives. The drive to assimilate practically useful knowledge from reliable conspecifics is an appropriate example of such motivation, which calls for action upon a concrete, physical end state derived from it. I suggest an account of the findings in which infants' imitative behavior (in both guises as emulation and imitation) diagnoses their drive to learn, activated by the

CEU eTD Collection

agent's communicative pedagogical intention. According to this extension, infants' pedagogical and teleological stance contribute to understanding actions and to their own goal-directed action via the same teleological reasoning mechanism; this is the single-mechanism multi-level teleological processing anticipated at the end of Chapter 1.

The second extension of the developmentalist theory brings to the fore the role of infants' capacity for planning. I suggest that on top of reasoning *toward* an interpretation (cf. van Lambalgen and Stenning (2008)) of the observed action context, infants' behavior is supported by planning computations, that is, by reasoning *from* goals towards executable actions. The theoretical whys and wherefores of planning are discussed in Section 2.3; I recap here for convenience the core aspects. Planning is a form of teleological reasoning, whose expected outcome is a temporally ordered sequence of means that attain the desired end-state. Computations grounded in a goal state s produce action rules of the kind *In order to bring about state* s *do action* a, whereby a has the teleological status of means action, and s – of goal. Action rules are the output of planning; they link goals to action execution. Planning is computationally tractable if constrained by closed-world assumptions. Closed-world reasoning provides a computational algorithm that promotes teleological interpretation, and that prescribes the means to be performed by the young reasoners. The elaboration with respect to planning is twofold, because of the twofold use of the algorithm. Planning fulfills, on the one hand, an explanatory function – during teleological interpretation of observed actions (i.e. in the demonstration phase of the experiment), and second, the function of prescribing the actions to be performed in the test phase. Similarity of inferential structure between the two computed functions, i.e. backwards reasoning, facilitates the formalization of teleological action explanation via a planning mechanism.

First, along the lines of the arguments presented in Sections 2.3 and 2.4, I propose that infants' capacity to plan efficiently actions for goals is recruited for an integrated teleological interpretation of the observed action context. Planning computations from the perspective of another agent, originated in the goals assigned to that agent, foster explanation. If an action is prima facie epistemically opaque, as is the case of the novel head action in the task, representing it as a component of the agent's plan to attain specific goals promotes comprehension. An action plan assigned to the agent integrates the opaque novel action in the dynamic context of ongoing actions, and thereby facilitates observers' making sense of the situation.

Furthermore, when allowed to interact with the funny looking light-box in the test phase, the young reasoners plan their actions online according to the retrieved teleological information from the experimental phase. The objective of formalization is to capture in a principled way the influence of the teleological interpretation of observed actions, onto planning infants' own, via the activated epistemic motivations. Online planning spells out the inferential path that links infants' motivation, to a concrete end-state, to the behavior performed. Planning computations terminate with basic actions. According to Chapter 2 such are those transitive actions for which the concrete role of the agentive relatum is specified to body parts like hand, head, torso.

CEU eTD Collection

To recap, infants' performance in the test phase is a manifestation of observational learning. As a result of the above theoretical elaborations, it is viewed as the result of two-staged inferential action processing: *teleological interpretation* of observations in the demonstration phase, and *online planning* of own actions in the test phase. The different forms of learning, manifested as different behaviors across the experimental setups, are a consequence of different contexts of observation, which activate differentially infants' epistemic drive. Representations of observations provide the inputs for processing (a formal characterization of which is available in Section 4.1.3). Along with Gergely et al. (2002); Király (2009), I suggest that the processes unfolding in young reasoners' minds upon presentation with the agent's instrumental actions onto the novel artifact construct a teleological interpretation of observations. An intrinsic epistemic desire to understand observations provides the motivation for processing them, and organizing them along a means – ends structure. Teleological interpretive computations are meant to make the observed situation intelligible by relating actions to goals that justify them contextually. For the current modelling enterprise I adopt an abductive approach to explanation (Shanahan 1993). Reasoning is like inference to the best explanation, where 'the best' is defined by closed-world considerations (as discussed in Section 2.4). Teleological processing activates the motivation to learn. Online planning is action-oriented, and grounded by this drive. It prescribes actions to be performed in order to bring about efficiently physical goals derived from such motivations. These computations are run in the test phase of the experiment, and manifest themselves in overt behavior.

Table 1 shows the reasoning steps and the knowledge databases over which they are run. Teleological action interpretation sets out with goal attribution to actions. The outcomes are conjectures of physical achievement goals. The goal conjectures require validation of the justificatory role. For one, infants generate an offline plan for the goal, based only on their prior knowledge. This is like counterfactually asking "what would I have done, assuming that I wanted to obtain what you seem to want?". Second, because the result of offline planning mismatches observations, the young reasoners verify whether the conjectured goal justifies its cause action in the holistic, dynamic context of performance. This is like factually asking "what did the agent do, assuming that she wants all the things that she seems to". I label this part of the explanatory process plan recognition. It is a planning process run by infants from the perspective of the observed agent. The database comprises the observed behaviors of the agent, along with infants' prior knowledge and the default plan recommended by offline planning. Depending on whether plan recognition provides a satisfactory explanation, different formats of epistemic motivations are triggered. They connect interpretation with action execution. In the test phase, reasoners compute a plan grounded by the drive(s) to learn active at the end of interpretation; they execute this plan.

CEU eTD Collection

Table 1. The reasoning processes involved in task performance.

Reasoning process	Database	Syntactic strategy	Outcome
Goal attribution	current observation	specialization	goal hypotheses
Offline planning	prior knowledge	bwd. chaining from main goal	action rule (own)
Plan recognition	processed current observation	bwd. chaining from assigned goals	action rule (agent's)
Online planning	processed current observation	bwd. chaining from main goal	action (own)

One final word to highlight the non-monotonicity of the goal-centered reasoning that underlies observational learning in the task. Goal inferences and planning are generally defeasible processes, even more so in the context of novelty (cf. Chapter 2). Offline planning constrained by closed-world assumptions produces default action rules, valid in typical circumstances. However, an analysis of the currently observed context may provide positive evidence of abnormality, which overrides the assumptions. If available this cancels the validity of the reasoners' offline plan for the observed context, and allows that teleological explanation of the observed action context be satisfactory despite mismatch with the default plan. Infants' making sense of novel behaviors by extracting their teleological structure and comparing it with their own default plan is inherently tied with non-monotonic processing.

Section 3.1 below justifies the claim that infants engage in action processing by capitalizing on their motives to sustain the computational effort. It details the first elaboration of the developmentalist theory. Up next, Section 3.2 describes the reasoning processing involved in each of the three stages of teleological interpretation. The planning computations ongoing in the test phase are presented in Section 3.3. At each step I emphasize the relevant differences between the computations occurring in the four experimental conditions.

CEU eTD Collection

3.1. Motivational factors for the teleological processing in action interpretation and action planning.

Motto: If active intervention in the world is necessary to infer its causal structure, then there needs to be some sort of motivational system to bring about such intervention. (Gopnik 1998, p.8)

I explained in Section 2.2 that human agents' goals, infants' own as well as the ones assigned to the agent, are hierarchically structured on a continuum from physical achievement goals, to more general physical maintenance goals, to abstract motivations and intentions. Infants want to understand, and this provides a primary motivation for processing. Furthermore they want to learn what is relevant about actions and their effects (and, conversely, disregard what is not), presumably because more practical knowledge grounds the ability to achieve more physical goals, more efficiently, and to understand more. They are motivated by an intrinsic (Gopnik 1998), potentially species-specific (Csibra and Gergely 2009) desire to assimilate relevant knowledge via understanding the surrounding world. The epistemic motivation is twofold: to understand – henceforth, the explanatory epistemic motive, and to learn from reliable conspecifics – henceforth, the assimilationist epistemic motive. The two aspects reinforce each other. Moreover, they are inherently connected to social cognition, since the knowledge providers are peer human agents. Social aspects play a crucial role in shaping observational learning. This is why a more appropriate label for these motives is socio-epistemic. They are represented in the model as maintenance goals, with a conditional formulation; the motives themselves (understanding, and learning) are expressed in the conclusion. Infants' reasoning from, and subsequent acting upon these motivations illustrates rational teleological behavior.

Gergely and Csibra (2003) describe teleological processing as a tripartite process. It is said to aggregate observed actions, relevant aspects of the context of observation, and goal states, into a coherent means – ends structure. In the head-touch task the first two elements provide the raw material over which computations are performed in order to assign goals to the agent's actions. They constitute the observational knowledge database for the computations in goal attribution (see Table 1).

I begin by discussing the contextual aspects that bear on teleological interpretation. More specifically, the formulation 'relevant aspects of the context of observation' in the description of the tripartite interpretive process may rightly raise questions about the kinds of aspects that infants take to be relevant[1]. The answer I suggest is along the lines of the first theoretical extension introduced in the preamble to this Chapter. The initial rationalist theory restricted the sphere of relevance for teleological processing to perceptually available (observable), or physical features of the context (Gergely and Csibra 2003; Gergely et al. 2002; Csibra and Gergely 2007), e.g., the spatial position of the agent's hands, the presence or absence of obstacles. Accordingly,

[1]The related problem, not of what *kinds* of aspects, but rather of how *particular* aspects of the situation are represented the current database for processing (i.e. the problem of coarseness of action representations) is touched upon in Section 5.1.

CEU eTD Collection

the notion of goals derived via teleological inferences referred to physical goals. First, I wish to call attention that relevant aspects for teleological processing of a situation as a whole, may involve more dynamic features, e.g., other physical goals pursued in the context. That is, the aspects of a situation that shape justificatory practices are not necessarily raw observations, they may be the result of inferential goal attribution. Furthermore, in the light of subsequent findings obtained in the experimental paradigm of the head-touch task, e.g., in Király et al. (2004); Király (2009); **?**); Király et al. (2013, to be submitted), it seems that restricting the contextual aspects with direct bearing on teleological processing of actions to overtly observable, physical features does not suffice to provide an account of observational learning. An agent's social communicative attitude has a crucial effect on imitative learning, triggered by its effect on the action interpretation that underpins it. Unless the agent exhibits ostensive cues (Király et al. 2004; Southgate et al. 2009; Király et al. 2013, to be submitted), or unless she conveys cues of knowledgeability or epistemic reliability (Chen and Gergely, in preparation), infants do not generalize the novel means actions by imitative performance in another context (than the one of the demonstration), they only emulate the function of a novel artifact. A related empirical fact is that whether the demonstrator agent is or is not present during the test phase, when infants presumably retrieve the stored interpretive outcomes and plan their responses based on them, also modulates their action upon the artifact Király (2009). It is thus crucial to note a point that Hacker (2009, p.78) makes explicit when discussing agency and action explanation:

> The circumstances of an act can be construed narrowly, so as to exclude such subjective factors as voluntariness, knowledge and intention [...] *or it can be stretched to include these*.

[added emphasis]. In the case of observational learning assessed with the head-touch experimental paradigm, factors such as the ones Hacker calls "subjective" are necessary preconditions for a certain kind of learning, namely, the acquisition from observation of novel means actions for which a default alternative is available in the observer's motor program. The model I set forth shows that learning such new skills is the upshot of action planning upon teleological explanation of observed actions, which includes the agent's pedagogical intention as a "subjective factor" in Hacker's sense. Consequently, it supports the idea that contextual aspects like behavioral manifestations of an agent's intentions and knowledge states, are relevant circumstances for the evaluation of the teleological structure of observed situations. If goals do not justify actions within a model of the physical context, the agent's ostensive behavior triggers observers' epistemic drive; means imitation thus compensates explanatory failure with respect to physical goals. The emergent theory is not excessively mentalistic, or psychologist (in the sense of Davidson's psychologism with respect to action explanation). For one, it is the directly observable, behavioral component of intentions, i.e. their physical goal content (what an action changes in the physical world), that plays the crucial role in reasoning. Second, as explained in Chapter 2, both the mentalistic and the physical representations of goals are handled by the same

CEU eTD Collection

mechanism of justificatory teleological processing. In so doing, I take seriously the proposal that

> even though mentalistic belief-desire explanations of actions are formulated as causal relations between intentional mental states and actions, their inferential structure always involves a teleological (reason-giving) element Csibra and Gergely (1998, p.256).

Goals represented at a higher level of abstraction that involve a mentalistic component in the form of intentions, are processed along the same lines of physical teleological reasoning. If the latter kind of computations fail to provide a satisfactory explanation then the observers capitalize on "subjective factors". I elaborate this idea below.

The interpretive processes are constrained by flexible use of closed-world assumptions. These require one to construct a minimal model[2] of the means – ends structure of the situation, which renders the observed context intelligible. A minimal model is physical, as mentalistic action construal is computationally more costly Csibra and Gergely (1998, p.258). If the physical contextual aspects, be those static (e.g., the position of the table) or dynamic (e.g., other simultaneous physical goals pursued by the agent), do not provide a satisfactory teleological explanation, the model is flexibly expanded (i.e. the reasoning world is opened). The agent's intentions (i.e. what the agent intended to convey to the observer), and not just the end-states that her actions bring about (e.g., lighting up a light-box), count as justificatory reasons for her actions. Learning despite failure to explain requires credulity, in the sense of taking for granted the epistemic authority (i.e. the knowledgeability) and social trustworthiness of the agent. Credulity is justified only if the agent behaves ostensively and communicates (Csibra and Gergely 2009), i.e. if her state of knowledge, and her intention to share it, have explicit behavioral manifestations. The infants' credulous take on the action context results in willingness to learn the bizarre instrumental behavior, despite not being optimal for the physical end-state conjectured as the function of the light-box. This is the core of the proposed account for the different findings between the Communicative and Incidental experimental conditions.

All in all the model accounts for the reasoning processes that support infants' learning of practical information in terms of a 'single-mechanism multi-level' teleological processing. Physical and social-psychological features are integrated in the interpretation of the action context. Furthermore, based on infants' motivational states, they infer physical goals to be attained in the current context and backwards planning from those results in action performance.

Action processing, both in the stage of interpretation and in that of planning, is driven by infants' epistemic motives introduced above: to understand (via explanatory teleological interpretation), and to assimilate practical knowledge. *Understand* and *learn* are expressed in the conclusions of maintenance goals. They are triggered by forward reasoning from observations

[2]I introduce the corresponding semantic notion in Chapter 4.

CEU eTD Collection

that their conditions obtain in the world. Throughout this Chapter the direction of reasoning is represented by arrows under the conditional formulation of the maintenance goals.

Based on the empirical evidence discussed in Section 1.1.3 I suggest that when infants observe the agent's ostensive signals at the outset of the demonstration in both conditions of the Communicative setting, they assign her a communicative pedagogical intention, and they are prepared to learn from her. It is an abstract intention whose content is a placeholder to be filled it as the demonstration progresses.

IF agent conveys ostensive signals THEN agent teaches x

$\longrightarrow$

Upon observation of ostensive signals infants' reactive attitude is to operationalize x, the content of the agent's pedagogical intention, to something concrete. Their genericity bias (Csibra and Gergely 2009) is also set off; that is, infants are prone to generalize over different contexts the episodic teleological information currently encountered. Such bias is the foundation of learning.

Young actors' natural motivation for action processing is to render observations intelligible – the explanatory epistemic drive to *understand*. Although it is active irrespective of the social features of the situation, it is pursued to a larger extent (e.g., by willingness to invest more computational effort) in a Communicative setup. In what follows I focus on its unconditional activation by observation of instrumental actions. It initiates the teleological explanatory process by triggering the application of the teleological stance. Infants' natural tendency to make sense of what they see encourages them to engage in an inferential explanatory process based on the teleological stance. The effect of the explanatory motivation on action interpretation is captured in two reasoning steps, one forward, and one backward, as follows.

IF observe instrumental behavior THEN understand

$\longrightarrow$

Observation of instrumental behaviors activates the maintenance goal *understand*. Action observation makes true its condition, and forward reasoning activates the motivation in the conclusion. It then becomes the reasoners' task to make the conclusion true.

teleological interpretation FOR understand

$\longleftarrow$

As discussed in Chapter 1, the model endorses the empirically justified claim that infants' understanding of instrumental actions is conditional on teleological action interpretation via application of the teleological stance (Gergely and Csibra 2003, 1997; Csibra and Gergely 1998). Teleological action interpretation is a goal-directed cognitive process, whose goal is understanding (see also Gopnik (1998)). Constructing a means – ends interpretation of the action context is meant to fulfill the drive to understand. I suggest that the teleological stance be seen as a subgoal of understanding (cf. discussion of goal representations in Section 2.2); it specifies the kind of inferences that are called for once the explanatory drive is activated by observation.

CEU eTD Collection

The epistemic motivation is reduced in one step of backward reasoning to application of the teleological stance to the observed context. It thus sets out teleological interpretation of the situation. The stance, in line with a Gricean-like quality maxim applied to the domain of actions, elicits the expectation that the observed agent's actions are performed in order to attain optimally physical goals in the current circumstances. It sets goal placeholders that must be filled in inferentially. Goal attribution is the first step of teleological interpretation; it is described in Section 3.2.1 below, and formalized in Section 5.3.

The assimilationist aspect of the epistemic motivation – to acquire practical knowledge, or to *learn*, grounds infants' online planning computations. In the test phase they retrieve the specific format of the motivation(s) activated during action interpretation, and plan actions for learning. By so doing, they situate themselves at the receivers' position of the pedagogical relation (Csibra and Gergely 2009). They act upon the pedagogical stance. Action production is the outcome of planning for learning, in the sense that it is driven by an intrinsic propensity to learn. This view of infants' behavior gives a clue as to why imitative behavior is more often that not labeled 'observational learning' in the psychological literature (Gergely and Csibra 2006; Gergely 2003; Heyes 2001; Csibra and Gergely 2006; Tomasello et al. 1993). The specific kind of learning, and correspondingly the particular action performed, are determined by the inferential outcome of teleological interpretation relative to physical goals, and by the social context of the demonstration. This way, the twofold manifestation of the motivation to learn connects action interpretation, to infants' planning of own actions, and eventually to action performance. The exhibited kind of observational learning is selective between the experimental conditions (as described in Section 1.2).

I now introduce the forward reasoning that activates the assimilationist epistemic motivation. The backward reasoning from the the assimilationist motivation toward action is part of online planning; I discuss it in more detail in Section 3.3.

The generic conditional formulation of the assimilationist epistemic motive *learn* regarding practical knowledge is:

IF observe novelty THEN learn

$\longrightarrow$

The learning mode is activated by forward reasoning from observation of novelty. The criteria for novelty are given in the seminal paper of Meltzoff (1988), and the following experiments in this paradigm have implicitly acknowledged them. I suggested in Section 1.2 that in the case of the head-touch task under scrutiny, novelty has a twofold origin. It stems from the artifact kind, as well as from the instrumental action that the agent performs on it. A more fine-grained expression of the conditional maintenance goal in the head-touch task is:

IF (observe novel artifact OR observe novel instrumental action) THEN learn

$\longrightarrow$

CEU eTD Collection

According to Section 2.6, the assimilationist epistemic motivation can take two objects: ends states achieved by an agent, and practical skills (instrumental actions). The difference resides in the teleological kind of the to-be-learnt object, i.e. the element assimilated to the reasoners' knowledge database. This is a goal in (1) and a means in (2). Learning that a particular goal state is associated with an artifact that makes it be the case, and subsequent expectation that a similar artifact produces the same goal, i.e. that the goal is enduring, justifies the expression '*function* learning'. According to labels from developmental and comparative psychology as discussed in the preamble to Chapter 1, (1) stands for emulative learning, while (2) is imitative learning. A feature of (2) specific to the head-touch task is that the skill that infants learn, i.e. the head action, is a suboptimal alternative to an instrumental action that is available in their motor repertoire, i.e. a hand action.

The categorization even in hypothetical form of actions *qua* means, and of states of affairs *qua* goals or functions, is the outcome of teleological interpretation of the observed context. The activation of a concrete assimilationist motivation is thus consequent on, indeed grounded by, the explanatory one. The drive to learn that the light-box can bring about state light-on is activated as soon as the state is conjectured as an achievement goal. Whether (2) is also activated although infants have knowledge of a default means action depends on the outcome of using the teleological stance for understanding. More specifically, only failure of explanation via the teleological stance with respect to physical goals activates imitative learning as represented in (2).

I set forth in Section 2.6 that the drive to learn artifact functions is primary. Indeed, in the head-touch task it is activated across the four experimental conditions as soon as a goal conjecture is available for the agent's artifact directed action.

IF (observe novel artifact OR observe novel goal-directed action) THEN learn artifact function

$$\longrightarrow$$

In order to exhibit the kind of behavior that evidences observational learning of functions, it is not necessary that reasoners be able to establish a satisfactory teleological interpretation of the situation in terms of that goal conjecture (cf. Section 3.2.3). That is, formulation (1) of the assimilationist epistemic drive is independent of the overall success or failure of the explanatory drive in terms of physical goals. It supervenes on the formation of an achievement goal conjecture upon observation of successful instrumental behavior. Learning is specialized to a concrete goal state as soon as a goal conjecture, presumably the artifact's function, is assigned to the transitive behavior. Consequently acquisition of functional knowledge is supported by lesser computational effort than (2), the acquisition of alternative practical skills.

However further research questions are raised as to why learning *that* takes precedence over learning *how*. It is possible that the ease of engaging in emulative learning has to do not with the object to be assimilated per se – learn function vs. learn means, rather with the kind of novelty involved in the function and the means action, respectively. That is, the light-box is completely

CEU eTD Collection

new to the infants, so its function is an element of 'radical novelty". In contrast, the head action is novel relative to what the infants know already, i.e. the default method to act on objects with the hand. Acquiring alternatives is more demanding than learning something completely new.

The current model of the reasoning involved in observational learning in the head-touch task pursues this direction. The fact of the matter remains that the motivation to learn novel means actions is tied to increased computational effort particularly when the observer knows a default way to bring about the end-state, relative to which the observed means is suboptimal. The vast majority of evidence shows that only human infants learn by imitating observed actions whose practical necessity for specific goals is not obvious, i.e. skills that are teleologically opaque (Gergely and Csibra 2006; Gergely and Jacob 2012; Király et al. 2013). The principled[3] motivation to acquire alternatives to default behaviors might require multiple level teleological processing abilities, including mental components of goal representations, that go beyond those of other species, but that are present *in nuce* in one-year-old humans.

The cited empirical evidence for the head-touch task recommends that there are two types of information that determine whether skill learning motivates action performance. One is the outcome of teleological inferences in a minimal physical model. More precisely, the outcome in question is failure of the teleological interpretation with respect to physical achievement goals to render the novel action intelligible in its context. The other one is provided by the social context of demonstration. The two constitute a conjunctive condition of the specific maintenance goal *learn goal-directed actions*. The assimilationist drive in the conclusion is activated only by co-occurrence of both conjuncts. They provide the preconditions of updating the practical knowledge database with means actions, as it manifests itself in Communicative – HandsFree conditions where reactive reasoning from ostensive signals triggers the expectation that the agent is teaching something.

When explanatory uncertainty occurs relative to the behavior of a communicative agent, it activates the drive to retain not only the novelty of what the light-box can be used for, but also the novel instrumental action. The conditional maintenance goal that represents infants' more demanding motivation *to learn practical skills* is:

IF observe novel transitive instrumental action AND teleological explanation of x AND agent teaches x THEN learn x.

$\longrightarrow$

Physical teleological explanation fails when the physical goal of an action does not justify the action that causes it in the given circumstances, i.e. when the action is not singled out as necessary for that goal in those circumstances. The explanatory process then appeals to the agent's pedagogical intention *to teach* x. This notion of intention does not require complex reasoning about mental states, because the focus is on its behavioral content – the observed action. The behavioral, physical, content of the agent's intention *to teach* x is specialized to

[3]In the sense of not-random.

CEU eTD Collection

the observed instrumental behavior not explained satisfactorily by the teleological stance in a minimal physical model. Hence the maintenance goal reads as

IF observe novel transitive instrumental action AND teleological explanation of instrumental action fails AND agent teaches goal-directed action
THEN learn goal-directed action.
$\longrightarrow$

The last conjunct is the conclusion triggered by the agent's initial ostensive cues, whose variable has been specified to the under-explained skill. It expresses infants' credulous stance with respect to the demonstrated action sequence. This credulous interpretation occurs only when the agent provides cues of social reliability, i.e. ostensive pro-social behavior.

To sum up, the model is developed on the empirically grounded assumption that reasoning processes over the observations task trigger activation of both formats of infants' epistemic motives, the explanatory and the assimilationist one. The activation is different across the four experimental setups. Infants' epistemic maintenance goal modulates behavior in the head-touch task through its twofold manifestations. The *understand* component motivates them to apply the teleological stance to interpret observations, in order to be able to make sense of the novel, thus explanatorily opaque, aspects. The *learn* element provides a conditional way out in case teleological assessment with respect to physical achievement goals does not provide a satisfactory explanation, and thus does not fulfill the drive to understand. The 'way out' is not absolute, but relative to the social constraints of the situation. It amounts to learning those actions that produce the conjectured achievement goal, but that are not justified by it. Infants thereby manifest themselves as rational actors motivated to understand novelty, and to acquire practical knowledge in the form of what to do, when, what for, from reliable sources (i.e. other rational actors) in the environment.

3.2. Teleological interpretation of the context.

As argued by proponents of the top-down cognitivist approach in developmental psychology, and backed by additional theoretical arguments discussed in Chapter 2, a satisfactory teleological explanation of actions requires inferential processing of raw observations.

Throughout this Section I present the reasoning processes meant to fulfill the explanatory epistemic drive activated by observing the agent's unusual behavior directed at the novel lightbox. The observed actions are parsed and interpreted with reference to goals, thereby constructing a teleological model that makes sense of the context. Teleological interpretation is an attempt to categorize observations (i.e. actions, activities, and resultant states of affairs) as means or goals. The disjunction in the above sentence is not exclusive, given the hierarchical nature of human teleology (discussed in Chapter 2).

Teleological interpretation has two stages. First, *goal attribution* outputs conjectures regarding the agent's physical achievement goals. Second, verification is conducted regarding the

CEU eTD Collection

justificatory role of the goals with respect to the observed instrumental actions in the context of occurrence. It is like checking whether the goal conjecture provides the best explanation (cf. Section 2.4). Verification is twofold – *offline planning* and *plan recognition*; the two processes are successive.

Infants observe subtly dissimilar sequences of events in the four experimental setups (as described in Section 1.2). Consequently the current knowledge representations over which interpretive processes are run, are also different. Focusing on the Communicative setting only, the differences between the HandsOccupied and HandsFree conditions take effect on how the action sequences are teleologically interpreted in the stage of plan recognition. Goal attribution and the 1st person validation procedure in offline planning are similar across all four experimental conditions, irrespective of the social context (Communicative vs. Incidental) and of the mode of presentation (HandsFree vs. HandsOccupied) of the novel action.

3.2.1. Action-to-goal reasoning: goal attribution. The initial interpretive reasoning process is concerned with assigning physical achievement goals to actions. Goal attribution sets out action processing via the teleological stance – the assumption that actions are *for* something.

A goal is a particular kind of effect of an action, namely, a state of affairs that is most efficiently caused by that particular action. The search space for goals is delineated by the foregoing causal individuation of events. The events that occur in the context are perceived as being causally linked, i.e. forming cause – effect sequences (Saxe and Carey 2006). Based on defeasible inferences over causally linked observations and relevant contextual aspects, particular states of affairs are represented as concrete goals that the agent's actions are meant to achieve. This is the process of *goal attribution*[4]. Because of the defeasibility of reasoning, the teleological status of goal is conjectured, and verified by further computations.

Csibra (2008) describes goal attribution in infancy as a twofold enterprise: identifying the agent of an observed behavior, and finding by means of action-to-goal inferences the goal that justifies the observation in its context. I suggest that in the experimental setup of the head-touch task the first process is rather trivial[5]: the agent is the observed experimenter.

I now focus on the action-to-goal inferences. The verification of the justificatory role of the main conjectured goal is described in Sections 3.2.2 and 3.2.3 below. Goal derivation starts with selecting the behaviors liable to teleological processing. The teleological stance activates goal expectations for instrumental behaviors. I propose that in the head-touch task the expectations are fulfilled by two strategies. On the one hand, the inferences refer to searching through the perceived effects of these behaviors, and selecting one salient effect that justifies the behavior in its context; this effect is represented as the achievement goal of the behavior. On the other hand, an instrumental behavior might activate in a quasi-automatic, straightforward manner a maintenance goal; more processing is then involved in deriving a subgoal of it, which counts as

[4]The terms 'attribution', 'derivation', 'assignment' (of goals) co-refer, and are used interchangeably.

[5]Though, see below in Section 3.2.2, and also in Chapter 5 discussion contrasting global agents (e.g., human beings) from local agentive relata (e.g., hands, head).

CEU eTD Collection

the sought achievement goal. The latter strategy is mostly suitable for instrumental behaviors with which the observer is highly familiar, and which are easily categorized as subgoals of a maintenance goal in a particular context. For example, were one to observe another agent opening the door of the fridge, while at the same time hearing his stomach growling, and knowing that the agent knows that the only (still) edible product in the fridge is a light yoghurt, one could infer that the agent's maintenance goal is to eat, and a plausible achievement goal to the effect "eat" in the context of observation is "eat light yoghurt".

I suggest that infants detect two goal-directed behaviors, and thus open two placeholders for achievement goals. The teleological stance is applied to the observation of the agent's shivering and frowning, and to her leaning towards and touching the light-box. The latter is the main train of teleological thought in the task because goal attribution is initiated by the target behavior. Teleological interpretation is aimed primarily to fulfill the two placeholders for physical achievement goals. This view of the task as calling for a twofold application of the teleological stance is an elaboration on the part of the current modelling endeavor, compared with the initial theoretical account in Gergely et al. (2002). According to the original theory, only the target novel behavior triggers teleological processing.

The two placeholders are fulfilled via the two inferential strategies mentioned above. In both cases goal attribution supervenes on a causal evaluation of events, i.e. on the the perceptual identification of the cause – effect relations that obtain in the observed sequence. Contact causality is the basic principle allowing for causal perception of physical events (empirical evidence is presented in Section 1.1.2). It allows the representation of the light-on state of the light-box as the effect of the light-box being contacted by the agent. The capacity for inertial causal reasoning promotes the perception of contact as the terminating event of the light-box directed activity, or the inference that the scarf should be held for a while onto the shoulders in order to change the distressing state.

The first observation that sets off the young reasoners' goal search is the agent's behavioral manifestation of distress for being cold. It triggers the expectation that she engages in changing the state of being cold. She is assigned the maintenance goal to change the distressing state.

In terms of action production, discomfort naturally directs an agent to change the state of the world to the effect that there is no longer a disturbing element. In terms of action interpretation, observers are also likely to assume that the agent's behavioral signs of discontent with the current situation trigger her attempt to change it for the better. This is like attributing the agent a 'reactive production rule' (Sadri 2011), which says that 'If she shows discontent for being cold she will do something to change it'. Such a rule relies on the expectation that human agents strive to maintain a stable harmonious relation with the world. Correspondingly, when this relation does not obtain, then they are expected to engage actively in pursuing it. In other words, an agent who is observed to be dissatisfied with her current state, is expected to to do *something* to change the unpleasant state. That is, infants expect her to act upon the assigned maintenance goal. Inferential processing specialize this *something* to a concrete behavior among the observed

CEU eTD Collection

ones, i.e. a behavior to the effect not cold. Such a behavior is the sought achievement goal, a subgoal, or an instance of the maintenance goal.

do something FOR change distressing state

$\longleftarrow$

Locating the scarf around her shoulders is the expected achievement goal derived from the desired change of distressing state. The derivation is grounded by the commonsense knowledge that a warm cloth placed for a certain period of time on the spatial location 'on an agent' causes her to not be cold anymore. The manipulated scarf is categorized as a token of a warm cloth, so the observed scarf manipulation fits this causal principle. Consequently, the agent is assigned the concrete achievement goal of placing the scarf around her in order to achieve the temporally extended goal of changing the state that caused her to shiver. Moreover, infants' background causal knowledge represent the physical relation between the thick material and an agent's sensation of warmth as an instance of continuous causation. The warming force of the scarf needs to be exerted over a temporal interval in order to produce the desired change of state via temperature transfer; the rule reads as 'continuous location of scarf on agent brings about warmth'. Furthermore closed-world reasoning licenses the inference that if the scarf occupies that place it will change the distressing state: unless (or until) some other relevant event takes place, the location around the agent's shoulders will continue to exert its warmth affordance, and after a certain non-negligible duration δ the distressing state will change. The closed-world assumption with respect to unbeknownst events permits this kind of inertial reasoning. All in all, in order to change the unpleasant state of cold, the agent should place the nearby scarf around her shoulders, and maintain the spatial location for time δ.

place scarf on for duration δ FOR change distressing state

$\longleftarrow$

The causal chain is not problematic to interpret teleologically because of familiarity. The behavior that triggers the teleological stance is recognizable; also the subsequent related sequence of actions maps onto prior practical knowledge. At 14 months, reasoners have some experience with overt, sometimes even overly emphasized, manifestations of adults being cold. Moreover, such behaviors are followed by adults' habitual courses of action meant to attain a desired pleasant change of state, e.g., cover up, get dressed. Therefore, constructing a teleological model via assigning a physiological maintenance goal to change an unpleasant temperature state, and interpretation of the scarf activity as an instance of a means behavior of the kind 'something to the effect change-of-shivering-state', fulfills the explanatory role of action interpretation. Justification comes in a pack with goal attribution. The scarf activity of the agent is backward-looking justified because of the behavior that precedes it and triggers it (shivering), as well as forward-looking justified by its causal efficacy in bringing about the changed state of the world (placing scarf on).

CEU eTD Collection

The target behavior in the head-touch task is an element of novelty. Novel behaviors are, at least prima facie upon observation, epistemically opaque. They call for interpretive schema to render them intelligible. The explanatory epistemic motive described in Section 3.1 triggers application of the teleological stance to the observed light-box directed behavior. Observation of the agent leaning down towards the light-box and touching it in a bizarre way, with the head, promotes infants' explanatory teleological inferences.

Goal derivation for the target novel action is done in a more straightforward manner because it does not involve the intermediary step of assigning a maintenance goal. It makes use of the first among the two strategies mentioned at the outset of the Section. Inertial perception of leaning down towards the light-box and contact causality permit the perception of the light-on state as the effect of the agent's artifact directed behavior. Because the light-on state is the only salient effect of the instrumental behavior, it is quite an easy inferential job to form the conjecture that the light-on state fulfills the teleological status of achievement goal. Multifinality does not hinder goal attribution in this particular case. Once the conjecture is formed, the categorization of 'light-on' as a goal state updates the database of current observations.

A behavior occurs, contacting the light-box, that produces an effect, light-on. The causal link between contact and the light-on state is perceived via the physical principle of contact causality (Saxe and Carey 2006). The principle is abstract, it does not specify the patient and the agent role of contact, i.e. the local agentive relatum. Consequently events are perceived as causally linked without the need to encode the particular entities that fulfill these roles. However, feature detection warrants perception that the local agentive relatum of the light-box directed behavior is the agent's head. This constitutes the novelty of the observed behavior. What infants see is the agent's head touching the light-box and the state light-on ensuing right after. The causal representation of what they see is something like 'contact causes light-on' (expressed above in conditional form).

Henceforth, the achievement goal light-on is referred to as the main or foreground goal in the task; the reasoning processes that it triggers constitute the main teleological train of thought. Correspondingly, the achievement goal to place the scarf on the agent's shoulders and to maintain the location state for duration δ is the background thread of the teleological interpretation.

A few comments are in order regarding the generalization of the suggested goal assignment procedure to other similar tasks. Goal derivation upon observation of failed attempts (Meltzoff 1995; Király 2009), accidental actions (Carpenter et al. 1998), or 'NoEffect' conditions (Király et al. 2013) involves more complex action-to-goal reasoning, i.e. processing of counterfactual development. Attributing a goal to an incomplete, unfulfilled, or failed action is based on counterfactual computations about how the action should have developed in order to result in an end-state that, nevertheless, does not obtain in the actual state of the world. The goal is thus the consequent state which is the most efficient causal effect that an observed action *should have* resulted in, had there been nothing abnormal to prevent it. The state which fills the goal

CEU eTD Collection

placeholder in the teleological stance is not itself part of the observational database; it is an inferential counterfactual construct. Therefore the procedure described for goal attribution in the head-touch task is applicable only to correspondingly simple cases of successful intentional actions.

Coming back to the task under focus, let me briefly locate the current inferential position of Little Johnny, the young reasoner. Upon observation of the bizarre head-banging lady who is shivering, he assigns her two achievement goals via different procedures. The first goal is to place and maintain for a while the scarf on her shoulders, in order to change the unpleasant state that caused her shivering. The other achievement goal is to turn the light-box on. The two goals can be rightly conceptualized as simultaneous although the attainment of the second is observed after the former already obtains. This is so because the achievement goal of having the scarf around her, given commonsense prior knowledge, is derived with the specification 'maintain for a while'. In order to change the distressing state, the scarf should be held onto the shoulders for some time δ. To turn on the light-box is a goal that must be achieved by the agent on the temporal interval that delineates duration δ, while she maintains her background goal to not be cold. Under this teleo-functional specification, the state light-on updates the database of raw observations.

Upon availability of the goal conjecture light-on, infants are in the position to specialize the default epistemic motivation to *learn novel artifact function* to the light-on state. Consequently, the main goal to turn the light-box on is stored as the goal that the novel artifact can be used for, and infants are driven to learn it. Because goal derivation computations are similar in the the four experimental setups, this version of the assimilationist epistemic drive is also activated across the board.

Once Little Johnny isolated a goal conjecture for the novel target behavior and is set on learning it, he pursues further the teleological stance by verifying the justificatory role of the goal with regard to its cause behavior. An action is normatively connected to its goal, as the action that should be done for that particular end state in the given context, if recommended by evaluation of its contextual efficiency. A goal is expected to justify the action in the given context, and provide a forward-looking reason for it (as suggested in Section 2.4). The ensuing computations verify if the achievement goal conjecture justifies the novel action, by being normatively linked with it as the optimal (most efficient) means action in that context. To this end, infants begin by computing offline a plan for its achievement.

CEU eTD Collection

3.2.2. Goal-to-action reasoning: offline planning. Offline planning initiates goal validation. This reasoning stage can be understood as an implementation of the principle that in a regular world, similar problems have similar solutions (Leake 1995). It is as if infants engaged in answering the query 'If the observed agent did ϕ for goal g, what would I have done?'. They seek a motor program available in their own motor repertoire for the goal attributed to the novel action of the agent. They attempt to reduce the main physical goal to an executable action that

they themselves would do typically in order to achieve the goal. 'Typically' refers to an abstract default context, where nothing abnormal interferes in choosing the thing to do. They do so offline. The database for goal reduction is constituted by a selection of the young reasoners' practical knowledge, i.e. default behaviors. Currently observed behaviors are disregarded.

Let us proceed with some remarks about the structure of the inferential path. It is insightful to view Little Johnny's offline planning as computations meant to answer the question 'Had I wanted to turn the light on, what would I have done?'. The verbal moods used in the question make one think about counterfactual processing of alternative courses of events. This is precisely what I intend to suggest, i.e. to view offline planning as computations about an alternative[6] world where the infant himself had the purpose to bring about the light-on state of the light-box. The counterfactual reasoning procedure is applied to a mentally simulated goal, previously assigned to the observed agent. Counterfactual simulationist planning serves as a basis for comparison between the reasoners' potential course of action assuming that nothing abnormal is the case, and the current observations. In Section 1.1.2 I presented empirical evidence for early availability in development of non-explicit counterfactual thought for use in various cognitive tasks, e.g., pretense play (Leslie 1994; Onishi et al. 2007). The evidence supports the inclusion of a reasoning step whose structure is similar with that of counterfactual processing in teleological interpretation.

I argued in Section 2.5 that planning is computationally tractable if framed by closed-world assumptions. Infants take over, as it were, the goal inferred for the bizarre action performed by the agent, and infer from it an action sequence that presumably achieves it in normal circumstances. The ultimate action must be basic, or executable. That is, its agentive relatum must be specified, so that a particular motor program can be activated for execution. Although computations are performed offline, this requirement specific to the functionality of planning is applicable. Infants plan their actions under the assumption that nothing abnormal is the case, hence the most habitual, prepotent, or default instrumental action is, to the best of their knowledge, optimal. They produce a sequence of ordered actions that they themselves would have done in order to turn the light-box on. They put themselves in the agent's teleological shoes (if I may be permitted the expression) only with respect to the main goal. Other aspects of the situation currently observed, for example, the agent's subsidiary simultaneous goal not to be cold anymore, are disregarded in counterfactual planning. The end state 'light-on' is the only observed element from the database for offline planning, in its teleologically processed form of a goal conjecture. It serves as a junction point between the observational database, and the prior knowledge that constitutes the database for offline planning. The goal grounds offline planning computations.

Under the assumption that nothing abnormal is the case, the database for offline planning contains actions that infants would have normally done, had they been set on the goal to light

[6]Presently alternative, at least, since at a subsequent moment in time the alternative might become, indeed does become, actual. See Section **??**.

CEU eTD Collection

up the light-box. Derivations are run over a restricted set of their practical knowledge about actions, the set deemed relevant for the goal 'light-on via artifact usage', e.g., contact light-box with the hand. The prepotent behavior is linked with the goal by prior experience[7]. It takes precedence over other causally appropriate motor responses in virtue of experiential primacy. The database also includes the causal principles that allow to connect in backwards processing resultant states to cause actions, e.g., light-on is caused by contact with the light-box.

The more firmly entrenched the association between the resultant state and its corresponding instrumental action, the more backwards planning from the goal is a direct, automatic selection of the means action. Planning requires little inferential processing. Activation of a default motor program for a certain kind of goal does rely on stored associations between actions and their causal effects, to the extent that it is precisely such associations that frame the database over which planning computations are run. But then an opponent from the direction motor resonance account of the empirical data (see Section 1.2.2) may question the utility of spelling out in a top-down manner the reasoning that leads to the choice of a default action for the effect 'light-on', when mere associationist activation of motor programs can do the work. I would reply that my suggested account is less parsimonious for a good reason, i.e. it paves the way towards understanding subsequent teleological explanatory processes in case the observed head touch does not overlap with the default. The prescription of the default action for the goal 'light-on' is based not only on a framed domain of planning, but also on making explicit during reasoning the assumption that nothing abnormal is the case in the context of offline planning. Should the counterfactual default mismatch actual observations, the assumption provides a way to continue the explanatory endeavor without canceling altogether the prepotent action rule. The assumption may be overridden in the observed action context, and this can be easily verified by searching a justificatory abnormality. Therefore, both here and in Section 5.3, I will elaborate the backward reasoning that starts from the foreground goal and prescribes an action that the infants would have performed.

$$\text{behavior } x \text{ FOR light-on}$$
$$\longleftarrow$$

The young reasoners seek a transitive instrumental action x to the effect light-on. Causal knowledge guides the search. The principle of contact causality prescribes contact as a subgoal of the desired effect. Further on, since planning is after executable actions, and executable means represented in a fine-grained manner, the derivation continues. The abstract notion of contact is specialized. Now, one way to specify instrumental actions expressed via transitive verbs like 'contact', is by mentioning the agent and the patient: what contacts what. Given the context where it has been derived, the goal 'light-on' makes inherent reference to the novel artifact at which the unusual action is directed, the light-box. The patient of contact is thus the light-box. Let us proceed now to specification of the element playing the agential role of contact,

[7]It is irrelevant for the current discussion whether it is about own experience with action production, or experience gathered from observations of others' performance.

CEU eTD Collection

i.e. the local agentive relatum (notion introduced in Chapter 2). In a global understanding of agency, the agent of contact is the very same agent who is engaged in offline planning, i.e. the reasoner itself. However, global agency is not specific enough, because an agent can contact an artifact using a variety of local agentive relata, e.g., whichever body part, another object (a tool). Assuming that no abnormality interferes contact can be transformed in a more fine-grained motor program, i.e. contact performed with the default local agentive relatum. As argued in Section 1.2, human agents' default way to act upon artifacts is with their hands. Young reasoners thus use a generic principle like 'act on artifacts with hand unless there is positive evidence of abnormality' in their offline planning computations to specialize the initial subgoal recommended by the principle of contact causality, to a default hand contact with the light-box.

The bottom line is that hand contact is the upshot of offline planning with closed-world assumptions over a restricted database of infants' practical knowledge. It is the most efficient with respect to the physical goal light-on, in a perfectly normal model of the world. The closed-world assumption regarding abnormalities facilitates to establish a connection between the goal and the default action. The fact that the goal light-on is achieved in normal circumstances via a hand contact with the light-box does not entail that the same goal must be similarly achieved in all models. Correspondingly, the outcome action rule is represented as a non-monotonic conditional of the form 'light-on if hand contact and nothing abnormal is the case'.

(do hand contact AND no positive evidence of abnormality) FOR light-on.

$\longleftarrow$

The conditional formulation shows that the action rule allows exceptions (see also discussion of non-monotonicity in Section 2.1). More specifically, justified exceptions to the rule occur when positive evidence of teleological abnormality is encountered. The default action rule is maintained in working memory at least during the subsequent interpretive steps as a counterfactual as 'the action that I would have normally done'. Moreover, the goal conjecture for the artifact directed action triggers the assimilationist drive to learn that the light-box causes the state light-on (as described in Section 3.1). Infants thus store the goal state.

As a result of offline planning, contacting the light-box with the hand is activated as optimally efficient with respect to the achievement goal 'light-on' in normal circumstances. It is *the* means action that should be performed in order to bring about the state 'light-on'. It mismatches the current observations: the agent performs head contact with the light-box. From the perspective of a rational actor (the infant) in a normal context, the agent could have done better, i.e. she could have achieved the desired state by contacting the light-box with the hand instead of the head.

The conjectured achievement goal does not justify the novel action as a means for its attainment, at least not from the standpoint of the observers' ideal, abstract context of offline

CEU eTD Collection

planning. If under normal circumstances a rational agent *would have done* better to attain light-on, the reasoners aim to understand why the observed agent did not actually *do* better. To this end, they verify whether a justificatory abnormality obtains in the context where the agent is observed acting.

3.2.3. Goal-to-action reasoning: plan recognition. According to Section 2.4, if a conjectured achievement goal does not justify the observed action because of mismatch with the outcome of offline planning, there are two potential solutions:

(1) goal revision;
(2) contextual re-evaluation aimed to find an abnormality in the context.

In the current task (to say the least), solution (2) is more convenient. I elaborate the positive argument for this suggestion in what follows. To begin with, the end state is a highly salient effect of the agent's transitive action, so the degree of confidence it its teleological status of a goal is high. Furthermore, the novel action does bring about the end state light-on (this causal perception is precisely what licensed the teleological categorization of 'light-on' as an achievement goal in the first place). It is not blatantly (i.e. causally) inappropriate for the goal. Rather, upon the offline planning validation procedure, it is found to be less optimal a cause of the end-state than an alternative, default means action. There are two remarks concerning this, one specific to the action context in the head-touch task, the other more general. For one, what makes the observed behavior 'less optimal' than the default, is a subtle difference among two agential specifications of the same action – head contact is observed, when hand contact is normally *the* way to act. The new agentive relatum of the demonstrated instrumental action is the fact that must be explained. More generally, the offline plan is computed for a context where nothing abnormal is the case. This may be an idealization relative to real-world, dynamic action contexts, and a rational offline planner is aware of this. Awareness manifests itself in the explicit proviso that accompanies the action rule, i.e. the conjunct 'no positive evidence of abnormality'[8]. Hence the rational explainer makes an attempt to safeguard the action rule by checking whether abnormality conditions obtain in the context, which justify the agent's overriding the closed-world assumption. All in all, the less radical solution (2) is preferred for arriving at a satisfactory interpretation.

Planning computations in the form of *plan recognition* are used to implement solution (2), as a second attempt to validate the conjectured goal by assessing the efficiency of the novel action in producing it in the observed context. Plan recognition is meant to accommodate the inconsistency between observation and infants' default action rule. It is meant to find among contextual features reasons for the novel action, which justify the difference from the default means action picked in offline planning. The key to inconsistency resolution is to emphasize that

[8]The conjunct used in the action rule is functionally and computationally identical with the 'unless proviso' discussed in Section 2.5.

CEU eTD Collection

the observers' context of offline planning is a closed-world model, i.e. a model of habitual circumstances, where nothing abnormal interferes in goal reduction. The closed-world assumption for abnormalities comes in a pack with a disclaimer – a definition of what constitutes abnormality and thus under what conditions the assumption is justifiably overridden. I suggest that the young reasoners proceed to a global, integrated teleological evaluation of the observed situation in order to find a justificatory abnormality in the observed context. As explained in Section 2.5, abnormality conditions provide a principled way to find exceptions to action rules. A contextual abnormality thereby justifies the agent's unusual approach to her main achievement goal. The teleological evaluation in search for abnormality is computationally implemented in backwards planning; it constitutes a plan assigned to the observed agent in pursuit of the two achievement goals. In offline planning only the action's inferred goal was expected to provide justificatory reasons for actions; explanation was focused on forward-looking reasons for the novel action. In the case of plan recognition, other contextual features than the action's goal may furnish backward-looking reasons to justify the action, for instance, an abnormality occurring prior to the action.

The explanatory procedure upon availability of goal conjectures, initiated by counterfactual planning, focuses again during plan recognition on the factual observations. The database is composed of the observed facts, the end states categorized as goals upon goal attribution, and the inferential outcome of offline planning. The two achievement goals initiate backwards chaining.

Teleological evaluation is global in the sense that it seeks ways in which the agent might satisfy both attributed goals simultaneously. Reasoners apply their background causal knowledge (i.e. contact causality and inertia) to the current observations in order to reduce both goals to executable actions, and thereby assign action plans to the agent. Plan recognition provides a satisfactory explanation if it can organize observations unambiguously in a global means – ends structure of the situation. Computations take into account potential interactions among the two causal chains corresponding to the two achievement goals. The background goal 'place scarf on and maintain for duration δ' is viewed as a dynamic contextual feature (according with the terminological suggestion from Section 2.1, relative to context-relativity of practical reasoning). The backward chaining processes are expected to yield a unique, unequivocal plan to satisfy the main goal, given the need to satisfy the background one.

These considerations lead me to introduce the conditions for teleological abnormality. In the case of teleological interpretation, something abnormal is the case with respect to a goal when the agent has an additional goal to be achieved over the same temporal interval in a different location. That is, attempting to bring about simultaneously two end states in two different locations constitutes an exception to the action rule which relates one of the goals with its default means action. The condition is made more specific by reference not to global agents, but to local agentive relata. There is something abnormal with respect to the involvement of an agentive relatum in bringing about one goal when the very same relatum is necessary at the same time for another goal in a different location. This abnormality condition is derivable from

CEU eTD Collection

the basic principle of contact causality that guides infants' causal understanding, and from from object individuation based on spatial location features. Spatial contact between two objects is the main ingredient of causal ascription at early ages (Saxe and Carey 2006; Leslie 1987). Since an object, e.g., an agent's body part, cannot be in two different spatial locations simultaneously (cf. the use of spatiotemporal continuity for object individuation – (Xu 1999)), the same object cannot exert simultaneous causal influence by contact in different locations. But if a goal is a particular kind of effect, then an object cannot be the local agent of two different goal states, if the difference involves different spatial location features of the end states. All in all, one object, an agentive relatum, cannot be located in two different points at the same time, and, consequently, it cannot produce effects, i.e. exercise its agency, via contact at these two locations. This means that a non-default behavior is contextually justified when called for by the fact that the default agentive relatum is necessary to bring about a different goal.

Abnormality condition:
IF agentive relatum a is necessary FOR goal s at time t in location l_1
THEN there is something abnormal with respect to a for goal r at t in l_2.

The abnormality condition assumes that $l_1 \neq l_2$. The necessary connection between a and s derives from the normativity of means – goals relations. If planning computations output that in order for state s to obtain an action by means of a should be performed, a is necessary for s. Contextual abnormalities thus can be encountered via backwards planning from goals.

As anticipated at the outset of Section 3.2, the differences in the two modes of presentation of the novel action performed by the agent while her hands are either free – HandsFree, or occupied – HandsOccupied, in the demonstration phase take effect on infants' action processing. Different action rules are computed for the background goal 'scarf on' in the two conditions. Because of this, the abnormality condition obtains in the HandsOccupied, and does not obtain in the HandsFree condition. The explanatory processes in the HandsFree, and HandsOccupied conditions are similar for both social modes of presentation in the Communicative and in the Incidental settings. Their consequences on the activation of the epistemic assimilationist goals, however, differ.

The first (in the order of observations) achievement goal assigned to the agent is to place the scarf around her shoulders, and hold it there for a while until the unpleasant state of being cold changes for the better. The young reasoners organize the agent's observed behaviors into a plan that plausibly helps her to attain the background goal in the given context. Causal backward chaining from the state 'scarf on' reduces it to the observed scarf manipulation. That is, in order to locate the scarf on her shoulders, the agent should engage in a hand activity with the scarf: reach for, grasp the scarf, then wrap it around her shoulders. The local agentive relata of the behaviors involving the scarf are the hands; in fact the hands are intrinsic elements of the scarf manipulation. Once the scarf reaches the desired spatial location, it must be maintained for duration δ (as argued in Section 3.2.1).

CEU eTD Collection

There are two possibilities to ensure that 'scarf on' obtains for duration δ: either a continuous force is exerted throughout δ, or an instantaneous event occurs that makes it obtain at t_1 of δ and inertia warrants its holding still across the desired interval. They are illustrated by the action sequences in the HandsOccupied and HandsFree conditions, respectively. In the HandsOccupied condition, no other behavior involving the scarf is observed besides its manipulation with the hands. Inertial reasoning leads infants to infer that the hand activity with the scarf must hold as long as the location state of the scarf must be 'on agent', i.e. for a duration δ. The prescription of the spatial location maintenance is fulfilled by the hands continuous exertion of force. The default agentive relatum is *occupied* maintaining circumstances expected to result in a desired state of affairs. The resultant action rule, which links the maintenance goal to observed behaviors, is 'in order to change state of cold, place scarf with hands around shoulders, and keep holding scarf with hands around shoulders', or

(place scarf on with hands AND keep holding scarf on with hands for duration δ) FOR change state of being cold.

$\longleftarrow$

In the HandsFree condition, on the other hand, infants observe a tying action at the end of the scarf manipulation. The fastening of the scarf by making a knot releases, as it were, the agent from manipulating the scarf in order to maintain its location, in order to change her state of being cold. An instantaneous action – knotting, institutes the effect of continuous spatial position of the scarf without the need for further exertion of the hands' force. The default agentive relatum is *freed* from involvement in bringing about the background goal. The reasoners' expectation of permanence of the spatial location is fulfilled by the action of knotting the scarf. The resulting action rule is 'in order to change state of being cold, place scarf with hands around shoulders, and knot scarf', or

(place scarf on with hands AND knot scarf) FOR change state of being cold.

$\longleftarrow$

The action rule assigned to the agent as her plan for the background goal becomes a situational constraint to be taken into account for the reduction of the main goal, to turn the light on.

In both HandsOccupied and HandsFree conditions backward chaining from the goal 'light-on' closely resembles the computational structure of offline planning. That is, via knowledge of contact causality applied to the context, the goal state is reduced to contact with the novel light-box. Given the observational database, the corresponding fine-grained representation as a basic action is head contact. The plan assigned to the agent in the given circumstances, HandsOccupied and HandsFree as well, is to light-up the lamp by contacting it with the head –

do head contact FOR light-on.

$\longleftarrow$

CEU eTD Collection

Consequently, given the upshot of offline planning, contact can be reduced to a fine-grained executable action in two ways: as contact with the hand, or as contact with the head. To the young reasoners' minds, it is as if the agent had a choice between the two, which she actualizes in the unique form of an unusual head contact. It is precisely this actualization that needs to be justified. Justification is furnished by an abnormality with respect to the default hand contact, i.e. by an exception to the default action rule. The exception constitutes a sufficient reason for the agent not to perform the hand action because it shows that the closed-world assumption is overridden.

The main goal state 'light-on' is achieved at an instant pertaining to the interval delineated by duration δ when the background achievement goal 'scarf_on' must hold. The two goals are simultaneous. Therefore according to the teleological definition of abnormality, in the HandsOccupied condition the hands' involvement in bringing about the agent's side goal, as prescribed by planning derivations, provides an exception to the reasoner's rule. It cancels the reduction of contact to the default hand contact, so the observed head action remains the only alternative – it should be performed. The agent justifiably resorts to it, since she could not have done otherwise. Therefore a new action rule such as 'for light-on do head contact' is acceptable as the plan assigned to a rational agent in her given exceptional conditions where she pursues the side goal by holding the scarf with her hands. The novel instrumental head action is no longer teleologically opaque; it is justified with respect to the physical goal 'light-on' in the abnormal observed context. Plan recognition fulfills the explanatory function. Little Johnny's explanatory motivation for processing is fulfilled. Only the default assimilationist epistemic drive concerning artifact functions is activated in both the Communicative and the Incidental social contexts of HandsOccupied demonstration. Little Johnny learns the function of the light-box. The motive for action is retrieved in the test phase of the experiment, upon encountering again the novel artifact.

The situation is relevantly different in the HandsFree condition. Here, none of hand contact or head contact is singled out as *the* way to reduce contact to a basic action by specification of the agentive relatum. Fastening the scarf by making a knot releases the agent's hands from involvement in bringing about the side goal. Her hands are free to act when the foreground goal is achieved, so teleological abnormality does not obtain. The main goal 'light-on' remains overdetermined with respect to means actions. The physical situational constraints do not provide an exception to the default action rule activated in offline planning. The agent could have performed the default hand contact. Although she could have done better (more efficiently), she achieves the end state 'light-on' in a suboptimal manner, with her head. Teleological evaluation in terms of physical causality provides no sufficient reason for the observed head action. It is not forced to occur by the goal 'light-on', and the physical constraints of the situation, including the background goal and the plan devised for its attainment. Teleological explanation of the novel action with respect to its physical goal conjecture, given the physical features of the context, fails. This is the outcome of interpretation in both Communicative and Incidental

CEU eTD Collection

social contexts, when the novel instrumental action is performed by the agent while her hands are not required to fulfill a different purpose.

However, Little Johnny is not forced to drop the hypothesis altogether. He may appeal to more abstract maintenance goals of the agent of which the physical end-state 'light-on' is a subgoal, and attempt to understand the action context in the light of those. If in a Communicative setting, he may interpret the odd head contact action as the behavioral content of the adult's pedagogical intention to teach a practical skill assigned at the very beginning of the action sequence. Formulation (2) of the epistemic motive acknowledges the possibility that his explanatory drive *to understand* remains unfulfilled after physical teleological explanation, and offers a potential solution. I restate it here:

(2) IF observe novel transitive instrumental action AND teleological explanation of instrumental action fails AND agent teaches instrumental action
THEN learn goal-directed action.

$\longrightarrow$

Conditional on the agent conveying communicative signals during her performance, an under-explained instrumental action is interpreted as the content of the agent's intention *to teach* x. But as explained in Section 3.1, this intention is assigned only upon ostensive communicative manifestations, that is, in the Communicative setup. When 'light-on' does not justify the observed head contact in the context, x is specialized to it. All the three conditions of the learning motive hold. Forward reasoning triggers the conclusion. In the Communicative–HandsFree condition, Little Johnny's epistemic drive to learn the novel means action is activated.

All in all, at the end of teleological action interpretation in the Communicative – HandsFree both assimilationist motivations are active. The result is that both the agent's under-explained instrumental head action, and the resultant state of the light-box are stored. The motivations are retrieved in the test phase by perception anew of the novel light-box, and guide online planning.

In the Incidental – HandsFree condition, no teaching intention is assigned to the agent. Consequently the third conjunct in the above formulation does not obtain, and neither does the conclusion. Little Johnny is not motivated to learn the novel action despite the fact that he does not understand it relative to the physical goal hypothesis. The odd head contact is merely disregarded. All in all, interpretation of the Incidental–HandsFree setup leads to similar results like in Communicative–HandsOccupied, that is, learning novel functions is the only epistemic motivation carried on beyond the context of observation.

CEU eTD Collection

3.3. Online planning.

These are the reasoning processes that go on in the test phase of the experiment, when infants are given the chance to interact with the novel artifact. Their artifact directed behavior is a manifestation of the action rule resulting from online planning. Because it is significantly

different from unconstrained exploratory behavior (as described in Section 1.2), this behavior qualifies as evidence that learning from the demonstration has taken place. It diagnoses learning.

The model developed here sets forth the idea that in the test phase infants act upon the maintenance goal *learn*, according to the formulation that has been activated during the interpretive stage. This idea is along the lines of Gopnik (1998), that

changing theories, generally involves active intervention in the world (p.106),

where 'changing theories' stands for learning. Teleological interpretation shapes infants' online action planning by specifying the assimilationist motivation to learning novel functions, or to learning new practical skills. Infants are motivated to perform the particular actions they perform in each condition by the correspondingly activated instance of the drive to learn. The derivations in online planning capture the influence of interpretation on the motivation to act, and consequently on infants' action performance. Their ultimate goal is to learn.

do x FOR learn x.

$\longleftarrow$

x in the above conditional stands for one of *artifact function*, or *goal-directed action*, i.e. the two objects that the assimilationist epistemic goal might take. Infants' action performance is the behavioral content of their intention to learn. In online planning infants reason backwards from it towards executable actions.

The inferential structure is as follows. Observation of the novel artifact in the test phase re-activates infants' assimilationist epistemic drive. First a concrete physical achievement goal is derived from the epistemic motivation to learn. This goal is further reduced to basic executable actions. The computational mechanism is backwards reasoning – from the abstract maintenance goal, to an achievement goal, and further to executable actions.

I focus first on reduction of the goal to learn that the light-box serves the purpose of achieving the state light on. According to Section 3.1, learning novel artifact functions is the default manifestation of infants' acquisition of practical knowledge from observations of novelty. It grounds infants' online planning across the four experimental conditions. Observation of novelty triggers the conclusion 'learn novel artifact function' once a goal conjecture is formed; it is re-activated upon encountering the novel artifact in the test phase. Reasoners make it obtain by computing a plan, and enacting the instrumental behavior from the resultant action rule. The conditional below translates the maintenance goal standing for their epistemic intention, into a concrete achievement goal.

do novel artifact function FOR learn novel artifact function.

$\longleftarrow$

Infants' achievement goal is thus to act such that light-on is the case. Backwards planning from it is computationally identical with offline planning described above in Section 3.2.2; the knowledge representations that underlie computations overlap as well. The outcome is the action rule 'in order to bring about light-on contact light-box with hand', or

CEU eTD Collection

do hand contact FOR light-on.

$\longleftarrow$

Under the closed-world assumption that nothing abnormal is the case when they act, infants push the light-box with their hands. The hand action represents an emulative manifestation of the acquired practical information. Because the same processing goes on throughout the four experimental setups, the resultant behavior is also expected to occur across the board.

Let us focus now on the case of the Communicative – HandsFree condition, the only one where the drive to learn an alternative practical skill is active. The epistemic drive to learn novel goal-directed action is translated into a concrete achievement goal by the conditional

do novel goal-directed action FOR learn novel goal-directed action.

$\longleftarrow$

In the test phase following the Communicative – HandsFree action demonstration, infants are motivated to learn the bizarre way of touching the light-box with the head. The motivation is operationalized in performance of the observed skill. That is, the maintenance epistemic goal is reduced to execution of a head contact. This outcome supervenes on infants' ability to map observations onto their own motor program. Although the issue of the correspondence problem in imitation is far from trivial (Nehaniv and Dautenhahn 2002a; Brass and Heyes 2005; Nehaniv and Dautenhahn 2002b), it is beyond the objectives of the current model. I merely assume that infants are able to do this mapping, and they thus re-enact the previously observed head action directed at the light-box. By so doing they exhibit imitative learning of the novel instrumental action. I wish to emphasize that, given the goal that grounds online planning computations and the structure of the goal reduction, the head action is not performed as a means for lighting up the light-box. In contrast, when infants press the novel artifact with their hands, they do that as a method to bring about state light-on (the achievement goal), and thereby learn it as an enduring function of the artifact (the maintenance goal). Hence when infants perform different light-box directed actions in the test phase, these are instrumental with respect to different goals.

All in all, online planning in the Communicative – HandsFree condition prescribes hand behaviors for the motivation to learn artifact goals, and head behaviors for the motivation to acquire new skills. Because the latter motivation is triggered in case interpretation with respect to physical goals of actions is not satisfactory, enactment of the novel head behaviors may be understood as compensating for explanatory failure (as explained in Section 3.2.3). In fact, the imitation of the under-explained instrumental action can be viewed as a further attempt for behavioral hypothesis formation with respect to the physical goal of the head action. Infants' touching the light-box with the head is an instance of 'behavioral abduction', in the sense of abduction introduced in Section 2.4. It is as if Little Johnny were asking the teaching agent 'I did it because you did it in order to teach me, but still... what is it good for, practically speaking?'.

CEU eTD Collection

It should be noted that the suggestion developed in this Chapter regarding the reasoning steps that underlie behavior, accounts for the observational finding that when the novel action is enacted, it is preceded by performance of the default motor response. I proposed that learning about states that can be reached via artifact utilization is the default manifestation of infants' pedagogical stance; therefore infants act primarily in order to bring about the physical goal 'light-on'. But action planning with closed-world assumptions requires performance of a default motor response, one that is already available in the their motor repertoire. Planning for learning novel means actions seconds the basic, default reduction of the light-box function to actions. Imitative learning of means actions comes after emulative learning of novel goal states brought about by action on an artifact.

In conclusion, the reader is asked to picture Little Johnny as an average participant in the head-touch task. As he observes new and exciting things, i.e. behaviors and objects, he is motivated by an intrinsic desire to become wiser by learning selectively what to do, when, what for. Imagine Little Johnny in the demonstration phase of the task. In the first place he represents the context sequentially, and sets up a database where the observed facts are represented. Motivated to understand observations and to learn from them, he constructs a physical teleological model of the situation. He hypothesizes that some resultant states among the observed ones qualify as concrete physical goals of the actions performed by the agent. He updates his observational knowledge database with representations of these end states *qua* goals, and of their corresponding means actions. In so doing he reasons towards an interpretation of the context (van Lambalgen and Stenning 2008). Finally, different formats of the assimilationist motivation are activated, depending on the outcome of teleological explanation and on the social context of presentation. The motivation to learn that the light-box can produce the light-on state is active in all conditions, because the goal derivation computations are identical. Explanation with respect to the physical goal conjecture light-on is satisfactory in the HandsOccupied conditions in both Communicative and Incidental settings, so the alternative instrumental action is disregarded; it fails in both HandsFree conditions. In the Incidental – HandsFree condition, lack of ostensive signals from the agent does not support expanding the minimal teleological model; Little Johnny acquires the same functional knowledge like the ones exposed to the HandsOccupied demonstration. It is only in the Communicative – HandsFree condition that the under-explained head contact is credulously taken to be the the content of the agent's intention to teach, and thereby the motivation to learn how to act on the light-box with the head is triggered. Imagine now Little Johnny in the test phase of the experiment. When allowed to act upon the previously observed novel artifact, he reasons toward a chain of actions that presumably bring about achievement physical goals derived from his own epistemic motivation to update the knowledge database with relevant practical information. What is relevant is retrieved from the stored result of action interpretation.

Finally, I list the high-level cognitive skills required by the reasoning steps described throughout the Chapter:

CEU eTD Collection

(1) use of the teleological stance for action interpretation (see Section 1.1.1);
(2) causal perception of events, with emphasis on inertia and contact causality;
(3) motor planning;
(4) counterfactual thought;
(5) role of functions in knowledge about artifacts.

Taken together such skills account for infants' capacity for selective (i.e. wise) observational learning; imitative behavior is a practical manifestation of this capacity. The proneness to learn from communicative demonstrations of actions upon objects has been shown in numerous studies and led to the theory of natural pedagogy; the theory and supporting empirical data have been outlined in Section 1.1.3. Empirical findings regarding the availability of abilities (1) to (5) for infants before the age of the subjects in the light-box tasks, i.e. 14 months, have been quoted throughout the Chapter, and reviewed in Section 1.1. The findings speak for the psychological well-groundedness of the proposed reasoning steps, and thus for the plausibility of the model.

CEU eTD Collection

CHAPTER 4

Some logic for a refined account of the behavioral phenomena

In his latest book *How to Be Artificially Intelligent* Robert Kowalski, a pioneer with respect to applications of computational logic to human reasoning, makes the point that agents

> may use Logic to represent their beliefs about the way the world is, and to represent their goals for the way they would like the world to be (Kowalski 2011, p.8).

A follow up claim can be made that the computational component of a logical theory allows principled transitions between these representations, thereby connecting beliefs, goals, and goal-directed actions. If 'beliefs' is taken to stand for representations of current observations, I use logic as a tool for precisely these purposes: to model how young reasoners' planning capacities are recruited to interpret the instrumental actions that they observe another human agent performing, and subsequently compute action plans to bring about their own achievement goals.

The solution I put forth in Chapter 2 to the computational concerns regarding teleological practical reasoning was to view it as framed by closed-world assumptions. The inferential processes involved in the head-touch task are a form of goal-centered closed-world reasoning. Negation as failure (NAF) is the basic formal manifestation of closed-world reasoning (van Lambalgen and Hamm 2005; van Lambalgen and Stenning 2008). NAF provides the general inference pattern; in Section 2.5 it was instantiated by the closed-world assumption about abnormalities in the absence of positive evidence for abnormalities.

A good framework to capture goal-based reasoning with closed-world assumptions is provided by a logical system called constraint logic programming (CLP), used as a computational tool over a representational environment constituted by the language of the event calculus. A logical formalism can be individuated through descriptions of three parameters: i.e. syntactic, semantic, and notion of validity. Manifestations of NAF are involved in each of the three parameters of CLP: it embeds closed-world reasoning at the syntactic, semantic, and notion of validity level.

The formalization developed in Chapter 5 spells out the reasoning steps that underpin infants' action responses along the lines of CLP derivations. The knowledge that underlies computations (i.e. observations and prior knowledge) is represented in the language of the event calculus. The model's working assumption is that the dynamic reasoning processes performed over representational structures expressed with the event calculus predicates, output by means of constraint logic programming derivation rules the behaviors (i.e. actions in the world) that

CEU eTD Collection

infants exhibit in their task performance and that operationalize practical learning. By showing how the transitions from raw observation, to action interpretation, to planning of infants' actions, and, finally, to the motor response exhibited in the test phase occur, an integrated and fine-grained explanation is to be provided for the behavioral findings as manifestations of observational learning.

The event calculus was introduced by Kowalski and Sergot (1986) as a formalism for representing events and their effects. Originally it was meant to address the concerns with respect to path planning in robotics, using logic programming as an inferential engine. Logic programming, and CLP as a particular version of it, is essentially a formalism that spells out an algorithm for planning Shanahan (1997, 1999). Constitutive features of the logical system furnish a priori reasons for its use in cognitive modelling of human behavior. First and foremost, it embeds closed-world reasoning under various forms:

(1) computations performed with the inferential rule for NAF;
(2) the semantics of NAF and of the conditional;
(3) the system's non-monotonic notion of validity.

Further, the language of the event calculus axiomatizes two notions of causality (instantaneous and continuous), both of which promote the planning computations suggested in Chapter 3 to be involved in the head-touch task. Also it has representational means that can accommodate the various goal representations introduced in Section 2.2. All in all the event calculus and CLP serve as a good representational and computational specification of the multi-leveled inferential processes under focus.

This Chapter emphasizes the manifestations of closed-world action processing in the formalism. In Section 4.1 I introduce the scaffolding of CLP and the event calculus formal language. I start with a a brief outline of the system's closed-world semantics, and its non-monotonic notion of validity (Section 4.1.4). I go on in Sections 4.1.1 to 4.1.6 with an outline of the syntactic component of the theory, i.e. of the inference rules. I illustrate the formal specifications with concrete examples of reasoning processes. I anticipate how the logical notions introduced throughout the Chapter apply to the computations involved in goal-based processing of actions (i.e. explanation, and planning) in the empirical study of interest. A concrete example of the formalization of the reasoning steps involved in a toy example is given in Section 4.1.7. After familiarization with the logical machinery Section 4.2 restates its cognitive relevance. Section 4.2.1 concludes by zooming in on the relevance of the logical tool for learning. I tackle the prima facie paradoxical capacity of a closed-world reasoning formalism to represent the kind of knowledge acquisition diagnosed by infants' action responses.

The description of the logical tool throughout this Chapter relies primarily on the version of the formalism employed in van Lambalgen and Hamm (2005) to model the use of planning capacities for a cognitive semantics of tensed speech. van Lambalgen and Stenning (2008) and Kowalski (2011) are also significant sources for the import of the formal details. However for

CEU eTD Collection

the present purposes, I abstract away from the temporal component of planning. I focus on the structure of the inferential mechanism and emphasize its being suitable for action planning and explanation, and consequently for learning about actions. The expected outcome is to persuade the reader that the reasoning processes described in Chapter 3 can be implemented in CLP terms.

4.1. Constraint logic programming and the event calculus for a theory of observational learning in preverbal infants.

In order to understand sequences of goal-directed actions in the environment, to represent them as another agent's plans of actions, and to compute one's own plans conditional on these representations, one needs a representational structure, and a computational mechanism. These amount to:

(1) a set of causal rules holding for all situations (norms for a causal understanding of the world),
(2) a finite set of relevant aspects of the currently observed situation and of the background knowledge,
(3) a structural definition of derivations,
(4) some general principles that guide inferences.

The first two items in the list above are necessary to construct a declarative representation of the current domain of interest. The last two make up the computational architecture for processing representations, or the procedural component of the theory. The procedural solutions for a reasoner's problems (e.g., finding a plan of actions for a certain goal, explaining observations) follow from the way its knowledge is represented.

The representational support made available by the event calculus language is introduced in Section 4.1.1. The logical formalism of choice makes (1) above available via the axioms of the event calculus. They provide the general theory of causality presumed to underly the teleological reasoning processes involved in action processing. The axioms are listed and explained in Section 4.1.2.

The facts expressing current observations and relevant prior knowledge constitute the 'raw' data upon which interpretive and planning computations are performed. The raw data is represented as sentences with a particular logical form; it constitutes the declarative component of the database called *scenario*. It has a twofold structure: a narrative of events, and a causal model (see Section 4.1.3).

The structure of CLP derivations is given by the inference rule of *query resolution*. The semantics of resolution, as well as the definition of validity in CLP are introduced in Section 4.1.4. The logical mechanism is described in Section 4.1.5. The basic derivation method in CLP is the reduction of goals to subgoals under the auspices of NAF.

CEU eTD Collection

The general principles, (4) in the list above, have normative status and are called *integrity constraints*. They are state conditions that must be satisfied by each state of the database (Mayol and Teniente 1999) over which computations are run. The infants' teleological with respect to observed actions, for instance, is expressed in the model as an integrity constraint. The logical notion is introduced in Section 4.1.3. Its crucial role in CLP derivations, given its different semantics from regular CLP sentences, is discussed in Section 4.1.6.

4.1.0.1. *Some necessary terminology.* Let us start with some logical terminology. The definitions below will gain weight through the upcoming exposition of the formal system throughout the chapter.

Terms are expressions constituted by constants, variables, or functions of the kinds of entities among the event calculus sorts. They fill in the argument slots of predicates. The terms that contain only constants are called *ground terms*. Constants are a particular kind of function symbols, namely, 0-ary functions.

Constraints are equations expressing mathematical relations, e.g., $=$, $<$, $\leq$, between terms of the language at stake.

Atomic formulas, or *atoms*, are n-ary predicate symbols applied to n terms. The list of the primitive predicates of the event calculus in Section 4.1.1.1 contains atomic formulas.

Literals are either atomic formulas (positive literals), negations of atomic formulas (negative literals), or constraints (equations between terms).

Clause is a conditional sentence, composed of literals and constraints. The antecedent is called *body* and the consequent is called *head*. Constants, variables, and functions are clause terms. The convention of universal quantification over variables in clauses is assumed, unless otherwise specified.

Positive clause is a clause whose body contains only positive literals.

Definite clause is a clause whose body contains also negative literals. The meaning of negation in definite program clauses is that of negation as failure.

Goal clause or *query* is the body of a program clause represented interrogatively. It sets out the CLP derivations. The head of a goal clause may include negation in the form of $\perp$.

Integrity constraint is a peculiar kind of goal clause, which imposes local norms on goal reduction relative to to a particular logic program.

Constraint logic program is a syntactic format for knowledge representation expressed in the language of the event calculus predicates ($\mathcal{K}$), augmented with the language of constraints ($\mathcal{L}$). It is a set of *clauses* in the sense defined above.

Definite logic program is a constraint logic program composed of definite clauses (i.e. a finite conjunction of definite clauses).

Resolution is the main derivation rule in CLP, based on backwards chaining of actions and states.

CEU eTD Collection

Completion of a logic program P ($comp(\mathcal{P})$) is a semantic restriction of P which provides only those models of P where closed-world assumptions are satisfied. It is obtained by introducing disjunction $\vee$ between all the bodies with the same head in P, and substituting implication $\rightarrow$ with equivalence $\leftrightarrow$ between the disjunctive body and the head (i.e. by rendering the clause biconditional).

4.1.1. The representational environment of the event calculus. The event calculus is a many-sorted first order logic, with sorts for the following:

(1) *fluent types*, represent time-dependent properties such as states (e.g., light being on, hands lying on a table), or activities (e.g., touching, manipulating);
(2) *event types*, represent happenings in the world (e.g., a light-bulb *lighting up*), as well as human actions (e.g., an agent's *contacting* an object), whose instantiations at a particular time (e.g., contact *now*) typically mark the beginning and end of time dependent properties (of activities or of states);
(3) *real numbers*, to represent time and other variable quantities;
(4) *individual objects*, represent entities such as human agents, light-boxes, scarfs.

The first two sorts express types of logical objects. Tokens are obtained by instantiation of the type to a particular time (instant, or temporal interval). Type representations of events and fluents allow that the same action or activity be performed in various contexts, i.e. have different concrete tokens.

The underlying ontology of the event calculus thus comprises events, properties, time points, and individual objects. The terms over which quantification is permitted pertain to these sorts. Predication is done over events, fluents, and times. Individual objects alone do not fill argument slots of the predicates; they are parameters occurring in terms that pertain to the sort of events or fluents. Throughout this section, I use the following notation: e denotes events (both human actions, and happenings in the world), f denotes fluents (both activities and states), t denotes temporal instants, δ denotes non-negligible durations. Parameters are expressed by lower-case letters x, y, z, etc.

There are at least two aspects of the event calculus ontological structure that might strike a reader. I briefly address them below.

Firstly, a classical logician might raise an eyebrow to view (time dependent) properties ontologically on a par with individual objects and time points. They do not function as predicates, as they do in classical logic. Time-dependent properties are reified in the event calculus, i.e. they are treated as terms and not as predicates. Rather, they are entities of the kind that fill argument slots of predicates. They may be quantified over. In AI this phenomenon is labeled 'reification' and the corresponding formalisms are 'reified logics' (e.g., Prior's 1968 Tense Logic, McDermott's 1982 temporal logic). The main justification for reification is that it allows a finer expression of facts about causality (Galton 2006).

CEU eTD Collection

> Cause-effect relations with complex factual effects will require complex propositional terms to refer to them (Hugh 1987 apud Galton 2006).

Furthermore, actions and activities pertain to different sorts, of events and of fluents respectively. Actions are said to *happen* or occur, whereas activities *hold* or obtain (see below discussion of predicates). Although the distinction is obscured by the natural language meaning of 'action' and 'activity', it is of primary importance in the suggested logical language. There is consensus among most logicians that a basic contrast between events and activities is essential for modelling time and causality ((Galton 2006)). Actions or events are instantaneous, that is indexed at a particular moment in time. Activities, on the other hand, are like processes, or forces exerted continuously over a temporal interval; they have a duration. Another relevant distinction that is obscured by the natural language meaning of the terms, is that activities by themselves do not have instantaneous causal power. They, however, may have a culminating point (see also definition of culminating events below) when an event occurs that terminates them, and causes a resultant state of affairs (either a state *per se*, or a subsequent activity). In this sense, activities bring about instantaneous causal changes in the world via a mediating causally relevant event. For example, consider the activity of pressing a light switch. Although it is a common linguistic practice to say that the activity causes the light to be on, in terms of the event calculus it is the event of contact between the agent and the switch that causes the instantaneous change of the state of the light in the room. Nevertheless activities do have some sort of causal power in and of themselves, without mediating events, as forces whose continuous exertion over temporal intervals causes change. In this sense the force exerted by the activity of moving a scarf on a certain trajectory for a duration δ is said to cause the new location state of the scarf at the end of δ.

Constants denote individuals, e.g., the touching of the table, the light-box, time 1. Variable terms stand for types or classes of entities of the listed sorts, e.g., touching a table, light-box, time t. Human actions are those events whose subject (the agent) is a human being. In general, when I use the word 'action', unless other specifications, I refer to human action. Instrumental entities from the category of events and fluents are goal-directed actions or activities. Culminating events are events that end an activity and bring about a resultant state.

Within the sorts of events and of fluents, there are parametrized entities. They are a function of a parameter x which, for a certain value $f(x)$ of x, yields a concrete object (event or fluent) in the model. A constant indexed with a variable parameter does not specify the value of the function; it is a variable in the model. Direct or indirect objects of verbs are good examples of parameters. In the current formalization endeavor, the local agentive relatum of an action or activity is represented as a parameter that specifies the event or the fluent. $\textit{Contact}_{hand}$ is a contact event performed with the hand, $\textit{activity}_{head}$ is an activity performed with the head. *Head* and *hand* are examples of parameters.

Attached to the ontological structure of the event calculus is the semantic assumption of uniqueness of names (UNA). The core idea is that different names of actions, states, activities,

CEU eTD Collection

or parameters refer to different entities in the world – i.e. to entities in the outside world, or to second-order entities in the minds of the reasoners who interpret the outside world. Conversely, if two entities bear the same name then they are identical. All in all, one constant stands for one and only one representation. Although in natural language one can say that the event of touching and that of contacting co-refer, in the event calculus if one introduces an event labeled 'touch' and an event labeled 'contact' then the two names stand for two different events. To take another example: when applied to parameters expressing local agentive relata, the assumption demands that $contact_x = contact_y \rightarrow x = y$.

4.1.1.1. *The predicates of the event calculus.* The predicates of the event calculus provide a format for representations of knowledge of what event happens when, of what fluent holds (is true) at what time, and of the causal relations that connect tokens of actions and their effects at certain time points. The predicates have two or more arguments, of which at least one stands for time (in the case of the first predicate in the list below, the time argument is implicit).

The event calculus is a logical environment that connects fluents, events, and their effects (resultant states), all indexed at particular time points. It formalizes causation as a relation between events and states of affairs in the world, similar to the way it is understood in physics. In fact it purports to capture two notions of cause, concerned with instantaneous, and with continuous change, respectively. An example of the former notion is the change of the light state in a room upon the event of pressing a switch. Properties with continuous variation, such as the acceleration of an object due to gravity, or the filling of a container with water under continuous flow of water from a tap, are instances of the latter notion of cause. Continuous variation is change without concomitant occurrence of events. Event occurrences are necessary to cause change of states in the former sense, but not in the latter. Consequently the two views of causation are prima facie in conflict. Cause as instantaneous change leads to one form of inertia: after the occurrence of the event marking the change, properties remain unchanged until another incidence of a relevant event. This conflicts with continuous causation under the influence of a force. The solution that renders the complex, twofold view of causation coherent is furnished by the set of primitive predicates.

The predicates are:

(1) *Initially(f)*: fluent f holds at time 0;
(2) *HoldsAt(f, t)*: fluent f obtains at time t;
(3) *Happens(e, t)*: event e occurs at time t;
(4) *Initiates(e, f, t)*: fluent f is the causal effect of event e at time t;
(5) *Terminates(e, f, t)*: the cessation of fluent f is the causal effect of event a at time t;
(6) *Trajectory* (f_1, t, f_2, δ): given that f_1 is a fluent expressing a force and f_2 is a property that changes under the influence of the force, if f_1 holds from t until $t + \delta$, then f_2 holds at $t + \delta$;
(7) *Releases(e, f, t)*: event e releases fluent f from the law of inertia;
(8) *Clipped(s, f, t)*: fluent f is terminated or released between times s and t;

CEU eTD Collection

(9) predicates expressing basic mathematical relations between time points and individual objects ('less than' and 'equals'): $\leq$, =[1].

Items (1) to (8) delineate the language of the primitive predicates of the event calculus, $\mathcal{K}$. *HoldsAt* expresses what states are true when, or what activities are ongoing at a certain time. It is the truth predicate in the event calculus. Truth applies to ground formulas only. 'Is true', 'obtains', and 'holds' are henceforth used interchangeably. *Initially(f)* is equivalent with *HoldsAt(f, 0)*. Or else, f is true at time point 0. (3) describes what happens when. It states that an event of type e occurs at time t. *Happens* and *HoldsAt* predicates provide the basic representational format for event narratives (see Section 4.1.3).

Initiates, *Terminates*, and *Trajectory* are the causal predicates. They represent the effects of events (the first two) and fluents (the latter). They express what actions do (Shanahan 1999). Predication of causal relations does not make commitments with respect to concrete occurrences of token events or fluents with the required causal power. The predicates establish causal relations between event types and fluent types. In other words, *Initiates*, *Terminates* and *Trajectory* do not entail *Happens* or *HoldsAt* at a given time point t. *Initiates* and *Terminates* are expressions of instantaneous causation whereby events are the cause of change. Predicate *Trajectory* expresses continuous causation

> where variable quantities may change their values without concomitant occurrences of events (van Lambalgen and Hamm 2005, p. 44).

Trajectory(f_1, t, f_2, δ) expresses gradual, continuous change by linking a fluent type f_2 to the force f_1 whose exertion from t to $t + \delta$ causes f_2 to be the case at $t + \delta$. It does differentiate between the kinds of fluents that may occupy the two fluent argument slots. Argument f_1 ranges only over activities, or forces (i.e. fluents with causal power), whereas f_2 ranges over state fluents as well as activities.

The convention (as set forth in van Lambalgen and Hamm (2005)) is that fluents start to hold *strictly after* the time of their initiating event. That is if *Initiates(e, f, t)*, then *HoldsAt(f, t+1)*. The opposite is true with respect to terminating events. If *Terminates(e, f, t)*, then *HoldsAt(f, t)* but $\neg$ *HoldsAt(f, t+1)*. The consequence is that time-dependent properties obtain on half-open intervals, open at their initial point. In other words, properties do not (yet) hold at the moment when an event causes them, but they do (still) hold at the moment when another event ceases them. This idea is implemented in the strict ordering between time point variables in the conditions of Axioms 1, 3, and 4 (see Section 4.1.2).

Predicate *Releases* accommodates the two notions of cause by exempting properties with continuous variation from the inertia of instantaneous change. It states that fluent f is not subject to the law of instantaneous inertia after the occurrence of e at time t. An illustration is a formula such as *Releases(start*$_{hand}$*, activity*$_{hand}$*, t)*, which expresses that after a starting

[1]The mathematical notation is simpler and hence more convenient. The very same relations might as well be expressed as dyadic predicates of the form *LessThan*(t_1, t_2), *Equal*(f_1, f_2)

CEU eTD Collection

event at time t the hands become the local agentive relatum of an activity. Before *start*$_{hand}$ at t, changes with respect to the activity$_{hand}$ (e.g, its holding or not) could be brought about only by explicit occurrences of events. However, from $t + 1$ onwards, it can change value continuously under the influence of a force (say, the fluent *will-power*) for as long as the force is in effect, without concurrent happenings. This continuous variation would be represented as *Trajectory(will-power, t, activity*$_{hand}$*,* δ*)* and would be governed by continuous inertia

Clipped(s, f, t) expresses the occurrence of a causally f-relevant event in a given temporal interval $(s, t]$. A fluent is clipped on the interval $(s, t]$ if it does not persist between s and t.

Predicates such as the ones in item (9) above pertain to a language of constraints between terms, $\mathcal{L}$, which is peculiar to the version of logic programming I use. The constraints are atomic formulas that may contain constants, as well as equations between variables. The basic relations expressed in equations are 'less than', and 'equals'. Some concrete examples are to be found in the axioms of the event calculus, more specifically, in the formulas that express relations between time variables, such as $t < s$ or $t^{'} = t + \delta$.

Constraint logic programming is concerned with the interplay between the language of constraints $\mathcal{L}$ and that of the primitive predicates of the event calculus $\mathcal{K}$. This interaction is spelled out in Section 4.1.5.

4.1.2. The axioms. One is now well enough equipped to grasp the event calculus axioms and their relevance for the causal action processing involved in teleological interpretation and planning. The list below amends the set available in van Lambalgen and Hamm (2005). The set of four axioms is abbreviated *EC*. All variables occurring in the axioms are assumed to be universally quantified.

The axioms spell out the conditions of properties being true about the world, and of properties ceasing to be true. These conditions are based on inertial reasoning, which is a a form of closed-world reasoning about causality (as explained in Section 2.5). Inertia is approached in the two senses introduced at the outset of Section 4.1.1, as a law that governs instantaneous and continuous change, respectively. The general format of the axioms is "when property f is caused to obtain, unless there is positive information about an f-relevant event on a given temporal interval, f obtains at the end of the interval".

CEU eTD Collection

AXIOM 1. *Happens*$(e, s) \ \wedge \ t < s < t^{'} \wedge$ (*Terminates*$(e, f, s) \ \vee$ *Releases*$(e, f, s)) \ \rightarrow$ *Clipped*$(t, f, t^{'})$

AXIOM 2. *Initially*$(f) \ \rightarrow \ HoldsAt(f, 0)$

AXIOM 3. *Happens*$(e, t) \wedge$ *Initiates*$(e, f, t) \wedge \ t < t^{'} \ \wedge \neg$*Clipped*$(t, f, t^{'}) \rightarrow$ *HoldsAt*$(f, t^{'})$

AXIOM 4. *Happens*$(e, t) \wedge$*Initiates*$(e, f_1, t) \wedge \ t < t^{'} \wedge t^{'} = t + \delta \wedge$*Trajectory*$(f_1, t, f_2, \delta) \wedge$ $\neg$*Clipped*$(t, f_1, t^{'}) \rightarrow$ *HoldsAt*$(f_2, t^{'})$

The first axiom defines the meaning of the *Clipped(t, f, $t^{'}$)* predicate with reference to the occurrence of an event that either causes the cessation of property f (a terminating event), or that alters the inertial status of f (a releasing event). It specifies the meaning of causally fluent-relevant (f-relevant) events to terminating or releasing events. One way to paraphrase the head of the clause is as *HoldsAt*(ab_f, r), $r \in (s, t]$. That is, something abnormal is the case with respect to fluent f on a given temporal interval $(t, t^{'}]$ if it is either terminated or released on that interval. A time-dependent property does not persist across a temporal interval if an event terminates it, or makes it such that it is no longer subject to the inertia of instantaneous change.

The other three axioms spell out the conditions under which certain properties are (made) true about the world, i.e. the conditions of *HoldsAt(f, t)* predicates.

Axiom 2 specifies the meaning of the $Initially$ predicate with respect to a fluent that obtains at a particular constant time, namely, the beginning of time, or time 0. It does not make any reference to the causal route by which the fluent is made true.

Axioms 3 and 4 are the causal axioms of the event calculus. They capture the principle of inertia (i.e. a property persists unless it is caused to change by the explicit occurrence of an event) if read in a closed-world reasoning framework, which assumes that no un-mentioned event occurs. The $\neg Clipped$ conjunct in the bodies gives a formal characterization of closed-world reasoning about abnormalities with respect to the temporal interference of events that alter the inertial status of fluent f. Reasoning with inertia amounts to expecting that a change of state occurs only upon an explicit causal influence. The two axioms express the idea that properties that are true at the beginning of a temporal interval are also true at the end, unless causally relevant events occur within the given interval, where the meaning of "causally relevant" is set forth in Axiom 1.

Axiom 3 is an expression of the principle of inertia with respect to instantaneous change. It states that a fluent f holds at a certain moment in time $t^{'}$ if it is started by an event e which takes place at an earlier moment t, and it persists (it is not *Clipped*) from the moment when it is started t until the moment when it is claimed to hold $t^{'}$. Or, conversely, unless there is information about a causally relevant event in a given temporal interval, the resultant state of an event continues to hold.

Axiom 4 expresses a similar principle with respect to continuous change, or causation exerted by a continuous force. It states that a fluent f_2 holds at $t^{'}$ if information is available about

CEU eTD Collection

some other fluent f_1, a force, which is brought about by the concrete occurrence of an event (i.e. instantaneously) at t prior to $t^{'}$, and whose continuous influence for a non-negligible duration δ causes the first fluent to hold at time $t^{'} = t + \delta$.

4.1.3. A logical format for knowledge databases. A database is the set of knowledge representations over which computations are run. It comprises a scenario, the axioms of the event calculus, and integrity constraints.

The scenario provides a description of a given situation in sentences with a particular logical format, i.e. clauses. Scenario clauses are composed of ground terms. They have the generic format

$$B_1, \ldots, B_n, \neg D_1, \ldots, D_m, c \rightarrow A.$$

The conclusion of a clause or the head is an atomic formula (A). Constraints in language $\mathcal{L}$ can also occur in the head. The conditions or the body may contain atomic formulas (B_i), negations of atomic formulas ($\neg D_i$), constraints (c), as well as tautologies ($\top$) and contradictions ($\bot$). Each of the constituents of the body is a *literal*. Commas between the literals and constraints are to be read as conjunction $\wedge$. The default quantification over clause variables is universal (exceptions arise in the case of integrity constraint clauses).

The meaning of the $\rightarrow$ symbol in conditional clauses is in accordance with the closed-world semantics of the logic programming conditional (for which see Section 4.1.4). In short, the body should be read as including an abnormality proviso of the kind 'unless something abnormal is the case'.

The *EC*, scenario clauses, and integrity constraints, share the logical format of conditional clauses. The difference between axioms and regular scenario clauses is that predicate arguments in the axioms are universally quantified variables, as opposed to constants in scenario clauses. If there are no negative formulas in the body of a conditional clause (i.e. if $m = 0$ in the generic format provided above), it is positive or a *Horn clause*. Among the *EC* Axiom 1 and Axiom 2 are instances of positive conditional clauses. The axioms that formalize inertia, Axiom 3 and Axiom 4, are definite clauses because they contain instances of $\neg Clipped$ in the body.

A scenario should be understood as a

> 'micro-theory' which states the specific causal relationships holding in a given situation, and which lists the events that have occurred in that situation (van Lambalgen and Hamm 2005, p.48).

This construal makes explicit the twofold structure of scenarios; they are composed of a narrative of events and a causal model. Narrative clauses contain the predicates *Happens* and *HoldsAt*. Observed facts expressed in the narrative have the clausal form

$$\top \rightarrow A.$$

This means that facts are the conclusions of tautologically true conditions. For simplicity I will omit the tautology write only the atom A. Causal model literals involve the causal predicates

CEU eTD Collection

Initiates, *Terminates*, *Trajectory*, and *Releases*. Clauses in the causal model whose atoms contain *Releases* and *Trajectory* predicates define the dynamics of the situation. The term 'micro-*theory*' draws attention to the fact that the set of individuated raw facts is implicitly influenced by cognitive aspects, namely, by considerations of relevance for further processing. The list of clauses in a database currently in use is regimented by closed-world reasoning in an informal sense: only those aspects of a situation that are relevant in a particular reasoning context, for a particular purpose (e.g., understanding current observations), are part of the micro-theory and expressed as scenario clauses.

Attached to scenarios are sets of local normative principles, which constrain the updates to the database; they are called *integrity constraints*. They express obligations and prohibitions that all states of the database must satisfy. The logical format of integrity constraints is similar to that of scenario clauses. Their degree of abstraction, and hence their scope of applicability, is lower than that of axioms (they are *local* normative principles), but higher than that of regular scenario clauses. This is manifested in that the terms in integrity constraint clauses may be, and most often are, variables, unlike regular scenario clauses where predication is done over constants. Combinations of existential and universal quantifiers may occur. This is related to the fact that causal predicates are customary constituents of the formulas in integrity constraints, and these predicates have a more complex argument structure. For instance, the clause $?Initiates(action, goal, t)$ `succeeds` is an integrity constraint that requires goal assignment to observed actions.

The integrity constraints are part of the database, but are distinguished from ordinary scenario clauses by their syntax. The main syntactic difference concerns quantification. The peculiar kind of normativity of integrity constraints (i.e. local normativity) manifests itself in the scope of quantification over variables. I get back to this in Section 4.1.5 when introducing the notion of goal clause (or query), of which integrity constraint clauses are a subspecies, and in Section 4.1.6, which deals with the role of integrity constraints in computations.

A *logic program with constraints* is a particular syntactic format expressed in the language of the event calculus predicates $\mathcal{K}$ augmented with the language of constraints $\mathcal{L}$, which can accommodate clauses in the *EC*, the scenario of a particular situation, and the integrity constraints. The *EC* are examples of logic programs used to characterize the primitive predicates of the event calculus. For instance, the first axiom is a single-clause logic program that defines the meaning of predicate *Clipped*, while Axiom 3 and 4 are logic programs that provide definitions for *HoldsAt*. If expressed in a similar format also the plan to turn on a light-box by means of a head action, or the interpretation of the light-on state as the goal of a head touch, or the explanation that a goal-directed head touch occurs because there is something abnormal with respect to hands' bringing about the goal light-on, are constraint logic programs. The formalization endeavor carried on in Chapter 5 spells out such reasoning threads as constraint logic programs. In order to get there one must first become acquainted with the derivational procedures that go along with constraint logic programs, as well as with the denotation of these procedures. The

CEU eTD Collection

next Section specifies the latter by discussing the semantics and the consequence relation of CLP.

4.1.4. Some relevant aspects of the system's semantics and notion of validity. Closed-world reasoning accounts for the specifics of the semantics and notion of validity parameters of CLP. In its guise as negation as failure, or as the completion of logic programs, the closed-world assumption for reasoning about abnormalities subsumes the relevant peculiarities of the logical system under focus.

Unlike the case of classical logic semantics where truth is skeptically evaluated with respect to all models[2] of the clauses under scrutiny, the semantics of CLP is best understood in terms of truth in minimal models. That is the semantic manifestation of closed-world reasoning as negation as failure, and the corresponding process of completion of a logic program, evaluate truth with respect to a minimal model of the clauses that constitute a logic program P. Among the truth connectives, negation and the conditional in particular are linked with the construction of, and further reasoning from minimal models (via what McCarthy (1980b) calls minimal entailment).

The notion of model construction in formal semantics is tantamount to interpretation. Interpretation is a psychological term that overlaps with the semantic notion. Action processing, and subsequent planning proceed via the construction of models for observations and relevant bits of background knowledge. Model construction, or the formal counterpart of reasoning towards an interpretation, had better be uniformly and efficiently computable[3]. This is ensured by the CLP semantics for a weak notion of negation, i.e. negation as failure (NAF) (Etzioni et al. 1997; Lifschitz 2002). CLP is a suitable computational format for practical reasoning as it manifests itself in teleological action interpretation and planning precisely because the semantics of negation as NAF allows computations to be performed in minimal models.

The semantics of closed-world reasoning sanctions the idea that computations with the CLP inference rules are performed in minimal models. A model of a situation is minimal in the sense that

> the occurrences of events and their causal influences are restricted to what is required by the scenario and the axioms (van Lambalgen and Hamm 2005, p. 53).

Such minimal models are closed-worlds. Those events and fluents that are not forced to occur given program P are assumed not to occur. The only events and fluents forced to occur are, according to negation of failure, the ones mentioned explicitly in $Happens$ or $HoldsAt$ literals in P clauses, and the the ones derivable from P in conformity with the EC. These

[2]The term 'model' is used here as Ôsemantic model'. It should not be confounded with its homonymous meaning of 'formal theory', e.g., the formal model that this dissertation is aimed at.

[3]Because of the need to deal with the concern of un-manageable size of the reasoning space, or the much discussed frame problem. See Section 2.5.

CEU eTD Collection

effects are computed with the rule of resolution (detailed below in Section 4.1.5). Hence a minimal model of P is the simplest, most parsimonious constructed interpretation of the program clauses. Constructing a minimal model of another agent's observed behavior is aimed only at the one interpretation presumably intended by the agent[4].

The CLP meaning of negation, NAF, is weaker than the one in classical logic. If there is no explicit reason for the truth of a fluent f holding, or of an event e occurring in a program P, its negation can be assumed (until further evidence to the contrary). Events/ fluents not forced to occur/ hold by the scenario are assumed not to occur/ hold. That is, if the database does not state that *HoldsAt(f, t)* or *Happens(e, t)*, nor can these be derived by the syntactic rules of CLP (introduced in Section 4.1.5) from P, then a reasoner can assume that ¬*HoldsAt(f, t)* or ¬*Happens(e, t)*, i.e. that no f is ongoing at time t or no event e occurs at t.

Once a minimal model of P is constructed, computations that reduce goal clauses may turn out to provide explicit information about events or fluents not included in the initial model. P is then minimally updated with these events or fluents, which thereby become part of a subsequent model of P. Allowing only a minimal amount of additional data to the initial minimal interpretation means, given derivations from it, in logical terms, an incremental computation of minimal models. This provides one reason why I set forth a formal treatment of practical reasoning in a CWR framework via logic programming and not circumscription, which does not permit such stepwise additions. They are important because they are related to the convergence toward a stable state of a neural network, which shows why model construction can proceed automatically (van Lambalgen and Stenning 2008, footnote 17, p. 189)[5].

4.1.4.1. *Formally on minimal models and minimization.* The formal Definitions and Lemmas are imported from van Lambalgen and Hamm (2005) with the authors' permission. The models to be considered have the form specified by Definition 1.

DEFINITION 1. *Let P be a positive program on a finite set L of proposition letters. An assignment $\mathcal{M}$ of truth-values {0, 1} to L (i.e. a function $\mathcal{M} : L \to$ {0, 1}) is a* model *of P if for $q \in L$,*

(1) $\mathcal{M}(q) = 1$ if there is a clause $p_1, \ldots p_n \to q$ in P such that for all i, $\mathcal{M}(p_i) = 1$
(2) $\mathcal{M}(q) = 0$ if for all clauses $p_1, \ldots p_n \to q$ in P there is some p_i for which $\mathcal{M}(p_i) = 0$.

(2) in the definition entails that for any q not occurring as the head of a clause in P, $\mathcal{M}$(q) = 0. Furthermore, if q occurs in a single formula of the form $p \to q$ and p does not occur as the head of a clause, also $\mathcal{M}$(q) = 0. The model $\mathcal{M}$ is minimal in the sense that any proposition not forced to be true by the program P is false in $\mathcal{M}$.

[4]Note the similarity with the Gricean pragmatics of discourse interpretation.
[5]This automaticity further reinforces the understanding of practical reasoning as not necessarily conscious, linguistically expressible inferences (as discussed in Section 2.1 at p. 36).

CEU eTD Collection

The set of models as defined in Definition 1 can be constructed iteratively. In formal terms, they are given by the fixed points of a consequence operator $\mathcal{T}_{\mathcal{P}}$[6], which is monotone in the case of positive programs. Monotonicity of $\mathcal{T}_{\mathcal{P}}$ is relative to truth values, and it should not be confounded with monotonicity of the system's definition of validity. Monotonicity is important because it implies the existence of fixed points of the operator $\mathcal{T}_{\mathcal{P}}$, with respect to which models of P can be iterated. The minimal model of a positive program P is the least fixed point of the monotone consequence operator associated to P.

DEFINITION 2. *The operator T_P associated to P transforms an assignment $\mathcal{V}$ (identified with the set of proposition letters made true true by $\mathcal{V}$) into a model $T_P(\mathcal{V})$ according to the following stipulations: if u is a proposition letter,*

(1) $T_P(\mathcal{V})(u) = 1$ if there exists a set of proposition letters C, made true by $\mathcal{V}$, such that $\bigwedge C \rightarrow u \in P$

(2) $T_P(\mathcal{M})(u) = 0$ otherwise.

DEFINITION 3. *An ordering $\subseteq$ on assignments $\mathcal{V}, \mathcal{W}$ is given by: $\mathcal{V} \subseteq \mathcal{W}$ if all proposition letters true in $\mathcal{V}$ are true in $\mathcal{W}$.*

LEMMA 1. *If P is a positive logic program, T_P is monotone in the sense that $\mathcal{V} \subseteq \mathcal{W}$ implies $T_P(\mathcal{V}) \subseteq T_P(\mathcal{W})$.*

DEFINITION 4. *A fixed point of T_P is an assignment $\mathcal{V}$ such that $T_P(\mathcal{V}) = \mathcal{V}$.*

LEMMA 2. *If T_P is monotone, it has a least and a greatest fixed point.*

LEMMA 3. *A fixed point is a model in the sense of Definition 1. Every such model is also a fixed point.*

CEU eTD Collection

However, it is rather definite logic programs that are involved in my use of CLP for modelling. Allowing negation in the body of clauses brings about some complications relative to the definitions above, because monotonicity fails, and hence so do the fixed points of T_P as defined above. The solution is to re-define a consequence operator that applies to definite programs in a three-valued logic, Kleene's logic, and thus called the three-valued consequence operator $\mathcal{T}_P^3$. The truth values are {0, 1, u}, where u is not a degree of truth intermediary between 0 and 1 (unlike in Lukasiewicz' three-valued logic, for example). Rather u is undecided and can evolve toward 0 or 1 as a computational upshot. Once a literal has been assigned value 0 or 1 by $\mathcal{T}_P^3$, it retains that value at all stages of the model construction; if it has been assigned value u, that

[6]The consequence operator is important because it provides the connection between minimal models and stable states of neural networks (van Lambalgen and Stenning 2008), and thus provides more indirect support for the cognitive relevance of the formalism. The consequence operator is mirrored by the activation flow in a neural network.

p	$\neg p$
1	0
0	1
u	u

p	q	$p \wedge q$
1	1	1
0	0	0
u	u	u
1	0	0
1	u	u
0	1	0
0	u	0
u	1	u
u	0	0

p	q	$p \vee q$
1	1	1
0	0	0
u	u	u
1	0	1
1	u	1
0	1	1
0	u	u
u	1	1
u	0	u

Figure 1. Three-valued connectives

value may mutate into 0 or 1 at a later stage. The truth values tables for Kleene's logic governing the three-valued consequence operator are given in Fig. 1. Equivalence $\leftrightarrow$ is defined in a two-valued logic by assigning 1 to $\varphi \leftrightarrow \psi$ if φ and ψ have the same truth value (in $\{u, 0, 1\}$), and 0 otherwise.

If P is a definite program the set of models of P is constructed as fixed points of $\mathcal{T}_P^3$, according to the following Definitions.

DEFINITION 5. *A three-valued model is an assignment of the truth values $u, 0, 1$ to the set of literals in P. If the assignment does not use the value u, the model is called* two-valued*. If $\mathcal{M}, \mathcal{N}$ are models, the relation $\mathcal{M} \leq \mathcal{N}$ means that the truth value of a proposition letter p in $\mathcal{M}$ is less than or equal to the truth value of p in $\mathcal{N}$ in the canonical ordering $\leq$ on $u, 0, 1$ (i.e. $u \leq 0$ and $u \leq 1$; 0 and 1 are incomparable in the ordering).*

DEFINITION 6. *Let P be a definite program.*

a. The operator $\mathcal{T}_P^3$ applied to formulae constructed using only $\neg$, $\wedge$ and $\vee$ is determined by the truth tables in Fig. 1.

b. Given a three-valued model $\mathcal{M}$, $T_P^3(\mathcal{M})$ is the model determined by

(a) $T_P^3(\mathcal{M})(q) = 1$ iff there is a clause $\varphi \rightarrow q$ such that $\mathcal{M} \models \varphi$

(b) $T_P^3(\mathcal{M})(q) = 0$ iff for all clauses $\varphi \rightarrow q$ in P, $\mathcal{M} \models \neg\varphi$

(c) $T_P^3(\mathcal{M})(q) = u$ otherwise.

Monotonicity with respect to truth-values is different for a three-valued operator than for the two-valued operator associated to positive programs according to Lemma 1. The following result spells this out.

LEMMA 4. *If P is a definite logic program, T_P^3 is* monotone *with respect to truth values in the sense that $\mathcal{M} \leq \mathcal{N}$ implies $T_P(\mathcal{M}) \leq T_P(\mathcal{N})$.*

Restricting the construction of a model for P to minimal models is called minimization. A minimal model of P is obtained by application of closed-world reasoning in the form of a syntactic operation called completion. Completion and minimization are co-referring processes.

CEU eTD Collection

Minimization has the effect that the only model of P which is of interest for CLP derivations does not consider information that is not explicitly stated as a clause in P. The model of the completion of the program P $comp(\mathcal{P})$ is a minimal model of P.

The term 'completion' covers both the method, and the effect of its application to a constraint logic program P. The notation $comp(\mathcal{P})$ expresses the latter. Computability of a minimal model is a consequence of applying syntactic restrictions to logic programs. $Comp(\mathcal{P})$ is obtained in two steps. First, disjunction is introduced between all bodies with the same head in the program clauses of P. Second, the implication $\rightarrow$ is replaced with equivalence $\leftrightarrow$. The result is that the head obtains *if and only if* the disjunctive antecedent does. Methodologically speaking, the syntactic operation of completion rejects any unforeseen or unexpected events are rejected from the model of P. It ensures that computations are performed in a minimal model of P, or a model of $comp(\mathcal{P})$.

DEFINITION 7. *Let P be a positive program. The* completion *of a positive program P is given by the following procedure:*

(1) take all clauses $\varphi_i \rightarrow q$ whose head is q and form the expression $\bigvee_i \varphi_i \rightarrow q$
(2) if q does not occur as a head, introduce the clause $\bot \rightarrow q$
(3) replace the implications ($\rightarrow$) by bi-implications ($\leftrightarrow$) (here, $\leftrightarrow$ is semantically interpreted by $\mathcal{V}(\psi \leftrightarrow \varphi) = 1$ iff $\mathcal{V}(\psi) = \mathcal{V}(\varphi)$, and 0 otherwise)
(4) take the conjunction of the (finitely many) sentences thus obtained; this gives the completion of P, which will be denoted by $comp(P)$.

Let us consider an example. Suppose P is a program formalizing an agent's observations that a lamp is switched on at time 4 – *Happens(on, 4)*, and switched off at time 7 – *Happens(off, 7)*. A conditional representation of the clauses may be given by introducing variable e to stand for an event that happens at time t; the program clauses are $Happens(on, 4) \rightarrow Happens(e, t)$ and $Happens(off, 7) \rightarrow Happens(e, t)$. Completion implements negation as failure by transforming the above clause into

Happens(on, 4) $\vee$ *Happens(off, 7)* $\leftrightarrow$ *Happens(e, t).*

Any other events but the ones explicitly mentioned are rejected from the model. Program P becomes $comp(\mathcal{P})$. A non-minimal interpretation would accept the possibility that, for instance, *Happens(off, 5)* and *Happens(on, 6)*, i.e. that e can be specialized to other values at other time points than on at 4 and off at 7.

CEU eTD Collection

The process of completion of a logic program P resulting in its transformation into $comp(\mathcal{P})$ is the fundamental technical tool to reformulate the semantics of CLP derivations (i.e. of resolution) as a classical semantics. Because the minimal model of P is in fact the[7] model of $comp(\mathcal{P})$, which is obtained from P by a closed-world reasoning process (completion), the utilization of the syntactic inferential rule of NAF is warranted for the resolution of goal clauses relative to P.

The procedure of minimization captures the fact that most often in default human reasoning conditionals are interpreted as biconditionals. For example, suppose that a reasoner's database represented in a program P contains the conditional action rule "light-on *if* hand contact and nothing abnormal is the case", and no other clause with 'light-on' in the head. Completion leads to representing it as an equivalence "light-on *if and only if* hand contact". It is licensed by the closed-world assumption that nothing abnormal is the case.

The following lemma provides a strategy for the construction of minimal models of definite programs, via the three-valued consequence operator.

LEMMA 5. *Let P be a definite logic program.*

a. *$\mathcal{M}$ is a model of the $comp(P)$ if and only if it is a fixed point of T_P^3.*

b. *The least fixed point of T_P^3 exists and is reached in finitely many steps ($n+1$ if the program consists of n clauses). The least fixed point of $\mathcal{T}_P^3$ is called the minimal model of P.*

c. *All models $\mathcal{M}$ of $comp(\mathcal{P})$ are fixed points of $\mathcal{T}_P^3$, and every fixed point is a model.*

4.1.4.2. *The CLP conditional and the consequence relation.* The relation between the body and the head of a conditional program clause is law-like, yet defeasible. Such laws hide an endless number of assumptions. When closed-world assumptions 'of normality' are overridden, exceptions to the laws arise in a justified manner. They do not cancel the validity of the law. The conditional has a stronger antecedent, enriched with a non-abnormality proviso (the 'unless proviso' introduced in Section 2.5). Should the antecedent be true, and the conclusion false, the conditional need not be considered invalid (as is the case for the classical logic conditional). Rather, the closed-world of reasoning is opened to admit that, after all, there is something abnormal that prevents the truth of the conclusion although the antecedent is true. Because of the semantics of the $\rightarrow$, the CLP conditional formulation provides an adequate translation in a formal language of human agents' action rules (as discussed in Chapter 2).

The meaning of the $\rightarrow$ serves as the foundation for understanding the system's consequence relation, or its notion of validity. It establishes a relation between the semantics of the formalism and its use in syntactic derivations. The success of a query resolution with respect to a logic program P is entailed by the completion of the program $comp(\mathcal{P})$. A query φ is successfully

[7] *The* implies uniqueness, which is indeed the case in the simple example presented here, but need not generally be so.

CEU eTD Collection

resolved with respect to P only under the assumption that P contains all the information that is relevant for derivations starting from φ. This is the case in a minimal model of P, or a model of $comp(\mathcal{P})$. Likewise, failure of a query means that the completion of the program entails the negation of that query. I introduce the symbol $\mid\!\approx$ for the CLP consequence relation to differentiate it from that of classical logic $\models$. If $P \mid\!\approx \varphi$, we say that φ follows from P by closed-world reasoning, or by negation as failure.

The classical logic definition of validity reads $P \models \varphi$ as 'φ is true in all models of P'. In contrast, $P \mid\!\approx \varphi$ means that φ is true in a minimal model of P, which is a model of $comp(\mathcal{P})$. Just like in the case of semantics, $comp(\mathcal{P})$ provides a link between the definition of validity in CLP and in classical logic. The relation is

$$P \mid\!\approx \varphi \textit{ iff } comp(P) \models \varphi.$$

The righthand side of the equivalence means that all three-valued models of $comp(\mathcal{P})$, which are given by the consequence operator according to Lemma 5, satisfy φ.

The closed-world assumption for abnormalities itself defines the system's non-monotonic validity (van Lambalgen and Stenning 2008). Think about a CLP clause of the form ***HoldsAt(p, t)***$\wedge$ $\neg$***HoldsAt***$(ab, t) \rightarrow$ ***HoldsAt***(q, t), in which the negative literal $\neg$***HoldsAt***(ab, t) is the unless proviso. Suppose then that the information available about the fluent ab is expressed in the following conditional clauses: ***HoldsAt***$(r_1, t) \rightarrow$ ***HoldsAt***$(ab, t), \ldots,$ ***HoldsAt***$(r_k, t) \rightarrow$ ***HoldsAt***(ab, t), and that this is *all* the information regarding ab. One can also add the trivial[8] implication $\bot \rightarrow$ ***HoldsAt***(ab, t). All in all, something abnormal is the case if and only if one of the fluents r_1, or $\ldots$, or r_k obtains. This can be expressed in a single equivalence clause ***HoldsAt***$(r_1, t) \vee \ldots \vee$ ***HoldsAt***$(r_k, t) \vee \bot \leftrightarrow$ ***HoldsAt***(ab, t). By classical validity we get that

HoldsAt(p, t) $\wedge \neg$***HoldsAt***$(r_1, t) \wedge \ldots \wedge \neg$***HoldsAt***$(r_k, t) \models$ ***HoldsAt***(q, t)[9].

If there is no information available about properties that lead to abnormalities, then the previous equivalence that defined abnormality reduces to $\bot \leftrightarrow$ ***HoldsAt***(ab, t), or $\neg$***HoldsAt***(ab, t). This is precisely a formal exposition of the closed-world assumption: unless positive/ nontrivial information is available, nothing abnormal is the case. However, if $HoldsAt(r_k, t)$ is added to P, then abnormality is the case and $HoldsAt(q, t)$ no longer follows from $HoldsAt(p, t)$. This means that the CLP definition of validity is non-monotonic. The addition of new information of the kind r_k, i.e. among abnormality conditions, leads to retraction or revision of previously held conclusions.

[8]Trivial because from an arbitrary contradiction everything follows

[9]Note that the closed-world inference is non-classical, because in classical logic nothing can be concluded about q from the conjunction of information about $p, r_1, \ldots, r_k$.

CEU eTD Collection

4.1.5. Computational strategies in CLP.. I proceed to introducing the computational ingredients of the theory, its 'inferential engine'. The following statement of the direction of CLP computations over the machinery of the event calculus gives a good starting point:

> The aim of a constraint computation is to express a programmed predicate symbol entirely in terms of constraints or [...] to find an assignment in the programmed predicate which satisfies a given constraint (van Lambalgen and Hamm 2005, p.52).

The constraints that constitute the satisfactory computational outcomes update the knowledge database. The update process is normed by integrity constraints, as described below in Section 4.1.6.

As mentioned at the outset of this Chapter, Logic Programming (of which Constraint Logic Programming is an extension) and the event calculus are formalisms developed in AI for planning agents. Computations licensed by CLP reflect the inferential structure of planning as described in Chapter 2. Planning is aimed at finding a sequence of actions to be performed in order to bring about a certain state of the world. This is like answering the question "given that such-and-such conditions obtain in the world, what actions should be performed in order to make such-and-such state hold". The outcome of planning inferences are action rules of the kind "In order for state f to be the case, do action a at time t", which connect a goal f to an event token a at t.

The notion of clause has been specified in Section 4.1.3 as the CLP format for knowledge representation in databases. From the computational standpoint, a particularly significant kind of clauses are *goal clauses* or *queries* (henceforth used interchangeably), which ground CLP computations. Integrity constraints are a peculiar sub-category of goal clauses.

Goal clauses are formal counterparts of the psychological notion of human goals understood in the framework set forth in Section 2.2. Both maintenance and achievement goals are expressed formally as queries. Integrity constraints provide a suitable formal expression of maintenance goals (following the proposal in (Kowalski 2011)), be those meta-level social-epistemic motivations, or object-level desirable states of the world that need be achieved and maintained through time. Achievement goals are represented as regular goal clauses. Just like a goal sets one about constructing a plan for the desired state of affairs, goal clauses initiate computations aimed at a temporally ordered sequence of events that presumably bring it about. The syntactic expression of the computational aim is that goal clauses are articulated interrogatively, as a query regarding the actions that should be performed in order to make a predicated property true about the world.

Queries are interrogative clauses represented by the notation

$$?c, G_1, \ldots, G_m,$$

where $c \in \mathcal{K}, G_k \in \mathcal{L}$. The predicate that occurs most often in goal clauses is *HoldsAt* since it expresses that a property obtains at a given time.

CEU eTD Collection

The terms in goal clauses are ground terms, or universally quantified variables or Skolemized functions[10]. The former is the case for regular goal clauses resolved in concrete action planning and action explanation. Variables occur in goal clauses that function as integrity constraints; hence quantification issues arise mainly with respect to integrity constraint clauses.

In describing the inferential procedures of goal clause reduction, I focus on the most common case where the head is a $HoldsAt$ predicate. The relevant peculiarities of integrity constraints processing are articulated in Section 4.1.6. Section 4.1.7 puts some flesh on the logical notions by providing a concrete example of resolving a goal clause by means of the CLP inferential rules.

> Operationally, one should think of a query $?q$ as the assumption of the formula ¬q, the first step in proving q from P using a reductio ad absurdum argument. In other words, one tries to show that $P, \neg q \models \bot$. In this context, one rule suffices for this, a rule which reduces a goal to subgoals van Lambalgen and Hamm (2005, p.230).

Goal clauses are resolved by reduction to an assignment of constraints. The reductive strategy is called *resolution*. Goal clauses are reduced to subgoals and, eventually, to basic actions. Basic actions are expressed as satisfiable constraints on the argument structure of the predicates occurring in the bottom-most atomic query of a derivation, after all possible reduction has been performed given information in the program P. A set of satisfiable constraints make up the *resolvent* or the *computed answer substitution* of a query[11]. The constraints express what must be assumed in order to derive the empty query.

DEFINITION 8. Resolution *is a derivation rule which takes as input a program clause* $p_1, \ldots, p_n \rightarrow q$ *and a query* $?q$ *and produces the query* $?p_1, \ldots, p_n$.

Backwards chaining is the reasoning strategy most congenial to resolution, just as backwards reasoning is to planning. Goal reduction is based on the fact that when evidence of truth is encountered for an atom in the query, it is deleted. Assuming that the head of the query obtains, the truth of its conditions can be inferred in a minimal model of P. The head is thereby reduced to the body. Furthermore, upon failure to provide evidence for the positive atom in a negative literal, the negative literal is erased from the query, which is thus reduced to a simpler clause. The strategy is licensed by negation as failure. The syntactic manifestation of the closed-world assumption is the derivation rule of negation as failure, NAF. Positive literals can be eliminated from the query upon encountering in the logic program P evidence of their truth. Variables and functions are specialized to concrete constants, which occur as arguments of predicates in P clauses. Upon unification (see below) with formulas in P, atoms are deleted from

[10]Skolemized functions express an existentially quantified variable in terms of the value of a function of another universally quantified variable (to be explained in Section 4.1.6). This justifies talk of universal quantification across-the-line.

[11]This is a crucial difference between CLP and standard logic programming, where successful derivations terminate in an empty clause.

CEU eTD Collection

the query. Forward reasoning can be used to integrate by conjunction different results obtained during backwards resolution, and to derive $HoldsAt$ or $Clipped$ atoms by application of the *EC*.

DEFINITION 9. *Let P be a definite program.*

a. *A* query *is a formula of the form $?L_1, \ldots, L_k, c_1, \ldots, c_m$ where the L_i are literals and the c_j are constraints (equations between terms).*

b. *Let $?L_1, \ldots, L_k, e_1, \ldots, e_m$ be a query, and suppose L_i is a positive literal, $L_i = L_i(t_1, \ldots, t_l)$, such that $K_1 \wedge \ldots \wedge K_n \rightarrow L_i(x_1, \ldots, x_l)$ is in P. Via resolution, one may then derive the new query*

$$?L_1, \ldots, L_{i-1}, K_1, \ldots, K_n, L_{i+1}, \ldots, L_k, x_1 = t_1, \ldots, x_l = t_l, e_1, \ldots, e_m.$$

c. *A derivation is* successful *if its last line consists of satisfiable equations between terms only.*

Let $y_1, \ldots, y_v$ be the variables occurring in the top query $?\varphi$. If the derivation from $?\varphi$ is successful, the equations $y_1 = s_1, \ldots, y_v = s_v$ derived from the bottom query are jointly called the computed answer substitution *or the* resolvent *for $y_1, \ldots, y_v$ in $?\varphi$.*

The initial goal clause and the derived subgoals are *nodes* in a *derivation tree* (illustrated in Fig. 3 of Section 4.1.7). A derivation tree may have just one, or more *branches*. For example, if a query is unified with the head of a clause with disjunctive body in P computations branch out. Each disjunct is pursued in a separate branch starting from the goal clause where the disjunctive program clause is called for. It is as if the goal clause *may be* reduced to either of the disjuncts. Reductive computations are further on performed independently for each branch. This is in line with the idea that

> alternative branches of the search space [for goals] are represented as disjuncts (Kowalski 1995).

The following definitions are more detailed follow-ups for the notions introduced in Definition 9.

DEFINITION 10. *A branch in a derivation tree is* successful *if it is finite and ends in a query of the form $?c$, where c is a satisfiable constraint. A branch in a derivation tree is* finitely failed *if it ends in a query $?c, B_1, \ldots, B_m$ such that either c is not satisfiable, or no program clause is applicable to the B_i. Otherwise the branch is called* infinite.

For the explanatory computations supporting infants' performance in the head-touch task Definition 10 is supplemented with an integrity constraint, which places an additional demand on the explanatory success of query resolution. Only derivations with a single successful branch are accepted as satisfactory, i.e. successful, explanatory computations (see discussion of the 'no disjunction in conclusion' constraint in Section 2.3).

CEU eTD Collection

DEFINITION 11. *A query* $?c, G$ *is* finitely evaluable *with respect to a program* P*, if its derivation tree is finite, i.e. if all branches in a derivation tree starting from* $?c, G$ *end either in success or in finite failure.*

According to Definitions 10 and 11, for resolution to be successful the top query must be transformed in a finite number of steps into satisfiable constraints. Basically the idea behind CLP derivations is to arrive at the simplest possible result that makes the query successful. This result is a set of constraints on relations between terms, which, if the query is finitely evaluable and successful, can be substituted for the variables in literals that occur in the nodes of the derivation tree. The constraints arrived at via backwards reasoning from the top query represent what the reasoner must assume in order to derive the empty query.

DEFINITION 12. *An assignment of constraints is* satisfiable *with respect to a program* P *if it specializes a variable term to a ground term occurring in* P*, or if it expresses mathematical relations between terms consistent with the set of primitive constraints in* P*.*

Definition 12 captures the semantics of successful query resolution.

Suppose a query *?HoldsAt(f, t)* is successfully reduced to $?c$. This means that c is satisfiable given the program P, and if the substitutions suggested by c are performed then f obtains at t. Using the normative language of planning, c states equations between terms with which the database *should* be updated if the goal in the top query is to be achieved. Conversely, if a branch fails finitely and ends in $?c, B_1, \ldots, B_m$, no clause of P can reduce the bottom clause any further: the assignment of constraints is not consistent with the relations between the terms of P, and there is no clause in P whose head can be unified with B_i. It is worth noting that the notion of consistency as used in the definiens of constraint satisfaction is to be understood via negation as failure. That is, a constraint is not consistent with P when it contradicts information in P; otherwise it can be assumed to be consistent. For example, should a derivation end with the temporal constraint $t < s$, and should P contain a clause of the form *Clipped(s, f, t)*, then the assignment of constraints is not satisfiable because the program clause entails $s < t$.

The event or fluent argument of the predicate in the bottommost atomic query in a derivation stands for a behavior that should be executed in order to bring about the property f in the top query. Instrumental human actions or activities are like behavioral manifestations of the conclusions derived from the background beliefs, current observations, and current goals represented in the reasoning knowledge database organized as a logic program P. The final constraint $?c$ is a conjunction of all the constraints accumulated as assumptions throughout reduction. It comprises an equation between the constant expressing an action or an activity to be performed, and the event or the fluent variable occurring in the atom before specialization. It must also give a concrete temporal order, established by the relations $<$ or $\leq$, between time variables. The constraint thus specifies what must happen/ hold when for f to obtain. If these specifications are not inconsistent with information in P, resolution of *HoldsAt(f, t)* is successful.

CEU eTD Collection

Resolution is computed by inferential rules applied to a query $?G$, in a closed-world framework established by the rule of negation as failure. Given a logic program P, composed of the scenario for a situation, and a set of integrity constraints (with their resolvents), the reasoner attempts to determine an assignment of constraints that make the query formula G succeed (i.e. a satisfiable constraint) or fail (i.e. a non-satisfiable constraint) finitely. Resolution reduces G in a stepwise fashion to subgoals, all the way to their simplest expression as constraints on relations between terms.

Backwards derivations reduce queries by application of the *EC*, and of clauses in P. The latter has been explained relative to Definition 8. Some remarks are in order here with respect to the usage of the axioms in backwards reasoning. The *EC* provide two ways by which a fluent f in a top goal clause can be shown to hold at a time t (different than 0), through Axiom 3 and Axiom 4. That is, there are two informative causal principles in the *EC* to be used in backwards reasoning for reduction of *HoldsAt(f, t)* goal clauses to the events or activities that result in f. Axiom 3 expresses the principle of inertia for instantaneous causality, so it is computationally useful when the causal model in P provides information about an initiating event of f. It reduces f to an event that brings it about at some previous moment. Axiom 4 reduces f in the goal clause to some other fluent, an activity or a force, whose continuous exertion brings about the state. Hence it is useful when information about such an activity is part of the causal model of the situation. The choice of either one of the axioms is dictated by the kind of causal information (either instantaneous, or continuous change) available in the descriptive database of P.

The syntactic expression of closed-world reasoning, *NAF*, is an inferential rule based on the definition of truth. As such *NAF* expresses reasoning with the completion of logic program P. It licenses the reduction of the head to its body, by applying the closed-world assumption that no abnormality interferes between the truth of the premises and that of the conclusion. To show that $\neg p$ is a true atom, one simply fails to show that p is true (cf. Kowalski (2011, p.285)), because, for example, one has no observational evidence for, and/ or has no background knowledge about the truth of p. Unless there is positive information with respect to the termination or change of inertial status of a fluent f on the interval $[t, t^{'}]$ along the lines of Axiom 1, *NAF* allows the deletion of a $\neg Clipped(t, f, t^{'})$ predicate from the derivation tree.

Another CLP syntactic strategy involved in goal reduction is *unification* of query literals with literals in the head of program clauses. Unification implements a 'search and substitute' strategy for resolving goal clauses. It is based on *specialization* of variables or functions to constants occurring in P clauses, and *substitution* of constants for variables (or of narrower for broader variables). A query literal is conveniently unified with the head of a scenario clause in P, thus licensing the inference of its truth. In a closed-world framework the head reduces to the body. Consequently the literal is deleted from the query and replaced with the body of the P clause. Throughout goal reduction, variables in the goal clause formulae are specialized to constants occurring at similar positions in the literals of P clauses. Unification with literals in

CEU eTD Collection

P clauses and its prerequisite, the specialization of variables, ensures that the set of constraints accumulated throughout goal reduction is satisfiable given P (in the sense of Definition 12).

Informally, specialization can be characterized as a method to derive specific from general information. One way to do so is based on knowledge of taxonomical relations. A syllogism such as *All A are B, a is A, Therefore a is B* may be taken to provide an example of specialization based on taxonomical knowledge. The major premise predicates property B over the whole category of entities that have property A. The minor premise states that constant a is an instance of the category of entities that have property A. In the conclusion of the syllogism constant a is substituted for the general term *All A* in the major premise. This is precisely how specialization works, allowing the derivation of informative specific conclusions from more general knowledge. The syllogism can be expressed as a positive logic program P, $\{A(x) \rightarrow B(x), A(a)\}$. Suppose now that given this program, one resolves query $?B(a)$. The required procedure involves the specialization of variable x to constant a. Based on the program clause expressing the major premise, one can derive derive from the top query *?A(a), a = x*. Further, based on the minor premise, one can infer the truth of $A(a)$ and delete it. One is then left with a line that consists only of the identity equation between terms, the constraint *x = a*. The constraint is satisfiable given P, because of lack of positive evidence of inconsistency with P. The topmost query $?B(a)$ succeeds in a finite number of steps, according to the criteria in Definitions 10, 11, and 12.

Unification of literals amounts to adding to the derivation conveniently chosen, i.e. satisfiable constraints. They are composed of equalities between terms, which are satisfiable given P. The unification algorithm provided in Schöning (2008, pp.84-85) is applied to two literals, one in the goal clause and one in P, of which at least one contains variables (functions are included in the category of variables). It determines the substitution to be performed between the two literals so that they become identical. P is scanned. When a term a is encountered in a literal in the head of a clause, at a similar position with variable x in a goal clause literal, the substitution $[x/a]$ is effectuated. The two literals are unified by the substitution $[x/a]$. The unified literal is deleted from the query. Constraints like $a = x$ are accumulated throughout the derivation, and they partake in the computed answer substitution. In the head-touch task, unification is the main strategy used for goal abduction (formalized in Section 5.3).

A CLP syntactic procedure for forward reasoning is the application of the *EC* to conjuncted atoms among computational outcomes, or clause heads in P. The atoms may contain descriptive, i.e. *Happens(e, t)*, or causal information, i.e. *Initiates(e, f, t)*, *Terminates(e, f, t)*, *Releases(e, f, t)*, *Trajectory(*f_1, t, f_2, δ*)* about actions and their effects. By forward inferencing the reasoner can derive information about specific fluents holding at a particular time, or about fluents being clipped. That is, if one has obtained independently the conditions of Axioms 1, 3, or 4, one can join them together and derive the conclusions of axioms by forward reasoning. This amounts to predicating certain properties about the world (for the case of Axioms 3 and 4), or blocking the use of NAF over specific temporal intervals by deriving information about the non-persistence

CEU eTD Collection

of a property on the interval (for the case of Axiom 1). Suppose, for instance, that the reasoner obtained by previous backward derivation from a goal that a certain event e occurring at time t_1 – $Happens(e, t_1)$ – brings about the spatial location on of a certain object – $Initiates(e, on, t_1)$. If P contains no other on-relevant information on the interval $[t_1, t_2]$ NAF licenses the inference that on persists between t_1 and t_2 – $Clipped(t_1, on, t_2)$. By forward application of Axiom 3 to these three atoms one could then infer that $HoldsAt(on, t_2)$.

To sum up, the general to-do list for goal reduction via resolution looks as follows. The concrete example of derivations provided in Section 4.1.7 will clarify the logical technicalities.

Suppose reasoning starts from a generic, universally quantified goal clause $?G$, given a program P (a narrative and a causal model), and a set of integrity constraints. $?G$ is recurrently reduced to subgoals.

(1) specialize by means of taxonomical knowledge variables in goal clause (if any) to more concrete terms (either constants or more restricted variables);
(2) reduce achievement goal clause to conditions of axiom(s) whose head it is, or to conditions ;
(3) specialize argument variables conveniently, to make unification possible with scenario clauses;
(4) unify atoms with corresponding atoms in the scenario, thereby inferring their truth;
(5) delete true atoms from the derivation tree;
(6) infer truth of $\neg Clipped$ literals by application of the *NAF* inference rule: unless a scenario clause, or the resolution of an integrity constraint provides information about the fluent-related conditions of Axiom 1 (namely, termination, or release from instantaneous inertia), $\neg Clipped$ may be deleted;
(7) repeat steps (2) – (6) until none of them is applicable anymore given P;
(8) conjugate all the constraints accumulated throughout (1) to (7), and call the conjunction query $?c$ expressed in the language of constraints $\mathcal{L}$.

Using goal resolution as a computational procedure one can attain several outcomes, e.g., assign the status of goals to certain effects of perceived actions, derive appropriate action sequences for one's (achievement or maintenance) goals, abduce causes of observed states of affairs and thus assign plans to other observed agents. I propose that relative to the conceptualization of the reasoning processes involved in the head-touch task (cf. Chapter 3) the planning of young reasoners' own actions, as well as the justificatory plan recognition for an observed agent, are finely expressed in terms of query resolution.

4.1.6. The role of integrity constraints in computations. CLP computations are modulated by integrity constraints, which dictate or prohibit certain updates to the database prescribed by resolution of regular goal clauses. Integrity constraints impose local norms on computations, either as obligations, or as prohibitions. They call upon, or prohibit, certain computational movements. The role of integrity constraints in computations can be understood by analogy

CEU eTD Collection

with the sudoku game, where the presence or absence of certain numbers on the rows, columns, and blocks of the table, dictate or prohibit filling in the gaps in certain ways. For example, the constraint *no row contains more than one instance of a number between 1 and 9* prohibits my adding 2 on R1C3 if digit 2 is on R1C7. If on R1C1 to R1C8, I have all the integers between 1 and 9, save 5, the constraint that *each row contains all numbers from 1 to 9* dictates that I fill in R1C9 with number 5.

Goal clauses that express integrity constraints have at least two distinctive features, both of which are related to their normative status. First, conditional integrity constraint clauses have a different semantics than regular conditional clauses. The semantics of integrity constraints is a thorny issue and there is no consensus in the literature as of yet (Godfrey et al. 1998; Kowalski 2011). I opt for the formal semantics of the classical logic conditional because I think it aligns well with their normative status. The truth of the body necessarily entails the truth of the head of the clause. The *IF . . . THEN* formulation is used in the case of integrity constraints in order to differentiate them from conditional program clauses. Second, as mentioned in Section 4.1.3, causal predicates are likely to occur in integrity constraint literals. Variables usually occupy the slots of these predicates. For instance, infants' teleological stance with respect to action interpretation is represented in the model as an integrity constraint clause that contains the predicate *Initiates(a, g, t)*. It states the norm that "for every individuated action a, there is a physical goal g", or "$\forall(a, \exists g)$". It combines the default universal quantification over action variables with existential quantification over goal state variables. The combination of quantifiers, more specifically the existential within the scope of the universal, is handled by skolemization (cf. Schöning (2008)). In short, the existentially quantified variable g within the scope of a universal quantifier is replaced with the value of a function of the universally quantified variable, $f(a)$. The function $f(a)$ is called Skolem function. Only one variable is left in the integrity constraint clause, namely a, and it is universally quantified. This way universal quantification remains 'the only rule in town' for goal clauses, which makes possible a smoother trajectory of resolution. The method for the substitution strategy of variables with concrete instances, or constants from P, is given by the unification algorithm in Schöning (2008, p.84) detailed in Section 4.1.5 above.

Kowalski (1995) sets forth the view that integrity constraints stand for

> passive goals that become active when they are 'triggered' by appropriate inputs (Kowalski 1995, p.7).

This is in line with my suggestion that integrity constraint clauses are an appropriate formal ingredient to express maintenance goals. They regulate the database update process resulting from regular query resolution. They represent norms with which the resolvents of regular goal clauses must comply. In Section 2.2 I emphasized the normative, guidance function that activated maintenance goals have in computations. This is why the abstract representations of

CEU eTD Collection

goals are well amenable to formalization as integrity constraints. Properly speaking, maintenance goals are the consequent of a conditional integrity constraint clause, while the antecedent expresses its triggering conditions. They are activated by reasoning forward from evidence that their conditions obtain. The antecedent requires activation of the consequent, which the reasoner must make true in integrity constraint resolution, and maintain true while resolving regular goal clauses with respect to P.

Computations starting from integrity constraints have the same inferential structure presented in the previous Section: the notion of resolvent or computed answer substitution is identical, and it is arrived at via resolution as described above. The activation, or 'triggering' of the the integrity constraint conclusion by its conditions is an instance of forward reasoning. The formula in the conclusion is resolved by backwards chaining.

There is one peculiarity of integrity constraint clauses that gives a flavor of their normative status. Success or failure of a query are not matters to be decided at the computational end point, depending on whether the assignment of constraints is satisfiable or not given P (according to Definition 10). Rather, success and failure are requirements that accompany explicitly the query. They prescribe that the final assignment of constraints be either satisfiable or not, and consequently they determine the reductive computational strategy applied in order to arrive at an assignment of either kind.

DEFINITION 13. *Let $R, R', R'' \ldots$ be a finite set of constants each denoting a reference time; these constants belong to the constraint language. An* integrity constraint *is a statement of the form*

($\dagger$) *IF $?\varphi$ succeeds THEN $?\psi(R, R', R'' \ldots)$ succeeds/fails, or*

($\dagger'$) *IF $?\varphi$ fails THEN $?\psi(R, R', R'' \ldots)$ succeeds/fails,*

where φ and ψ are formulae of the event calculus.

The operational meaning of ($\dagger$) is that if the scenario satisfies φ, the goal $?\psi(R, R', R'' \ldots)$ must be made *to succeed, or to fail finitely. To determine whether the scenario satisfies φ, one has to investigate whether the goal $?\varphi$ succeeds.*

An integrity constraint required to succeed dictates certain updates to the database. It has the form *?G* `succeeds`, and it reads as "make it be the case in the situation model that G". A prohibitive clause looks like *?G* `fails`, and it reads as "make it be the case in the situation model that G is not the case". The former are represented in conditional form as

$$G \rightarrow \top,$$

and the latter as

$$G \rightarrow \bot.$$

Integrity constraint resolution is the computational way to "make it be the case" in the sense introduced in Definition 13. Resolving

$$G \rightarrow \top$$

CEU eTD Collection

amounts to finding an assignment of constraints that make the top query G successful. It is noteworthy that the requirement of `success` works like an existential quantifier. Resolution is called to show that there exists such an assignment of constraints, satisfiable given P, that make G obtain.

> The requirement that a query such as *?Happens(e, R), R* $<$ *s* `succeeds`, entails that the scenario to which the constraint is added must be updated in order to ensure success. Together, the scenario and the axioms of the event calculus strictly regiment the update process, so that always a minimal extension is chosenvan Lambalgen and Hamm (2005, p.118).

Resolving prohibitive queries like

$$G \rightarrow \perp$$

is equivalent to finding a set of constraints applicable to some of the formulae in G which make it that G does not obtain irrespective of how the other terms are grounded. The requirement that query $?G$ `fails` thus works like a universal quantifier over G.

For the task at hand integrity constraints express infants' maintenance goals, whether higher-order socio-epistemic motivations and intentions, or simple physical goals that need to be maintained throughout time; they are spelled out in Section 5.2.1. Integrity constraints play a crucial role in choosing (in the case of obligations) or eliminating (in the case of prohibitions) candidate explanations for observations (during context interpretation), or candidate actions (during planning). An example of an epistemic integrity constraint on interpretation is the teleological stance, which sets children on assigning goals to actions.

4.1.7. A concrete example of CLP derivations. Let us now apply the representational format and the derivational mechanism spelled out throughout this chapter to the example introduced in Chapter 2. Imagine an agent who desires financial well-being. It is a maintenance goal that she is committed to at all times. It motivates her to derive its concrete content, i.e. a goal to be achieved at a particular time, and subsequently to act upon it.

Let P be the definite constraint logic program that formalizes the knowledge database for this situation. It is composed of a scenario, goal clauses to be resolved, and integrity constraints.

The scenario contains a narrative of events, and a causal model. The narrative specifies the fluent variables $well_being$ and PhD, and the time variable $deadline_{submission}$. The causal model states that the performance of a writing activity for a duration of three months results in thesis submission, and that an event type *start_writing* brings about a writing activity at a variable time t. The maintenance goal for financial well-being is represented as an integrity constraint $i.c._1$. $I.c._2$ prohibits that, if P contains a clause such as $HoldsAt(writing, t)$ (as a result of computed updates), it is updated with information about a different property f holding

CEU eTD Collection

simultaneously, at t[12]. The second constraint prescribes that writing is something of an "all or nothing" activity. Variable t is universally quantified within the integrity constraint clauses: financial well-being is sought at all times, whenever *writing* obtains f does not.

The knowledge database looks like:

- Narr 1: *well_being = PhD*;
- Narr2: *PhD = thesis_submitted*;
- Narr 3: $deadline_{submission}$ = *May*.
- CM 1: *Trajectory(writing,now,thesis_ submission,3m)*;
- CM 2: *Initiates(start_writing,writing,t).*
- I.c. 1 *?HoldsAt(well_being, t)* `succeeds`;
- I.c. 2 *?HoldsAt(writing, t), HoldsAt(f, t),* $f \neq$ *writing* `fails`.

Computations start with the evaluation of the integrity constraint $i.c._1$ relative to the scenario. It seeks satisfiable constraints to be added to P, so that the success of the query is granted. Resolution outputs an achievement goal, which is further reduced to executable actions. Fig. 2 and fig. 3 illustrate the processing steps. The two figures are examples of derivation trees used as a graphical depiction of CLP computations.

Given information in the narrative, a concrete instance of the agent's maintenance goal is to submit the PhD thesis in May (fig. 2). It becomes her achievement goal. She computes a plan that presumably achieves it (fig. 3) to the best of her knowledge.

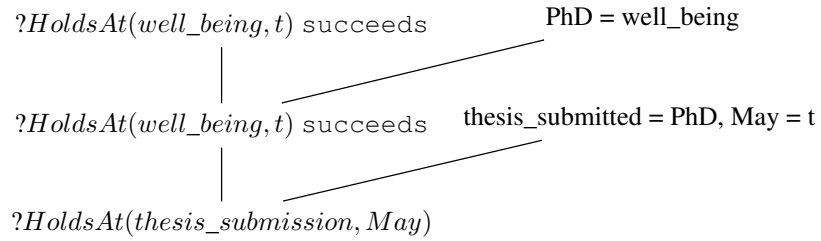

Figure 2. Reduction of maintenance goal represented as an integrity constraint.

The integrity constraint $i.c._1$ is resolved by using the equation between terms in the narrative. Forward reasoning helps the agent to move forward towards action performance by replacing an abstract desire with an undefined, universally quantified time window, with a concrete instance of it, namely, the submission of the thesis in May. The constraint is satisfiable given P, and resolution of the integrity constraint is successful if the corresponding formula *HoldsAt(thesis_submission, May)* is added to P. The integrity constraint to be wealthy at all times reduces to the obligation to update the database with a concrete achievement goal of submitting

[12]The reader is asked to accept this for the sake of the argument. Of course $writing$ and $thinking$, for instance, are not just compatible but even advisable. The integrity constraint can be made more specific by specifying f as *activity incompatible with writing*, where an instance of f is $going_out$.

CEU eTD Collection

the thesis in May. Thesis submission in May is a subgoal of a lifetime of financial well-being. It is a goal clause to be resolved.

The agent is after a chain of actions that make the state *thesis_ submission* true about the world at time point *May*. To this end she engages in resolving the query *?HoldsAt(thesis_ submission, May)*.

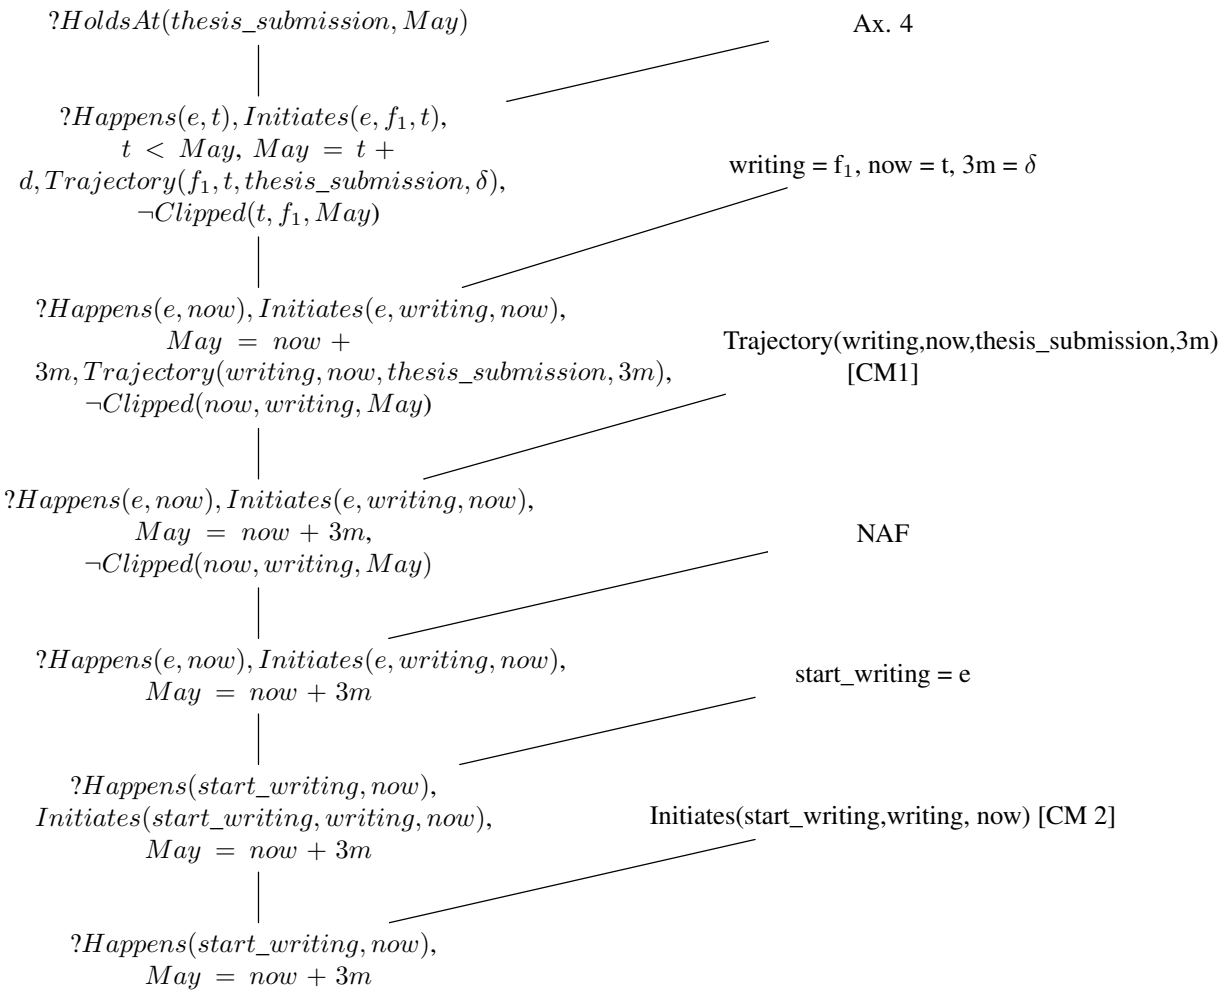

Figure 3. Reduction of achievement goal represented as goal clause.

Because the knowledge database contains information about the state *thesis_ submission* coming about due to continuous exertion of a force such as *writing* (CM 2), the reasoner uses Axiom 4 to reduce the query. Adding the constraints *writing* = f_1, *now* = *t*, and *3m* = δ grants the substitutions $[writing/f_1]$, $[now/t]$ and $[3m/d]$ performed in step (2) of the derivation. Upon unification with the atomic formula in the causal model clause *Trajectory(writing,now,thesis_ submission,3m)* (CM 1), the *Trajectory* literal is deleted from the derivation tree in step (3). Further on, *NAF* licenses deletion of the $\neg Clipped$ predicate in step (4). That is, under the closed-world assumption that the writing activity is continuous for a duration of three months between now and May, the truth of $\neg Clipped$ on the temporal interval $(now, May]$ can be

CEU eTD Collection

inferred. Hence it is deleted it in step (5). The clause CM 2 in the causal model provides information about a type of action that brings about writing, something like a starting event (be it turning the computer on, or pressing the computer key corresponding to the first letter of the text, etc., depending on how fine grained a plan I am after). The event variable e is specialized to *start_writing* in step (6). This makes possible unification of the *Initiates* formula in the query with the clause CM 2, and thereby its deletion. The clause *?HoldsAt(thesis_ submission, May)* is successful if P is updated with a start writing event at time now.

It is worth noting that reasoning forward with Axiom 3 from the 4th node in the derivation leads to *HoldsAt(writing, t)*, $t \in (now, May]$. This means that the literal *HoldsAt(writing, t)*, $t \in (now, May]$ updates P; let us call the extended program P'. From then onward $i.c._2$, *?HoldsAt(writing, t), HoldsAt(f, t), f* $\neq$ *writing* `fails`, prohibits the addition to P' of another literal such as *HoldsAt(going_out, t)* on the same temporal interval $(now, May]$. Suppose the agent had another maintenance goal, to have fun at all times – *HoldsAt(have_fun,t)*. This goal cannot obtain, given P', for any $t \in (now, May]$. $I.c._2$ prohibits it.

Let us summarize the computations meant to find a way to satisfy the agent's desire for financial well-being given P. The query *?HoldsAt(well_being, t)* is finitely evaluable with respect to P. In fact the derivation tree must be finite, as the query is required to `succeed` (see Definition 11). It has only one branch, which is successful according to Definition 10: it ends with a set of satisfiable constraints with respect to P. The computed answer substitution is $start_writing = e$, $now = t$, and $now + 3m = May$. The constraints are satisfiable because the equalities have been selected by a 'search and substitute' strategy over P. The upshot of the resolution of *?HoldsAt(success, t)* `succeeds` is an action rule. It states that for financial well being, the agent should start writing the thesis immediately and not cease the activity (much as she also wanted to have fun, for instance) for the next 3 months.

4.2. Cognitive relevance of CLP.

In this Section I recapitulate the information presented so far throughout the Chapter. I thus highlight the applicability of CLP as a modelling tool to human cognition.

Constitutive features of the formalism composed of the event calculus language and CLP, furnish background reasons for its use in cognitive modelling. First, the expressive capacity of the event calculus is fairly good. Its causal predicates formalize two notions of cause, instantaneous and continuous (van Lambalgen and Hamm 2005). $Initiates$ and $Terminates$ as expressions of instantaneous change are just what is needed to form causal representations with the principle of contact causality – a form of instantaneous causation. $Trajectory$ expresses continuous causation, by linking a fluent to a force whose exertion for a certain duration brings about the fluent. It is a necessary ingredient to represent observations of temporally extended change (e.g., activities that cause changed states of the world). The two notions of cause are essential for infants' causal processing (according to the review in Section 1.1). Therefore it is

CEU eTD Collection

justified to believe that the formalism can capture the understanding of the world's dynamics underlying performance in the head-touch tasks.

Further on, the fact that CLP embodies closed-world reasoning provides a number of reasons for its cognitive relevance more generally.

For one, the non-monotonic consequence relation specific to logic programming with the closed-world assumption is computationally efficient (van Lambalgen and Hamm 2005). It is less complex than other formalisms (e.g., classical logic) which have been shown to be inappropriate for formal descriptions of some forms of human reasoning (e.g., Byrne (1989); Wason (1968)). Moreover, logic programming is also less complex than other non-monotonic logics, such as Reiter's default logic (Reiter 1980) because it does not involve consistency checks (van Lambalgen and Hamm 2005). The logical notion of reduced complexity implies that, as a computational tool, it places lower demands on the processing capacities of working memory. Consequently it serves well cognitive modelling endeavors.

Relatedly, the system's semantics of negation as failure provides a good formalization strategy for human agents' reasoning toward an interpretation of observed behaviors because it demands that the computed model of the observations be minimal, defined by a suitable completion (as detailed in Section 4.1.4). Minimal models provide an appropriate environment to represent interpretive and planning computations of entities with limited processing and storage resources. The computations themselves impose minimal demands on reasoners because CLP's main inferential strategy, resolution, is grounded by the syntactic implementation of negation as failure.

Finally, the cognitive relevance of the event calculus with CLP has been demonstrated explicitly from a variety of perspectives: successful applications to discourse interpretation (van Lambalgen and Hamm 2005), formalization of cognitive tasks (Baggio et al. 2008), which facilitated the derivation of predictions with respect to autistic subjects' performance on those tasks (Pijnacker et al. 2009), or appealing implementations in neural networks (van Lambalgen and Stenning 2008). The first example is particularly relevant because it exploits precisely the use of human cognitive capacities for planning in the construction of temporal meaning; thereby a cognitive formal semantics of tensed speech is provided based on a formalism for planning. In Pijnacker et al. (2009), by developing ideas introduced in van Lambalgen and Stenning (2008), it was shown that different reasoning patterns between people with autism and normal controls are to be expected upon abstracting the logical form of the reasoning task using such logical formalization methods. The hypotheses generated by formal models in CLP were confirmed in behavioral studies.

I now zero in on the relevance of logic for the human phenomenon of observational learning.

4.2.1. Logic as a formal expression of learning in the head-touch task. By learning I mean the generalization of instrumental actions and achieved end-states from a teleologically structured observed context (the definition set forth in Section 2.6).

CEU eTD Collection

There are three main reasons why CLP and the event calculus are appropriate tools for a principled model of observational learning. First, the ontological commitment of the event calculus to types of events, activities, and states, fosters generalizability of their episodic representations (i.e. in the current context of observation), and hence learning about them.

Second, the motivation to acquire practical knowledge from observed agents can be represented as an assimilationist epistemic maintenance goal (as argued in Section **??**). In Section 4.1.6 above I have shown that CLP fares well in expressing maintenance goals as integrity constraints on the processing of observations, given a certain set of background beliefs. Section 5.2.1 sets forth the concrete proposal of representing maintenance goals, among those the epistemic motivations, in the format of integrity constraints. The suggestion aligns with the view expressed in Reiter (1988) that integrity constraints are epistemic in nature, standing for true statements about what the database to which they apply knows.

Third, in the proposed conceptualization of the reasoning that underpins learning (set forth in Chapter 2), planning algorithms are central. Action explanation is computed by means of planning derivations. More, the measure of learning is infants' behavior in the test phase. It is the result of backwards planning from the epistemic maintenance goal activated upon interpretation. All in all the processing that underlies learning recruits the observers' planning capacity to fulfill their explanatory drive. It is thus crucial that the chosen formalism provides an accurate description of the planning capacity recruited for explaining observations, as well as for bringing about infants' own goals. CLP fares best in this respect.

The three reasons recommend the formalism of choice as a suitable representational and computational environment for learning practical information from observation of other agents' actions, based on inferential processing of these observations in the context.

Observational learning exhibited by infants in the head-touch task has a twofold manifestation: as emulative learning of novel artifact functions (i.e. end-states that can be achieved through action upon artifact), or as imitative learning of novel means actions. Each of them is an instance of the epistemic maintenance goal. Maintenance goals are expressed as conditional integrity constraint clauses, whose antecedents are the set of triggering conditions of the motivation stated in the conclusion. The representation of infants' epistemic motivations as conditional integrity constraints provides a principled way to discern between the two: different conditions that obtain activate different goals. The conditions refer to the outcome of action interpretation, and to the socio-communicational interpreted context. Consequently, different outcomes of the explanatory processing of observed behaviors in the particular communicational context of demonstration, activate either the learning about functions, or the learning about means actions. Integrity constraints ensure consistency of the learnt information (either a novel function, or a novel action) with the background knowledge.

The representation of the epistemic goal *to learn* as an integrity constraint on reasoning ensures the link between learning and action performance. As such it justifies from a formal perspective the interpretation of infants' behavior in the task as manifestations of learning. The

CEU eTD Collection

information represented in P either as observed facts, or as updates based on interpretive inferences, makes true the antecedent of the integrity constraint. Once this is the case, it becomes the task of the reasoner to make true that learning as represented in the conclusion obtains, i.e. that *HoldsAt(learn, t)* `succeeds`. This is done by deriving an achievement goal and further reducing it to executable actions along the lines described in Section 4.1.5. Hence the reasoner acts upon learning, and ultimately exhibits learning in action performance.

The conceptualization of learning as a fluent that obtains in the conclusion of a conditional integrity constraint in CLP is beneficial for yet another reason. That is, its conditions can formalize all in the same key information about observed states of affairs, be those physical (e.g., change of physical position) or mentalistic (e.g., communicative attitudes, emotions), inference-laden observations (e.g., goals), as well as outcomes of previous computations (e.g., abnormalities, success or failure of explanatory attempts). All these properties read as fluents in the language of the event calculus. This allows a uniform handling of a wide range of conditions whose variability might make one skeptical with respect to the possibility of logic to model the complexity of factors that lead to fulfillment of the human epistemic motivation. Moreover, parameterization of the fluent *learn* permits to differentiate the two instances of epistemic goals that I suggest to be involved in the task under scrutiny.

I emphasized throughout this chapter the manifestations of closed-world reasoning in the formalism. Now, it is easy to see how such reasoning is called for when planning or explaining in terms of goals, given the need to overcome the combinatorial explosion of thoughts in order to decide how to act upon goals (according to Chapter 2). However, formalization of learning in a closed-world system appears at least prima facie paradoxical. The closed-world assumption is, after all, an auto-epistemic assumption (cf. Kowalski (2011)). The reasoner sets out interpretative computations assuming that (s)he knows everything there is to know about the context where the observed agents perform default actions in order to achieve their goals. In other words, nothing abnormal is the case. Should an agent be observed acting differently, interpretation proceeds by searching for contextual abnormalities that prohibit her default action. The beauty of a non-monotonic formalism is precisely that it allows for flexible use of closed-world assumptions. After all, paraphrasing Kowalski (2011), an evolving agent who learns needs to have an open mind, at least to some extent. The assumptions are endorsed initially in order to support reasoning towards a minimal interpretation, or the construction of a teleological model in which observations can be accounted for. Positive evidence about abnormalities overrides the assumption, providing the condition of possibility for updating the knowledge database with new facts.

Abnormalities can be described at several levels. For one, there can be physical abnormalities that block the conditional connection between actions and their goals. For instance, my having to stir the frying pan now (because otherwise the food would get burnt) may prevent my answering the phone now. Or, related to the example in Section 4.1.7, my having to use all my time for writing the thesis (because otherwise I would not be successful in the longed-for

CEU eTD Collection

academic career) may prevent me from going out as much as I feel like. Such abnormalities can justify odd behaviors, such as my not answering the phone, or my rejecting all party invitations for a certain period of time. In these cases, abnormality can be accurately represented formally as the fluent *abnormality*, and its conditions can be expressed in the consequent of an integrity constraint as conditions of $HoldsAt(abnormality, t)$. My contention is the following: failure to find physical contextual abnormalities that might justify odd behaviors is one of the preconditions of the more computationally complex learning of alternative instrumental actions. Lack of positive evidence with respect to abnormalities may trigger learning of contextually unjustified behaviors, despite availability of the learner's default action for the physical goal.

As mentioned in Section 4.1.1.1 the standard abnormality predicate of the event calculus is $Clipped$. Axiom 1 defines a fluent f being *clipped* by reference to the occurrence of e, which is either a terminating event, or an event that changes the inertial status of f. The closed-world assumption (represented at the syntactic level of CLP as the rule of *NAF*) licenses the inference that no such (abnormal) event occurs in the time interval of interest, and hence that $\neg$ *Clipped($t^{'}$, f, t)*. However, the Axiom provides a principled way in which the assumption can be over-ridden: the application of *NAF* may be blocked by explicit evidence of f-relevant events. There is something abnormal about fluent f, i.e. it does not persist, when it is either terminated or released from the inertia of instantaneous change. However the event calculus provides a practically unlimited number of ways to express abnormalities than the predicate $Clipped$. Different kinds of abnormalities (e.g., teleological abnormalities, such as the impossibility to pursue several goals that require physical presence at different locations simultaneously) can be expressed as fluents under the scope of $HoldsAt$ predicates.

The suggestion set forth in Chapter 2 was that action rules be represented by a non-monotonic conditional of the kind "*behavior*$_{known}$ and *nothing abnormal is the case* FOR *goal*". The logical expression is given by a CLP conditional clause like

$$\textit{HoldsAt(behavior}_{known}\textit{, t)} \wedge \neg \textit{ HoldsAt(ab, t)} \rightarrow \textit{HoldsAt(g, t)}.$$

This means that, unless something is abnormal, a goal-directed behavior known to be appropriate for achieving the goal, will in fact achieve it. Now, suppose that one's reasoning database represented as a program P contains information that the default behavior $behavior_{known}$ is terminated by the occurrence of some other event before it brings about the goal. Forward reasoning with Axiom 1 leads to the conclusion that *Clipped($t^{'}$, behavior*$_{known}$*, t)*, where t is the time when the goal should be achieved, and $t^{'}$ is a previous moment. This reads as "there is something abnormal of the behavior relative to its bringing about the goal at t", or *HoldsAt(ab, t)*. Still, the reasoner wants to attain the goal. Given that there are abnormalities with respect to the known behavior *behavior*$_{known}$ in the current context, she needs to recompute how to achieve the goal and come up with an alternative action rule. That is, she needs to plan for the goal on a different route, which does not involve a clipped behavior, and for which *NAF* can be safely applied: if the behavior *behavior*$_{new}$ obtains, presumably so does the goal. If her

CEU eTD Collection

database of motor knowledge does not contain any such *behavior*$_{new}$ then P might be updated to P' based on teleologically processed observations.

Finally, I briefly address the possibility to choose an alternative logic for formalizing learning behavior. The applicability of Inductive Logic Programming, for instance, to belief and goal revision as an

> essential activity for a real agent interacting with the real world (Kowalski 2006)

is well documented in machine learning approaches to areas such as molecular biology, bioinformatics, or natural language (Muggleton 1999; De Raedt 1997). On the other hand the domain of relevance of CLP has not, so far at least, covered explicitly belief revision or learning. So then why not, one may pertinently ask, opt for the former, as an acknowledged tool for formal models of learning.

I do not appeal to the machinery of Inductive Logic Programming because currently induction is beyond the scope of the model. Learning is diagnosed by action performance in a single occasion – in the test phase, which is the consequence of a single observation during the experimental phase. The context of performance is identical with that of observation, save the actions performed by the agent. Notwithstanding its importance the suggested formal model in the current format does not address head on the issue of generalization of the acquired bit of practical knowledge, the extent or the conditions of generalization. It may well be that a system for inductive reasoning is insightful in this respect. Consequently it is a highly relevant direction for future research the extent to which a move towards Inductive Logic Programming, or Bayesian probabilistic inferences may yield new predictions with respect to infants' acquisition of generalizable practical knowledge from observation of other agents. For instance, once CLP provides a clear, straightforward formal understanding of how the inferential processes of interpretation and planning may result in knowledge updates, Inductive Logic Programming should be called for to model the appropriate way to frame the scope of applicability of those behaviors. Similar considerations apply to probabilistic Bayesian formalisms (Rao et al. 2004).

CEU eTD Collection

CEU eTD Collection

CHAPTER 5

The Model II: The Reasoning Formalized

This Chapter depicts the full-fledged formal model of the reasoning steps that result in learning about the functionality of, and about the instrumental actions onto the novel artifact introduced to infants in the head-touch task. It draws on the logical tools described in Chapter 4. I focus on the Communicative context of presentation because it calls for more complex reasoning strategies. It provides a good illustration of the 'single-mechanism multi-level' kind of teleological processing that I set forth as the underlying cause of infants' observational learning. The discussion is led off in Section 5.1 by introducing the notation to be used in the formal model. The notation is implemented in the three databases that reasoners use for teleological interpretation (Section 5.2), offline planning (Section 5.4), and for online planning (Section 5.6). The interpretive reasoning processes are performed over the first two of these knowledge databases. The queries resolved are questions along the lines of *Why does she do* ϕ*?*, *What should she do in order to* φ*?*, or *What should I have done in order to* φ*?*. They are formalized in Sections 5.3, 5.4, and 5.5. The crucial functional aspect of teleological interpretation is that it is meant to render the agent's light-box directed behavior intelligible to the observer. Interpretation influences young reasoners' online planning computations in the test phase of the experiment by differential activation of the their assimilationist epistemic drive. Section 5.6 shows how planning computations prescribe actions that evidence different kinds of learning in the HandsOccupied and HandsFree conditions. In Section 5.7 I explain why in the absence of ostensive communicative cues from the observed agent in the Incidental setting, infants exhibit only learning about the novel artifact's functional properties, but not about its action affordances. Finally I give a concrete example of related empirical findings that the current formalization explains. Section 5.8 speaks for the generalizability of the model.

CEU eTD Collection

5.1. The ontology of the head touch task.

The list below comprises the terms used in the logical model of reasoning. The terms are notation conventions for the actions, activities, and states observed in the demonstration, or that can be inferred from aspects of it. In previous Chapters I referred interchangeably to the actions and behaviors performed by infants or by the experimenter, and I hardly ever used 'actvity'. In the context of the ontological structure of the event calculus, however, the distinction between actions (events) and activities is relevant in respect to the kind of causal relations that they are part of (as explained in Section 4.1.1). A plan is composed of instantaneous actions and continuous activities; when the distinction between the two sorts does not matter, I use 'behavior' as

an umbrella term. 'Actions' and 'events' are used interchangeably. Two actors perform behaviors: the experimenter in the demonstration phase (referred to as 'the agent'), and the infants in the test phase (for simplicity of expression, henceforth throughout this Chapter I use the generic singular). The agent fluents and events occur as arguments of the predicates that represent infants' observations, both in raw (e.g., agent touches light-box with head) and processed (e.g., agent teaches me head-banging) form. They constitute the observational database used for teleological interpretation of the context. The infant fluents and events are used for planning own actions, in offline and online planning computations. They form the database of actions to be performed by the infant.

Because the model formalizes an imitation study, the target novel behavior (i.e. the light-box directed bodily movement, and its culminating point – contact) has a twofold, symmetrical notation, when performed by the agent or by the infant. The agent activity fluent a corresponds to the infant activity fluent c. The agent contact action l corresponds to the infant contact n. As stated at the outset in Chapter *Introduction and road map* the reasoning processes are modelled under the assumption that they supervene on the mapping of the agentive relata of observed behaviors onto the observer's corresponding motor program; for instance the observed activity a_{head} is mapped onto the infant's motor program to c_{head}.

Actions and fluents may be represented with different degrees of coarseness. I explain this for the case of contact, with the specification that the same applies to the light-box directed bodily movement. A contact event has at least two arguments specifying the agent and the patient of the action – *contact(who, what)*, but it can have more, for instance *contact(who, what, how)*. Henceforth I will omit from notation the *what* slot. In the given context the action is artifact directed, so the direct object argument is assumed to be filled with the individual object constant *light-box*. Moreover, because I use different notations for the experimenter's and child's contact actions (l and n, respectively), the *who* slot is implied in notation. The additional indirect object argument *how* stands for the particular manner in which the action is performed; for our purposes, a more accurate expression is *with what*. It refines the agentive *who*; it is the local agentive relatum slot. A fine-grained representation of contact is *contact(global agent, patient, local agent)*. The local agent is represented by notation as a parameter attached to the contact action – e.g., $contact_{head}$. In Csibra (2008), identifying the agent of an action is the initial step in the process of goal attribution. This identification can be coarse, when the agent is global, i.e. the observed human being, or fine-grained, when the agent is a local agentive relatum, e.g., hand, head. Fine-grained representations are basic behaviors in the sense of 'basic actions' coined by Danto and Morgenbesser (1963) (introduced in the preamble to Chapter 2). Planning computations are aimed at such fine-grained representations of behaviors. On the last line of planning derivations, the agentive relatum must be specified, in order to render behaviors executable by activation of the corresponding motor program.

CEU eTD Collection

(1) temporal constant now: *temporal index for fluents occurring in top queries that ground one's own planning computations – 'the fictitious now of planning'(van Lambalgen and Hamm 2005)*;
(2) agent fluent p: *agent exhibits object-directed ostension, and child-directed communicative behavior*;
(3) agent fluent d: *agent exhibits behavioral discomfort for feeling cold*;
(4) agent fluent $change_d$: *agent is engaged in changing distressing state*;
(5) agent activity fluent m_{hand} : *agent manipulates scarf*;
(6) agent action k: *agent knots scarf*;
(7) agent activity fluent a: *coarse representation of the agent performing light-box directed bodily movement*;
(8) agent activity fluent a_{head}: *fine-grained representation of the agent performing light-box directed movement with the head*;
(9) agent contact action l: *coarse representation of the agent touching light-box*;
(10) agent contact action l_{head}: *fine-grained representation of the agent touching light-box with head*;
(11) agent contact action l_{hand}: *fine-grained representation of the agent touching light-box with hand*[1];
(12) infant activity fluent c: *coarse representation of own light-box directed bodily movement*;
(13) infant activity fluent c_{hand}: *fine-grained representation of own light-box directed movement performed with hand*;
(14) infant activity fluent c_{head}: *fine-grained representation of own light-box directed movement performed with head*;
(15) infant contact action n: *touch of the light-box*;
(16) infant contact action n_{head}: *fine-grained representation of own touching of the light-box with head*;
(17) infant contact action n_{hand}: *fine-grained representation of own touching of the light-box with hand*;
(18) light-box fluent e: *light state of the light-box – on (e) or off ($\neg e$)*;
(19) spatial location fluents are expressed as a combination of two individual objects assumed to be linked by spatial preposition on – the first object is on the second one: e.g., *light-box_table, hand_table, scarf_on* stands for *scarf_agent*;
(20) initiating event of an activity $start_f$: *the 'first motor' of activity f*;
(21) goal state variable g^{f}: *resultant state of an instrumental behavior f that justifies f*;
(22) teleological abnormality fluent $ab_{ag.g}$: *unusual state of affairs in the current context with respect to the involvement of agentive relatum ag in goal g*;

[1]Note that in the experimental setting of Gergely et al. (2002) the experimenter does not actually contact light-box with hand.

CEU eTD Collection

(23) agent teaching parameterized fluent *teach*$_f$: *agent intends to teach* f;

(24) infant learning parameterized fluent *learn*$_f$: *learn* f *from observed agent*;

(25) unsatisfactory explanation fluent $failure_{f.g}$[2]: *teleological explanation of behavior* f *with respect to physical goal* g *fails because the goal does not justify the action in its context of occurrence*.

Most of the terms in the notation list are constants standing for observed actions, activities, and states. They partake in the ground formulas of the scenarios. Temporal constant now is used as an index of the infant's planning its own actions, both offline and online. When planning from the perspective of the agent (i.e. in plan recognition), the temporal index of the goal queries is the actual moment of observation represented in the scenario by real numbers. Variables are listed at the end – items (20) to (25). Variable $start_f$ represents the initiating event of one's own activity, or of an observed activity. It occurs in scenario clauses that should be understood as existential claims with respect to a type of initiating events of activities, derived from commonsense inertia. No claim is made that reasoners actually represent the concrete event token. For this reason the temporal index is given as a variable ranging over a certain interval – all activities have a 'first cause' type of action, whose moment in time is not strictly fixed to an instance. The other four g^{f}, $ab_{ag.g}$, $teach_f$, and $learn_f$ occur as arguments of the predicates in integrity constraint formulas. Their meaning is specified by resolution of the constraints. They are parameterized fluents, and the parameters themselves are variables. Variable g^{f} stands for interpretation-laden resultant state inferred to be the physical achievement goal of behavior f (an action or an activity) performed by the agent. It reads "goal g of behavior f". Variable $ab_{ag.g}$ stands for teleological abnormality; according to the notion of 'abnormality' introduced in Section 2.5 it constitutes positive evidence of violation of the closed world assumption CWA_{ab}. It has two parameters – ag standing for local agentive relatum, and g standing for goal. It reads "abnormality of an agentive relatum ag with respect to goal g". The fluents *teach*$_f$ and *learn*$_f$ are also variables. Parameter f is the thing taught, or to be learnt, respectively; it is the object of the diadic pedagogical situation. According to the conceptualization of learning in Section 2.6, f may be either a behavior performed by the agent and observed by the infant, or a resultant state achieved by the agent and observed by the infant; its sort is not fixed.

CEU eTD Collection

5.2. Database for teleological interpretation in a Communicative context.

The database takes the infant's perspective of observations in the HandsOccupied and Hands-Free conditions, respectively, in the Communicative setting. I assume that scenario clauses represent what the infant observes and what he knows about the causal structure of the world; integrity constraints express the generic principles that guide his interpretive endeavor. The representations in the database are used for goal attribution (Section 5.3), and for explaining

[2] *Failure* is a meta-fluent; its reference is the result of explanatory computations.

observations through the assignment of action plans to the agent (Section 5.5). The database for offline planning is listed in Section 5.4.

5.2.1. Integrity constraints on teleological interpretation. As explained in Section 4.1.3 the predicates in integrity constraints (i.c.'s) normally take variables and Skolemized functions as arguments. The first two i.c.'s in the list, as well as constant $hand$ in $i.c._4$ and $i.c._5$, are exceptions. The variables and functions are specialized to the appropriate constants among representations of current observations and prior knowledge, allowing unification of literals in the i.c.'s with ground literals in the scenario, and thereby their application to the current context. x is a variable ranging over observed instrumental (goal-directed) actions, y over instrumental activities, z over states, loc_k over location states, and g over physical achievement goals. Goals g need not be ultimate goals, they can also denote subgoals which are further reduced to basic motor acts. When the distinction between activities and actions is not relevant I use variable f, meaning action or activity (non-exclusive disjunction). Variables s, t, and r stand for time points, and δ for temporal duration.

Save for the first one, the norms expressed by integrity constraints are also used in the Incidental conditions; they govern computations in all four experimental setups described in Section 1. Sections 5.3 to 5.6 describe their role in the two Communicative conditions – HandsOccupied and HandsFree. Section 5.7 elaborates on what changes, given these constraints, when agent fluent p is not observed in the Incidental experimental setups.

(1) IF $?Initially(p)$ `succeeds` THEN $?HoldsAt(teach_f, t)$ `succeeds`;

(2) IF $?HoldsAt(d, 1)$ `succeeds` THEN $?HoldsAt(change_d, t), 1 < t$ `succeeds`;

(3) $?Happens(x, t), Initiates(x, f(x), t), g^{x} = f(x)$ `succeeds`
OR $?HoldsAt(y, t), Initiates(y, f(y), t), g^{y} = f(y)$ `succeeds`;

(4) Abnormality condition:
(IF $?HoldsAt(y_{hand}$, t), $Trajectory(y_{hand}, s, g_1, \delta), s < t, t = s + \delta, Happens(x, r)$,
$Initiates(x, g_2, r), r < t, g_1_loc1 \neq g_2_loc2$`succeeds`
THEN$?\neg HoldsAt(ab_{hand.g2}, t)$ `fails`) AND
(IF $?Happens(x_{hand}, t), Initiates(x_{hand}, g_1, t), Happens(x', t), Initiates(x', g_2, t)$,
$g_1_loc1 \neq g_2_loc2$ `succeeds` THEN $?\neg HoldsAt(ab_{hand.g2}, t)$ `fails`) AND
(IF $?HoldsAt(y_{hand}, t)$, $Trajectory(y_{hand}, s, g_1, \delta), s < t, t = s + \delta$,
$HoldsAt(y, t), Trajectory(y, t, g_2, \delta), g_1_loc1 \neq g_2_loc2$ `succeeds`
THEN $?\neg HoldsAt(ab_{hand.g2}, t)$ `fails`) AND
(IF $?Happens(x_{hand}, t), Initiates(x_{hand}, g_1, t)$,
$HoldsAt(y, t), Trajectory(y, t, g_2, \delta), g_1_loc1 \neq g_2_loc2$ `succeeds`
THEN $?\neg HoldsAt(ab_{hand.g2}, t)$ `fails`);

CEU eTD Collection

(5) Meaning postulate for default behaviors:
$?\neg HoldsAt(ab_{hand.x}, t), Happens(x, t), x \neq x_{hand}$ `fails`
OR $?\neg HoldsAt(ab_{hand.y}, t), HoldsAt(y, t), y \neq y_{hand}$ `fails`;

(6) Well-formedness constraint for teleological explanation:
(IF $?Initiates(x_1, g, t), Initiates(x_2, g, t), x_1 \neq x_2$ `succeeds`
THEN $?HoldsAt(failure_{x1.g}, t)$ `succeeds`) AND
(IF $?Trajectory(y_1, t, g, \delta), Trajectory(y_2, t, g, \delta), y_1 \neq y_2$ `succeeds`)
THEN $?HoldsAt(failure_{y1.g}, t)$ `succeeds`);

(7) Assimilationist epistemic maintenance goal with respect to goal states (artifact functions):
IF ?*HoldsAt(z,t),* g^f=z `succeeds` THEN ?*HoldsAt(learn$_z$,t)* `succeeds`;

(8) Assimilationist epistemic maintenance goal with respect to practical skills:
IF ?*HoldsAt(teach*$_f, t)$, $HoldsAt(failure_{f.g}$*,t)* THEN ?*HoldsAt(learn$_f$,t)* `succeeds`.

The first three clauses in the list ground the infant's formation of concrete achievement goal conjectures. While the first two concern the way in which observations activate maintenance goals by reactive forward reasoning, the third concerns the direct derivation of achievement goals from observations (not mediated by previous maintenance goal assignment).

$I.c._1$ expresses the assignment of an abstract maintenance goal, i.e. a pedagogical intention, to the observed agent conditional on her object directed ostensive cues, and infant directed communicative behavior.

$I.c._2$ formalizes the specific interpretive norm that manifestations of distress trigger the agent's attempt to change it for the better. The desired change of state $change_d$ is an instance of the maintenance goal.

$I.c._3$ is the main CLP expression of the teleological stance with respect to physical achievement goals of behaviors. It demands assigning an achievement goal to an observed goal-directed behavior, i.e. finding a value for the Skolem function $f(x)$. According to discussion of quantification in Section 4.1.6, the requirement `succeeds` makes an existential claim with respect to a physical goal, i.e. there is such a state as g, which is a function $f(x)$ of an action x or an activity y. Given the definition of g in the notation list, $f(x)$/ $f(y)$ is the state optimally brought about by action x/ activity y. The drive to unify the value of function $f(x)$/ $f(y)$ with a constant picked out among the fluents from the scenario is an application of the teleological stance with respect to physical goals (i.e. in its original conceptualization from Gergely and Csibra (2003, 1997); Csibra and Gergely (2007, 1998), etc.). The antecedent of the constraint is tautology $\top$, which means that the explanatory drive that it expresses is unconditional. It formalizes the teleological stance introduced in Section 1.1.1, that infants are prone to assigning

CEU eTD Collection

goals to instrumental behaviors in order to understand them. The possibility that goal-directed behaviors be actions or temporally extended activities determines the disjunctive formulation of the constraint. A successful unification is sought by specializing the function $f(x)$ to a constant fluent from the narrative of events, given an evaluation of the causal relations available in the causal model.

$I.c._4$ formalizes the conditions for teleological abnormality, and thereby it spells out the conditions when the CWA_{ab} for goal centered action interpretation is justifiably overridden. The fluent that expresses it has two parameters: one stands for goal state g, the other one for the default agentive relatum $hand$. The antecedent stipulates the conditions for something abnormal to be the case regarding the instrumental use of the hand for goal attainment. Performance of a goal-directed behavior with the agentive relatum $hand$ constitutes an abnormality with respect to the simultaneous necessary involvement of the same agentive relatum in bringing about another physical goal state in a different spatial location. The constraint is used in the process of goal validation via plan recognition, in case the infant's offline planning outputs a different basic behavior than the observed agent's. Both goals g_1 and g_2 can be brought about either by an activity (via continuous causation) or by an action (via instantaneous causation), so the i.c. is composed of four conjuncts. Only the first two will be used in the plan recognition derivations modelled here. If both goals are caused by actions, simultaneity means that both actions are indexed at exactly the same time t. If one is caused by an action and the other by an activity, simultaneity means that the moment of the action t belongs to the temporal interval when the activity is ongoing $(s, t + \delta)$, where $s < t$.

$I.c._5$ specifies the default fine-grained representation of instrumental behaviors in terms of the local agentive relatum. The former disjunct concerns actions, the latter concerns activities. It is used in planning in order to reduce a behavior coarsely represented to a basic behavior, which can be directly executed. It guides all instances of planning the reasoner's own behavior, i.e. offline planning for teleological explanation, and online planning for action. According to Section 2.5, the CWA_{ab} prescribes performance of a prepotent, default response available in the reasoner's repertoire. Because the constraint is used in planning computations, i.e. in goal reduction, x and y are behaviors to be reduced to subgoals; they are immediately superior goals to the expected basic behaviors performed for their attainment. This is the meaning of the $x = g$ and $y = g$ equations. The constraint stipulates that the prepotent, default local agent for goal-directed behaviors is the hand. It reads as "unless something abnormal is the case, perform goal-directed behavior with the hand". Under the assumption that no teleological abnormality obtains in a given context ($\neg HoldsAt(ab_{hand.g}, t)$), the constraint prescribes performance of a default behavior. The assumption is overridden when the other two conjuncts in the i.c. obtain. Default representations of behaviors are thus defeasible. CWA_{ab} is justifiably overridden if the conditions spelled out in the antecedent of $i.c._4$ obtain and provide concrete positive evidence of abnormality, or if the context to be explained is trustworthy enough so that the abnormality need not be specified to a concrete instance.

CEU eTD Collection

$I.c._6$ provides a well-formedness standard for plan-based explanatory process of the other agent's actions. It does so by formalizing the condition 'no disjunction in conclusion', which expresses reasoners' expectation of normative means – ends relations (introduced in Chapter 2). An action should be executed by an observed agent assumed to be rational if it is the only one required to happen by the database which comprises the current knowledge representation (i.e. the narratives, the causal models), and by a general theory of causality (i.e. the EC and the integrity constraints). A one-to-one relation between the conjectured achievement goal and the observed action is expected for a satisfactory explanation. Goal reduction is explanatorily successful when depicted by a non-branching derivation tree. The constraint defines the meaning of $failure_{f.g}$ (f is expressed in the constraint as action x in the first conjunct, or activity y in the second): teleological explanation of f with respect to g fails when the goal is reducible to more than one instrumental behavior in a certain context. Like in $i.c._2$ and $i.c._4$ the disjunctive form of the antecedent captures the fact that goal states may be brought about via instantaneous or via continuous causation.

The epistemic maintenance goal is represented formally in the conclusions of two integrity constraints, $i.c._7$ and $i.c._8$. The i.c.'s define variable *learn*$_f$ by specifying parameter f in the conclusion. They provide an interface between the infant's processing of observed actions and the planning of his own action performance. They warrant the effect of action processing during teleological interpretation on the reasoner's online planning by having in the antecedent teleologically interpreted observations, and in the consequent the reasoner's own assimilationist epistemic maintenance goal *learn*$_f$. It is reduced to basic actions in online planning, given the integrity constraints in Section 5.6. The infant plans his actions in the test phase motivated by at least one of these forms of learning from novelty, activated during teleological interpretation.

$I.c._7$ is the formal expression of the maintenance goal to learn about physical goals attainable via artifact directed action. In the antecedent fluent z is equated with an observed end state that has been teleologically categorized as the physical goal g^{f} of some transitive behavior f. The scope of f is transitive behaviors, i.e. artifact directed behaviors. z need not be new per se (i.e. a state of affairs never seen before), rather it is novel because the artifact used to attain it is novel to the observer. Learning about goals is conditional only on goal assignment to instrumental behaviors upon novel objects, irrespective of whether the goal conjecture is validated as a proper justification for the behaviors. It is the default kind of practical knowledge acquisition from observation of artifacts that afford instrumental action (as explained in Section 3.1). The epistemic motivation in the conclusion is activated when a physical goal conjecture is available for an artifact directed behavior. Inferring the goal of a light-box directed behavior g^{f} prescribes learning *that* the new light-box can be used for g^{f}.

$I.c._8$ expresses the reasoner's maintenance goal to learn new practical skills. The antecedent formalizes the preconditions of imitative learning of a goal-directed behavior demonstrated by an agent, i.e. assignment of pedagogical intention, and failure of teleological evaluation to explain the skill f by means of goal g. Because the second conjunct refers to the outcome of

CEU eTD Collection

explanatory computations, i.e. $HoldsAt(faiIure_{f.g}, t)$, the constraint can be activated only in the final step of plan recognition, after it has been decided based on $i.c._6$ whether the physical goal conjecture justifies observed behaviors in the context. The agent's pedagogical intention, assigned according to $i.c._1$ above, triggers the infant's epistemic trust in her teaching of the observed instrumental behavior f despite its not being teleologically justified in the physical context of occurrence. The constraint is a formal expression of learning from a communicative agent *how* a novel artifact can be used. Its behavioral manifestation is evidenced only in the Communicative–HandsFree context.

5.2.2. The scenarios. The set of clauses in the scenario represent the actions and states that the infant individuates ('what happens when', the narrative of events), and the basic causal relations holding between them ('what actions do', the causal model).

I list two scenarios for the two experimental settings in Gergely et al. (2002), HandsFree and Hands Occupied in a Communicative context. Scenarios consist of a narrative and a causal model. Each clause in the causal model is justified with respect to the causal mechanism that presumably allows its formation (i.e. contact causality, or inertia), and the database over which the mechanism operates (i.e. current observation and prior knowledge). The justification is provided between square brackets after each clause. Within the scope of predicates in the scenario clauses there are only constants for the individuated events, activities, and states. The time argument is expressed by real numbers, or as temporal intervals over which quantification is universal in the case of activities or states (time variable t ranges over the intervals). γ is a constant denoting a certain non negligible duration. The real numbers are chosen as such given the convention that fluents start to hold after initiation (i.e. their first moment is strictly later than the moment of the cause event), and they hold until the moment when they are terminated (i.e. their last moment is equal with the moment of the terminating event) (van Lambalgen and Hamm 2005).

Before listing the clauses, I wish to address a distinction between the narrative and the causal model of observations. Actions and fluents may have different representations relative to coarseness in the narrative and in the causal model (according to the discussion in Section 5.1 above). Cognitive economy considerations are responsible for these differences. Because of perceptual availability of all the elements for fine-grained representation, contact appears as a parameterized event – l_{head} or l_{hand} in the narrative; the parameter specifies the local agentive relatum. However, it bears a stable interpretation in the causal model upon coarse encoding as a non-basic action l. It fulfills a causal justificatory role regarding its potential resultant states without the fine-grained agentive relatum. Identifying a global agent that contacts the object suffices for a causal understanding of events; contact causality (Leslie 1987; Saxe and Carey 2006) recommends the coarse representation in the causal model (without specification of the agentive relatum). Behavioral novelty is made available by non-causal perception – 'head contact occurs, and afterwards the light is on'. It is thus a feature of the fine-grained representation

CEU eTD Collection

in the narrative of the events currently observed. In contrast[3], the scarf manipulation activity m_{hand} fulfills a causal role in virtue of the local agentive relatum *hand*. Consequently, it is represented in the causal model as well as in the narrative as m_{hand}, a basic behavior.

- Hands Occupied

  HO-Narr: Narrative

  1. agent exhibits artifact-directed ostensive signals, and child-directed communicative behavior: *Initially(p)*;
  2. light-box is initially off: *Initially(¬ e)*
  3. agent expresses through gestures and language discomfort for being cold: *HoldsAt (d, 1)*;
  4. agent performs with hand scarf activity: *HoldsAt (m_{hand}, 4)*;
  5. scarf is located around agent's shoulders: *HoldsAt ($scarf_on$, 6)*;
  6. agent is engaged in light-box directed activity with the head: *HoldsAt*$(a_{\text{head}}, 7)$;
  7. agent's head contacts light-box: *Happens(l_{head}, 8)*;
  8. light-box is on: *HoldsAt(e,9)*;
  9. scarf is located around agent's shoulders: *HoldsAt ($scarf_on$, 9)*;
  10. light-on state obtains on the table, where the light-box is located: *HoldsAt(e_table, 9)*.

  HO-CM: Causal model

  1. scarf manipulation brings about the scarf location state on the agent's shoulders at some later moment: *Trajectory*$(m_{\text{hand}}, 3, scarf_on, 6)$ [by inertial perception of the hand activity resulting in a different location state];
  2. definition of $change_d$ with respect to scarf location - scarf location on the agent brings about no-distress state at some later moment: *Trajectory*$(scarf_on, 6, change_d, \gamma)$ [by inertial perception of the location state of the scarf as a continuous force that causes a change in the temperature state];
  3. contact is the culminating point of the light-box directed bodily movement: *Terminates(l, a, 8)* [by inertial perception of the ongoing activity];
  4. light-on is the causal effect of contact: *Initiates (l, e, 8)* [by contact causality].

- HandsFree

  HF-Narr: Narrative

[3]Yet another different representational strategy is presented for the experiment described in Section 5.8.

CEU eTD Collection

1. agent exhibits artifact-directed ostensive signals, and child-directed communicative behavior: *Initially(p)*;
2. light-box is initially off: *Initially(¬e)*
3. agent expresses through gestures and language discomfort for being cold: *HoldsAt (d, 1)*;
4. agent performs with hands scarf activity: *HoldsAt (m_{hand}, 4)*;
5. agent knots scarf: *Happens(k,5)*;
6. the location state of agent's hands is on the table: *HoldsAt($hand_table$,7)*;
7. scarf is located around agent's shoulders: *HoldsAt($scarf_on$,6)*;
8. agent is engaged in light-box directed movement with the head: *HoldsAt (a_{head},7)*;
9. agent's head contacts light-box: *Happens*(l_{head}, 8);
10. light-box is on: *HoldsAt(e,9)*;
11. scarf is located around agent's shoulders: *HoldsAt ($scarf_on$, 9)*;
12. light-on state obtains on the table, where the light-box is located: *HoldsAt(e_table, 9)*.

HF-CM: Causal model

1. knot is the culminating event of scarf manipulation: *Terminates(k, m_{hand}, 5)* [by inertial perception of the ongoing activity];
2. the scarf location state is the causal effect of the knot action: *Initiates(k,$scarf_on$,5)* [by prior causal world knowledge];
3. definition of $change_d$ with respect to scarf location - scarf location on the agent brings about no-distress state at some later moment: *Trajectory*$(scarf_on, 6, change_d, \gamma)$ [by inertial perception of the location state of the scarf as a continuous force that causes a change in the temperature state];
4. contact is the culminating point of the light-box directed movement: *Terminates(l, a, 8)* [by inertial perception of the ongoing activity];
5. light-on is the causal effect of contact: *Initiates (l, e, 8)* [by contact causality];

In the Incidental contexts of presentation, the scenarios are the same, save for the first clause in the narratives, $Initially(p)$, which is absent.

CEU eTD Collection

5.3. Goal derivation.

Computations have the same outcome across the HandsFree and HandsOccupied conditions; the same goals are picked in both Communicative conditions.

Given the database above, the infant derives three goals. For one, the clause in HO-Narr1 and HF-Narr1 triggers by reactive forward reasoning the conclusion of $i.c._1$ – $HoldsAt(teach_f, t)$. The agent is thereby assigned a pedagogical intention to teach f. The intention is an abstract maintenance goal, whose behavioral content is not yet specified. The parameter f will be rendered concrete with respect to the infant's corresponding maintenance goal – his drive to learn, which occurs in the conclusions of $i.c._7$ and $i.c._8$.

The expected outcome of the derivations based on the integrity constraints $i.c._2$ and $i.c._3$ are conjectures about the physical achievement goals of observed behaviors. The different directions of reasoning have relevant consequences for the degree of certainty in the conjectures. The achievement goal based on $i.c._2$ is inferred backwards as a subgoal of the maintenance goal $change_d$ in the conclusion. The strategy is to find in the scenario a behavior that the desired change of state justifies, given the infant's world knowledge. Once found it is categorized as an achievement goal with high degree of certainty.

It is worth noting that the assignment of an abstract intention via $i.c._1$ and of a physical maintenance goal via $i.c._2$ is done by the same inferential mechanism: forward reasoning from factual observations, given the world knowledge available. This illustrates the 'single-mechanism multi-level' teleological processing suggested in Chapter 2 to be at work in infants' interpretive and planning processes that result in practical knowledge extension.

I zero in on the computations generated by each constraint in turn, starting with $i.c._2$:

IF $?HoldsAt(d, 1)$ `succeeds` THEN $HoldsAt(change_d, t), 1 < t$ `succeeds`.

Upon observation of behavioral manifestation of distress (d) at time 1, reactive forward inferencing assigns the agent the physical maintenance goal to change the distressing state ($change_d$) at some later moment. The observer expects that the agent act upon it. $change_d$ is thereby reduced to a concrete behavior meant to bring it about in the current context, i.e. an achievement subgoal represented by a concrete observation in the scenario.

The backwards processing in the reduction of the maintenance goal is depicted in Fig. 1. The clause standing for the maintenance goal is resolved by application of Axiom 3 or Axiom 4 depending on whether an action or an activity is encountered in the current causal model that brings about $change_d$. Resolution is based on the 'search and substitute' strategy of unification introduced in Chapter 4. A scan of the scenario encounters a relevant clause in the causal model HO-CM2 and HF-CM3 – $Trajectory(scarf_on, 6, change_d, \gamma)$. It defines $change_d$ by appeal to the location fluent $scarf_on$, whose continuous exertion of force for a non-negligible

CEU eTD Collection

duration γ brings it about at time $6 + \gamma$. The corresponding clause in the narrative is HO-Narr5 and HO-Narr7 – $HoldsAt(scarf_on, 6)$.

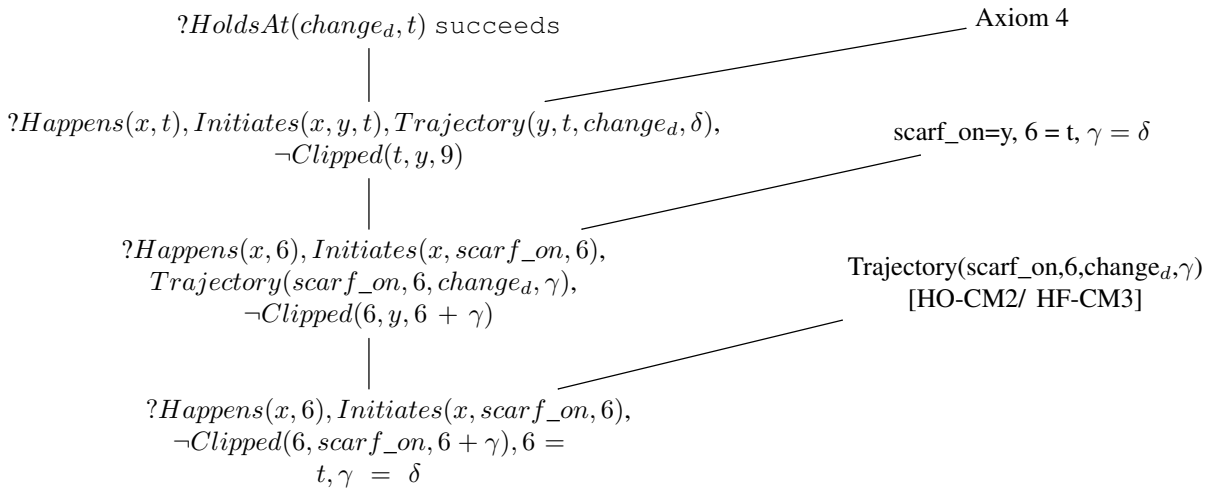

Figure 1. Achievement goal derivation from the maintenance goal $change_d$.

Fluent $scarf_on$ is thus substituted for variable y – [*scarf_on/ y*]. Temporal substitutions $[6/t]$ and $\gamma/\delta]$ permit unification with the literals in HO-CM2 and HF-CM3. In the next step the $Trajectory$ predicate can be deleted from the derivation tree. The temporal substitutions accumulate as constraints with which the database must be updated to ensure success of the top goal clause.

Forward closed world reasoning sums these results and derives via Axiom 3 that $HoldsAt(scarf_on, 6+\gamma)$.

$$Happens(x, 6) \wedge Initiates(x, scarf_on, 6) \wedge \neg Clipped(6, scarf_on, 6 + \gamma) \rightarrow HoldsAt(scarf_on, 6 + \gamma)$$

Location fluent $scarf_on$ is the subgoal of the maintenance goal $change_d$. The scarf location state must be maintained for duration γ in order to bring about the desired change of state – scarf is on the agent's shoulders at time 6 and maintained for duration γ. Now, HO-Narr9 and HF-Narr11 represent the observed state $HoldsAt(scarf_on, 9)$. Consequently, the most parsimonious way to make the goal clause successful is by setting γ to 3. This way, the observations in clauses 9 and 11 in the two narratives are teleologically categorized as the achievement goal of the agent who shows distress for being cold at time 1. The observational scenario is updated with the constraint $scarf_on = g_1$. The observed fact $scarf_on$ at time 9, represented in the scenarios as HO-Narr9 and HF-Narr 11 makes the conclusion of $i.c._2$ successful; the clause $HoldsAt(scarf_on, 9)$ ensures that the agent fulfills her maintenance goal $change_d$. Because of familiarity with behavioral manifestations of distress for being cold, the corresponding desire not-to-be-cold-anymore, and subsequent putting-on-something-warm to this end [4] the infant's

[4]According to criteria for confidence in goal status in Section 2.4, at p. 54.

CEU eTD Collection

confidence in the assigned action rule (i.e. 'keep scarf on FOR change distressing state') makes it so that the goal does not require validation.

I now turn to the resolution of $i.c._3$,

?Happens(x,t), Initiates(x, f (x), t), f(x) = $g^x{}_2$ succeeds
OR *?HoldsAt(y,t), Initiates(y,f(y),t), f(y) =* $g^y{}_2$ succeeds.

Resolution aims to spell out the concrete goal conjecture, by substituting a constant from the scenario for $f(x)$ or $f(y)$, which is a function of, and is expected to justify observed behavior x or y. Fig. 2 presents the graphic depiction of the resolution. In the first step unification is grounded by specialization of the action variable x to an occurring event (i.e. an event that $Happens$) or activity y (i.e. an activity that $Holds$) from the narrative. Reasoners then seek a value of fluent $f(x)$ that witnesses the existential claim made by the i.c., that there is such a fluent which justifies an action x or an activity y by recourse to a causal physical appraisal of the context. The value of the function is the achievement goal placeholder, to be specialized to a constant from the causal model. Following unification, in order to make the constraint successful the knowledge database must be updated with the equation between the goal variable g and the substituted constant. The expectation that g_2 justify the behavior must be validated by further computations, to be described in the next Sections.

In the head-touch task, the artifact-directed behavior consists of an activity a_{head}, and its culminating event, contact l_{head}. Given contact causality, it is contact that bears causal power; the information is represented in clause HO-CM4/ HF-CM5. The behavior that engages the infant in goal attribution is action l_{head}, the agent's contacting the light-box with the head. The reasoner resolves the action disjunct of $i.c._3$. The occurrence of contact is represented in clauses HO-Narr7 and HF-Narr9 $Happens(l_{head}, 8)$. The action variable x in the i.c. is specialized to the head contact action, and the substitution $[l_{head}/x]$ is effectuated. Upon unification with the narrative clause, $Happens$ is discarded from the derivation. Computations proceed by specializing the goal function, now $f(l_{head})$, to a state caused by contact with the light-box.

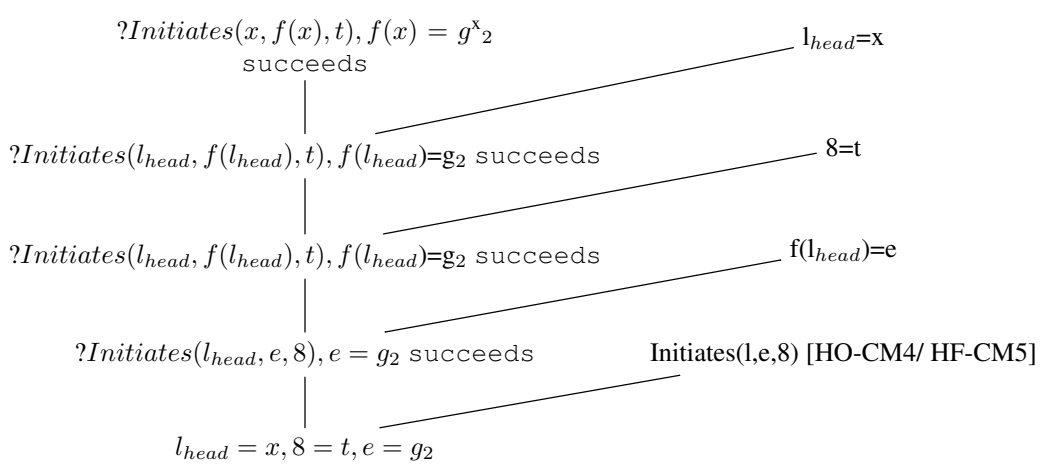

Figure 2. Goal derivation for the agent's contact action.

CEU eTD Collection

Upon observation of a novel instrumental action followed by a perceptually salient change of state, as is the case in the head-touch task of Gergely et al. (2002), specialization of the goal to a concrete action function involves limited computations. Contact causality fosters the infant's perception of the state of the light being on as an effect of the contact between the agent and the artifact. This causal understanding is represented in HO-CM4/ HF-CM5 as $Initiates(l, e, 8)$. The young reasoner substitutes the resultant state for the goal function, $[e/f(l)]$, and time point 8 for variable t. This allows unification of the formula in the i.c. with the corresponding literal in the causal model, and hence its deletion from the derivation tree. The integrity constraint is resolved, providing a goal conjecture for the target novel action. The goal fluent 'light-on' is an interpretation-laden causal effect of the agent's exerting force by means of a light-box directed bodily movement culminated by contact. The interpretive load manifests itself in the teleological categorization of the 'light-on' state as the observed agent's achievement goal. It is the foreground goal in the head-touch task.

Availability of a goal conjecture involving the novel artifact allows activation of the assimilationist epistemic drive formalized in the conclusion of $i.c._7$:

IF ?*HoldsAt(z,t), g^f=z* `succeeds` THEN ?*HoldsAt(learn$_z$,t)* `succeeds`

The substitutions performed in the integrity constraint are $[e/z]$, $[l/f]$, and $[9/t]$. The antecedent is made true by observations (NarrHO-8/HF-10 – $HoldsAt(e, 9)$) and by the database updated with the physical goal conjecture $e = g_2$ obtained in Fig. 2. At time 9 the infant is set on learning that the light-box can be used to achieve e, i.e. to make light, irrespective of whether the expected justificatory effect of the goal with respect to the observed behaviors obtains. He is thereby committed to making such that ?*HoldsAt(learn$_z$,t)* `succeeds` at a not yet specified time t. He has the chance to do so in the test phase, when allowed to interact with the light-box.

Until then, the reasoner pursues further on his explanatory epistemic drive. To do so, he verifies whether the foreground goal justifies the novel action as optimal for in the current circumstances.

5.4. Offline planning.

This is a formal depiction of the reasoning processes detailed in Section 3.2.2: computing a temporally ordered sequence of behaviors that would have lead the reasoner to attain the foreground goal by using the novel object.

I spell out below the database for offline planning. Clauses in the database represent the infant's prior knowledge, applied to a fully normal context of instrumental transitive action upon the novel artifact. The artifact is the observed entity that triggers the infant's relevant prior knowledge, which constitutes the constraints, the narrative, and the causal model. Because offline planning is for the infant's own actions, predicates take as arguments the infant actions and activities listed in Section 5.1.

CEU eTD Collection

I. Integrity constraints
   1. $?\neg HoldsAt(ab_{hand.x}, t), x = g, Happens(x, t), x \neq x_{hand}$ `fails`;
   2. $?\neg HoldsAt(ab_{hand.y}, t), y = g, HoldsAt(y, t), y \neq y_{hand}$ `fails`.

II. Narrative
   1. light-box is initially off: *Initially(¬ e)*;
   2. a light-box directed bodily movement obtains: *HoldsAt(c,1)*;
   3. contact with the light-box occurs: *Happens(n,2)*.

III. Causal model
   1. light-on is the causal effect of contact with the light-box: *Initiates(n,e,3)* [by contact causality];
   2. contact ends light-box directed bodily movement: *Terminates(n, c, 3)* [by inertial conceptualization of a light-box directed motion].

IV. Goal clause to be resolved
   1. the light-box should be *on* at time point *now*: ?*HoldsAt (e, now)* `succeeds`.

Only one integrity constraint among the ones introduced in Section 5.2.1 is actually involved offline planning, namely, $i.c._5$. I reproduced it in the database above, and, for convenience, I split the two conjuncts into two clauses: one for actions and one for activities. The two clauses instantiate the effect of the CWA_{ab} on the reduction of coarsely represented behaviors to basic behaviors, by specification of a prepotent agentive relatum: the hand. They prohibit update of the database with a non-default basic behavior, unless positive evidence of abnormality. That is, they stipulate reduction of transitive behaviors to behaviors performed with the hand. $I.c._4$ in the database for teleological interpretation of the current context (Section 5.2.1) is not considered unless the need arises upon goal reduction, i.e. unless the goal clause cannot be made to succeed given the scenario for offline planning. The CWA_{ab} can be overridden in an exceptional context, but the context of offline planning is not such. Offline planning is reasoning in an idealized default context, where there is no positive evidence of abnormality.

The scenario is composed of a descriptive narrative for the context involving the novel artifact, and a causal model. The light-box is the only element furnished by current perception. Clauses about artifact directed behaviors are recruited from the infant's existent practical knowledge concerning instrumental behaviors on objects (i.e. transitive instrumental behaviors). The narrative contains clauses about the initial state of the artifact, and artifact directed behaviors: a temporally extended activity, i.e. the movement towards the light-box, and instantaneous contact with the light-box. In contrast to the perceptually available narrative used for goal derivation (see Section 5.2 above), behaviors are represented coarsely, without specification of the local agentive relatum. They are reduced to fine grained representations of basic executable behaviors by planning computations. The causal model is a selection of infant's causal knowledge relevant for goal reduction. An inertial understanding of a temporally extended activity (i.e. of a movement towards an object) and contact causality connect the light-box directed movement

CEU eTD Collection

c with the resultant state of the light-box e via a culminating point of the activity – the causally efficient event n.

The goal clause ?*HoldsAt(e,now)* `succeeds` requires as a resolvent a temporally ordered sequence of concrete behaviors that lead to the attainment of the achievement goal *light-on* at time *now* (the fictitious now of planning). Basic motor responses are expected in the final line of the derivations because they are executable.

The goal that e be the case now is reduced via Axiom 3 for instantaneous change due to an event with causal power. The first integrity constraint prescribes performance of a default means action onto the artifact.

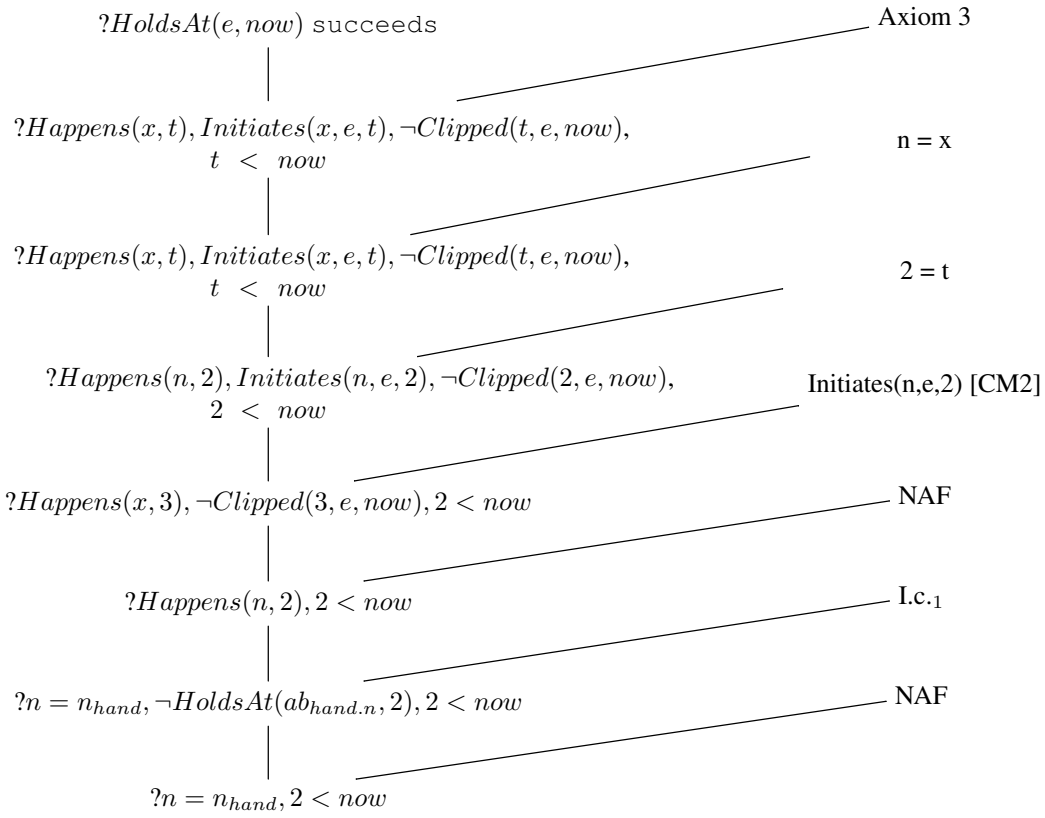

Figure 3. Offline planning: Derivation from the goal *light-box is on now*.

Fig. 3 shows the derivation from the achievement goal ?*HoldsAt(e,now)* `succeeds`, using the principle of inertia for instantaneous change as expressed by Axiom 3. The Axiom reduces the goal state to an initiating event that takes place before now. Substitutions $[n/x]$ and $[2/t]$ permit unification of the $Initiates$ atom with CM1, and thereby its deletion from the tree. Because there is no information available in the database about another causally e-relevant event on the temporal interval $[2, now]$, NAF deletes the $\neg Clipped$ atom without much ado. The contact event is then specialized to hand contact, in order to ensure failure of the $n \neq n_{hand}$ conjunct within integrity constraint 1 in the database. In the last step of the derivation NAF permits deletion of the $\neg HoldsAt(ab_{hand.n}, 3)$ conjunct from the antecedent of the clause, for

CEU eTD Collection

lack of positive evidence of teleological abnormality with respect to contact at time point 2. The top query succeeds in a normal context if the scenario is updated with a hand contact event prior to the fictive now at which the *light-on* state is required to hold.

Contact is represented causally in CM2 as terminating a movement towards the object, i.e. it is the culminating point of the activity. The planning derivations that connect contact at a time prior to now to the light-box directed movement are depicted in Fig. 4. In the second step of the derivation Axiom 1 makes the activity $Clipped$ because it is terminated by contact at time 2. However, it can only be so after contact occurs, because of information in Narr2 about a light-box directed movement at time 1, so $Clipped(2, c, 3)$ and not $Clipped(1, c, 3)$. Further on $i.c_1$ warrants a light-box directed movement of the hand at time 1. Axiom 3 used under the auspices of NAF ultimately prescribes that the hand movement be started and performed uninterruptedly until the occurrence of contact. NAF is warranted in the ninth step and $\neg Clipped(t, c_{hand}, 1)$ can be discarded because the artifact directed movement is $Clipped$ at a moment after 2 (according to the third step where we have $Clipped(2, c, 3)$).

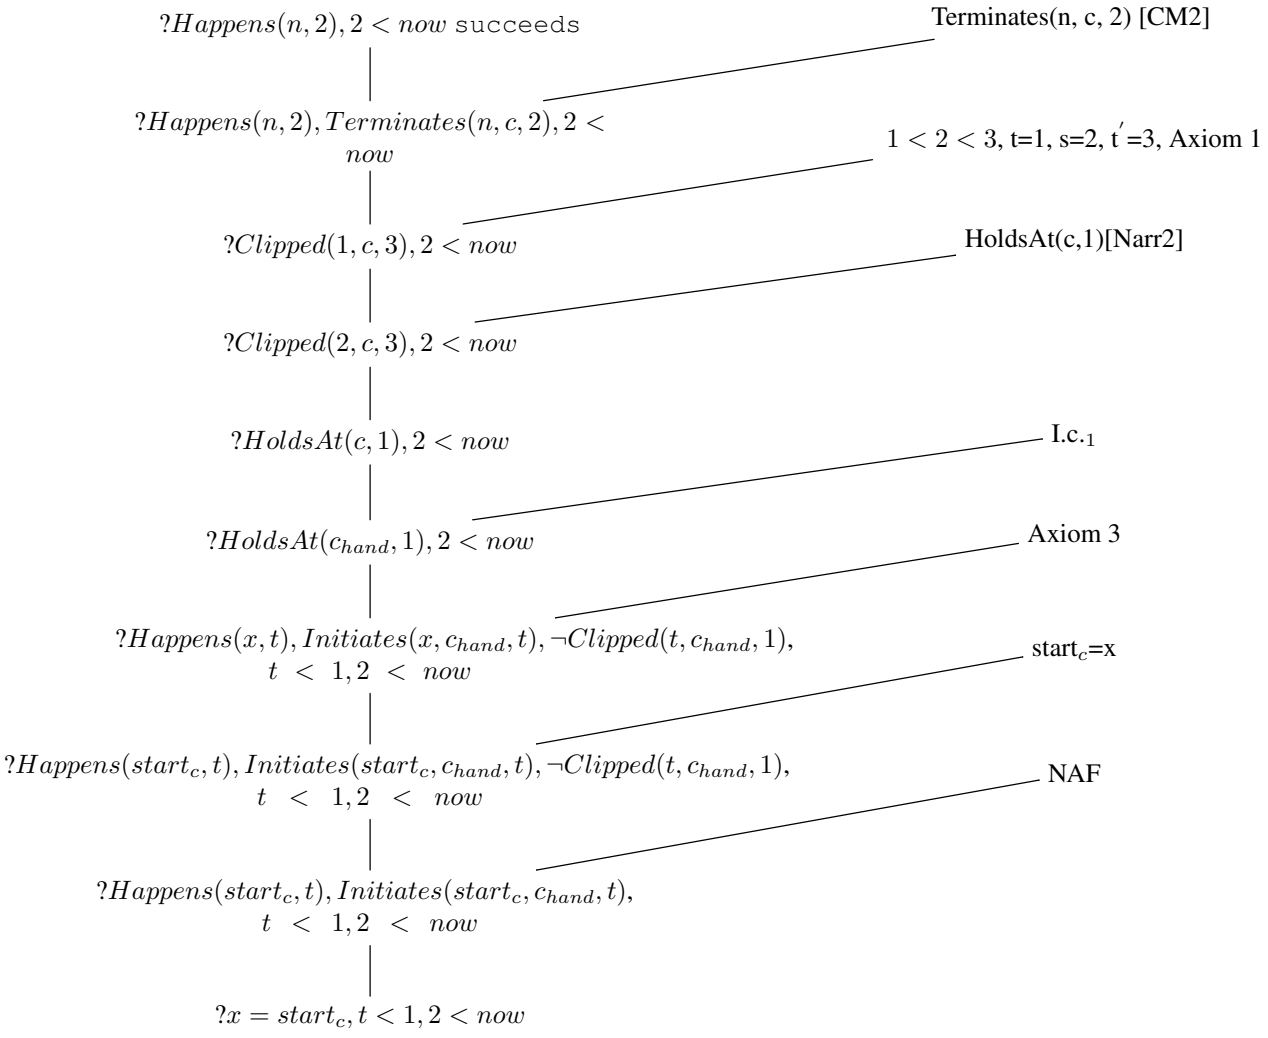

Figure 4. Offline planning: Resolution of the contact event query.

CEU eTD Collection

Starting from the requirement that the state *light-on* holds now, by application of the principle of inertia for instantaneous change (Axiom 3), the scenario must be updated with the occurrence of an event that sets out a light-box directed activity with the hand, terminated by hand contact, which initiates the desired consequent state of the light at time 2 prior to now.

$$HoldsAt(e, now) \leftarrow Happens(start_c, s), Initiates(start_c, c_{hand}, s), s < t, Happens(n_{hand}, t), Initiates(n_{hand}, t), t < now$$

All in all, in offline planning, based only on the default motor knowledge about bodily agentive relata of instrumental actions, and on the general causal principles of inertia and contact, the infant reduces the achievement goal *light-on* to starting a light-box directed activity with the hand and performing it uninterruptedly until its culminating event, i.e. a hand contact, which causes the desired state of the light-box. Contact appears as the causal mediator of the transition between a light-box directed movement and the consequent state of the light-box, i.e. the goal *light-on*. In a normal context the infant would perform the basic action of contacting the object by hand in order to light it up. This creates the expectation that the state *light-on* is the resultant state of a hand contact.

But clauses HO-Narr7 and HF-Narr8 in the database of observed facts look like *Happens (l_{head},8)*. The agentive relatum of the default action plan recommended by offline planning (the hand) mismatches observations of a head contact.

5.5. Plan recognition.

The computations to be described here implement the interpretive reasoning processes from Section 3.2.3.

The mismatch between the outcome of offline planning and observations engages reasoners in further validation of the foreground goal *light-on* as a suitable explanation for the novel behavior. Because the effect categorized as an achievement goal of the agent's unusual instrumental behavior is the only effect, and because it is highly salient, the goal conjecture is entertained with high confidence. Consequently the observer re-evaluates the agent's context of action performance, rather than re-computing the goal. The strategy amounts to plan recognition, i.e. interpreting the actions performed by the experimenter in the given context as her plan to attain the two goals that she has been assigned.

The reasoners' objective is understanding the inclusion of the unusual behavior in the agent's global action plan, all relevant teleological aspects of the situation taken into account. More concretely, this amounts to the aim of canceling out one of the actions tied to the goal 'light-on', i.e. hand or head contact. The way to go is to look for contextual abnormalities in the observed context that rule out performance of the default hand action. If found, they re-establish expected normativity between means and goals: although in a normal context something else would have been optimal, in the current context something abnormal calls for performance of an exceptional

CEU eTD Collection

means. Explanatory exhaustivity of the whole observed sequence via plan recognition under the constraint 'no disjunction in conclusion' (cf. Chapter 2) is operationalized as a twofold aim:

(a) obtaining for each query a single successful derivation, i.e. a unique sequence of observed actions and consequent states that leads to the top query, and
(b) ensuring that the two derivations are consistent.

The former, (a), is ensured by integrity constraint $i.c._6$ in Section 5.2.1. The latter, (b), is secured by the fact that derivations grounded in the background query take effect on reduction of the main query, because $i.c._4$ involves both. The background goal is a relevant teleological aspect, a dynamic feature of the situation as a whole[5].

The achievement goals $scarf_on$ and e have been assigned to the agent, indexed at the very same moment in time. Consequently derivations start from the two queries expressing that both hold simultaneously. The content of HO-Narr8 and HF-Narr10 provide the top query for the main thread of derivations – ?*HoldsAt(e, 9)* `succeeds`, which stands for the foreground achievement goal. The background query is the achievement goal $scarf_on$ also holding at time 9, ?*HoldsAt(scarf_on, 9)* `succeeds`. Background reasoning starts from it because its cause behavior m_{hand} is observed first (cf. HO/HF-Narr4 $HoldsAt(m_{hand}, 4)$).

The computational strategy is the same across Communicative conditions, for both Hands-Free and HandsOccupied. However the differences between the behaviors demonstrated by the agent lead to different outcomes. I present the formalization in two separate subsections.

Once the explanatory endeavor via plan recognition from the assigned goals is finalized, depending on its outcome, a particular version of the infant's assimilationist epistemic drive $learn_f$ – as represented in the conclusion of $i.c._7$ or $i.c._8$ – is activated. Parameter f of $learn_f$ is specialized either to an instrumental behavior, or to a goal state, depending on whether teleological explanation relative to the goal conjectures is satisfactory. This influences what is stored from observations, and, furthermore, what guides online planning in the test phase.

5.5.1. Plan recognition in the HandsOccupied condition. Axiom 4 is used to reduce the background achievement goal $HoldsAt(scarf_on, 9)$. Usage of the Axiom for continuous change is justified because the causal model contains information about $scarf_on$-relevant activities, and not actions. The resultant fluent in the Axiom is specialized to the goal to be reduced, $scarf_on$, and the time variable is taken to be 9.

Specialization of the activity variable y to constant m_{hand}, of time t to 3, and of duration δ to 6 allows unification with the literal in clause CM1 of the causal model, and thereby deletion of $Trajectory(m_{hand}, 3, scarf_on, 6)$ from the tree. For lack of information about a terminating or releasing event, the NAF rule warrants the closed-world assumption that the manipulation activity is not clipped, and hence the literal can be discarded from the tree. The substitution $[start_m/x]$ is effectuated. What is left is the conjunction of a $Happens$ and an $Initiates$ predicate relative to a starting event for the scarf activity performed by the agent's hand – $start_m$.

[5]The distinction between static and dynamic contextual features is introduced in Section 3.1.

CEU eTD Collection

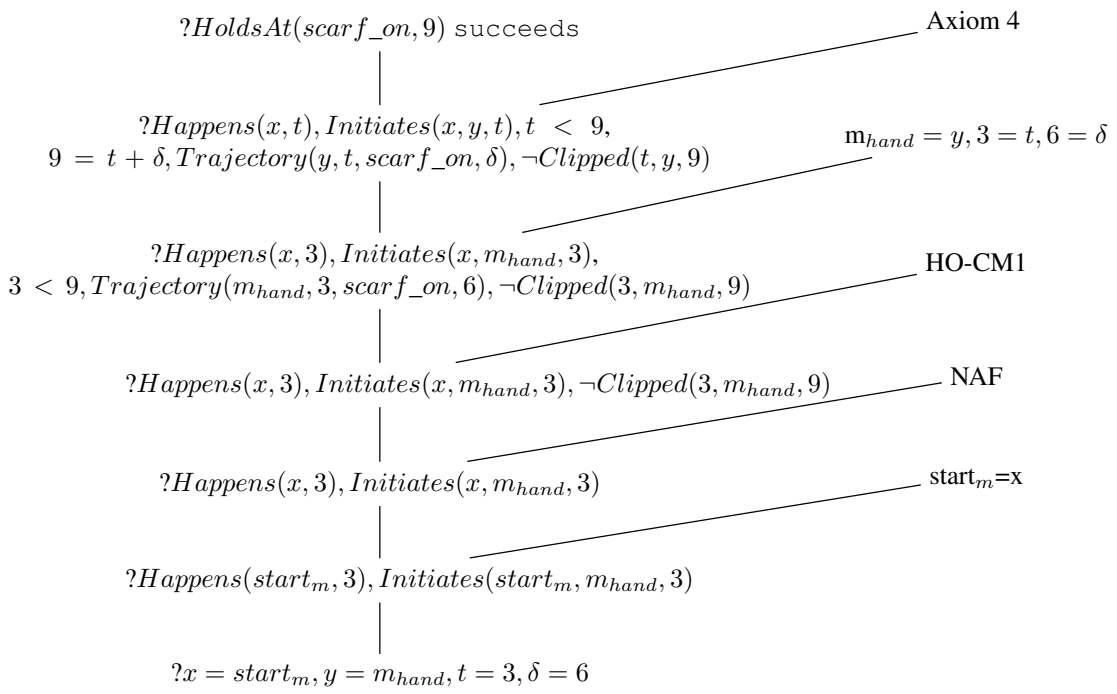

Figure 5. Plan recognition: Reduction of the background achievement goal $scarf_on$ in HandsOccupied.

Information regarding the scarf manipulation is the bottommost query necessary to make the achievement goal state $scarf_on$ succeed. Backward reasoning with Axiom 4 reduces the achievement goal $scarf_on$ to a scarf manipulation activity m_{hand}. For the top query to succeed, the database must be updated with the constraints from the last line of the derivation. They represent information about an initiating an event $start_m$, which occurs at time 3; also, scarf manipulation must be performed for duration 6.

The NAF rule allows addition of $\neg Clipped(3, m_{hand}, t)$ for whichever $t > 3$ in a temporal interval where there is no positive evidence of a m_{hand} relevant event. The reasoner takes $t = 9$. Forward closed world reasoning derives via Axiom 3 that $HoldsAt(m_{hand}, 9)$.

$$Happens(start_m, 3) \wedge Initiates(start_m, m_{hand}, 3) \wedge \neg Clipped(3, m_{hand}, 9) \rightarrow HoldsAt(m_{hand}, 9)$$

All in all the database must be updated with the occurrence of an initiating event at time 3 (as a result of backwards goal reduction), and with clause $HoldsAt(m_{hand}, 9)$ (as a result of forward use of Axiom 3) in order to make the query successful. The scarf is on the agent at time 9 if the scarf activity is initiated by the occurrence of a $start_m$ event at 3, and still holds at time 9. A plan consisting of a hand activity with the warm scarf, initiated by a starting event, is assigned to the agent whose achievement goal is to locate and maintain the scarf onto herself.

The resolution of the foreground query ?*HoldsAt(e,9)* `succeeds` follows the strategy used in offline planning (Fig. 3 and Fig. 4). Axiom 3 reduces e to the occurrence of contact at time

CEU eTD Collection

8, as depicted in Fig. 6. Clause HO-Narr7 contains information about contact with the head, so l is specialized to the basic action l_{head}. The constraints in the last line make the top query successful: action x from the Axiom is specialized to contact with the head (l_{head}), and time t of its occurrence is 8.

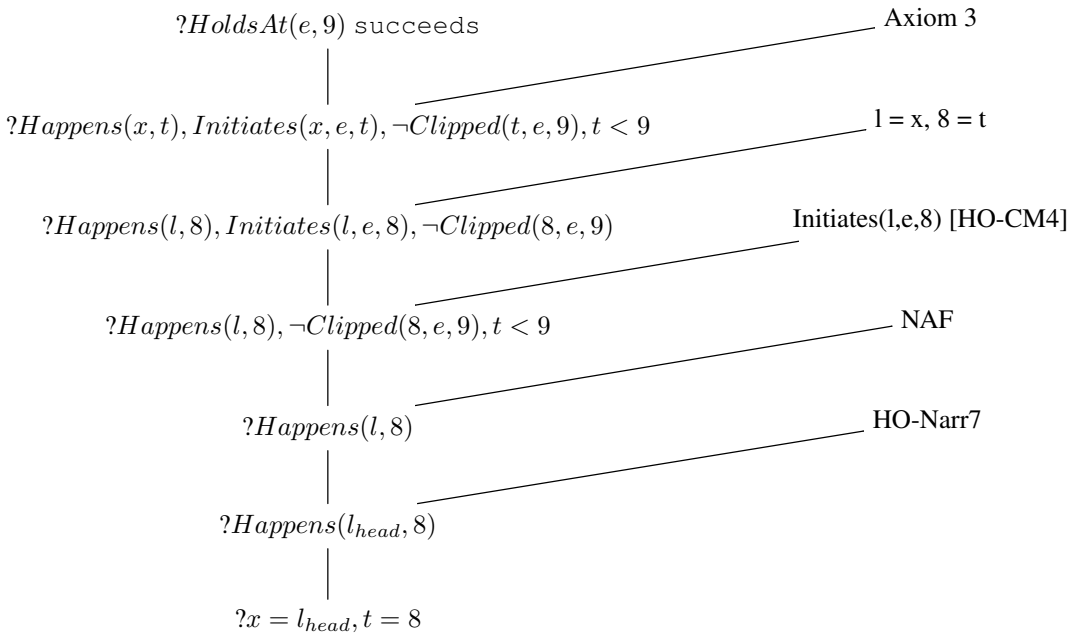

Figure 6. Plan recognition: Reduction of the foreground goal *light-on* in HandsOccupied.

Further derivations grounded in ?$Happens(l, 8)$ succeeds (obtained in the fourth step of the previous derivation) are presented in Fig. 7. The contact event terminates a light-box activity, according to HO-CM3. This clause makes the light-box directed movement $Clipped$ between an earlier and a later moment, say, time points 7.5 and 9. HO-Narr6 furnishes information about a head movement, a_{head} holding at time 7, so this is fed into the derivation, and the constraint $a = a_{head}$ is accumulated. Because NAF is warranted for times prior to 7[6], Axiom 3 reduces the head activity to the occurrence of an initiating event before time 7. An activity performed with the head must be started before time point 7.

CEU eTD Collection

[6]Given that $Clipped(7.5, a, 9)$, NAF cannot be applied to times later than 7 and earlier than 9.

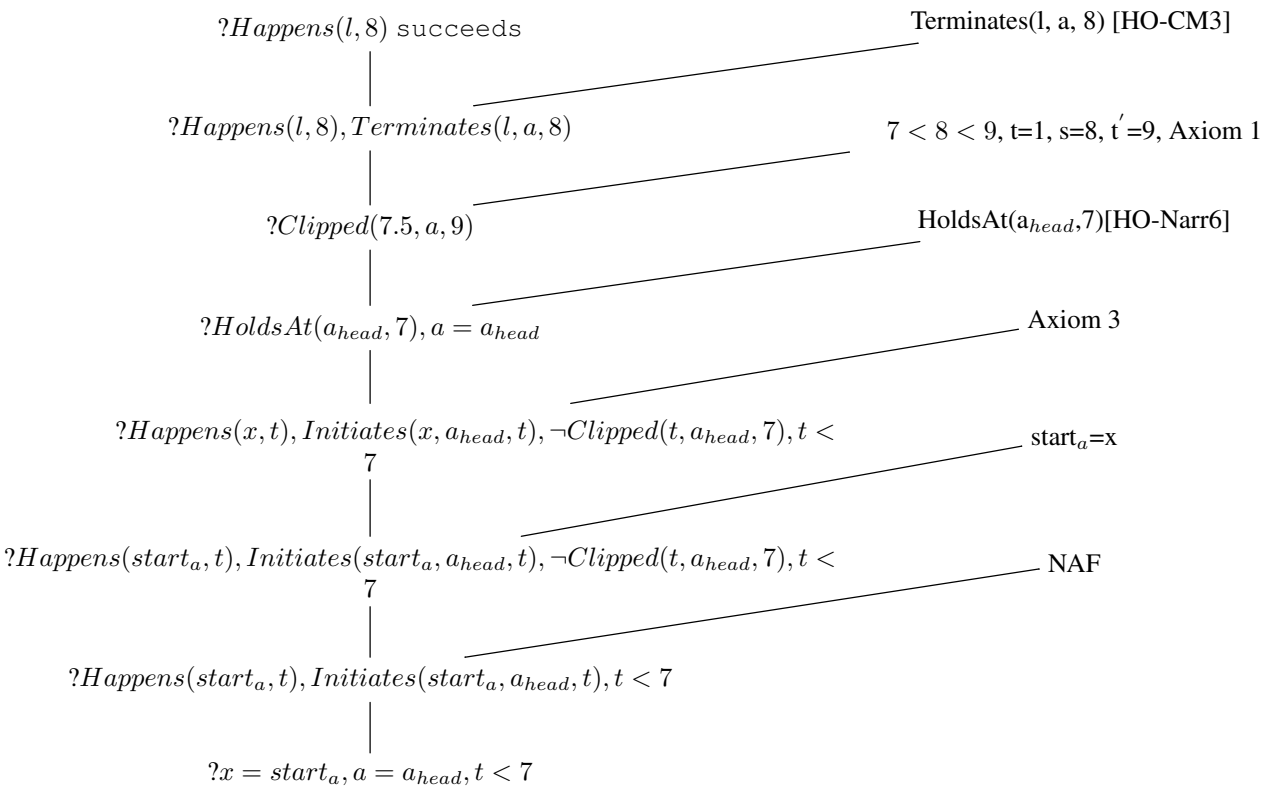

Figure 7. Plan recognition: Resolution of the contact query in HandsOccupied.

Summing up the derivations in Fig. 6 and Fig. 7, the goal clause that e be the case at time 9 prescribes the agent's performance of a head contact with the light-box at time 8, which culminates a head movement initiated at a time earlier than 7. The database must be updated with information about the occurrence of a starting event, all the rest are given in the scenario.

The plan assigned to the agent prescribes performance of a non-default action (head contact) which terminates a non-default activity (head movement towards the light-box). It constitutes an alternative to the infant's default plan, computed during offline planning. In the observed circumstances, the physical achievement goal *light-on* reduces to an action differently specified in terms of the agentive relatum. This can be expressed as a disjunctive action rule of the form

(do hand contact OR do head contact) FOR light-on,

$\longleftarrow$

and laid out in a branching derivation tree grounded by the top query $?HoldsAt(e, t)$ `succeeds`[7].

The left branch reduction relies on the CWA_{ab} that no abnormality occurs in the context. Reduction of the goal e in the observed context (the right branch of the tree) conflicts with the

[7]Correspondence between one's own and another agent's bodily ensures transformation of the infant's own motor program for a default action n obtained in offline planning into the corresponding program of another agent – l, and its use from a 3rd person perspective in the action rule and the derivation tree in plan recognition. As mentioned in the introductory Chapter, and restated in the preamble to this Chapter, the model developed here assumes correspondence and does not attempt to account for it.

CEU eTD Collection

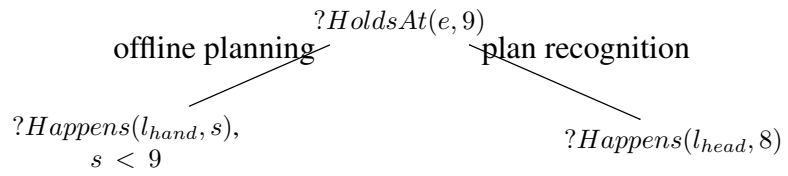

Figure 8. Branching derivations from the goal *light-box is on at time 9*.

default representation of basic instrumental behaviors specified by $i.c._5$. The integrity constraint requires that

$$?\neg HoldsAt(ab_{hand.x}, t), Happens(x, t), x \neq x_{hand} \texttt{ fails}$$
$$\text{OR } ?\neg HoldsAt(ab_{hand.y}, t), HoldsAt(y, t), y \neq y_{hand} \texttt{ fails}$$

i.e. in normal conditions goal-directed behaviors must be performed by hand (nothing else but the hand qualifies as normal agentive relatum). The causally relevant behavior in the reduction of the goal e is contact l_{head}, so it is the first disjunct in the i.c. that is of interest. The substitution $[l_{head}/x]$ makes it such that $l_{head} \neq l_{hand}$ `succeeds`. Since $Happens(l_{head}, 8)$ is represented in HO-Narr7 so $Happens(l_{head}, 8)$ `succeeds`, the only way to make the query fail is by having $?\neg HoldsAt(ab_{hand.l}, t)$ `fail`.

This shows formally that the agent's performance of a non-default instrumental behavior, which overrides the CWA_{ab}, is justified if a contextual abnormality that cancels the default plan for a normal context is found. In order for the teleological conjecture with respect to state *light-on* to be explanatorily useful for the agent's odd behavior, there must be something abnormal concerning the involvement of the agentive relatum $hand$ in an action that brings about e (namely, l) at time 9. The reasoner proceeds to checking if this is the case.

$?\neg HoldsAt(ab_{hand.l}, t)$ `fails` is the consequent of $i.c._4$, the abnormality condition constraint. The reasoner focuses on the first conjunct of the integrity constraint because the foreground physical goal e has been reduced to an occurring event, i.e. contact l, whereas the background achievement goal $scarf_on$ has been reduced to the activity m_{hand}. The conjunct in question is:

$$\text{IF}$$
$$?HoldsAt(y_{hand}, t), Trajectory(y_{hand}, s, g_1, \delta), s < t, t =$$
$$s + \delta, Happens(x, r), Initiates(x, g_2, r), r < t, g_1_loc1 \neq g_2_loc2 \texttt{ succeeds}$$
$$\text{THEN } ?\neg HoldsAt(ab_{hand.g2}, t) \texttt{ fails}$$

The constraint links failure of the CWA_{ab} in the consequent to the conditions when overriding it is legitimate. As explained in Section 5.2.1, the conjuncts in the conditions refer to simultaneous involvement of an agentive relatum in two differently located goals. The conditions obtain if the involvement is prescribed by planning computations, and is thus necessary in the given circumstances. The reasoner attempts to find an assignment of constraints that

CEU eTD Collection

specializes the variables and allows unification of the non-causal literals ($HoldsAt(y_{hand}, t)$ and $Happens(x, r)$) with formulas obtained previously in plan recognition, and of the causal predicates ($Trajectory(y_{hand}, s, g_1, \delta)$ and $Initiates(x, g_2, r)$) with formulas in the HO-CM.

Given goal derivation outcomes, the substitutions for goal variables are $[scarf_on/g_1]$ and $[e/g_2]$. They have different spatial locations since HO-Narr9 states the location of the scarf on the agent, HO-Narr10 states the location where e obtains is $table$, and $on \neq table$. The substitutions for instrumental behaviors are $[m_{hand}/y_{hand}]$ and $[l_{head}/x]$. They have been obtained in plan recognition as the means to be performed for the two goals. If $[9/t]$, $[3/s]$, $[6/\delta]$, and $[8/r]$ then ***HoldsAt(*m_{hand}*, 9*** succeeds (obtained by forward reasoning in the reduction of background goal $scarf_on$), ***Trajectory(*m_{hand}*,3,scarf_on,6)*** succeeds (HO-CM2), ***Happens(*l_{head}*,8)*** succeeds (obtained by backward reduction of goal e), and ***Initiates(l,e, 8)*** succeeds (HO-CM4). This assignment of instrumental behaviors, goals states, and temporal constraints provides an instance of the successful antecedent of $i.c._4$. For this assignment the conclusion is triggered – $?\neg HoldsAt(ab_{hand.e}, 9)$ `fails`. The constraints provide positive evidence for teleological abnormality of a hand action relative to achievement goal e at time 9, so the CWA_{ab} is overridden. Consequently performance of a light-box contact with the hand at a time when the scarf manipulation is required for the goal *scarf_on*, e.g., time 9, is blocked.

The left branch in Fig. 8 shows the offline planning reduction of the light-on state to contact performed with the default agentive relatum hand. The effect of $i.c._4$ is that this branch fails (Fig. 9) for the suggested assignment of constraints. The top query $?HoldsAt(e, 9)$ is thus univocally reduced to the head action obtained in plan recognition (the right branch).

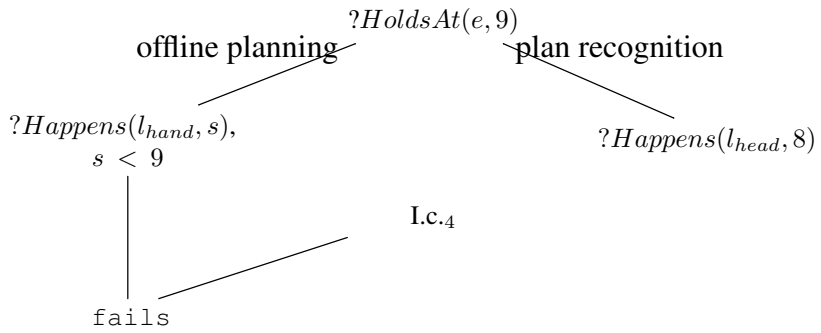

Figure 9. Failing default branch in reduction of goal *light-box is on at time 9* in HandsOccupied.

Overall, in the HandsOccupied condition all events are justified with respect to the main goal conjecture *e*, the side goal *scarf_on* and its subgoal, the scarf manipulation, m_{hand}, within the physical constraints of the situation. Each of the two achievement goals attributed to the agent reduces univocally to a basic observed behavior. The mismatch between the result of offline planning for *light-on* and observations is accounted for because the reduction of the agent's background goal *scarf_on* cancels out the use of the default agentive relatum. Because hands must exert continuous causal influence onto the blanket in order to bring about the side goal,

CEU eTD Collection

there is something abnormal with respect to a hand activity for the goal state *light-on*. The plan for the goal *scarf_on* provides positive evidence of abnormality by application of $i.c._4$. Since hands cannot be the fine-grained agentive relatum of contact with the light-box, planning for *light-on* reduces it unambiguously to the observed head activity. A well-formed plan prescribes that the agent should engage in a light-box directed movement culminated by head contact in order to bring about the conjectured goal. The interpretation of the observed novel head action as *the* means for the goal *light-on* within the constraints of the agent's situation, including the side goal *scarf_on*, , is supported. A univocal holistic plan is assigned to the agent, thereby providing an exhaustive means – ends structure of the situation.

$$HoldsAt(scarf_on, 9) \leftarrow$$
$$Happens(start_m, 3), Initiates(start_m, m_{hand}, 3), (HoldsAt(m_{hand}, 9)$$

$$HoldsAt(e, 9) \leftarrow Happens(start_a, t), Initiates(start_a, a_{head}, t), t <$$
$$7, Happens(l_{head}, 8), Initiates(n_{head}, 8)$$

Because the novel behavior composed of a movement of the head and a head contact with the light-box is contextually justified, and because a default alternative is available in the observer's motor program, cognitive economy considerations do not recommend its storage for use in other situations. The novel instrumental behavior demonstrated by the agent (the head touch) is discarded as irrelevant perhaps unless something abnormal is the case. However the goal that the artifact is used for, i.e. to make light, does count as relevant practical information to be retained, as explained at the end of Section 5.3. The infant who observed the demonstration in the Communicative–HandsOccupied condition is set on learning only that the light-box can be used for goal e, i.e. to make light. Online planning in the HandsOccupied condition is grounded by this specification of the assimilationist epistemic drive – $HoldsAt(learn_e, t)$.

5.5.2. Plan recognition in the HandsFree condition. The background goal $scarf_on$ to be achieved at time 9 is reduced to basic actions. The strategy is similar to the one used to resolve the foreground goal in the HandsOccupied condition. First, Axiom 3 for instantaneous causation is used for query resolution because HF-CM2 provides information about the action of tying the scarf, k, which brings about the desired end state $scarf_on$. The goal is reduced to performance of a knot action with the hand at time 5 (Fig. 10). The database need not be updated because this action is observed and is part of the narrative of events – $Happens(k, 5)$ in HF-Narr5.

Action k brings about the goal state upon culminating the instrumental manipulation of the scarf m_{hand}. $?Happens(k, 5)$ is the next query to be resolved. The reasoner asks what must be the case so that $?Happens(k, 5)$ `succeeds` (Fig. 11). Axiom 1 is used to derive that m_{hand} is clipped between some prior time 4.5 and time 6 – a moment after the occurrence of the knot action terminates it. Axiom 1 then reduces the m_{hand} activity, which holds at time 4 according

CEU eTD Collection

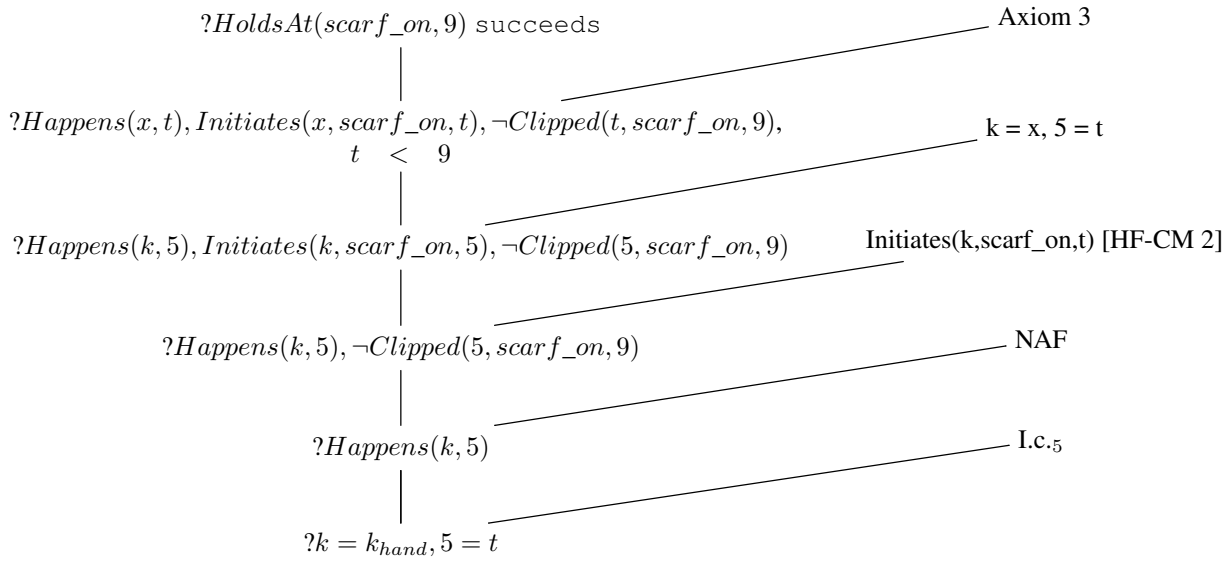

Figure 10. Reduction of background goal $scarf_on$ in HandsFree.

to HF-Narr4, to a starting event that brings it about. NAF in the final step of the derivation, for the temporal interval $(3, 4)$, is warranted because m_{hand} is clipped only later than 5.

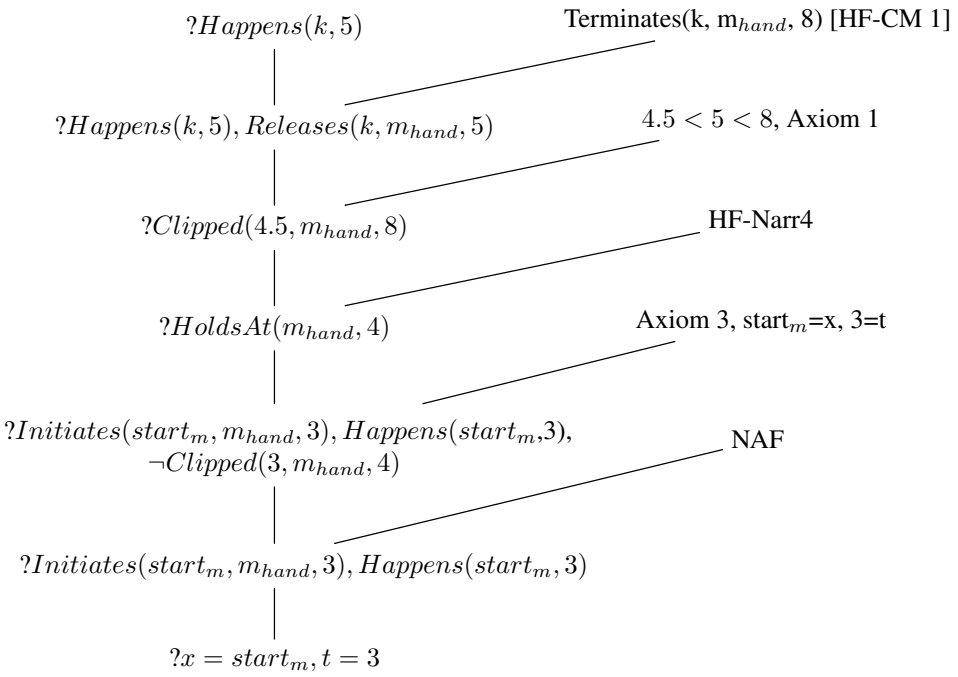

Figure 11. Resolution of the knot query in HandsFree condition.

The background goal $scarf_on$ is reduced in the HandsFree condition to scarf manipulation to be performed until the knot event terminates it. The database must be updated with a starting action for the scarf manipulation behavior. Unlike in the HandsOccupied condition, the scarf manipulation activity m_{hand} no longer holds at the time when the goal state of the scarf being

CEU eTD Collection

on must obtain; it is previously terminated by the knot action k that causes instantaneously the scarf location state at time 5.

The query regarding the main goal ?*HoldsAt(e,9)* `succeeds` is resolved exactly as in the HandsOccupied condition, and yields identical outcomes (see Fig. 6 and Fig. 7 above). That is, it demands the agent's performance of a head contact with the light-box at time 8, which culminates a head movement initiated at a time earlier than 7. The database must be updated with information about the occurrence of a starting event of the light-box directed movement. Because the plan assigned to the agent prescribes performance of a non-default action (head contact) which terminates a non-default activity (head movement towards the light-box), it is depicted by a branching derivation tree like the one in Fig. 8.

The reasoner must verify whether there is positive evidence for contextual abnormality, which justifiably overrides the CWA_{ab}, and thereby cancels the default disjunct. An assignment of constraints is sought, which satisfies the antecedent of $i.c._4$. Both goals have been reduced via Axiom 3 to cause actions, so the antecedent of the i.c. conjunct referring to two goals brought about by instantaneous causation is relevant.

IF

$$?Happens(x_{hand}, t), Initiates(x_{hand}, g_1, t), Happens(x', t), Initiates(x', g_2, t),$$

$$g_1_loc1 \neq g_2_loc2 \texttt{ succeeds}$$

THEN $?\neg HoldsAt(ab_{hand.g2}, t)$ `fails`

The light-on state e is brought about by contact at time 8. Unless planning computations prescribe hand's involvement in another goal-directed behavior performed at a different location precisely at time 8, then the antecedent does not succeed. This is indeed the case, because the knot action k that achieves goal $scarf_on$ takes place and causes the desired state before that, at time 5. Moreover, it terminates the hand's involvement in the inertia of continuous change (cf. HF-CM1), so they are in fact free to be the local agentive relata of an instrumental action afterwards. The condition for teleological abnormality in $i.c._4$ does not hold, thus nothing follows relative to the no-abnormality proviso in the conclusion. There is no positive evidence of abnormality, so the CWA_{ab} is not overridden. No contextual abnormality justifies performance of a non-default behavior for the goal e at time 9. None of the two branches of the derivation tree in Fig. 8 can be shown to fail. Teleological explanation of the agent's novel instrumental behavior is not satisfactory relative to the physical goal conjecture *light-on*.

The idea is rendered formal by the integrity constraint $i.c._6$ that defines the failed explanation of a behavior f with respect to an assigned goal g – $failure_{f.g}$. Given reduction of goal e with Axiom 3, it is the first conjunct that is relevant:

IF $?Initiates(x_1, g, t), Initiates(x_2, g, t), x_1 \neq x_2$ `succeeds` THEN $?HoldsAt(failure_{x1.g}, t)$ `succeeds`

The fact that both branches of the reduction of the foreground goal e are successful reads formally as updating the database with $Happens(l_{head}, 8)$ and $Happens(l_{hand}, 8)$ upon backwards

CEU eTD Collection

planning from e. In previous steps of derivations literals $Initiates(l_{head}, e, 8)$ and $Initiates(l_{hand}, e, 8)$ had been obtained and deleted upon unification with clauses in the respective causal models. The relevant fact is that both $Initiates(l_{head}, e, 8)$ and $Initiates(l_{hand}, e, 8)$ obtain in the reduction of e over the observational and the offline planning databases. Moreover $l_{head} \neq l_{hand}$ obtains, warranted by the Uniqueness of Names Assumption (see Section 4.1.1). Effectuating substitutions $[l_{head}/x_1]$ and $[l_{hand}/x_2]$ in the antecedent makes it true, given the derivations in plan recognition. Hence forward reasoning draws the conclusion that $HoldsAt(failure_{lhead.e}, 9)$ `succeeds`. Teleological explanation in terms of physical causality fails to make the novel behavior intelligible.

The light-box does afford attainment of a desirable state, i.e. the light-on state e. According to Section 5.3, learning that the novel artifact brings about e has been prescribed already. However the particular basic action, i.e. l_{head}, used by the agent to make it work is an unforced sub-optimal one. It is under-explained with respect to e as a goal conjecture; the physical change of state *e* does not justify the non-default instrumental behavior composed of a head movement a_{head}, and contacting the light-box with the head l_{head} at a time when hands are not required to attain another goal in a different location.

$$HoldsAt(scarf_on, 9) \leftarrow$$
$$Happens(start_m, 3), Initiates(start_m, m_{hand}, 3), Happens(k_{hand}, 5)$$

$$HoldsAt(e, 9) \leftarrow Happens(start_a, t), Initiates(start_a, a_{head}, t), t < 7, Happens(l_{head}, 8), Initiates(n_{head}, 8)$$

Because $HoldsAt(failure_{lhead.e}, 9)$ `succeeds`, and $HoldsAt(failure_{f.g}, t)$ is a conjunct in $i.c._8$, the integrity constraint is brought to the fore by the substitutions $[e/g]$ and $[l_{head}/f]$. I restate it here:

IF ?***HoldsAt(teach***$_f, t$***)***, $HoldsAt(failure_{f.g}$***,t)*** THEN ?***HoldsAt(learn***$_f$***,t)*** `succeeds`

$HoldsAt(teach_f, t)$ has been assigned as the agent's pedagogical intention in the Communicative – HandsFree context of presentation (cf. Section 5.3)). Given the upshot of plan recognition relative to the failure of the explanatory attempt, f has been specialized to the instrumental head contact l_{head}. This renders concrete the pedagogical intention, as the intention to teach the under-explained action. t is set to 9 because this is the time that indexes evidence of failure. Under these substitutions the antecedent of the integrity constraint `succeeds`. It follows that the reasoner should learn the head contact with the lamp.

The bottom line of plan recognition is that both the goal derived for the agent's artifact directed behavior, and the novel instrumental behavior itself, are prescribed for learning in the Communicative–HandsFree condition. They motivate action performance in the test phase.

CEU eTD Collection

5.6. Reasoning in the test phase: online planning.

I begin with a review of the outcomes of derivations carried out in teleological interpretation of the Communicative action contexts. Forming the goal conjecture *light-on* activates the reasoner's drive to learn that the novel object achieves it. The result of computations in offline planning can be summarized as 'In order to turn on the light, I would have started hand movement and continued until hand contact with the light-box, assuming that nothing abnormal is the case'. It conflicts with the observation that the agent's head movement culminated by head contact brings about the *light-on* state of the light-box. However in the HandsOccupied condition the CWA_{ab} is justifiably overridden because, based on the conditions of $i.c._4$, hands' bringing about the background goal $scarf_on$ qualifies as abnormality relative to the main goal. The abnormality justifies exceptional performance, i.e. the novel head behavior, for the goal *light-on*. Because something abnormal is the case, the default branch of the derivation computed in offline planning fails. Teleological explanation relative to a physical achievement goal succeeds. In the HandsFree experimental setup on the other hand, both branches of the reduction of goal *light-on* are successful: it is achieved via head touch, and it might as well have been achieved by the default hand touch. No abnormality that leads to the failure of the default branch set forth by offline planning obtains. This means that the novel action sequence is not prescribed as a means for the goal *light-on*, in a minimal model of physical situational constraints. Consequently $i.c._6$ dictates that teleological explanation fails. $HoldsAt(failure_{lhead.e}, t)$ obtains in the HandsFree, but not in the HandsOccupied condition. As a consequence, upon teleological interpretation in HandsOccupied the reasoner is motivated to learn only the goal attained through the light-box, and store it as the light-box function. In HandsFree the reasoner stores also the new instrumental behavior – the head contact, as the concrete content of the agent's pedagogical intention.

I suggested in Chapter 3 that the derivations in online planning, which occur in the test phase of the experiment, are grounded by the infant's intrinsic motivation to learn. When faced again with the novel light-box the reasoner retrieves the information prescribed for learning during teleological interpretation. He engages in backwards reduction of the epistemic maintenance goal to learn only the function of the light-box (in HandsOccupied), or both the function and the means that it affords (in HandsFree). Goal reduction proceeds via the integrity constraints below. They provide achievement goal clauses to be resolved in online planning.

(1) ?$HoldsAt(learn_f,t)$, $HoldsAt(f',t)$`succeeds`;
(2) ?$HoldsAt(learn_f,t)$, $Happens(f',t)$`succeeds`.

$I.c._1$ and $i.c._2$ flesh out the connection between the epistemic maintenance goal $learn_f$ activated during teleological interpretation, and a corresponding achievement goal f' that the infant may act upon by reduction to basic executable behaviors. f in $i.c._1$ is a fluent (a state or an activity) that must be made to hold; f in $i.c._2$ is an action. The integrity constraints operationalize the abstract assimilationist epistemic drive into action performance. They capture the idea that

CEU eTD Collection

the ultimate reason for practical knowledge acquisition is instrumental action. Practical knowledge acquisition is substantiated in action. The i.c.'s cover both kinds of learning exhibited in the task, i.e. of goal states or artifact functions (emulative learning), and of means (imitative learning properly speaking). Parameter f may stand for end states attributed to the agent as achievement goals (if in the conclusion of $i.c._7$), or for the agent's instrumental behavior (in the conclusion of $i.c._8$). $f^{'}$ is an infant fluent or action that is mapped onto the agent fluent or action f. When f is specialized to the derived foreground goal state (in both HandsOccupied and HandsFree) $f^{'}$ is the achievement goal appropriated by the infant in the test phase. When f is specialized to the novel artifact directed action (only in HandsFree) $f^{'}$ is the infant's corresponding motor program and agentive relatum.

The epistemic motivation $HoldsAt(learn_f, t)$ to learn from observations is re-activated by perception anew of the novel object, the light-box. This means that ?$HoldsAt(learn_f$,now) `succeeds`. Hence it is for the young actor to make $HoldsAt(f^{'}, t)$ successful at $now = t$. Maintenance goal $learn_f$ is transformed via one of $i.c._1$ or $i.c._2$ into a concrete physical goal $f^{'}$ to be achieved now. The infant reduces it to executable basic actions.

* * *

Online planning in the HandsOccupied condition. Computations reduce the epistemic maintenance goal in the conclusion of $i.c._7$, triggered upon assignment of a physical achievement goal to the agent's light-box directed behavior – ?$HoldsAt(learn_e, 9)$`succeeds`. It is an instance of the first conjunct in the integrity constraint $i.c._1$ above, obtained by substitution $[e/f]$, and by adaptation to the now of planning computations. The other conjunct provides the subgoal to be achieved now – $HoldsAt(e, now)$. Hence online planning in the HandsOccupied condition amounts to resolving the query ?$HoldsAt(e, now)$ `succeeds`. Resolution seeks an assignment of constraints for basic behaviors from the scenario, which make such that ?*HoldsAt(e,now)* `succeeds`.

The scenario and the integrity constraint for online planning from the achievement goal e are identical with the ones in offline planning (p. 154). Goal reduction is shaped by $i.c._5$ in Section 5.2.1, which specifies the closed-world default representation of basic behaviors. Derivations and their corresponding outcomes are also identical with the ones depicted in Fig. 3 and Fig. 4.

Under the CWA_{ab} that nothing abnormal is the case, implemented by $i.c._5$, the infant reduces the physical achievement goal *light-on* to starting a default instrumental hand movement c_{hand} and performing it uninterruptedly until its culminating event, a hand contact n_{hand}, which initiates the desired state of the light-box.

$$HoldsAt(e, now) \leftarrow Happens(start_c, s), Initiates(start_c, c_{hand}, s), s < t, Happens(n_{hand}, t), Initiates(n_{hand}, t), t < now$$

Goal reduction prescribes performance of the default action sequence. In other words, what infants *would have done* had they wanted to turn on the light-box during action interpretation

CEU eTD Collection

(i.e. in offline planning) is precisely what they actually *do* when they commit themselves to bringing about the goal by using the light-box.

All in all, in the experimental setup where action explanation is successful (i.e. the HandsOccupied condition), learning amounts to appropriation of the achievement goal whose pursuance by the agent has been observed. Planning in a minimal model prescribes default behavior upon this goal, i.e. using the hand to make the light-box work. This is precisely what the majority of infants exposed to action demonstration in the HandsOccupied condition do.

* * *

Online planning in the HandsFree condition. The upshot of teleological interpretation in the HandsFree condition is that the conclusions of both i.c.'s that express the knowledge assimilation drive, i.e. $i.c._7$ and $i.c._8$, are active: the infant must learn that e, as well as the method to e. In the test phase both learning motivations are triggered by observation of the light-box. Online planning computations in the HandsFree resolve two queries: $?HoldsAt(learn_e, now)$ and $?Happens(learn_{lhead}, now)$. The infant plans in order to bring about both maintenance goals.

The resolution of $HoldsAt(learn_e, now)$ is identical with the process that goes on in online planning in the HandsOccupied condition. $I.c._1$ derives the subgoal $HoldsAt(e, now)$. Backwards reasoning prescribes the same default instrumental behavior with the hand – initiating a hand motion, and performing it until finalized by hand contact the light-box.

I elaborate on the computations for learning the observed light-box directed instrumental behavior. If $HoldsAt(learn_{lhead}, now)$ `succeeds`, then $i.c._2$ specifies the subgoal to be achieved by the infant, namely, that he instantiates means learning in his own instrumental action towards the light-box. The subgoal $Happens(l'_{head}, now)$ sets out online backwards planning computations. It is the achievement goal to be acted upon.

The database is by and large similar with the one for offline planning. According to the notation conventions listed in Section **??** the infant action that corresponds to the agent action l is n. The infant maps his own contact action n onto the agent's l based on correspondence between motor programs. The narrative of events contains an additional clause relative to the one used for offline planning and online planning in HandsOccupied, which makes correspondence explicit – $l_{head} \approx n_{head}$. Fig. 12 shows transformation of the l'_{head} action into infant's own head contact, based on the mapping of motor programs and of agentive relata.

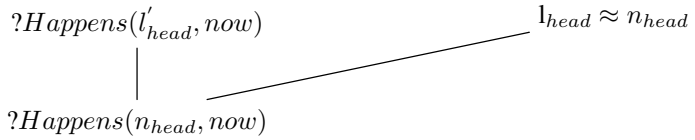

Figure 12. Reduction of goal *instantiate learning in own instrumental behavior*.

From this point onwards, computations are similar with the reduction of the contact event in offline planning, depicted in Fig. 4. Head contact occurs upon performance of a starting event

CEU eTD Collection

of the head movement c, and continuous performance of c until culminated by head contact n_{head}.

$$HoldsAt(learn_e, now) \leftarrow HoldsAt(e, now) \leftarrow Happens(start_c, s), Initiates(start_c, c_{hand}, s), s < t, Happens(n_{hand}, t), Initiates(n_{hand}, t), t < now$$

$$HoldsAt(learn_{lhead}, now) \leftarrow \text{Happens}(\text{l}'_{head}, now) \leftarrow Happens(start_c, s'), Initiates(start_c, c_{head}, s'), s' < t', Happens(n_{head}, t'), Initiates(n_{head}, t'), t' < now$$

All in all, computations prescribe hand behaviors for the epistemic drive to learn artifact goals, and head behaviors to fulfill the motivation to acquire new skills. Because the latter motivation is triggered in case interpretation with respect to physical goals of actions is not satisfactory, enactment of the novel head behaviors may be understood as a practical means to compensate for explanatory failure (as explained in Section 3.2.3).

Performance of the default hand action in the HandsFree test phase is instrumental with regard to the default learning motivation, to learn goal states brought about by new objects. Because both the motivation and the action that it calls for are default, the default instrumental behavior is activated faster and hence its performance is expected before the newly computed one, i.e. $s < s', t < t'$. The behavioral manifestation of the acquisition of novel skills for which a prepotent alternative is available, is delayed by simultaneous learning a goal of the novel artifact and enacting a default motor program for its attainment.

The current model expects performance of both default actions and newly observed head touches directed at the light-box in the Communicative – HandsFree condition. Moreover it expects that performance is temporally ordered as stated above. Yet more, it expects that this condition is the only one where head touches are performed by a significant number of infants. All these 'retrospective expectations' are the case in the experimental findings reviewed in Section 1.2.

CEU eTD Collection

5.7. 'What if she doesn't talk to, nor look at me, nor try to show me anything fun?'

The considerations below provide an answer to the first question set forth in Chapter 1, namely:

> **Question 1:** Why is learning of novel instrumental actions by imitation manifested in the Communicative – HandsFree condition, given that the adult could have used a more efficient means for the goal? Why learn what is not fully comprehensible?

In the Incidental mode of presentation where the experimenter does not engage in communicative interaction with the infant and does not ostensively signal the novel artifact, the scenarios for the HandsOccupied and HandsFree conditions do not contain information about

fluent p. Action bservation is incidental. The narratives are similar to the ones in the Communicative setup, save HO-Narr1 and HF-Narr1, $Initially(p)$. The causal models are identical for Communicative and Incidental contexts, because p does not influence the causal perception of actions and resultant states.

In fact, the only integrity constraint where a formula containing p occurs explicitly is $i.c._1$:

IF $?Initially(p)$ `succeeds` THEN $?HoldsAt(teach_f, t)$ `succeeds`

Since $Initially(p)$ is not part of the scenario, nothing can be concluded relative to the conclusion that refers to the agent's pedagogical intention. The assignment of a pedagogical intention to the agent is not supported by observational evidence.

For the other achievement goals the goal derivation processes are identical with the ones undergone in the Communicative settings. The same physical goal is conjectured for the novel behavior – state light-on. Consequently $i.c._7$ prescribes that the infant learn e. Offline planning is also identical with the Communicative condition. Plan recognition in the Incidental – HandsOccupied condition is analogous with the processes described for the Communicative – HandsOccupied.

The conclusion of $i.c._1$ is a condition for activation of the infant's assimilationist epistemic goal with respect to instrumental actions in $i.c._8$:

IF ?***HoldsAt(teach***$_f, t$***)***, $HoldsAt(failure_{f.g}$***,t)*** THEN ?***HoldsAt(learn***$_f$***,t)*** `succeeds`

In the Communicative – HandsFree context the conclusion was activated because both conjuncts in the antecedent obtained. Now, by making substitutions $[l_{head}/f]$ and $[e/g]$, the second conjunct $HoldsAt(failure_{lhead.e}, t)$ holds in the Incidental – HandsFree condition as well. Teleological action interpretation with respect to physical goals does not differ. The goal conjecture e does not justify the agent's performance of n_{head}. But the first conjunct, the conclusion of $i.c._1$, does not obtain. Forward reasoning does not trigger the infant's epistemic goal to learn the under-explained instrumental behavior with the head. Absence of evidence about p blocks the effect of $HoldsAt(failure_{f.g}, t)$ on computations. Consequently the infant's online planning and corresponding action responses in the Incidental setting in both HandsOccupied and HandsFree conditions are grounded only by the $i.c._7$. The only kind of learning that takes place in the absence of communicative referential signals is the default learning about goal states, or artifact functions. If observation of p offered a way out of explanatory dissatisfaction with the physical goal e via triggering the assignment of a pedagogical intention as the agent's maintenance goal in the Communicative – HandsFree context, in the Incidental – HandsFree condition all that the infant can rationally do is to disregard the odd head behavior.

5.8. Generalizability of the formalization.

Aside from providing a detailed account of the findings obtained for the head-touch experiments of Gergely et al. (2002) and Király et al. (2013) the result of modelling may be

CEU eTD Collection

generalized to explain other related phenomena that evidence learning about actions, artifacts, and their desirable effects.

Finally, I show how the model developed in this Chapter accounts for the results of Pinkham and Jaswal (2011) for an experiment in the head-touch paradigm. The sub-optimal head action was presented as instrumental for the resultant light-on state in two experimental conditions. 25 18-month-old infants in the Demonstration Only condition were presented with an action sequence similar with the Communicative – HandsFree setting, but simpler: the agent did not show behavioral signs of being cold, nor did she manipulate a scarf. She just placed her visibly free hands on the table next to the artifact, and demonstrated the instrumental head behavior. Other 22 infants were assigned to the Self-Discovery condition; they were allowed to explore the light-box by acting upon it for 20 seconds prior to the demonstration. All 22 used their hands to touch the light-box, which was thereby activated. The test phase was identical for the 47 participants – they were presented with the light-box, and the experimenter prompted them to act upon it by exclaiming enthusiastically "Your turn now! Can you show me how this works?". Across conditions all infants achieved the light-on state at least once during the test phase. However, only the infants in the Demonstration Only condition learnt the new instrumental head action, i.e. 72% used their head to illuminate the light. In contrast, in the Self-Discovery condition 14% of the infants did so. Prior experience of successful instrumental action with the hand resulted in disregarding the alternative head method for lighting up the lamp, despite its presentation in an ostensive communicative manner. The authors concluded that in the Self-Discovery condition,

> infants generally disregarded the ostensive demonstration in favor of their own more efficient method, suggesting that efficiency can "trump" pedagogy (Pinkham and Jaswal 2011, p.539).

They interpret the results as showing that by 18 months of age the principle of rationality, i.e. third-person teleological interpretation of instrumental actions, weighs heavier than the principle of relevance, i.e. second-personal understanding of relevant information conveyed communicatively (Gergely and Jacob 2012) on infants' selection of information to be learnt.

I would not agree with the interpretation of efficiency "trumping" pedagogy. The current model explains why it is that all infants exhibit learning of the new artifact function, while performance of the default hand action prior to the demonstration reduces the rate of learning the new skill, as a more nuanced interplay between infants' desire to be practically efficient, and their epistemic drive to learn from reliable adults.

Infants' teleological interpretation in the Demonstration Only condition differs from the computations detailed for the Communicative – HandsFree condition in one respect: it is centered on the assignment of one goal for the only instrumental behavior performed by the agent, i.e. the critical light-box directed movement and the head contact. The integrity constraint $i.c._2$

CEU eTD Collection

in Section 5.2.1 does not affect computations. Upon goal derivation, infants are set on learning that the light-box serves to bring about goal state e. Offline planning is identical with the derivations in Section 5.4. Since no abnormality obtains (the condition of $i.c._4$ does not obtain), plan recognition outputs failure of teleological explanation of l_{head} with respect to e. Hence the conclusion of $i.c._8$ is activated, and infants are prone to learn the opaque head action. There are no substantial differences from the computational structure of the original Communicative – HandsFree condition.

Things are different though in the Self-Discovery condition. I elaborate on the relevant distinctions. In the exploration phase, infants learn by doing that the novel light-box is for making light, which is thus stored as the goal of a light-box directed activity. When the pedagogical demonstration begins, these infants observe an instrumental behavior directed at an artifact that is *not* novel in the sense in which it is for the ones in the Demonstration Only condition, or for the ones in the head-touch tasks of concern. They know what it is for, and how to make it work.

Recent prior experience with the light-box adds to the observational database concrete information about occurrence of infants' default means for the artifact-mediated goal 'light-on'. The scenario for Self-Discovery contains a fine-grained representation of the contact action as hand contact n_{hand} in the narrative and in the causal model. The fine-grained default representation of contact is available from the beginning of the demonstration session; it is expressed in clauses SD-Narr1 and SD-CM1. It is as if the initial database for teleological interpretation includes the one for offline planning described in Section 5.4. The negative temporal index of infants' hand contact n_{hand} expresses retrieval of the information from (short term) memory.

The narrative of events looks like:

SD-Narr: 1. performance of hand contact: *Happens(n_{hand}, -1)*;
2. agent exhibits artifact-directed ostensive signals, and child-directed communicative behavior: *Initially(p)*;
3. light-box is initially off: *Initially($\neg e$)*
4. agent is engaged in light-box directed movement with the head: *HoldsAt (a_{head},2)*;
5. agent's head contacts light-box: *Happens*(l_{head}, 3);
6. light-box is on: *HoldsAt(e,4)*.

The causal model looks like:

SD-CM: 1. light-on is the causal effect of hand contact: *Initiates (n_{hand}, e, -1)* [effect of self-discovery];
2. contact is the culminating point of the light-box directed movement: *Terminates(l, a, 3)* [by inertial perception of the ongoing activity];
3. light-on is the causal effect of contact: *Initiates (l, e, 3)* [by contact causality].

CEU eTD Collection

Integrity constraint $i.c._3$ prompts infants to assign goals. A goal is attributed to the agent's contacting the box. The representation of contact that generates goal derivation is the coarse-grained one in clause SD-CM3. The process is similar with the one involved in the Communicative – HandsFree context, described in Section 5.3.

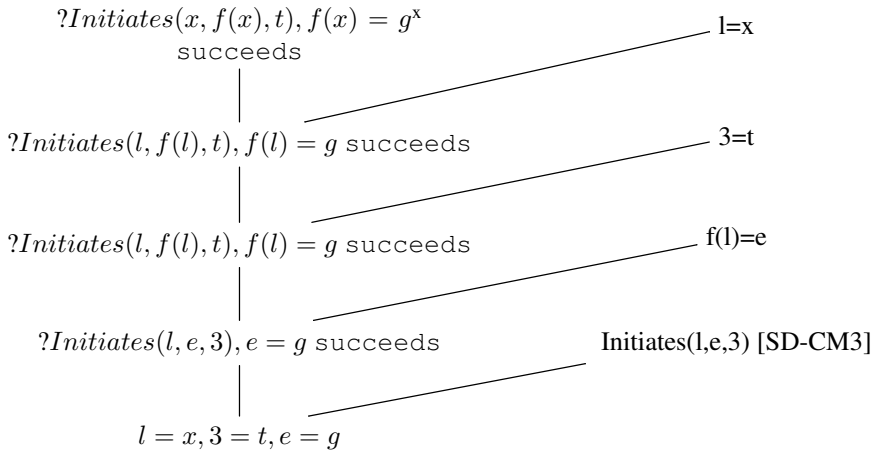

Figure 13. Self-Discovery: Goal derivation for the agent's contact action.

The light-on state e is the goal conjecture assigned to the adult's contact action l. Infants' degree of confidence in the conjecture is very high given their own experience with the light-box (according to criterion (b) in Section 2.4, p. 54). Relatedly, high confidence also characterizes the teleological categorization of the performed default instrumental action as a means for the goal light-on.

Upon specialization of the agent's teaching intention z to e, the conclusion of $i.c._7$ is activated – ?*HoldsAt(learn$_e$,t)* `succeeds`. Infants are motivated to learn the resultant state that the light-box affords.

Clause SD-CM1 replaces the counterfactual offline planning computations. It expresses the recent memory trace of infants' factual light-box directed action. Such processing arguably requires less sophistication, since the reasoners do not have to compute what they would have done habitually; they simply retrieve the fine-grained specification of the action that they actually performed, namely, n_{hand}.

Mismatch between previous performance of hand contact and observation of head contact engages infants in plan recognition. Computations from the goal e branch out: the agent performed l_{head}, while infants did n_{hand}. It is noteworthy that all infants do actually contact the light-box by hand during the exploration phase, while none of them does so by head. The basic action prescribed by recent experience is indexed to its concrete moment of performance, unlike plan recognition as modelled in Section 5.5 where its counterfactuality manifests itself as variable temporal index.

CEU eTD Collection

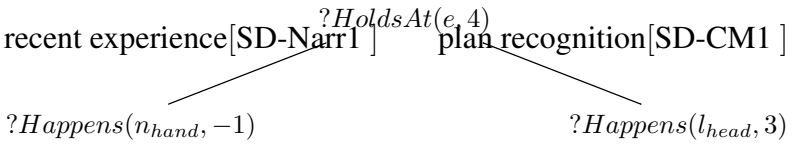

Figure 14. Self-Discovery: Branching derivations from the goal.

SD-Narr5 represents an observed instrumental action whose agentive relatum is different from the hand. Given the first disjunct of the meaning postulate for instrumental default behaviors in $i.c._5$, ?¬ ***HoldsAt(ab***$_{hand.l}$***,t)*** `fails`. The CWA_{ab} is overridden in the agent's context. In order for the teleological conjecture e to be explanatorily useful for the agent's odd behavior, something abnormal concerning the involvement of the agentive relatum *hand* in the action l that brings about e at time 3 must be the case. Infants' high degree of confidence in the goal conjecture and the corresponding default means action makes it such that they need not verify whether there is such an assignment of constraints that makes the conditions of $i.c._4$ obtain, given current observations. The abstract 'something abnormal is the case' does not require specification to a concrete contextual aspect. Infants do not know at this point precisely what is abnormal; however the failed conjunct of $i.c._5$ constitutes sufficient evidence for some unspecified abnormality. It cancels the default action, and thereby justifies the agent's performance of a head action. The derivation tree is in fact similar with the one presented in Fig. 9 for the canceling of one disjunctive means action in the Communicative – HandsOccupied condition. In this case however the rationale for the failing branch is $i.c._5$ (the meaning postulate for default behaviors) plus the acceptance of an *unspecified* abnormality, and not $i.c._4$ or the encounter of a *particular* abnormality like in Fig. 9.

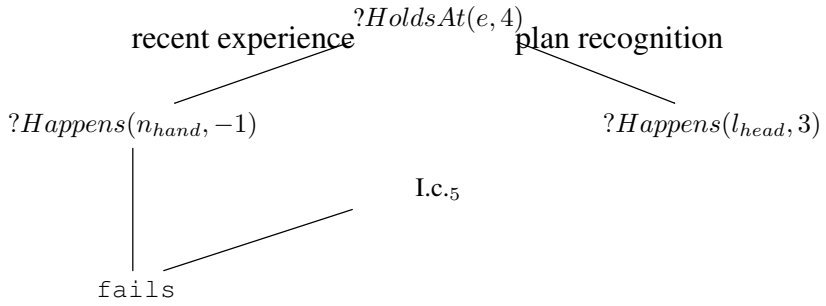

Figure 15. Self-Discovery: Failing default branch in plan recognition.

Since the conditions of $i.c._6$ do not obtain, neither does the conclusion. That is, explanatory failure is not the case. Consequently, upon teleological interpretation in the Self-Discovery condition, the assimilationist epistemic drive with respect to practical skills in the conclusion of $i.c._8$ is not activated. Infants in the Self-Discovery condition are not motivated to learn the head action. Indeed, only 14% of them show evidence of skill learning by performing the new instrumental action with the head.

CEU eTD Collection

Efficiency and pedagogy considerations prescribe infants' learning of the artifact function, but not of the instrumental action. In fact these results provide evidence for how the interaction of both principles shapes infants' understanding of action contexts, and subsequent observational learning. I suggest that the influence of the pedagogical setting on infants' interpretive computations manifests itself as trust that something abnormal is the case in the agent's context, without the need to check furthermore whether the abnormality condition obtains. However, social trust is not sufficient by itself (otherwise, the same would have been the case in the original Communicative – HandsFree condition); confidence in the teleological categorizations involved in own goal-directed action performance, i.e. object-level practical reasoning concerned with physical efficiency, is also required in order to disregard unsatisfactory explanations and learn from observed actions. This is why I stated above that I do not endorse the authors' formulation according to which efficiency "trumps" pedagogy. Both principles are rationally joined together to provide a maximal cognitive benefit in terms of understanding and learning.

CEU eTD Collection

CEU eTD Collection

CHAPTER 6

Concluding remarks

The formal framework presented infants' learning what to do, when, what for in the head-touch task of Gergely et al. (2002) and Király et al. (2013), as supported by two motivational factors: the assimilationist and the explanatory drive with respect to practical knowledge. I suggested that infants act upon these motivations, their maintenance goals, by reducing them to concrete achievement goals and then to actions via the same algorithm used for teleological interpretation of actions. A planning algorithm subserves multi-level teleological explanation of the observed context. The assigned goals are represented at different degrees of abstractions: as physical end-states of actions, and as second-order goals that have a psychological component (intentions or motivations). This sort of analysis gave rise to the multi-level single-mechanism theory of the teleological processing involved in the task. The mechanism is twofold: abductive goal attribution, and planning computations by backwards reasoning from goals. I showed how the processes can be formalized in Constraint Logic Programming with the language of the event calculus. Given the input variables, i.e. observations and background knowledge, represented in the event calculus, CLP computations prescribe the behavior exhibited by the majority of infants in each of the four experimental setups of focus. The model speaks from the infants' perspective in the sense that it expresses the implicit reasoning processes that lead from observations to corresponding action performance in the test phase. It is like using a software that allows one to play videos in very slow motion, and to zoom in on sequences not visible with naked eyes; the 'video' contains the processes undergone in Little Johnny's mind.

In the realm of practical knowledge, one can assimilate information about what things can do (i.e. functions), or about methods (i.e. practical skills). I proposed that the former is the default manifestation of practical knowledge acquisition, given pragmatic arguments, as well as empirical evidence regarding human agents' general focus on teleo-functional information. The end-state is learnt *qua* goal of an action sequence involving an unknown object; the attempt to bring it about by acting upon the object in a subsequent context points to storing it as an enduring function of the object.

Learning novel instrumental actions is underpinned by more complex processing; it supervenes on failed teleological explanation with respect to physical goals for which infants have an alternative available in their motor repertoire. The observed agent's pedagogical manifestations are constitutive of infants' teleological evaluation; the evaluation is meta-representational. Despite its complexity it is plausibly accessible to 14-month-olds. First, the empirical evidence

CEU eTD Collection

discussed in Section 1.1.3 suggests that even younger infants' behavior is guided by the attribution of pedagogical intentions to adults. Second, the model developed here captures reasoning about such intentions via the same teleological mechanism used for physical goals; the latter is also evidenced from earlier on, as discussed in Section 1.1.1.

Infants' two epistemic motivations, to understand and to learn, are intertwined. Learning a practical skill alternative to a prepotent action is based on unsatisfactory fulfillment of teleological understanding. That is, the assimilationist drive, supported in a pedagogical framework established by ostensive communicative signals, compensates for the inability to produce a well-formed teleological explanation with respect to a physical goal. Learning by doing again, i.e. by imitating, is infants' way out of explanatory failure at the object-level. It deems the novel instrumental action optimal with respect to the agent's intention to teach, and the infants' motivation to learn.

In the modelling endeavor I assumed the consensus view in developmental psychology, according to which actions' goals are central with respect to infants' observational learning manifested in imitation and emulation. I adopted a top-down view of the mechanisms by which goals are processed, along the lines of Gergely and Csibra (1997, 2003); Csibra and Gergely (2007). Nonetheless I used insights from bottom-up approaches. For instance, I presented goal attribution in the head-touch task as fostered by perceptual processes, such as recognition of the salience of the actions' results. Furthermore the offline planning stage of teleological processing of observations is a simulationist idea. I suggested that the outcome of planning is selected among infants' motor programs most strongly associated with the goal that grounds planning, which aligns with the theory of action – effect association. One may wonder why then, is an inferential account to be preferred. First and foremost, inferences perform an executive function. They allow integration of various sources of information in an explanatorily satisfactory picture of observations: the counterfactual goal simulation is compared with actual observation, and the mismatch is handled in a principled manner – the search for justificatory contextual abnormalities. The fact that when these are available, i.e. in the HandsOccupied contexts, infants disregard the novel behavior and act upon the goal using their own most efficient action, points to a flexibility of processing that exceeds the explanatory power of low-level accounts. More, the inferential model sets forth a finer-grained account of the phenomena. Where bottom-up theorists speak in terms of direct activation of motor programs (e.g., via motor resonance with observed actions), the closed-world reasoning model elaborates on the actual processes that connect an end-state with the prepotent choice of an action. This detailed description brings along the continuation of the interpretive endeavor, i.e. the search for abnormality, which explains the different findings for the different event structures in the HandsFree and HandsOccupied conditions.

Under these specifications, the answers recommended by the modelling endeavor to the questions raised in Chapter 1 are as follows.

CEU eTD Collection

Question 1:

Why is learning of novel instrumental actions by imitation manifested in the Communicative – HandsFree condition, given that the adult could have used a more efficient means for the goal? Why learn what is not fully comprehensible?

Because the suboptimal instrumental action is recommended by the agent's object-directed ostension and infant-directed communication. The cues frame the action context as a pedagogical setting. The model represents this formally as assigning the agent a maintenance goal that stands for her intention to teach. It raises infants' epistemic trust in the relevance of the subsequent action sequence. In the Communicative – HandsFree condition the transitive instrumental head action is not explained by the physical goal hypothesis 'light-on' entertained with a presumably high degree of confidence. The physical goal placeholder corresponding to the head action remains unfilled upon object-level teleological interpretation. The physical content of the intention is assigned to the unusual action with the head. It triggers infants' corresponding drive to learn it. This is why they exhibit imitative learning of instrumental actions. Despite the action's being suboptimal relative to physical efficiency standards, learning is rational because it is prescribed by planning upon the drive to learn, which is justified by the agent's communicative behavior. The model suggests that imitation be viewed as a form of 'behavioral abduction', whereby infants in the test phase attempt to fill the physical goal placeholder, perhaps by attempting to trigger additional cues from the teacher – agent.

Question 2:

Why is the prepotent hand action always used to act upon the light-box? Why is it used first, i.e. why does imitation follow goal emulation? In particular, if the hand action does achieve the desired end-state, why do infants produce the novel action?

Because in the Communicative – HandsFree condition infants act upon two goals: one is to learn the function of the funny looking light-box, the other is to learn what is otherwise incomprehensible, namely, using the head instrumentally (no metaphor involved). The former is primary because it is prescribed by the default epistemic motivation to learn what artifacts are good for (what is their function). The physical resultant state of the transitive action is stored as the function of the artifact. When the opportunity arises to interact with the artifact, it is reduced by backwards chaining with the closed-world assumption for abnormalities CWA_{ab} to infants' default or prepotent behavior. This answer has been sketched in Section 5.7.

CEU eTD Collection

Question 3:

What shapes the interaction between the two kinds of preconditions of imitation – one social (i.e. communicative ostensive presentation), and one physical-causal (i.e. availability of hands)? How are the two factors weighed in infants' decision what to do, when, what for?

Imitation is a behavioral diagnostics for the acquisition of instrumental actions, which provide an alternative means to infants' default method for achieving a physical goal. Under the assumption that infants favor one-to-one mappings between means and ends, alternative means are worth learning insofar as they are useful for some other, unbeknownst yet, desirable end-state (see response to Question 1). Therefore the extent to which the primary explanatory drive is fulfilled by the physical evaluation of actions with respect to goals determines whether infants learn a new skill from the agent, or simply learn about the artifact kind. The talk in terms of 'the extent to which . . .' points to the generalizability potential of the current model subsequent empirical manipulations in the head-touch paradigm (as described in Section 5.8).

* * *

In a recent paper Gergely and Jacob (2012), along the lines of the distinction between instrumental and communicative agency introduced by Gergely (2010), suggest that observational learning is underpinned by action interpretation with respect to two principles: the rationality principle (Gergely and Csibra 1997, 2003) and the principle of relevance (Sperber and Wilson 1995). The former is teleological reasoning with respect to physical goals; it leads to a teleological structure of action contexts according to criteria of physical efficiency. The latter, brought to the fore by the observed agent's ostensive communicative behavior, supports across-context generalization of information, i.e. learning, even when it remains teleologically opaque upon application of the rationality principle. By and large I sympathize with this approach. The main difference from the theory I pursued in modelling is due to my proposal of a single mechanism that integrates both principles: teleological reasoning with respect to a broader scope of goals. A comparative in depth discussion goes beyond the limits of a concluding chapter; I will emphasize just a few points that recommend my proposal. Given Ockhamian considerations an explanation subsumed by a unique principle is preferable, all other things being equal. This is a rather weak argument, particularly because of the *ceteris paribus* proviso. Focusing on the issue of generalization of information, the theories are indeed equal. I accounted for across-context generalization via infants' epistemic intention to learn, which corresponds to the pedagogical intention assigned to the agent. The epistemic intentions, both infants' and the observed agent's, are depicted as maintenance goals, which are processed via closed-world teleological reasoning. Maintenance goals have been introduced in the AI literature as goals universally quantified over time, i.e. states of the world desirable at all times upon activation of their triggering conditions. Once the content of the intention to learn is specified to a concrete object (e.g., a function,

CEU eTD Collection

an action) given the physical teleological aspects of the situation, it is generalized in virtue of the universal quantification relative to time. Hence teleological reasoning with maintenance goals that stand for epistemic motivations has resources to account for across-context generalization as a diagnostic condition of learning. Furthermore, the two-mechanism account does not fare well explaining in a principled manner how relevance and rationality are integrated in the cognitive phenomenon of learning. A closed-world reasoning view of teleology, on the other hand, states that action interpretation in terms of an agent's maintenance goal to teach new skills is called upon only when the minimal teleological explanation relative to physical goals is not satisfactory. This last point has been operationalized in Section 5.8 by showing how the single-mechanism model accounts for the findings of (Pinkham and Jaswal 2011).

6.0.1. Sketching an answer to 'why model?'. The usefulness of a model grounded on a dataset of empirical results is twofold. It stems from:

(1) the potential of the model to yield further empirical predictions: by modification of the input variables, different output behaviors are to be expected given the cognitive mechanisms spelled out formally in the model;
(2) the generalizability of the model, i.e. its applicability over related experimental findings: a model is applicable to a set of data to the extent that it provides a satisfactory account of the data.

The two are inter-related, since in the non-vicious circularity of science new explanations generate further testable hypotheses, which contribute to theory refinement. The logical model developed here in the form of CLP computations over knowledge represented in the language of the event calculus, exhibits both. It raises further hypotheses that are worth investigating empirically, and it explains subsequent experimental manipulations to the head-touch paradigm used by Gergely et al. (2002) and Király et al. (2013). I now take turns to elaborate on each of the two.

I begin with the issue of deriving testable hypotheses. The current model of learning what to do, when, what for involves two stages subserved by backwards chaining algorithms: explaining another agent's actions, and planning one's own. The memory encoding of observations depends on explanation, and it influences subsequent interaction with the previously observed objects. These memory traces modulate the effect of interpretation on planning, as described towards the end of Section 3.2.3 and Section 5.5. The model has yielded predictions relative to teleological explanation of novel instrumental actions via non-monotonic processing, and its presumed effect on subsequent stimuli encoding and memory consolidation. It fostered setting forth an experimental proposal for an electrophysiological study on adults. The experimental design and the predictions are available in the Appendix; the study has been conducted in the lab of ELTE University in Budapest under the supervision of Ildikó Király. It speaks for the practical utility of the modelling endeavor. Electroencephalography (EEG) is a neural investigation technique that records electrical activity at scalp level. It is used for cognitive science research

CEU eTD Collection

because the electrical activity indexes various sensory, motor, and cognitive processes. The processing of stimuli elicits electrical signatures. Event-related potentials (ERP) is a particular EEG method whereby the averaged electrophysiological responses are stimulus time-locked. This allows to capture accurately the processes generated by critical stimuli, e.g., a contact event that bears causal power in bringing about an end-state. Furthermore, the electrical oscillatory activity can be assessed with EEG methods. This is relevant because electrical oscillations in theta and gamma frequencies have been shown to index memory encoding processes (Ward 2003; Klimesch 1999). A crucial benefit of EEG techniques is the high temporal resolution, in the order of milliseconds. Because of this they warrant conclusions about the fine-grained causal arrangement of neural events elicited by stimulus presentation, and thereby about the corresponding cognitive phenomena. Given that the event calculus is a formal language that allows fine-grained representations of temporal structures of action sequences, EEG is an appropriate methodology to test predictions derived from the model developed in the thesis. Recent ERP studies have confirmed hypotheses regarding the non-monotonic teleological recomputations supported by the brain during discourse processing (Baggio et al. 2008) and explicit conditional inferencing (Pijnacker et al. 2011). We used this method of analysis to test predictions about human adults' action interpretation underpinned by teleological reasoning. The trend obtained in the pilot results confirmed the predictions. The data analysis for the experiment described in the Appendix is current work in progress.

More, the event calculus as a formalism able to express temporal relations in a fine-grained manner seems particularly suited for the EEG method of investigation. For instance, by a strict manipulation of the time variables for stimuli presentation as used in the model depicted in Chapter 5, we could derive fine grained predictions with respect to the causal sequence and temporal distribution of the underlying neural processes that support interpretation.

I now turn to the generalized applicability of the model. I modelled 14-month-olds' acquisition of practical knowledge from observation of adult agents acting upon novel artifacts by unusual methods, as supported by a single mechanism of teleological reasoning used for multi-level action explanation and action planning. In Section 5.8 I showed how the model accounts for the findings obtained in the study of Pinkham and Jaswal (2011). Some generic follow-up issues concerning the generalizability of the model are:

(1) To what other kinds of learning is the model applicable?

A vast array of existing empirical findings shows that functional information and teleo-functional information processing, are drawn on from early ages for artifact categorization (Kelemen and Carey 2007; Futó et al. 2010; Tráuble and Pauen 2007). Practical information, such as the functions of objects, subserves artifact categorization. Knowledge of functions is useful to the development of theoretical knowledge about these objects. The proposed model suggested that learning about functions is underpinned by teleological explanatory evaluation of observations. It is likely then that the reasoning processes involved in the interpretation of transitive instrumental actions

CEU eTD Collection

are recruited to learn about artifact kinds. The development of theoretical knowledge about artifacts, i.e. artifact categorization, may be grounded by counterfactual simulationist planning of own actions upon a new artifact observed while being used by an adult agent. The teleological explanatory mechanism proposed here is based on the assumption that explanation is satisfactory insofar as a one-to-one mapping between an artifact kind and its function can be established. Therefore the essentialist construal of artifacts is likely accountable by such mechanism.

Conversely, I wish to emphasize that in the present format the model does not account for the acquisition of radically novel[1] instrumental actions, i.e. goal-directed actions for which an alternative default is not available in the observer's motor repertoire. In this case the counterfactual simulationist plan of the observer's would not be of help in the explanatory endeavor – she would not be able to provide any answer to the question 'What would I have done had I wanted to achieve ψ?'. A structurally and functionally identical alien who lives on a planet where artifacts do not exist would not learn the head action because her knowledge database does not contain a normal, habitual skill for closed-world reduction of the assigned goal 'light-on'.

(2) To what other kinds of action processing is the model applicable?

The experiments that ground the model are concerned with instrumental actions upon objects, or tools. But an agent's body parts are themselves goal-directed tools. Tool use is in fact secondary instrumental behavior: a body part is used to activate an object, which is used to bring about a certain state of the world. The model may be extended to account for the processing of direct instrumental actions, where the bodily agentive relatum produces a desired effect without artifact mediation. That is to say, nothing in the proposed theory ties it to transitive actions. The model is applicable generally to teleological processing of goal-directed actions. The experiment described in the Appendix does in fact test the generalizability of the interpretive mechanism for goal-directed actions that do not involve tool use (see the four triads of object – unusual action – resultant state described in Section 1.1.).

(3) How do the teleological processes involved in observational learning differ for other age groups and/ or for populations with specific impairments?

I suggested that learning what artifacts are good for is the default kind of practical knowledge acquisition from observation. On the other hand, learning novel instrumental actions relies on some sort of meta-representational capacity; it allows assigning the observed agent a teaching intention, which, in turn, triggers infants' drive to learn teleologically opaque actions with respect to physical goals. The theory can be in principle applied to account for observational learning in populations with certain disorders that

[1]Radical novelty is introduced in Section 3.1, p. 81.

CEU eTD Collection

involve meta-representational deficits (e.g., autism), or in younger infants whose capacity to understand intentional action is not yet in place. It would be interesting to investigate empirically whether the two capacities, suggested here to be underpinned by the same inferential mechanism, can be dissociated for special populations such as the above mentioned ones. More concretely, the theory according to the suggested model would predict that in an imitation paradigm, populations whose meta-representational capacity to reason in terms of intentions is impaired would exhibit only emulative learning, and thereby perform only the actions that are most efficient with respect to physical achievement goals.

(4) For what other kinds of planning could the model be used?

The formalization deals with motor of planning of basic actions, which is not accessible to the agent's awareness; the temporal distance between the achievement goal and the executed actions is very small. However, especially because of the complex sphere of goals that can be reasoned over with the suggested mechanism, I expect that it be applicable to conscious deliberation about what to do, when, what for. Because the focus has been on the assimilationist epistemic intention, I expect that a primary direction of research may be the field of education. A fine-grained description of the processes that underpin learning in a formally established pedagogical framework may be used to enhance the knowledge acquisition capacities of students. Taking into account the cognitive make-up of particular age groups, and the corresponding processing abilities, one may apply the model to right-tailor the emphasis in teaching and learning practices.

Finally, I wish to acknowledge the possibility that the formal model developed here lead to more nuanced, fine-grained predictions by joining forces with other methods for modelling the cognitive phenomenon of observational learning. Furthermore this may contribute to increased explanatory power, i.e. to the generalizability potential. Explicit pitting against each other the logical approach to learning through planning algorithms, with statistical Bayesian models of learning such as (Gopnik et al. 2004; Rao et al. 2004), is a further theoretical endeavor that is called for in the near future. More concretely, given the conceptualization of teleological interpretation in the head-touch task presented in Chapter 3, it could be that a model of infants' degree of confidence in the attributed goals benefits from use of probabilistic methods. Logics and probability theory working side by side may refine the model and improve its degree of generalizability. Aside from theoretical concerns regarding Bayesian models (for which see van Lambalgen and Stenning (2008) and the reply of van Lambalgen and Stenning to Oaksford and Chater (2009)) these could provide a more firm justification for the pursuit of either goal revision, or contextual re-evaluation upon mismatch between observations and offline planning. According to Sadri (2011) such combinations of logical and probabilistic tools for modelling have been used for intention recognition, e..g., Demolombe and Fernandez (2006) drew on the situation calculus and probability theory frameworks, Pereira and Anh (2009) used Causal

CEU eTD Collection

Bayes nets and logical plan generation. Furthermore, some of the assumptions made in modelling, although recommended by the existing findings, require further empirical validation. I assumed that infants' normative understanding of means – ends relations imposes on teleological explanation the standard that actions and goals are linked in a one-to-one bijective mapping. Also, I assumed that the condition for teleological abnormality is the necessary involvement of a given agentive relatum in bringing about another end-state. Both assumptions call for further empirical testing in order to validate their applicability for theorizing upon the role of teleological reasoning in action interpretation, action production, and learning.

* * *

The methodological upshot of the formal model presented in the thesis is that logic and the empirical sciences is, after all, a marriage made in heaven. Despite Wason's popular statement that logic is irrelevant a field of study for the psychology of reasoning (Wason 1968), it appears that a well-chosen formal system should be included in investigations of information processing by human agents. It helps to provide a more fine-grained understanding of the phenomena, and thereby suggests novel directions of research that might confirm or falsify this understanding. Logic contributes to the non-vicious circularity of science: from theory, to empirical tests, to models which refine the theory, and to corresponding new empirical tests. This path of science is analogous, in fact, to the structure of teleological information processing by infants as conceptualized in the proposed model.

CEU eTD Collection

CEU eTD Collection

Bibliography

M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Specification and Verification of Agent Interaction Using Social Integrity Constraints. *Electronic Notes in Theoretical Computer Science*, 85(2):94–116, 2004.

G. Baggio, M. van Lambalgen, and P. Hagoort. Computing and recomputing discourse models: An ERP study. *Journal of Memory and Language*, 59:36–53, 2008.

C.L. Baker, J. Tenenbaum, and R. Saxe. Action Understanding as Inverse Planning. *Cognition*, 113(3):329–349, December 2009.

R. Barr and H. Hayne. It's not what you know, it's who you know: Older siblings facilitate imitation during infancy. *International Journal of Early Years Education*, 11(1):7–21, 2003.

K. Begus and V. Southgate. Infant pointing serves an interrogative function. *Developmental Science*, pages 1–8, 2012.

H. Bekkering, A. Wohlschlager, and M. Gattis. Imitation of Gestures in Children is Goal-directed. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology*, 53(1):153–164, 2000.

M. Brass and C. Heyes. Imitation: Is cognitive neuroscience solving the correspondence problem? *Trends in Cognitive Sciences*, 9(10):489–495, October 2005.

M.E. Bratman. *Intentions, Plans, and Practical Reason*. Cambridge, Mass.: Harvard University Press, 1987.

M.E. Bratman. Practical Reasoning and Acceptance in a Context. *Mind*, 101(401):1–16, 1992.

M.E. Bratman, D.J. Israel, and M.E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, (4):349–355, 1988.

A. Brugger, L.A. Lariviere, D.L. Mumme, and E.W. Bushnell. Doing the right thing: Infants' selection of actions to imitate from observed event sequences. *Child Development*, 78:806–824, 2007.

D. Buttelmann, M. Carpenter, J. Call, and M. Tomasello. Rational Tool Use and Tool Choice in Human Infants and Great Apes. *Child Development*, 79(3):609–626, May/ June 2008.

R.M.J. Byrne. Suppressing valid inferences with conditionals. *Cognition*, 31:61–83, 1989.

J. Call and M. Carpenter. Three sources of information in social learning. In C.L. Nehaniv and K. Dautenhahn, editors, *Imitation in Animals and Artifacts*, pages 211–228. MIT Press Cambridge, MA, USA, 2000.

J. Call, M. Carpenter, and M. Tomasello. Copying results and copying action in the process of social learning: Chimpanzees (Pan troglodytes) and human children (Homo sapiens). *Animal*

CEU eTD Collection

Cognition, 8(3):151–163, 2005.

S. Carey. *The Origin of Concepts*. Oxford: Oxford University Press, 2009.

M. Carpenter, N. Akhtar, and M. Tomasello. Fourteen- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, 21(2): 315–330, 1998.

K. Casler and D. Kelemen. Reasoning about artifacts at 24 months: The developing teleofunctional stance. *Cognition*, 103(1):120–130, 2007.

E. Charniak and D. McDermott. *Introduction to Artificial Intelligence*. Addison Wesley, Reading, MA, 1985.

N. Chomsky. *Language and Mind*. Cambridge University Press, 3 edition, 2006.

L.J. Claxton, R. Keen, and M. McCarty. Evidence of Motor Planning in Infant Reaching Behavior. *Psychological Science*, 14(4):354–356, 2003.

G. Csibra. Teleological and referential understanding of action in infancy. *Philosophical Transactions of the Royal Society for Biological Sciences*, 29:447–458, 2003.

G. Csibra. Action mirroring and action understanding: An alternative account. In P. Haggard, Y. Rossetti, and M. Kawato, editors, *Sensorimotor Foundations of Cognition (Attention & Performance)*, pages 435–480. Oxford: Oxford University Press, 2007.

G. Csibra. Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, 107: 705–717, 2008.

G. Csibra and G. Gergely. The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, 1(2):255–259, 1998.

G. Csibra and G. Gergely. Social learning and social cognition: The case for pedagogy. In Y. Munakata and M.H. Johnson, editors, *Processes of Change in Brain and Cognitive Development. Attention and Performance, XXI*, pages 249–274. Oxford: Oxford University Press, 2006.

G. Csibra and G. Gergely. 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60–78, January 2007.

G. Csibra and G. Gergely. Natural Pedagogy. *Trends in Cognitive Sciences*, 13(4):148–153, April 2009.

G. Csibra, G. Gergely, S. Biró, O. Koós, and M. Brockbank. Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition*, 72(3):237–267, October 1999.

G. Csibra, S. Biró, O. Koós, and G. Gergely. One-year-old infants use teleological representations of actions productively. *Cognitive Science*, 27(1):111–133, 2003.

J. Dancy. *Practical Reality*. Oxford: Oxford University Press, 2000.

A. Danto and S. Morgenbesser. What We Can Do. *The Journal of Philosophy*, 60(15):435–445, July 1963.

K. Dautenhahn and C.L. Nehaniv. The Agent-Based Perspective on Imitation. In C.L. Nehaniv and K. Dautenhahn, editors, *Imitation in Animals and Artifacts*, pages 1–40. MIT Press Cambridge, MA, USA MIT Press Cambridge, MA, USA MIT Press Cambridge, MA, USA,

CEU eTD Collection

2002.

D. Davidson. Actions, Reasons, and Causes. *The Journal of Philosophy*, 60(23):685–700, November 1963.

L. De Raedt. Logical settings for concept-learning. *Artificial Intelligence*, 95(1):187–201, August 1997.

D. Dennett. Cognitive Wheels: The Frame Problem in Artificial Intelligence. In Z.W. Pylyshyn, editor, *The robot's dilemma: The frame problem in artificial intelligence*, pages 41–65. Norwood, NJ: Ablex, 1987.

I. Douven. Abduction. `http://plato.stanford.edu/archives/spr2011/entries/abduction/`, 2011.

S. Duff, J. Harland, and J. Thangarajah. On proactivity and maintenance goals. In *Proceeding AAMAS '06 Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 2006.

B. Elsner. Infants' imitation of goal-directed actions: The role of movements and action effects. *Acta Psychologica*, 124(1):44–59, 2007.

B. Elsner and G. Aschersleben. Do I get what you get? learning about the effects of self-performed and observedactions in infancy. *Consciousness and Cognition*, 12(4):732–751, December 2003.

O. Etzioni, K. Golden, and D.S. Weld. Sound and efficient closed-world reasoning for planning. *Artificial Intelligence*, 89(1–2):113–148, January 1997.

A.J. Fodor. Modules, frames, fridgeons, sleeping dogs and the music of the spheres. In Z.W. Pylyshyn, editor, *The robot's dilemma: The frame problem in artificial intelligence*, pages 139–151. Norwood, NJ: Ablex, 1987.

H.G. Frankfurt. Peirce's Notion of Abduction. *The Journal of Philosophy*, 55(14):593–597, July 1958.

J. Futó, E. Téglas, G. Csibra, and G. Gergely. Communicative function demonstration induces kind-based artifact representation in preverbal infants. *Cognition*, 117(1):1–8, October 2010.

Antony Galton. Operators Vs. Arguments: The Ins and Outs of Reification. *Synthese*, 150(3): 415–441, June 2006.

G. Gergely. The Development of Understanding Self and Agency. In U. Goswami, editor, *Blackwell Handbook of Childhood Cognitive Development*, pages 26–46. Blackwell Publishing Ltd., 2002.

G. Gergely. What should a robot learn from an infant? mechanisms of action interpretation and observational learning in infancy. *Connection Science*, 15:191–209, 2003.

G. Gergely. Kinds of Agents: The Origins of Understanding Instrumental and Communicative Agency. In U. Goswami, editor, *Blackwell Handbook of Childhood Cognitive Development*, pages 76–105. Blackwell Publishers: Oxford, 2010.

G. Gergely and G. Csibra. Teleological reasoning in infancy: The infant's naive theory of rational action: A reply to Premack and Premack. *Cognition*, 63(2):227–233, May 1997.

CEU eTD Collection

G. Gergely and G. Csibra. Teleological reasoning in infancy: the naive theory of rational action. *Trends in Cognitive Sciences*, 7(7):287–292, July 2003.

G. Gergely and G. Csibra. Sylvia's Recipe: The Role of Imitation and Pedagogy in the Transmission of Cultural Knowledge. In N. J. Enfield and S. C. Levinson, editors, *Roots of human sociality: culture, cognition and interaction*, pages 229–255. A C Black Publishers Ltd, United Kingdom, 2006.

G. Gergely and P. Jacob. Reasoning about instrumental and communicative agency in human infancy. In F. Xu and T. Kushnir, editors, *Rational Constructivism: Advances in Child Development Behavior*, volume 43, pages 59–93. Academic Press, 2012.

G. Gergely, Z. Nádasdy, G. Csibra, and S. Biró. Taking the intentional stance at 12 months of age. *Cognition*, 56(2):165–193, August 1995.

G. Gergely, H. Bekkering, and I. Király. Rational imitation in preverbal infants. *Nature*, 415: 755, 2002.

G. Gergely, K. Egyed, and I. Király. On pedagogy. *Developmental Science*, 10(1):139–146, 2007.

P. Godfrey, J. Grant, J. Gryz, and J. Minker. Integrity Constraints: Semantics and Applications. In J. Chomicki and G. Saake, editors, *Logics for Databases and Information Systems*, volume 436 of *The Kluwer International Series in Engineering and Computer Science*, pages 265–306. Springer US, 1998.

A. Gopnik. Explanation as Orgasm. *Minds and Machines*, 8(1):101–118, 1998.

A. Gopnik and H.M. Wellman. Why the child's theory of mind is a theory. *Mind & Language*, pages 145–171, 1995.

A. Gopnik, C. Glymour, D.M. Sobel, L.E. Schulz, T. Kushnir, and D. Danks. A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111(1): 3–32, January 2004.

P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics: Speech Acts*, volume 3, pages 41–58. New York: Academic Press, 1975.

P. Hacker. Agential Reasons and the Explanation of Human Behavior. In C. Sandis, editor, *New Essays on the Explanation of Action*, pages 75–93. Palgrave Macmillan, 2009.

M.M. Haith. Who put the cog in infant cognition? is rich interpretation too costly? *Infant Behavior and Development*, 21(2):167–179, 1998.

G.H. Harman. The Inference to the Best Explanation. *Philosophical Review*, 74(1):88–95, January 1965.

J. Haugeland. An Overview of the Frame Problem. In Z.W. Pylyshyn, editor, *The robot's dilemma: The frame problem in artificial intelligence*, pages 77–94, 1987.

F. Heider. *The psychology of interpersonal relations*. New York: Wiley, 1958.

C.G. Hempel. Rational Action. In *Proceedings and Addresses of The American Philosophical Association*, volume 35, pages 5–23. American Philosophical Association, 1961–1962.

CEU eTD Collection

C. Heyes. Causes and consequences of imitation. *Trends in Cognitive Sciences*, 5(6):253–261, June 2001.

K.V. Hindriks and M.B. van Riemsdijk. Satisfying maintenance goals. In *Proceedings of the 5th international conference on Declarative agent languages and technologies V*, 2007.

B. Hommel, J. Ḿusseler, G. Aschersleben, and W. Prinz. The Theory of Event Coding (TEC): a framework for perception and action planning. *Behavioral and Brain Sciences*, 24:849–937, 2001.

V. Horner and A. Whiten. Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal Cognition*, 8(3):164–181, 2005.

P. Jacob and M. Jeannerod. The motor theory of social cognition: a critique. *Trends in Cognitive Sciences*, 9(1):21–25, January 2005.

S.P. Johnson, D. Amso, M. Frank, and S. Shuwairi. Perceptual Development in Infancy as the Foundation of Event Perception. In T.F. Shipley and J.M. Zacks, editors, *Understanding Events: from Perception to Action*, pages 65–95. Oxford: Oxford University Press, 2008.

A.C. Kakas, R. Kowalski, and F. Toni. The Role of Abduction in Logic Programming. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in AI and Logic Programming*, volume 5, pages 235–324. Oxford: Oxford University Press, 1998.

D. Kelemen. Function, goals and intentions: children's teleological reasoning about objects. *Trends in Cognitive Sciences*, 3(12):461–468, December 1999a.

D. Kelemen. Why are rocks pointy? children's preference for teleological explanations of th natural world. *Developmental Psychology*, 35(6):1440–1452, November 1999b.

D. Kelemen. Are children "intuitive theists"? reasoning about purpose and design in nature. *Psychological Science*, 15(5):295–301, May 2004.

D. Kelemen and S. Carey. The essence of artifacts: Developing the design stance. In E. Margolis and S. Lawrence, editors, *Creations of the mind: Theories of artifacts and their representations*, pages 212–230. Oxford: Oxford University Press, 2007.

I. Király. The effect of the model s presence and of negative evidence on infants the effect of the model's presence and of negative evidence on infants' selective imitation. *Journal of Experimental Child Psychology*, 102:14–25, 2009.

I. Király, G. Csibra, and G. Gergely. The role of communicative-referential cues in observational learning during the second year. Poster presented at the 14th biennial International Conference on Infant Studies, Chicago, 2004.

I. Király, G. Csibra, and G. Gergely. Beyond rational imitation: Learning arbitrary means actions from communicative demonstrations. *Journal of Experimental Child Psychology*, in press, 2013.

I. Király, K. Egyed, and G. Gergely. Relevance or Resonance: Inference based selective imitation in communicative context. to be submitted.

W. Klimesch. EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Research Reviews*, 29:169–195, 1999.

CEU eTD Collection

L. Kotovski and R. Baillargeon. Reasoning about collisions involving inert objects in 7.5-month-old infants. *Developmental Science*, 3(3):344–359, 2000.

A. Kovács, T. Tauzin, E. Téglas, G. Csibra, and G. Gergely. Infants' preparedness to be taught: 12-month-olds point to learn new and relevant information. in preparation.

R. Kowalski. The Logical Way to Be Artificially Intelligent. In *Lecture Notes in Computer Science*, volume 3900, pages 1–22, 2006.

R. Kowalski. *Computational Logic and Human Thinking: How to be Artificially Intelligent*. Cambridge University Press, New York, 2011.

R. Kowalski and M. Sergot. A Logic-based Calculus of Events. *New Generation Computing*, 4 (1):67–95, 1986.

Robert Kowalski. *Meta-logics and logic programming*, chapter Using meta-logic to reconcile reactive with rational agents, pages 227–242. MIT Press Cambridge, MA, USA, 1995.

D. B. Leake. Abduction, Experience, and Goals: A Model of Everyday Abductive Explanation. *The Journal of Experimental and Theoretical Artificial Intelligence*, 7(4):407–428, April 1995.

A.M. Leslie. Spatiotemporal continuity and the perception of causality in infants. *Perception*, 13(3):287–305, 1987.

A.M. Leslie. A Theory of Agency. Rutgers University Center for Cognitive Science, Technical Report TR-12, 1993.

A.M. Leslie. "Pretending" and "believing": issues in the theory of ToMM. *Cognition*, 50(1-3): 211–238, Aprile–June 1994.

A.M. Leslie and S. Biró. Infants' perception of goal-directed actions: development through cue-based bootstrapping. *Developmental Science*, 10(3):379–398, 2007.

V. Lifschitz. Computing Circumscription. In *Proceedings IJCAI 85*, volume 1, pages 121–127, 1985.

V. Lifschitz. Answer set programming and plan generation. *Artificial Intelligence*, 138(1–2): 39–54, June 2002.

P. Lipton. Is the Best Good Enough? *Proceedings of the Aristotelian Society*, pages 89–104, 1993.

Y. Luo and R. Baillargeon. Can a self-propelled box have a goal? psychological reasoning in 5-month-old infants. *Psychological Science*, 16:601–608, 2005.

Y. Luo and R. Baillargeon. Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105:489–512, December 2007.

D.E. Lyons, A.G. Young, and F.C. Keil. The hidden structure of overimitation. *PNAS*, 2007.

D. Marr. *Vision: a computational investigation into the human representation and processing of visual information*. New York: Freeman, 1982.

E. Mayol and E. Teniente. A survey of current methods for integrity constraint maintenance and view updating. In P. Chen, D. Embley, J. Kouloumdjian, S. Liddle, and J. Roddick, editors, *Advances in Conceptual Modeling*, volume 1727 of *Lecture Notes in Computer Science*,

CEU eTD Collection

pages 62–73. Springer Berlin / Heidelberg, 1999.

J. McCarthy. Circumscription – A form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980a.

J. McCarthy. Circumscription – A form of non-monotonic reasoning. *Artificial Intelligence*, 13 (1–2):27–39, April 1980b.

J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 27(1):89–116, January 1986.

J. McCarthy and P. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In *Machine Intelligence*, pages 463–502. Edinburgh University Press, 1969.

D. McDermott. We've Been Framed: Or, Why AI Is Innocent of the Frame Problem. In Z.W. Pylyshyn, editor, *The robot's dilemma: The frame problem in artificial intelligence*, pages 113–122. Ablex Publishing, 1987.

J. McDowell. What Myth? *Inquiry*, 50(4):338–351, August 2007.

A. N. Meltzoff. Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, 24(4):470–476, July 1988.

A. N. Meltzoff. Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5):838–850, September 1995.

A. N. Meltzoff. Imitation as a mechanism of social cognition: origins of empathy, theory of mind, and the representation of action. In U. Goswami, editor, *Blackwell Handbook of Childhood Cognitive Development*, pages 6–25. Malden, MA: Blackwell, 2002.

A. N. Meltzoff. Imitation and other minds: the "like me" hypothesis. In S. Hurley and N. Chater, editors, *Perspectives on imitation: from neuroscience to social science*, volume 2, pages 55–77. Cambridge, MA: MIT University Press, 2005.

A. N. Meltzoff. The 'like me' framework for recognizing and becoming an intentional agent. *Acta Psychologica*, 124(1):106–128, 2007.

G.A. Miller. The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences*, 7 (3):141–144, 2005.

E.T. Mueller. *Commonsense Reasoning*. Morgan Kaufmann Publishers, 2006.

S. Muggleton. Inductive Logic Programming: Issues, results and the challenge of learning language in logic. *Artificial Intelligence*, 114(1–2):283–296, October 1999.

C.L. Nehaniv and K. Dautenhahn. The correspondence problem in social learning: what does it mean for behaviors to "match" anyway? Poster presented at Perspectives on imitation: from cognitive neuroscience to social science, Royaumont Abbey, France, May 2002a.

C.L. Nehaniv and K. Dautenhahn. The Correspondence Problem. In C.L. Nehaniv and K. Dautenhahn, editors, *Imitation in Animals and Artifacts*, pages 41–61. MIT Press Cambridge, MA, USA, 2002b.

A. Newell. The knowledge level. *Artificial Intelligence*, 18:87–127, 1982.

N. Nielsen. Copying actions and copying outcomes: Social learning through the second year. *Developmental Psychology*, 42:555–565, 2006.

CEU eTD Collection

M. Oaksford and N. Chater. Precis of Bayesian Rationality: The Probabilistic Approach to Human Reasoning. *Behavioral and Brain Sciences*, 32(1):69–120, February 2009.

T. O'Connor. Reasons and Causes. In T. O'Connor and C. Sandis, editors, *A Companion to the Philosophy of Action*, pages 129–138. Blackwell Publishing Ltd., 2010.

K.H. Onishi and R. Baillargeon. Do 15-month-old infants understand false beliefs? *Science*, 308:255–258, 2005.

K.H. Onishi, R. Baillargeon, and A.M. Leslie. 15-month-old infants detect violations in pretend scenarios. *Acta Psychologica*, 124(1):106–128, January 2007.

E. Pacherie. The content of intentions. *Mind & Language*, 15:400–432, 2000.

M. Paulus, S. Hunnius, M. Vissers, and H. Bekkering. Imitation in infancy: Rational or motor resonance? *Child Development*, 82:1047–1057, 2011a.

M. Paulus, S. Hunnius, M. Vissers, and H. Bekkering. Bridging the gap between the other and me: The functional role of motor resonance and action effects in infants' imitation. *Developmental Science*, 14(4):901–910, 2011b.

J. Pijnacker, B. Geurts, M. van Lambalgen, C. Kan, J.K. Buitelaar, and P. Hagoort. Defeasible reasoning in high-functioning adults with autism: Evidence for impaired exception handling. *Neuropsychologia*, 47:644–651, 2009.

J. Pijnacker, B. Geurts, M. van Lambalgen, J.K. Buitelaar, and P. Hagoort. Reasoning with exceptions: An event-related brain reasoning with exceptions: An event-related brain potentials study. *Journal of Cognitive Neuroscience*, 23(2):471–480, February 2011.

A.M. Pinkham and V.K. Jaswal. Watch and Learn? Infants Privilege Efficiency Over Pedagogy During Imitative Learning. *Infancy*, 16(5):535–544, September/ October 2011.

J.L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press Cambridge, MA, USA, 1995.

J.L. Pollock. A Recursive Semantics for Defeasible Reasoning. In G. Simari and I. Rahwan, editors, *Argumentation in Artificial Intelligence*, pages 173–197. Springer US, 2009.

Z.W. Pylyshyn. *The robot's dilemma: The frame problem in artificial intelligence*. Alex Publishing Corporation, Norwood, New Jersey, 1987.

W.V. Quine. On What There Is. In *From a Logical Point of View*, volume 1–19. Harvard University Press, 1953.

W.V. Quine. *Actions and Events: Perspectives on the Philosophy of Davidson*, chapter Events and Reification, pages 162–171. Blackwell, 1985.

R.P.N. Rao, A.P. Shon, and A. N. Meltzoff. A Bayesian Model of Imitation in Infants and Robots. In K. Dautenhahn and C. Nehaniv, editors, *Imitation and Social Learning in Robots, Humans, and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press, 2004.

R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 119–140. Plenum Publ. Co., New York, 1978.

R. Reiter. A logic for default default reasoning. *Artificial Intelligence*, 13(1/2):81–132, 1980.

CEU eTD Collection

R. Reiter. On integrity constraints. In *Proceedings of the 2nd conference on Theoretical Aspects of Reasoning about Knowledge, TARK'88*, pages 97–111. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1988.

R. Reiter. What should a database know. *The Journal of Logic Programming*, 14:127–153, 1992.

G. Rizzolatti. The mirror neuron system and imitation. In S. Hurley and N. Chater, editors, *Perspectives on imitation: from neuroscience to social science*, volume 1, pages 55–76. Cambridge, MA: MIT University Press, 2005.

G. Ryle. Knowing How and Knowing That. In *Proceedings of the Aristotelian Society*, volume 46, pages 1–16. The Aristotelian Society, 1945–1946.

F. Sadri. Logic-based approaches to intention recognition. In N.-Y. Chong and F. Mastrogiovanni, editors, *Handbook of Research on Ambient Intelligence and Smart Environments: Trends and Perspectives*, volume 1, pages 346–375. IGI Global, May 2011.

R. Saxe and S. Carey. The perception of causality in infancy. *Acta Psychologica*, 123(1-2): 144–165, September–October 2006.

R. Saxe, J. Tenenbaum, and S. Carey. Secret agents: 10- and 12-month infants inferences about hidden causes. *Psychological Science*, 1001(16):995, December 2005.

R. Saxe, T. Tzelnic, and S. Carey. Knowing Who Dunnit: Infants Identify the Causal Agent in an Unseen Causal Interaction. *Developmental Psychology*, 43(1):149–158, January 2007.

Uwe Schöning. *Logic for Computer Scientists*. Birkhäuser Boston, 2 edition, 2008.

L.E. Schulz, C. Hooppel, and A.C. Jenkins. Judicious Imitation: Children Differentially Imitate Deterministically and Probabilistically Effective Actions. *Child Development*, 79(2):395–410, March/April 2008.

S. Schütz-Bosbach and W. Prinz. Perceptual resonance: action-induced modulation of perception. *Trends in Cognitive Sciences*, 11(8):349–355, August 2007.

S. Schwan and B. Garsoffky. The role of segmentation for perception and understanding of events. In *Understanding events: From perception to action (Oxford Series in Visual Congnition)*. Oxford: Oxford University Press, 2008.

C. Schwier, C. van Maanen, M. Carpenter, and M. Tomasello. Rational imitation in 12-month-old infants. *Infancy*, 10:303–311, 2006.

J.R. Searle. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press, 1983.

J.R. Searle. *Rationality in Action*. Massachussets Institute of Technology, 2001.

S. Sehon. Teleological Explanation. In C. Sandis and T. O'Connor, editors, *A Companion to the Philosophy of Action*, pages 121–128, 2010.

A. Senju and G. Csibra. Gaze following in human infants depends on communicative signals. *Current Biology*, 18:668–671, 2008.

M. Shanahan. Explanation in the Situation Calculus. In *Proceedings IJCAI 93*, pages 160–165, 1993.

CEU eTD Collection

M. Shanahan. Event Calculus Planning Revisited. In *Lecture Notes in Computer Science*, volume 1348/1997, pages 390–402, 1997.

M. Shanahan. The Event Calculus Explained. In *Lecture Notes in Computer Science*, volume 1600/1999, pages 409–430, 1999.

M. Shanahan. Perception as abduction: Turning sensor data into meaningful representation. *Cognitive Science*, 29(1):103–134, January – February 2005.

M. Shanahan. The Frame Problem. `http://plato.stanford.edu/archives/win2009/entries/frame-problem/`, 2009.

J.A. Sommerville and A.L. Woodward. Pulling out the intentional structure of action: the relation between action processing and action production in infancy. *Cognition*, 95:1–30, 2005.

V. Southgate and G. Csibra. Inferring the Outcome of an Ongoing Novel Action at 13 Months. *Developmental Psychology*, 45(6):1794–1798, 2009.

V. Southgate, C. van Maanen, and G. Csibra. Infant Pointing: Communication to Cooperate or Communication to Learn? *Child Development*, 78:735–740, 2007.

V. Southgate, C. Chevallier, and G. Csibra. Sensitivity to communicative relevance tells young children what to imitate. *Developmental Science*, pages 1013–1019, 2009.

V. Southgate, C. Chevallier, and G. Csibra. Seventeen-month-olds appeal to false beliefs to interpret others' referential communication. *Developmental Science*, 2010.

S.E. Spelke, G. Katz, S.E. Purcell, S.M. Ehrlich, and K. Breinlinger. Early knowledge of object motion: continuity and inertia. *Cognition*, 51:131–176, 1994.

D. Sperber and D. Wilson. *Relevance: Communication and Cognition*. Blackwell Publishers: Oxford, 2 edition, 1995.

B. Streumer. Practical Reasoning. In C. Sandis and T. O'Connor, editors, *TheA Companion to the Philosophy of Action*, pages 244–251, 2010.

P. Thagard. *Mind: Introduction to Cognitive Science*. MIT Press Cambridge, MA, USA, 2 edition, 2005.

M. Tomasello. Cultural transmission in the tool use and communicatory signalling of chimpanzees? In S.T. Parker and K.T. Gibson, editors, *'Language' and Intelligence in Monkeys and Apes. Comparative Developmental Perspectives*, volume 1, pages 274–311. Cambridge University Press, 1990.

M. Tomasello. First step toward a usage-based theory of language acquisition. *Cognitive Linguistics*, 11(1/2):61–82, 2000.

M. Tomasello, S. Savage-Rumbaugh, and A. C. Kruger. Imitative Learning of Actions on Objects by Children, Chimpanzees, and Enculturated Chimpanzees. *Child Development*, 64 (6):1688–1705, December 1993.

M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, 28:675–735, 2005.

CEU eTD Collection

J. Topál, G. Gergely, Á. Miklósi, Á. Erdöhegyi, and G. Csibra. Infants' perseverative search errors are induced by pragmatic misinterpretation. *Science*, 321(1831–1834), 2008.

B. Tráuble and S. Pauen. The role of functional information for infant categorization. *Cognition*, 105(2):362–379, November 2007.

M. van Lambalgen and F. Hamm. *The Proper Treatment of Events*. Blackwell Publishing, 2005.

M. van Lambalgen and K. Stenning. *Human Reasoning and Cognitive Science*. Bradford Book, The MIT Press, Cambridge, Massachussets, 2008.

C. von Hofsten. An action perspective on motor development. *Trends in Cognitive Sciences*, 8 (6):266–272, 2004.

C. von Hofsten and S. Fazel-Zandy. Development of visually guided hand orientation in reaching. *Journal of Experimental Child Psychology*, 38(2):208–219, 1984.

L.M. Ward. Synchronous neural oscillations and cognitive processes. *Trends in Cognitive Sciences*, 7(12):553–559, 2003.

P. Wason. Reasoning about a rule. *The Quarterly Journal of Experimental Psychology*, 20(3): 273–281, 1968.

A. Whiten and D. Custance. Studies of imitation in chimpanzees and children. In P. Hayes and B.G. Galef, editors, *Social Learning in Animals: The Roots of Culture*, pages 347–370, New York, 1996. Academic Press.

A. Whiten and R. Ham. On the nature and evolution of imitation in the animal kingdom: reappraisal of a century of research. In P.J.B. Slater, J.S. Rosenblatt, C. Beer, and M. Milinski, editors, *On the nature and evolution of imitation in the animal kingdom: reappraisal of a century of research*, pages 239–283. Academic, New York, 1992.

B. Williams. Internal and External Reasons. In B. Williams, editor, *Moral Luck*, pages 101–113. Cambridge: Cambridge University Press, 1981.

A.L. Woodward and J.A. Sommerville. Twelve-month-old infants interpret action in context. *Psychological Science*, pages 73–77, 2000.

F. Xu. Objectindividuation and object identity in infancy: The role of spatiotemporal information, object property information, and language. *Acta Psychologica*, 102(2-3):113–136, September 1999.

N. Zmyj, M.M. Daum, and G. Aschersleben. The development of rational imitation in 9- and 12-month-old infants. *Infancy*, 14:131–141, 2009.

CEU eTD Collection

CEU eTD Collection

APPENDIX A

Electrophysiological signatures of teleological action processing and memory encoding of novel means actions

The study investigates the differential processing, and correspondingly differential memory encoding, of novel goal-directed actions embedded in two contexts that differ with respect to relevance for the goal. To this end, behavioral (questionnaires) and electrophysiological (ERP, and oscillatory activity) analyses are conducted. The research explores the electrophysiological signature of defeasible teleological action processing, and its presumed effect on memory encoding, exhibited in neural oscillatory activity.

The study assumes elaborative encoding, by which memory encoding of the presented action stimuli is influenced by the computational load of action processing. We hypothesize that the processing of novel goal-directed actions is more computationally demanding when embedded in a goal relevant context compared to a goal irrelevant context, and that the increased processing effort takes effect on stimuli encoding. In order to test this hypothesis we record neural activity while participants observ action sequences. Sequences consist of 2 actions: a successful unusual goal-directed action, preceded either by a goal-relevant, or by a goal-irrelevant action.

On the electrophysiological level, processing of goal relevant and irrelevant first action of the pairs was compared. Specific ERP components are expected as indices of various computational demands imposed by teleological action processing. Oscillatory activity is expected to indicate differential memory effects.

The study started with testing healthy adults; an extension to action processing and novelty encoding in infancy is planned.

CEU eTD Collection

A.1. Method.

Within subjects design. All subjects watch 4 types of unusual goal-directed transitive (i.e. artifact-directed) actions, performed by 2 agents (1 male and 1 female), embedded in 2 action contexts (Goal Relevant Vs Goal Irrelevant).

Stimuli. The 4 triads of object – unusual action – goal are:

(1) light-box – head contact – light on (Gergely et al. 2002);
(2) transparent box with a lamp inside, covered with a lid, on top of which is the lamp switch – sit on box – light on;
(3) 2 boxes, one inside the other – contact inner box with a magnet hanging from a ball – inner box removed (adaptation of Királyi 2009);

(4) small toy suitcase – act on the suitcase button with a beer opener – open suitcase.

Novelty of means action stems from either usage of an alternative bodily agentive relatum instead of a prepotent action with the default agentive relatum (i.e. hand) (1. and 2.), or tool use instead of prepotent hand action (3. and 4.).

The novel stimuli are embedded in 2 action contexts:

Goal Relevant: Failed Prepotent – Novel Successful: Model performs 2 actions onto the artifact. 1st action is a failure of the prepotent means action (with the hand) onto the artifact. 2nd action is novel/ unusual, and it is successful. A trial is composed of the 2 actions.

Goal Irrelevant: Prepotent Irrelevant – Novel Successful: Model performs 2 actions onto the artifact. 1st action is irrelevant for the goal (back and forth movement). 2nd action is novel/ unusual, and it is successful. A trial is composed of the 2 actions.

Experimental design.

Familiarization: Subjects watch videos of novel successful action. The aim is to familiarize them with the goal state, and with the unusual means action used to achieve it. 4 times X 4 Novel Successful videos (2 with male agent, 2 with female agent, randomized). In between every 2 videos a blank screen with a fixation cross appears for 1000 ms.

Post-familiarization questionnaire: Participants are asked whether they would perform the novel action (referred to as *the unusual action*) for each of the 4 tasks. They are asked to give reasons for their answer.

Would you do the same unusual action if you wanted to . . . (each of the 4 end-states requires a Yes, because . . ./ No, because . . . answer)

Test phase (EEG recording): Subjects watch in randomized order the two 2-action trials for the 2 conditions Goal Relevant (failed prepotent – novel successful) and Goal Irrelevant (prepotent irrelevant – novel successful). One trial is composed of 5 pictures: wave hello (attention grabber), critical action (Failed Attempt or Prepotent Irrelevant), neutral face, novel action, goal attained + happy face. The 3rd picture (neutral face) stays on screen longer than the others (+ 2100 ms). In between every 2 trials a blank screen is presented for 500 ms.

Only the 2nd picture varies across the 2 conditions. In a FA-NS or PI-NS conditions, respectively, actors perform a hand action in a goal relevant (failed attempt) or irrelevant manner (prepotent irrelevant), respectively; no effect ensues on the manipulated object. This procedure allows to compare the two conditions while the external stimuli (showing the actor with a neutral face between the following the first action and before the goal achievement) in the two condition are identical.

Test phase has 8 blocks, with small breaks in between (length of break decided by the participant – button press when they want to go on). Each block consists of 24

CEU eTD Collection

trials (3 X 4 FP-NS + 3 X 4 PI-NS) in randomized order. Overall, test phase consists of 84 FP-NS trials (in the Goal Relevant condition), and 84 PI-NS trials (in the Goal Irrelevant condition).

ERP's are measured following the onset of the 1st action. Amplitude of sustained oscillatory activity is measured in 6 frequency bands (delta 0.5-4 Hz; theta 4-8 Hz; lower alpha 8-10 Hz; upper alpha 10-13 Hz; beta 13-30 Hz; gamma 3-45 Hz) following 2000 ms the offset of the goal relevant action is compared with the activity induced by the goal irrelevant action.

hello	FA	neutral face	NS	goal state	
2000 ms	2000 ms	2000 ms + 2400 ms	2000 ms	2000 ms	1000ms
		recording			

hello	PM	neutral face	NS	goal state	
2000 ms	2000 ms	2000 ms + 2400 ms	2000 ms	2000 ms	1000ms
		recording			

CEU eTD Collection

Post-recording questionnaire: Participants are given a longer questionnaire, which includes the question in the post-familiarization stage.

1) Why do you think the agents performed the unusual action? (open question)

2) Do you think the agents' choice of the unusual action was well justified? (Yes, because . . ./ No, because . . ./ Sometimes (please specify when), because . . .)

3) Would you do the same unusual action if you wanted to . . . (each of the 4 end-states requires a Yes, because . . ./ No, because . . . answer).

A.2. Objectives.

Results in the pilot experiment. The pilot study was identical, except that videos of the actions were used as stimuli.

FA (failed attempt of the prepotent goal relevant means action) compared to the PI (prepotent goal irrelevant action) evoked response: negative potential around 400 ms after the end of the videos, only in the frontal brain areas. The increased negativity can observed only for the left fronto-temporal, but not for the left temporal electrode.

Delta, Theta and Gamma absolute power tends to increase in the frontal brain areas of the left hemisphere after the Failed Prepotent, but not after the Prepotent Irrelevant actions. Increased delta power indexes increased attention. Increased theta and gamma power indexes subsequent memory effects, and increased maintenance of the information encoded.

Hypotheses. Action processing is hypothesized to be subserved by defeasible teleological inferences (i.e. analysis in terms of goal states and means actions). Semantic teleological action processing is more demanding when the unusual goal-directed action follows a failed attempt to achieve the goal with the usual means, because it requires recomputation of the teleological structure. Moreover, such teleological recomputations have been shown to be indexed by sustained negative electrical activity, observable particularly in left anterior regions (Baggio et al. 2008; Pijnacker et al. 2011). The Goal Relevant context calls for teleological recomputation: the novel action is categorized as means to the detriment of the habitual one (the prepotent, or default, response).

We expect enhanced event related response around 400 ms after critical picture of the failed prepotent actions in the Goal Relevant condition, compared to the response evoked by prepotent irrelevant actions, as an index of increased teleological processing (recomputation). We expect the negative shift in the left anterior sites to be temporally sustained, and to lack a peak.

With respect to memory, we rely on the widely accepted view that encoding stimuli, and further consolidating memory traces in the absence of the external stimuli is indexed by electrical oscillations (in line with Hebb's proposal about consolidation of memory traces). Based on that, we compare the amplitude of the oscillatory activity in 6 frequency bands following 2000 ms after the offset of the goal relevant actions (failed attempts), with the activity induced by the goal irrelevant actions. We expect topologically distinct increased activity for the enhanced

CEU eTD Collection

memory encoding processes in the Goal Relevant condition, indexed by oscillations in the theta and gamma bands.

We expect enhanced event related response around 400 ms expected in Goal Relevant condition compared to the response evoked by prepotent irrelevant actions, as an index of increased computational processing. We expect the negative shift to be more sustained and to lack a peak, as an index of computational effort demanded by teleological recomputation in the Goal Relevant condition.

CEU eTD Collection