Combinatorial interpretations of frameworks of emergence offer novel ways to study synergy, self-organization and redundancy

By

Mihaly Hanics

Submitted to
Central European University
Department of Network and Data Science

In partial fulfillment of the requirements for the degree of
MSc in Social Data Science

Supervisor: Federico Battiston

Vienna, Austria

May 2025

Author’s Declaration

I, the undersigned, **Mihaly Hanics**, candidate for the MS degree in Social Data Science declare herewith that the present thesis is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography. I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes on any person’s or institution’s copyright. I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, 22.05.2025

Mihaly Hanics

Copyright Notice

Abstract

The study of complex systems, such as the human brain, the weather, or societies, has been a core scientific challenge in their respective fields. These systems, despite their incomprehensible nature, have similarities and prompted researchers to establish the field of complexity science, which studies complex systems systematically, aiming to provide general tools for analysis. Such works include studying features that commonly appear in these systems, such as emergent behaviour (entities acting collectively as a whole). A common way to study emergence, the presence of emergent effects and patterns, is through statistical approaches, built on tools from information theory. While such frameworks are more applicable on real world systems than other ones, they do not typically account for the algorithmic component of the dynamics, and only capture the effects over a large aggregated scale. For this reason, approaches that focus on the algorithmic components of the dynamics may offer novel insights. In this work, I formulate the simple concept of dynamical impact, a way of tracking causal influences at the micro-level, which I show is related to both synergy and redundancy. Then, I reformulate two major statistical frameworks of emergence, the concept of synergy in the partial information decomposition (PID) framework and its generalizations, and the dynamical independence (DI) framework. Examples on cellular automata are shown. In particular, I categorize different types of interactions in Conway's Game of Life based on dynamical impact and measure how synergistic or redundant these types of interactions are when viewed through an information-theoretic lens. Using this concept of dynamical impact, I develop new combinatorial notions of synergy, redundancy, and dynamical independence, explore some of their properties, and discuss connections between synergy and dynamical independence. Experiments on simulated systems support the alignment between these new combinatorial definitions and their classical counterparts. I argue that this approach, while primarily applicable to systems where the algorithmic aspects are known (typically simulations), may still serve as an important link towards a more unified framework of emergence and a deeper theory of complex systems by better connecting these major concepts.

Acknowledgements

To my mother and my girlfriend, who have supported me through this process all along, pushing me and rooting for me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Complex systems are all around us – from the workings of the human brain, through ecology, to the financial markets, societies and ecosystems. They stand out from other non-complex systems by producing surprising, chaotic, large-scale behaviours from the interactions of many smaller parts. Understanding how these ”emergent” patterns come about is a central goal of complexity science. While we have novel tools and frameworks to study these systems, particularly those based on information theory it seems we still haven’t even come close to discovering the hidden secrets of these systems

Existing research tends to look at different aspects only on coarse-grained level, or study components such as structure, dynamics, or initial conditions separately. No such approaches have given us a comprehensive picture, and possibly they may never be able to provide that, because complex systems are sensitive to small deviations too, making such approaches too ”harsh”. Instead, we shall find approaches more fine-grained, incorporating every aspect, to better understand the general theory of these systems.

This thesis tackles this challenge by exploring a more direct, combinatorial approach to understanding how emergent patterns are built from the ground up. This is possible and done in cellular automata, as those are artificially defined systems. The central problem I address is the need for a framework that can better capture the causal unfolding of dynamics from micro-level interactions.

To address this, I introduce and investigate a concept I call ”dynamical impact.” This is a simple idea: a component has an impact if its state directly makes a difference to the state of another component in the next step, given the context of everything else happening. The main contribution of this thesis is to develop this notion of dynamical impact into a basic combinatorial tool and to explore how it can help us analyze emergent phenomena. My objective is not to overturn existing frameworks, but rather to offer a complementary, causally-grounded perspective that can potentially bridge different views of emergence, such as those focusing on synergistic information versus those emphasizing self-controlling, dynamically independent subsystems. I aim to show that by looking at these fundamental impacts, we can tackle complexity with a clearer, simpler, and more intuitive toolset.

The data used is manually generated. I generated systems based on certain parameters, and I generated randomized dynamics on top of these systems, these make up the data that was used for analysis. The simulations were run on cellular automata and graph models, and allow for other networked models, which is important for future work. I compare findings with established information-theoretic measures popular in the field. The analysis is therefore both conceptual – developing the idea of impact – and computational – applying it to concrete examples.

The thesis is structured as follows:

Chapter 1 lays down the core concepts of complexity science and discussing key ideas and existing frameworks related to emergence. This provides the context for the work that follows.

Chapter 2 details the software implementation developed for this research, explaining the chosen software design.

Chapter 3 forms the core of the thesis. Here, I formally introduce the concept of "dynamical impact." I then explore how this combinatorial approach can be used to re-examine and potentially provide new insights into established frameworks of emergence, such as PID and DI, and demonstrate its utility through experiments on simulated data. Finally, I conclude the thesis by summarizing the findings.

Chapter 2

Concepts of complexity science and related work

This chapter serves as a combination of explanation of the fields' concepts and literature review, as concepts related to the thesis are vast and need lengthy explanation.

As this thesis lies within the field of complexity science, it is essential to firstly introduce the goals and boundaries of the field, and the relevant concepts, before summarizing the literature relevant to this work.

2.1 Comprehensive description and discussion of concepts and ideologies in the field of complexity science

Complexity and complexity science: Complexity science is a field that studies systems such as the human brain, the weather, societies, markets, living organisms, condensed matter – systems that we call "complex" as we have not captured their entirety in simple equations – and makes attempts to deal with them. It provides tools to better understand, model, and predict systems that are hardly interpretable, i.e., complex [1].

What counts as "complex" in this context ? The most frequent notion is that a system is considered complex if it does not have a feasible description with mathematical formulas that describe the state(s) of the system, at a given point in time or over time. This definition is closely related to Kolmogorov complexity, which conceptualizes the shortest description length of an object as a measure of its complexity [2]. The field makes a clear distinction between complexity and complication. Complicated systems may be lengthy or detailed, but are straightforwardly interpretable. A good example is a large decision tree: despite its size, it has a straightforward representation of the outcomes based on the inputs. In contrast, a relatively small system, such as the neural connectome of the C. elegans worm ($\tilde{3}00$ nodes), is hard to compress into interpretable, explicit equations, and thus considered complex. Often, such systems have surprising and important patterns at various scales, making them hard to concisely describe completely. These arise from the interactions between entities in the system, which introduce different patterns than the entities alone would produce, making the dynamics chaotic. Typically, it is hard to describe the dynamics of a complex system over time even with a few parameters in a simple, explicit, and compact manner – as an illustrative example, the 3-body problem only has three parameters (positions of objects), yet despite that, no closed-form analytical solution exists using elementary functions and integration, and infinite series expressions found are only slowly converging.

A natural question to then ask is: What is the value in (systematically) studying complex systems if they are difficult to model and seemingly uninterpretable? Since many real-world systems that we care about are complex (as is the case with the weather), the field could offer general tools to study any system that is complex. A core motivation is that many complex systems share universal patterns, such as the presence of emergent phenomena, self-organizing structures, sensitivity to initial conditions, and more, suggesting there are hidden regularities across these systems. For example, if we could discover the causal effects of these universal patterns, we might better understand how to capture them, model the system with regard to such patterns, or even reveal mechanisms underlying complexity itself.

Below I explain the key concepts of the field, including common phenomena in complex systems:

Entities and interactions: Most complex systems consist of smaller, distinguishable, interacting parts. It would be hard to consider a system as complex that does not have parts that interact with each other; any such system could be perfectly described by the juxtaposition of its components (if multiple) and would be simply interpretable as the mere collection of the parts. (The system would not be more complex than its most complex component, just more complicated.) The interacting parts together make up networks, modelled mathematically with graphs. These can be of various types – simple graphs, hypergraphs (particularly important in for studying higher-order relationships [3] [4] [5] [6]), temporal graphs, multilayer networks, and even combinations of distinct types. These components can be called nodes, elements, entities. This thesis uses the term entities. Entities have states in any classical complex system, and these states may change over time. Interactions may induce the changes in states. As an example, if a person can support a project for an initiative, but due to social influence from his friends, ends up supporting another project (the project that is chosen to be supported here is considered as a state). Connections (called edges in graphs) define the channels where interac-

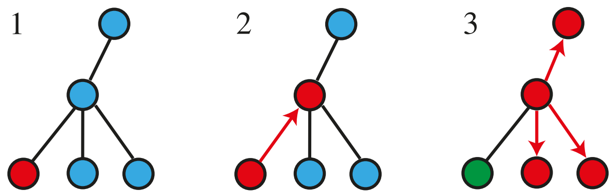

Figure 2.1: An abstraction of a disease spreading as an illustration for a dynamic system. The circles represent humans (entities), which interact through the edges (personal contacts) with each other. In the second timestep, an infected person (red node) interacts with another person, and infecting them, who in the next timestep infects all other contacts. The dynamics encapsulate the interactions and how the disease spreads. We can measure different observables at different scales: on entity level, we can analyze when a person got sick and for how long, whereas on macroscale we can work with aggregated measures, such as on average how many people are infected in a timestep.

tions can or do happen.

Mathematically, the mentioned network types differ only slightly (they are special cases with additional constraints, but fundamentally represent basic subclasses of graphs), but software implementations often treat them as separate cases for simpler implementation, which has its drawbacks. In this work I developed a unified structure class (encapsulating any object with entities and the connections facilitating interactions) to derive these from this shared representation.

Dynamics: Refers to the changes in the system over time, particularly the states of the entities, driven by its underlying mechanism (set of rules). It encompasses the evolution of the states, typically represented as a temporal sequence. Dynamics can be continuous, but in this work I focus on discrete time systems.

Behaviours, patterns, features: The observable outcomes of the dynamics. In this context, higher-order behaviours or features refer to emergent properties (collective behaviour from interactions) in the system [7]. Common patterns include presence of modular (seemingly independent) structures, states synchronizing in groups. Using novel methods described below, such patterns can be statistically analyzed and detected across various levels (scales, sizes).

Different scales: Patterns can appear on different spatial (or temporal) scales [8]. (Space in this notion may not have a geometric interpretation, but can refer to any dimension along which entities are organized, e.g. a feature space.) The scale of the quanta, the scale of individual entities (e.g. neurons or particles), is denoted as local or microscale. If the scale encapsulates the entire system including all entities and their connections, it is denoted global or macroscale (e.g. the brain organ). In between the two ends, patterns can appear on the level of groups of entities (e.g. modular communities), this is called the intermediate or mesoscale. Some literature refers to emergent mesoscale subsystems as macroscale, denoting everything emergent simply as “macro”. The term “higher-order level” is also used on any scale other than the smallest possible (micro)scale, i.e. scales with multiple entities.

Complex systems often display different patterns on different scales. In economics, the different scale patterns is the reason why microeconomics and macroeconomics are fundamentally different fields not only in their scope of analysis but also in the laws and observables.

Mechanisms: The mechanism of a system is the set of rules driving its dynamical evolution [9]. In this work, I assume mechanisms to be described and to operate on micro-level, meaning that the rules driving a system are described by (separate) formulas for each entity and apply individually to entities. This does not have to be the case, there are macro-level mechanisms used in statistical physics, e.g. the Ising model and other models described by Hamiltonian functions [6] [9], having one general rule encompassing all entities, and mechanisms with mesoscale rules are also possible.

When mechanisms are defined on the micro level, we make use of the so-called most fine-grained causal structure: studying the properties of mechanisms and how variation in the rules (and initial states) changes the outcomes, we can find the true causal effects via derivations (which are, as a side effect, also explanatory). This is a core ideology underpinning the methodology of the thesis; I advocate for frameworks that combinatorially causally derive patterns through mechanisms (and initial conditions), over statistical inference within specific systems. (This position is in line with Karl Popper's emphasis on deductive reasoning and rigorous derivations over experiments for reasoning, described in *The Myth of the Framework* [10]). Moreover, studying mechanisms can also enable analysing the joint effects of structure and dynamics, rather than studying them in isolation. Much of the existing literature focused on either the effect of the structure (topology) [11] or the dynamics (e.g. different models of contagion on the same network structure) [9], by isolating one aspect and varying the other – but such approach cannot capture the synergistic (emergent) effect of the two.

Sensitiveness to initial conditions: Complex systems often display chaotic behaviour. Even small differences in the starting states of a system can result in widely divergent outcomes. (The beauty of it is that these outcomes can contain clear, unique patterns and that are not merely random, as often imagined.) This concept is discussed below in the context of difficulty in (long-term) prediction of complex systems. Chaotic behaviour makes systemic deviations of a model and the real-world system more apparent over time, as it amplifies inaccuracies.

Emergence: "Emergence" is likely the most widely discussed and defining complex system feature, alongside complexity itself; some authors even define complex systems as systems exhibiting emergent features. Emergence typically refers to the presence of arising macro level (dynamical) patterns, structures, behaviours otherwise not apparent on microlevel. This means that entities themselves do not possess these patterns, these arise from the collective behaviour of entities, through interactions. An emergent pattern is thus any pattern that exhibits some collective behaviour. Interactions are necessary for patterns to exist that appear at a collective level.

Emergence has been a debated concept of science and philosophy, as it is argued that emergent patterns cannot be fully understood on the level of entities, examining entities alone, but only on an appropriate level of collections of entities. For example, one cannot understand society completely even if they understand every individual, as the interactions between individuals also shape society. The seminal paper of physicist Philip Warren Anderson titled *More is Different* [12] formulated this concept within science, stating that some laws of systems with many entities can differ from what is derivable from the combinations of laws governing their parts. In the lens of physics, this also suggests that the laws of quantum physics alone are insufficient to describe chemistry and biology.

This concept is closely related to the idea of holism (with a modern interpretation) or constructivism, which believes that real world systems possess laws or properties that cannot be decomposed into elementary (quanta) parts without a loss of (crucial) information. This can occur if some laws are only apparent to groups or combination of the parts, therefore are not simply universal laws holding in any setup. This goes against the reductionist perception, which argues that all phenomena can ultimately be derived from the lowest-level fundamental rules governing fundamental components, and generally seeks to break down systems to their minimal concepts. Often, a system with emergent patterns is described as more than sum of its parts, whereas reductionism believes that systems are exactly the sum of their parts (the validity of these statements may depend on the context). It is worth to note that statements being true (e.g. a phenomenon happening in a complex system) not being provable in the axiom system (microscale rules) despite being true is a phenomenon described by Gödel in 1931 [13], known now as Gödel's incompleteness theorems, although probably there is no connection between the incompleteness of an axiom system and the incompleteness of microscale laws.

It is important to clarify that antireductionism, particularly holism, does not by any means state the existence of "magical" higher-order laws (in dynamical complex systems) that cannot be derivable. It rather argues that many observable patterns in complex systems, that are not explicit in the axiom system (set of fundamental laws), could depend on the initial states of the system too. In the notable cellular automaton, Conway's Game of Life – a famous complex system with a simple ruleset [14] – so called "gliders" can appear. These are group of cells exhibiting repeating patterns while propagating across the grid to infinity. This is not typical behaviour of objects in Game of Life, mobile periodic patterns (called spaceships) are rare. The observable (deterministic) behaviour is thus limited only to particular dynamical subsystems, which may or may not appear based on the initial conditions. Therefore, they are not solely a

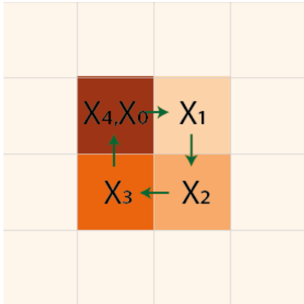

Figure 2.2: A periodic pattern apparent from the (collective) interactions defined by the mechanism, and is ”deducible to micro level”, despite being a mesoscale pattern.

product of the atomic laws of the system, but of the laws combined with the particular initial states, making the fundamental laws incomplete to describe these patterns without looking at specific states of entities to apply the laws on.

From my understanding, I've found that the statement of collective patterns cannot be manifested on the level of entities holds *except for the extraordinary case, when a collective pattern appears in every setting regardless of initial states, and is derivable from the microscale rules on entities* – for example, think of the simple dynamics of every cell in a cellular automaton (a system where entities are square cells on a grid) where cells iteratively copy the state of their left, down, right, and up neighbour, repeating in this order. This simple rule creates the collective behaviour of all 2x2 cells are repeating at periods of 4, 2 or 1, regardless of what was the initial state, and this emergent behaviour is provable from the rules (e.g. take the value stored in the cell X, after the first iteration, this value will be stored in the right neighbour of X, and following the dynamics, its value will be stored in X again after 4 steps). Any other case, where initial states make a difference on the presence of a collective pattern, means that the microscale rules themselves are not enough to deduce the behaviour.

Proponents of both of the two ideologies however do agree, that in a deterministic system, given complete knowledge of the laws driving its dynamics, and the initial states of the system, then theoretically one could derive any "theorem" of the system (or gather any phenomenon) by just simulating the dynamics starting from the initial states, observing the complete evolution of the system. (This is a concept likely without a name.) Some thinkers in complexity science extend this view to phenomena like human consciousness or the evolution of the universe, regarding them as systems that emerged very intricately, but are deterministic and theoretically simulatable if we "play the tape again" (from the initial states, using the laws). However, according to Stephen Wolfram, even if the ruleset of the dynamics is compact, the actual trajectory of the system shaped by the dynamics applied on the initial conditions are too extensive to be written into formulas (with too many operations), therefore the most computationally efficient way to gather an outcome of the system at some time period is to run a direct simulation and "watch the unfolding of the rules on the system". This would mean, that in such a system, some things cannot be determined faster than actually taking the system and observing it evolve and this is what Wolfram refers to as **computational irreducibility** [15], stating that consciousness in his view is that human consciousness is so complex, that it is computational irreducible and appears as if it was free, but actually follows the underlying physical rules.

Lastly, the idea of emergence is also closely related to the gestalt psychology [16], but the idea has existed long before and is often attributed to Aristotle. The deducibility of properties to microlevel causes (between the entities via their states) is referred to as ontological reducability, meanwhile (total) causal dependence on microlevel rules is called supervenience (a function of the entities is supervenient on the microlevel, as it cannot vary from any external factors independent of the entities and their states). The microlevel is considered *causally complete* [17] because it encompasses all causal factors necessary to determine the system's subsequent states and behaviour.

Hierarchy of laws, necessity and sufficiency: Previously, I described that higher-order laws may emerge that are not derivable just from the microscale rules, as they may emerge only in certain initial states. On the other hand, they are supervenient on the microscale rules, depending purely on them and their states. This creates a hierarchy of laws, where the microscale rules are fundamental, and macroscale laws need to obey the microscale rules, but macroscale principles may not be fundamental as they depend on the initial conditions. This is what David Krakauer refers to as "necessary, not sufficient". I give an example from software engineering: The programming languages of both C (a language that is procedural, and has race conditions) and Haskell (a functional programming language with no race conditions) are built on the same Assembly code (a lower-level programming language to which the compilers GCC and GHC compile C and Haskell code to, respectively). When writing a program, we may choose to write it in C or Haskell, which will be the deciding factor of whether the program may contain race conditions or not. The presence or absence of race conditions is a phenomenon, and this occurs on the level of the programming language. We can be sure that the decision whether it happens does not happen on the level of Assembly code, as this is a "common ancestor" of the two programming languages. Assembly code level is not the level of abstraction where this phenomenon "gets triggered to happen or not", such a phenomenon can only be decided on a more coarse-grained level.

Self-controlling emergent layers: Some authors refer to a stronger effect as emergence (see for example Mark Bedau's notion of "weak emergence" [18]), describing emergence as the presence of substructures, collections of entities which are not only exhibiting collective behaviour, but also have a degree of autonomy from external entities and primarily act independently. A good example are live cells in biology. Such structures have self-sufficient dynamics (at least for a portion of time), and not only do they possess higher-order interactions but also higher levels of organization.

These organized, self-controlling emergent structures – which are often abstract structures in nonspatial dimensions, have been discovered in different disciplines, and described by multiple authors with different flavours: David Krakauer calls them ”layers” or ”levels of abstraction” (putting more emphasis on the hierarchical form); biologist Stuart Kauffman uses the term ”Kantian wholes”; physicists David Pines and Robert Laughlin name the effect of not depending on external entities ”quantum protectorates”; and Douglas Hofstadter identifies them as ”self-referential loops” (with strong interpretation on humans) or ”strange loops.” For example, in chemistry, while all chemical reactions are fundamentally based on quantum physical interactions, chemistry as a discipline is effectively described by its own set of laws and principles, operating with a degree of independence from the explicit invocation of quantum physics for many phenomena. These chemical dynamics can be considered self-sufficient at their own level. The ”structure” here refers to the stable patterns of interaction and organization (like molecules, reaction pathways) that define the chemical level of description.

Seemingly, all complex systems have evident self-controlling parts, and self-referential dynamics, as described in this paper [19]. (It appears that this is the only literature that focuses on the connection between complexity and self-controlling dynamics, outside sociology) This work also mentions how self-referentiality also appears in the mechanism commonly, bringing up examples even from logic (all paradoxes such as the Russell-paradox or “this sentence is a lie” arise from self-referring logical statements), which seems to be a phenomenon on its own. It suggests that self-referential (think of recursion in this notion) rules are necessary for complexity arising from simple dynamics. I believe these findings point to a hidden principle possibly for a system to have simple rules that create complex dynamics, these rules have the most “impact” in the shortest description if they are described with self-referential formulas (recursively impacting their own states collectively at every timestep), along with noncontinuous functions, thinking of these from the other way around: some complex dynamics can be compressed into simple rules, and the best “compressors” are self-referential and nonlinear formulas. (This is true for the mechanisms of all complex discrete cellular automata models described, e.g., Game of Life and Rule 110) This is an interesting direction for future research.

Order from disorder: It is worth to mention that whilst most complex systems are algorithmic [20], recent and past research demonstrate that order can naturally forge from large chaos, seemingly and truly random patterns, This was first shown in the 1950s by Turing on Turing patterns, discovering that a (uniformly) randomized initial state can develop ordered patterns under the right mechanisms, and such is the case with Game Of Life too. Recently, it has been shown that order can arise from randomized dynamics as well, for example via self-replicating code [21].

It seems, that in a system with a mechanism that can create a diverse portfolio of patterns (by definition, Turing-complete systems can theoretically build. . . any digital computing system;

those would be plausible), then with a high probability, self-sufficient structures will arise in the diverse chaos. This is mostly beyond the scope of the thesis, but a follow-up work estimating the probabilities of self-sufficient structures arising in various artificial systems will follow from this work.

Symmetry breaking: Loosely speaking, the term symmetry breaking can refer to either the inverse of ”order from disorder” (i.e., patterns ceasing to exist), or, rather similar to *order from disorder*, a state in a system ”collapses” from a set of possible values to one specific value. The latter notion is popular in quantum physics (wavefunction collapse) - in complexity science, for example, it can refer to the system gradually developing stable, emergent orders from ”unstable” states, collapsing the set of possibilities to the set that ”actually happened”, with this process being **irreversible**. (This notion may be misleading, as one may rather argue that a different type of symmetry arises, and also emergent complex systems are predominantly computational/algorithmic in nature (except arguably at the level of biology), therefore, theoretically, there is no literal collapse, there is only algorithmic evolution.) Emergent patterns may break when confronted with external forces, that disrupting the entities that take part in constructing the pattern. Typically, these entities are closely tied and have a ”**boundary**”, that ideally regulates the external effects around the subsystem. When subsystems meet other subsystems and disrupt each other, typically a complex set of interactions happen over time and new patterns arise. I call this process ”complex mixing”.

The field includes many other concepts not in the scope of this thesis. These include self-organized criticality, complex system memory/history, evolution and adaptation, adjacent possible, artificial chemistries, scaling theory, ergodicity, phase transitions, broken symmetry, edge of chaos, leakage, feedback loops. Some of these concepts stem from different sciences, such as biology, physics, chemistry, social sciences and engineering (e.g. control theory or dynamical systems).

The books *Introduction to the Theory of Complex Systems* by authors Stefan Thurner, Rudolf Hanel, and Peter Klimek from the Complexity Science Hub in Vienna, and *The Complex World: An Introduction to the Foundations of Complexity Science* by David Krakauer (Santa Fe Institute) provide deeper insights into these topics.

2.2 Frameworks of emergence

The major frameworks of emergence build on one of two notions of emergence:

- “Synergistic” behaviour: building on the idea of some information containing informa-

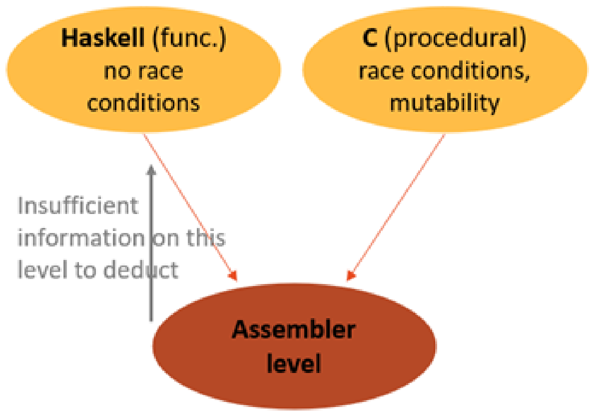

Figure 2.3: The Haskell-C example: the languages differ in patterns, but they compile to the same lower-level language, Assembly. From the "laws" of Assembly code, there can not be any derivations on a program having race conditions or not, this is a higher level pattern.

tion not captured by single entities,

- Self-organized behaviour: thinks of entities that make up "emergent", self-sufficient structures as emergence

The synergy-redundancy framework originates from Williams and Beer [22], establishing the Partial Information Decomposition framework to study emergence. It offers a decomposition of mutual information, for more than two variables. Rather than analyzing on the level of entities, it looks at the collections of entities that cooperate to create some effect. This takes inspiration from the "more than sum of parts" notion. The framework splits higher-order information into synergistic, redundant and unique parts, with synergistic information being an indicator of emergent phenomena. See the figures above.

There is also a distinction among these measures: those that quantify synergistic information contributed by a group of variables on a target variable's states are called directed, and those that capture the overall internal synergy of a set of variables are referred to as undirected measures [23].

The Partial Information Decomposition framework and its generalizations [24] [25] [26] still build on this concept for studying emergent patterns. There is no agreement on how to split mutual information into synergistic, redundant and unique information components.

A personal critique of this framework is that it is not evident that redundant and synergistic (or rather emergent) effects are disjoint, opposing in their functional roles. In the information theoretic sense, it may hold that redundant information is completely separable from synergistic information, but many emergent functionalities also include a significant amount of redundancy in them. In the DI framework, discussed below, emergence is understood as a

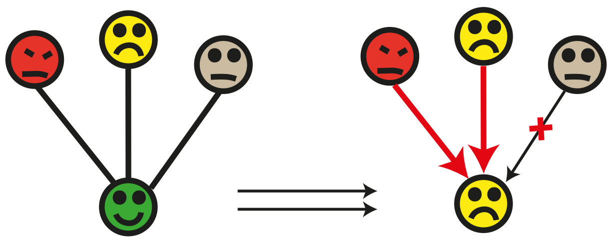

Figure 2.4: Illustration of synergy and redundancy. The left and middle actors together influence the person at the bottom to be unhappy; if it would have been just one of them unhappy, there person at the bottom wouldn't change status. The person on the right would also reduce the happiness level of the influenced individual, however it makes no impact after the two others – making it "redundant". These are the effects the framework is trying to measure.

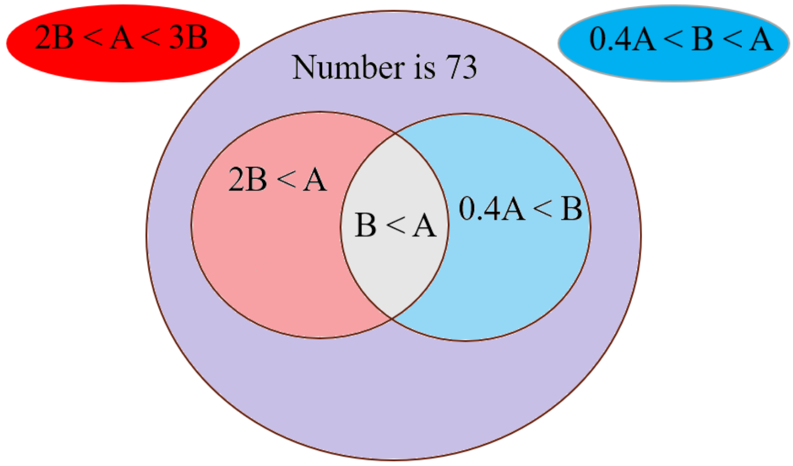

Figure 2.5: An example of redundant, unique, and synergistic information as defined by the PID framework. We get two packets of information about a 2-digit number, with the first (decimal) digit denoted by A, and the second digit denoted by B: 1) $2B < A < 3B$ 2) $0.4A < B < A$. We can categorize different pieces of information using these inequalities. The information that B is smaller than A is clear from both sources, therefore is redundant, and there are unique information. Combining these, one can deduce that the number is 73 (no other number fits the conditions), which is synergistic information.

group of entities with high self-control, which doesn't exclude redundant effects. In fact, often effects in complex systems are redundant, such is the case in the brain. Figure 2.6 shows a "still life" (static) cellular automata in Conway's Game of Life, namely the barge, that holds itself together through collective stabilizing interactions. Yet, almost all cells (a barge can be of arbitrary size), except the 4 cells at the ends can be individually removed, and the system would not lose any more cells. These cells only behave redundantly, yet they make up a large collective. Redundancy induces robustness, e.g., neural circuits in the brain can exhibit graceful degradation after damage, rather than catastrophic failure. Synergistic systems also seem to have some degree of redundancy, and there is a possibility that some redundancy can be positively related to synergy (contradicting this approach's viewpoint of it being only negatively related). Moreover, comparing the two measures (optionally against a target variable) may be misleading when there is a significant amount of both in the system, which is the case especially in large self-controlling entities as both are apparent and necessary for ideal self-controlling dynamics. Comparing one effect against the other in such systems may not actually quantify how much synergy/emergence there is in the (sub)system, as the two quantities may cancel out when compared. Lastly, it is not evident if decomposition to such independent parts is always possible. A single function may very well be both synergistic and redundant, for example an OR logic statement on many variables (only one is needed to satisfy the logical condition, every other variable will be redundant), that are themselves synergistic (e.g. the XOR function is synergistic), and it is not evident that it breaking up into just synergistic and just redundant parts is always possible.

The other approach to emergence are self-controlling subsystems, with high level of independence from external effect. The dynamical independence (DI) framework described by Barnett and Seth [17]. It is based on the idea of entities of a structure acting mostly within the structure itself, "independently" from entities outside the group – the subsystem behaves in a way that is internally governed and relatively autonomous from the larger system. Seemingly, all complex systems exhibit such self-controlling parts [19], so this is also a very common phenomenon. This is measured by assessing how much of the dynamics of the subsystem is explained by the subsystem's own previous states (typically using information-theoretic measures like transfer entropy.

Both of these frameworks, as typical in the field, build on the information theoretical notion of entropy, a measure of uncertainty or disorder and also a measure of expected information of an information theoretical message, which is central to the analysis of emergent behaviour. Entropy is a measure of disorder as well as a measure of expected information, when information is defined as the base 2 logarithm of the inverse probability of the event happening.

At this point, iit is worth noting that in complex systems, interactions and entity states often co-develop [20], The specific geometry, the inherent nature of its dynamics, and the initial states may not only matter, but codevelop and complicate analysis. To overcome this, we may tackle the problem from the direction of the mechanisms, as those can give a fine-grained picture, whereas statistical analysis is coarse-grained.

2.3 Related works

I will work with synergy and self-controlling structures primarily with The Game of Life cellular automaton [14], a famous complex system from a few simple rules. Every entity is placed on a grid, the states of entities depend on how many live (value of 1) neighbours they have. In

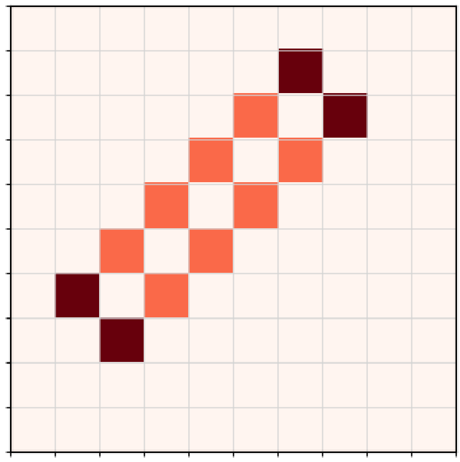

Figure 2.6: The barge; a static collective of living cells that stay alive by their collective impact, yet includes many redundant cells.

each step:

- Any dead cell with exactly 3 live neighbours will have 1 as next state value
- Any living cell with more than 3 live neighbours or less than 2 dies.
- Any living cell stays living if they have 2-3 live neighbours.

Boolean expressions on variables are commonly represented graphically by Karnaugh-maps, a frequently used diagram from digital circuit design and electrical engineering. (Typically, such maps are drawn for 3-5 variables, as more variables require more 2D plots, which is why Game of Life is not typically represented with them.)

Fernando Rosas et al. have showcased how emergent patterns may form generally in (artificial) complex systems, providing an example on cellular automata [27]. On cellular automata, a more explored concept is self-organization, typically by researchers of artificial life [28] [21]. This is directly related to what I discuss in Chapter 3.3.

I use the measures of Oinfo [29] and GradientOinfo [30], defined by Scagliarni et. al. for quantifying how synergistic a system is, and use the minimum mutual information measures [31] to analyze synergy or redundancy separately. Other measures such as Granger causality [31] and dOinfo [9] are also often used for such analysis. For implementation of many of these calculations, I used the HOI Python library.

For extending to applied uses and nonsynthetic data, one can take inspiration from methods of complexity science applied to other disciplines, which include neurology [32], physics [6], and even artificial intelligence [33].

Chapter 3

Software implementation, data

This work only uses simulated data for simplicity, as it requires the driving rules to be known, which real world systems typically do not have. Various graph and cellular automata (in particular, the Game of Life cellular automata) simulations were implemented that contain the evolution of a system from initial conditions via the rules of the dynamics, generating emergent patterns from simple rules, that as a form of generated data I took for experiments, As such, the data is synthetic, and does not contain sensitive information.

The software implementation is broad enough to enable various set of rules (defined mechanisms) on any defined structure as long as that structure implementation has entities and connections (this includes grids, graphs, other types of networks e.g. hypergraphs). This separates the structure from the dynamics and enables mechanisms to be defined regardless of the underlying structure (i.e., the implementation allows for using the dynamics of Game of Life on graphs, instead of cellular automata). The current implementation includes grid (cellular automata) and graph structures, and various mechanisms of dynamics. Implemented in Python, in order to keep the structures implementation and mechanisms implementation separate, I used a design pattern that I would call "structures and dynamics" (likely not a software design pattern that has ever been published): inherit concrete structures from an abstract structure class for structure, having general methods that return the entities and connections in the structure, and implementing the dynamics as a single class that takes in a structure (inherited from the abstract class) as input, and a function containing the logic of updating the entity states each step (the rules of the dynamics). The implementation uses discrete time steps.

Both the implementation and code is available on GitHub[cite].

Chapter 4

Combinatorial frameworks of emergence and the novel concept of dynamical impact

In the first chapter, I discussed how complex systems are challenging to predict partially due to sensitivity to initial states. The initial states in this case do not have to mean just the states of entities, but could also mean other components of the setup, for example, parameters of the structure. As discussed, aside from the mechanism, the initial states also play a crucial role in the outcomes, with observables having high sensitivity to initial states. There have been studies on the effect of the topology of the structure on the dynamics [11], which can also cause outcomes to be intricately changing. All of these components – mechanism, structure, and initial state of entities – matter, and they matter concurrently: the three components often give rise to emergent features cooperatively; therefore, any isolated analysis focusing on only one of these components risks missing key information about their synergistic interplay.

For example, I present a scenario where joint conditions have to be satisfied for a feature to (not) happen: Take a discrete SIS (susceptible – infected – susceptible) contagion model, where the states of individuals take exactly one timestep and infecting happens exactly one step after having become infected. The feature that nodes have a nonzero probability at every timestep to be infected after sufficient time (ergodicity), can only not hold in one specific scenario: if the structure is a bipartite graph (i.e. there are no odd-length cycles), and the initial infected individuals were in the same part of the graph. In this case, the infection will alternate between the two levels. For this to happen, the structure cannot have odd length cycles (thus it has to be a bipartite graph) and furthermore the infected nodes have to be on one side of the two parts of the graph. This pattern emerges from the combination of the structure, the initial state, and the given mechanism.

The fact that the different components jointly impact the dynamics makes research even

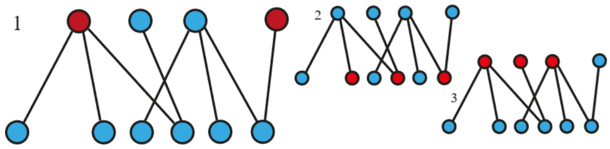

Figure 4.1: In this scenario, in each step, only one part of the graph can have infected nodes. For this pattern to happen, the structure and the initial states had to satisfy specific conditions.

more challenging, as it seems that the emergent features that we'd like to capture are also the result of dynamics, which is itself emergent from the interplay of the structure, the mechanism, and the initial states (a so-called Catch-22). We have to come up with new ways to measure emergent patterns, that bypass these issues. The idea, that we can "play the tape" (simulate the dynamics from the initial states, knowing the mechanisms) can guide us for better solutions.

In this chapter, I introduce a new concept to tackle this problem. While it is not a measure, it can be measured simpler than emergent effects, is easier to interpret, and serves as a bridge between various other concepts. The concept of dynamical impact may make the study of artificial complex systems easier, which can deepen our understanding of the theory of complex systems, but in its current form is rarely applicable to real-world systems as it relies on knowing what impact certain actions made, which is usually not available information due to only being able to observe what happened, and not what would have happened, for a comparison. This is a recurrent problem of causal inference – trying to estimate how much difference a different action would have made, to evaluate the impact of a decision. Potentially, ideas from the field of causal inference can offer tools to estimate the dynamical impact in real-world systems, and uncover not yet seen truths about them.

4.1 Unfolding the dynamics: Dynamical impact

Relying on the ideas discussed in the first chapter, we could theoretically capture a deterministic system fully by taking its initial values (including states, structure, etc.) and the rules that describe how it evolves and simulating what happens – in ideal case, we get a "recording" of the system. It begs for an approach that takes this into account, representing an imprint of the dynamics. This is the motivation behind a new concept I call dynamical impact: keep track of all of the situations in the dynamics, when one entity causally impacted the other entity's next state. Dynamical impact is thus a collection of causal impacts over time, compressing the dynamics into a dynamic causal diagram. I refer to "dynamical impact" to this collection for the rest of the paper.

But how should we measure causal impact? I introduce a simple, yet powerful measure, to establish this direction of research, based on the idea of causality, as understood in the field of

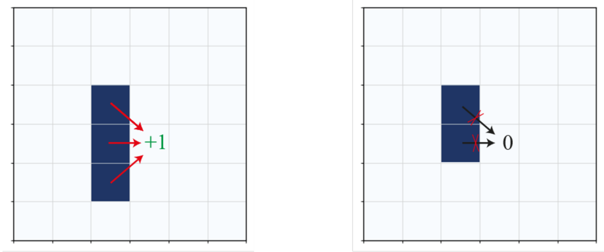

Figure 4.2: In Game of Life, a dead cell having exactly 3 neighbours results in it becoming a live cell (“reproduction/birth”), the three cells collectively “give birth” to the dead cell. Had any of these nodes inactive (not living), the empty cell would have not been born. Therefore, these three cells each had a changing effect, alias an impact, which we store for further analysis. On the right plot however, no matter if the one of the active cells would be inactive, the outcome would be the same, therefore there is no impact. In essence, on the left, the three cells all made an impact, whereas on the other picture, both nodes are redundant.

causal inference. In the thesis, I refer to this furthermore as “impact” (analogous to the term impact or effect in causal inference, where it denotes the extent to which an action alters an effect or an outcome [34]): a variable X at any time step directly impacts variable Y, if the next value of Y, conditional to all other variables’ current values (except X), is determined by what value X takes. In another sense, this means that even with fixing down the values of all other variables, with their current values, varying the current state of X still can vary the next state of Y. This resembles a “decision action” – choosing the value of X . Storing these impacts can recover what interactions in the dynamics made a change in the system, and for example, we can also uncover the chain reactions using the dynamical impact. This is a highly indulgent definition of impact, as it includes every possibility where varying X can cause any difference in Y conditional to all other states, the only cases not included as impact are where Y’s value stays the same regardless of what value X takes. This “anticase” will be important for the notion of redundancy too..

A mathematical formulation of impact could thus look like this (note, that I talk about entities and their states at a fixed timestep, however I did not include time in the notation for simplicity):

Assuming $Z = f(X_i, X_{-i})$, X_i impacts Z, denoted as $(X_i \rightarrow Z)$, if and only if exist two different values of $\exists x, x' \in \mathcal{X}_i$, such that $f(x, X_{-i}) \neq f\left(x', X_{-i}\right)$, or equivalently, $Z|x, X_{-i}$ does not equal $Z|x', X_{-i}$. Here, $\mathcal{X}_i$ represents the set of possible values of X_i, in our case, 0,1 - basically in this setting, X has to have some values for which Z equals 0, and some for which Z equals 1. Dynamical impact is just the set of impacts over time. This is the definition that I will use in this work.

The core idea here is not the introduction of a causal inference approach, but to capture the dynamics ”like if we could play the dynamics all over again” and use this rich information. That

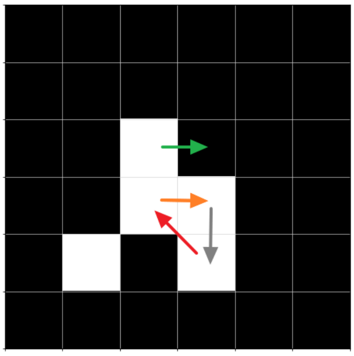

Figure 4.3: Different types of impacts in Game of Life: green arrow shows the impact one cell has on the empty cell (that 2 neighbouring cells also have), as it is key in making it a live cell by it having exactly 3 neighbours. If the cell where the impact comes from was not active, the empty cell would stay empty, therefore this cell made an impact by being active. This type of impact I call “birth”. The red arrow shows a “kill” impact, as it plays part in the middle cell dying, which would not have happened had the source cell been empty. The grey arrow shows an impact which is an essential impact in keeping the target node alive. The yellow arrow shows not an impact, but a redundancy. No matter the value of the middle (source) cell being 0 (empty) or 1 (living, active), the target cell would stay active.

is what dynamical impact attempts to do, and offers many benefits. Firstly, dynamical impact can capture joint effects in initial states and topology, as it is an unfolding of the dynamics, it varies as do the dynamics by varying the components. We lose less (if any) of the emergent effects of the dynamics by studying dynamical impact “isolated”, unlike with isolating the structure. It is itself an emergent construct, just as are the dynamics, but its main advantage is that it is easier to analyze and observe.

Secondly, this notion offers a fitting definition redundancy, which I define as:

(Momentary) redundancy: Entity X in general could have impact on Z (i.e., they are connected), however the current state of X with the other variables’ states fixed does not change Z: $\forall x \in X_i, Z = f(x, X_{-i})$ is constant, i.e. there are no two values $\exists x, x' \in X_i$ such that $f(x, X_{-i}) \neq f(x', X_{-i})$. This means that all the other states already determine the next value of Z. These definitions will be helpful moving forward.

4.1.1 Characterizing different types of interactions based on impact: Game of Life

As also seen on Figure 4.2 and 4.3, we can characterize different types of interactions in Game of Life based on whether they generate impact or are redundant, and their functionality.

Let’s consider the different type of interactions in Game of Life that create impact. For simplicity (since most of the cells on the grid are empty, from now on I only consider the impact of entities that are living, which I call 'active impact'.

The 4 types of interactions with X impacting Y, conditional on Z:

- "birth": Y has a current state of 0, with 2 active neighbours from Z – here X taking a value of 1 makes Y become a live cell, otherwise Y stays 0,
- "kill": Y has a current state of 1, with 3 active neighbours from Z – here X taking up the value of 1 will deactivate Y, otherwise it stays active,
- "live": Y has a current state of 1, and has 1 active neighbour from Z– for Y to survive, X has to equal 1,
- "no birth": Y equals 0 and has 3 active neighbours from Z – if X is active then Y will not be active, otherwise it will be in the next timestep.

These are seen evidently from the rules of Game Of Life. All of these are synergistic effects, as they require the combination of the nodes to be active at the same time. There are also 2 possible cases of redundancy:

- "redundancy kill": Y is 1, but has at least 4 active neighbours from Z – no matter the value of X, Y will be deactivated,
- "redundancy live": Y is 1 with two active neighbours from Z, no matter the value of X, Y stays alive.

I will analyze these in the next subchapter, where I measure also their synergy and redundancy based on the measures of the PID framework.

4.2 A combinatorial approach to the PID framework based on causal impact - directed synergistic causality

To produce a combinatorial version of the PID framework, the measure of synergy and redundancy need to be defined, which I will do in this subchapter. Redundancy has already been defined above, it is the "opposite" of impact under the condition of Z values, X does not impact the value of Y in a step regarding the other. We have thus a different interpretation of redundancy, not based on information theory, but based on lack of causal impact of an interaction.

We can plot how often each keyword appears (2400 simulations)

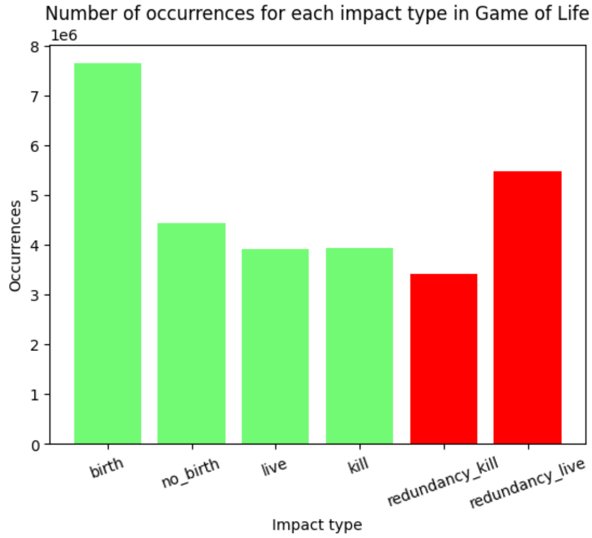

Figure 4.4: Impact type occurrences- Surprisingly, most impacts are births, despite requiring 3 active and 1 free nodes, and the births are significantly more than the deaths combined.

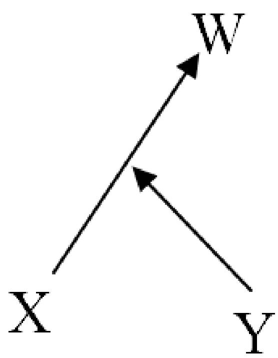

Figure 4.5: Y causally impacts the way X causally impacts on W. This is a sufficient case of synergy, the other possibility would be if X impacts the way Y impacts W, or both.

Let's try to create a similar definition of directed synergy. Synergy requires multiple entities acting cooperatively. We can measure whether the impact of two entities, X and Y, is synergistic to a target entity W, if the impact of let's say X on W, also varies by Y. This means, that through the way X impacts W, Y also impacts W, making this impact synergistic. This, of course, requires both X and Y to be able to vary W.

Proposition: Entity X is synergisticly impacting entity W with entity Y, if its impact varies by the values of Y, implying that the set of $F(X=1,Y=y,\boldsymbol{Z})-F(X=0,Y=y,\boldsymbol{Z}) = \Delta F_X|(Y=y,\boldsymbol{Z})$ values across all possible values of Y, conditional to Y and the fixed values in Z, contains at least two different values (otherwise there would be no varying). This is a measure of the causal impact of X being active (controlling for Y, and other parameters, if X does not impact W, this will equal 0 for any combination), and if this measure, the causal impact of X, differs for different values of Y, then this effect is a joint synergistic impact of X and indirectly Y too. Here, Y impacts W through a "backdoor".

Therefore, in this case, Y impacts the way X impacts W. This is certainly a sufficient case of synergy, and it is a question whether this is necessary, whether any synergistic impact on W through X and Y would imply that Y has an impact on X's impact on W. This follows, if "impact on impact" is symmetric, i.e. goes both ways: if Y impacts the way X impacts W, then, assumably, X also impacts the way Y impacts W. I show that this hypothesis is true in our simple framework:

Assume that we have X impacting W, and Y impacting the impact of X on W, conditional to fixed values of set of variables Z. Then, by definition, conditional on these $\boldsymbol{Z}$ values, there must be two values of Y, where $\Delta F_X|(Y=y,\boldsymbol{Z})$ differs, that is, $\Delta F_X|(Y=y_1,\boldsymbol{Z}) \neq \Delta F_X|(Y=y_2,\boldsymbol{Z})$. For simplicity, we can assume $y_1 = 1$ and $y_2 = 0$, as we work with binary variables - but the proof can be modified to generalize for cases when the set of possible values for each variable can have more than two values.
Indirect proof: assume that X does not impact Y's impact on W. But then also $\Delta F_Y|(X = 1,\boldsymbol{Z}) = \Delta F_Y|(X=0,\boldsymbol{Z})$. This would also imply $F(Y=1,X=1,\boldsymbol{Z}) - F(Y=0,X=1,\boldsymbol{Z}) = F(Y=1,X=0,\boldsymbol{Z}) - F(Y=0,X=0,\boldsymbol{Z})$, which after rearranging, would imply $F(X=1,Y=1,\boldsymbol{Z}) - F(X=0,Y=1,\boldsymbol{Z}) = F(X=1,Y=0,\boldsymbol{Z}) - F(Y=0,X=0,\boldsymbol{Z})$. This cannot be true, as we started with the assumption that they cannot be equal in the case of these Y values and $\boldsymbol{Z}$ values, therefore we arrived at a contradiction and the disproved that X can possibly not impact Y's impact on W. We made use of the set of possible values only including 0 and 1, but this proof shall be generalizable for any finite set of possible values for each variable.

This way, one can be sure that if Y has an impact on the impact of X on W, then X also has an impact on the impact of Y on W. There is no one without the other. If none of them impact the other's impact on W, there is no synergy, only "uniqueness" (in the PID framework). Impact of a variable on the impact of another variable on an outcome variable is therefore, in this framework, "undirected"; it goes both directions, as previously stated. (Impact, generally speaking, has a direction.)

It would be interesting to know how we can find/define joint synergies of higher order (i.e., when the values of Y altogether are also synergisticly impacting W, and not just pairwise with X), or what is the distinction between three variables jointly impacting an outcome, versus only the three pairs impact the outcome (no "triangle effect"). I have derivations for pairwise synergies, but none imply a joint, common synergy across all entities. In this work, this is not a problem, because I implement experiments on Game of Life and we know that the effects in Game of Life are synergistic, as they require the joint presence of multiple active cells for any active impact (reproduction, death, staying active, these only can happen based on the rules

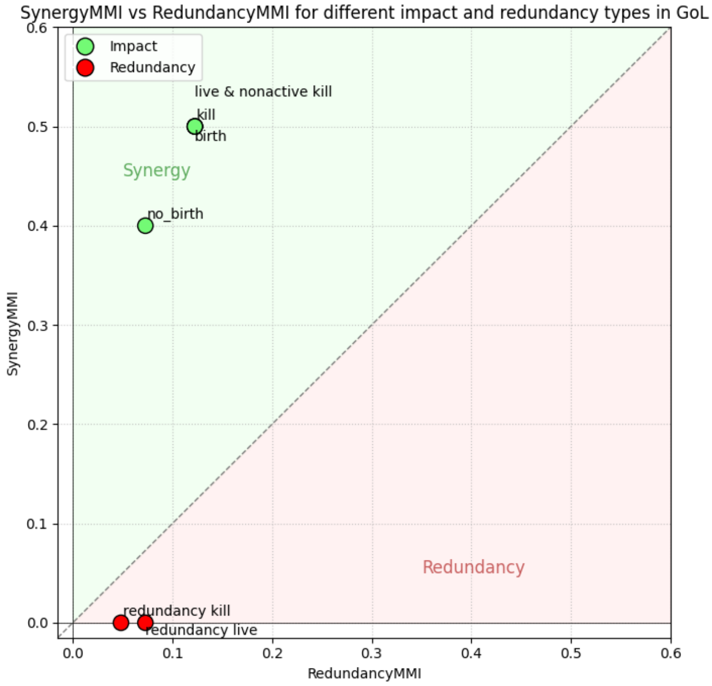

Figure 4.6: Synergy vs. Redundancy - defining based on impact yields the correct results.

when 2, 3 or 3+ neighbours are jointly active).

Based on a directed measure of synergy, one can provide an undirected one by choosing a target entity in the group of analyzed entities and measuring directed synergy, iterating through the entities as targets, and averaging the sum of the results. Such measures are not used in this work.

The last definition to provide is uniqueness: An impact is unique (nonsynergistic, nonredundant) when an entity X has an impact on W, which is not impacted by any other entity Y.

4.2.1 Measuring impact synergy and redundancy with information theoretic measures

To quantitatively show that the definitions of impact and redundancy rightly represent synergy and redundancy in Game of Life, I used the popular information theoretical measures of synergy and redundancy, checking if the two evaluations overlap:

Total correlation is a measure associated with redundancy, against dual total correlation, that encapsulates more synergistic information. Therefore, for synergistic interactions, we expect DTC to be higher than TC, and vica versa for redundancy.

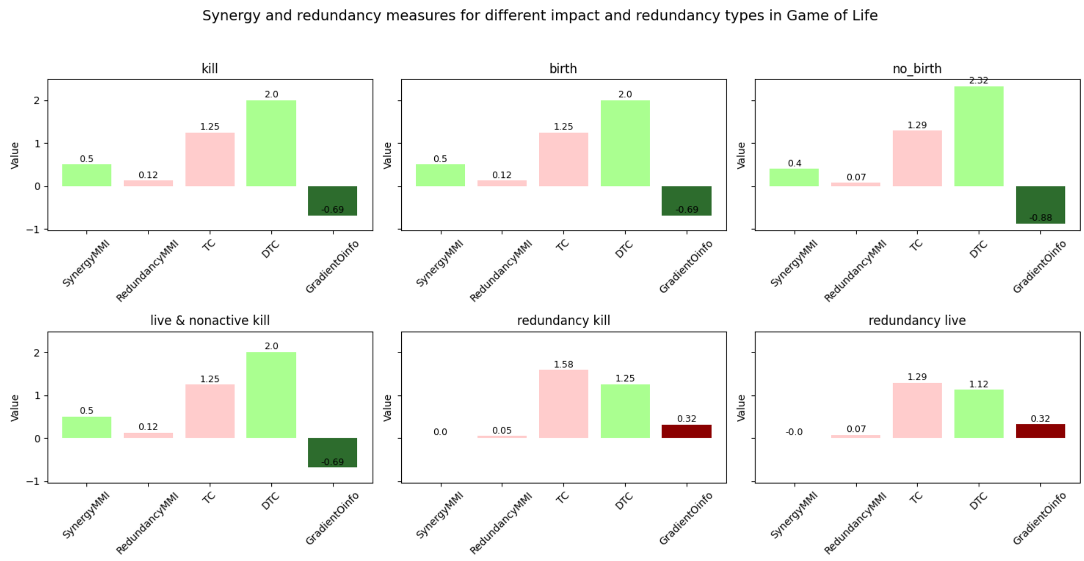

Figure 4.7: Synergy vs. Redundancy - all measures classify the same way as impact does - great for such a simple measure

Indeed, where there is impact (in the first 4 types of impact on the second plot), there is statistical synergy in all measures and it is higher than redundance, whereas for the two types of redundancy, the opposite is true, and the minimum synergy measured by SynergyMMI is 0. This is evidence that dynamical impact consistently measures synergistic effects also measured by information theoretic measures, and its counterpart, redundancy, also aligns with information theoretic redundancy.

4.3 Combinatorial approaches to the DI framework

4.3.1 Group theory as a possibility

Group theory provides the toolset to study periodic patterns (or symmetries). It has been widely used in physics, such as for crystal structures [35] and particle physics [36], but it gains increasing popularity in applied sciences. Recently the field of geometric deep learning was founded on group theory [37], as groups such as Lie groups have special geometric properties too, enabling to connect two seemingly unrelated fields.

Groups with certain periodicities appear frequently in real life, and also in cellular automata, For example, in a cellular automaton with 5 cells, after at most $2^5 + 1 = 33$ steps, some global state must repeat. In practice, many CA patterns settle into periodic or stable behaviors. In Conway's Game of Life, for example, we see oscillators and mobile repeating structures like gliders. For mobile patterns, some adjustments are needed to account for translation when

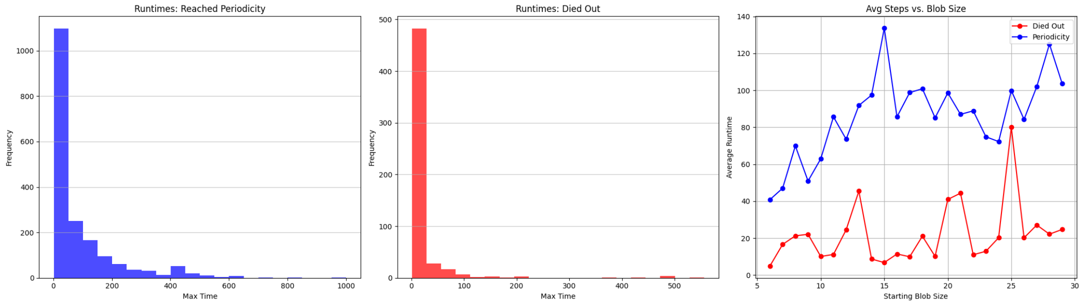

Figure 4.8: Runtimes till dying out / reaching periodicity

detecting repetition.

This work won't make use of group theoretical tools, but it borrows the idea of looking at periodic patterns. To explore this, I ran simulations from various initial states, tracking whether they became periodic or died out over time.

Interestingly, we can see on the left plot, that there are sudden bumps near time 400 in frequency. This has to be due to the mechanism, somehow the probability of grids that get periodical at that time is significantly higher - likely the mechanism causes certain periods be more frequent. From the right diagram (every point is aggregated over 100+ runs) shows that the two diagrams despite showing data for unrelated cases, share similarities, and these similarities in the plot appear quite irregular, very nonlinear. This could be an example of a surprising pattern as a result.

This approach has advantages: group theory naturally handles symmetry, making it easier to spot repeating or moving patterns, in particular I could implement this for cellular automata. But on arbitrary networks, finding equivalent patterns (e.g. under isomorphism) becomes hard, as there's no general efficient algorithm for graph isomorphism. Moreover, it's rare that an object in real life arrives at the exact same state two times in real life.

4.3.2 Dynamical impact as alternative to the DI framework

Dynamical independence (DI) considers emergent structures as collections of entities and their internal interactions if they tend to converge to be informationally complete about their future, and no external or lower-level entities provide significant extra information. We can recreate a similar notion using the dynamical impact: emergent structures defined by most of the causal impact related to any entity in the group takes place on both ends inside the group itself, independent of external entities. This can be measured by the fraction of the dynamical impact flowing inside the group (from sources to targets), against the combined impact in and out of each node in the group. This can be insightful in algorithmic systems, where proper probability distributions are often nonexistent by the nature (requiring more advanced methods, at least Markov chains), which are needed for entropy.

To find groups with the highest self-controlling effect, I use the mentioned measure: the ratio of impacts of the group that flows in the group, versus all of the group's impacts, compromising of impacts happening inside, flowing in, and flowing out. I defined this ratio, and aimed to find groups with high such value.
However, the highest such value group would be the group of all nodes, as in that case the ratio equals 1. Therefore, it is better to compare to a null model of the same size (a graph with randomized directed edges representing impact, where we calculate an expected value of inner impact in a group of fixed size). One can derive that for a random graph, the expected ratio at group size k is $m\frac{k-1}{n-1}$, where m, n are the number of edges and nodes in the whole graph. (This formula can be obtained by either double counting on edges: every edge appears in q). This is the formula I compare against.

Creating a simple heuristic algorithm that keeps growing the group size until the ratio of the real value and the null value grows. Here is a plot of two of them:

The figure shows that both groups have high ratios when they "exist", which means our heuristic worked. Running this algorithm for 20 other times, all of the groups with the highest dynamical impact were only consisting of at most 4 nodes. This makes sense, because many small patterns are completely repetitive, and likely such isolated patterns have the highest impact, especially compared to the null model.

Combined with previous framework, periodic groups will have periodic time series.

A theoretical condition on dynamically independent groups can be formulated.
Theorem: Any dynamically independent group of variables has to contain loops in its aggregated subgraph of the dynamics.

This can be seen by establishing that impacts (or information, interactions) have to flow circularly in order to have a self-sufficient dynamical subsystem, which can only be achieved via loops-in-time, otherwise all interactivity would flow out of the subsystem.

If we now focus on systems, where any interaction regarding determining the next state of a variable is "impactful" (which is reasonable, otherwise interactions wouldn't make sense if the outcome can never be changed), we can further state that the dynamical impact has the contain loops for the same reason. Any dynamically independent subsystem has to have dynamical loops of interactions, with periodicity length at most as high as the periodicity of the subsystem, unless the system is "moving" (variables in the system change) such as the glider in Game of Life.

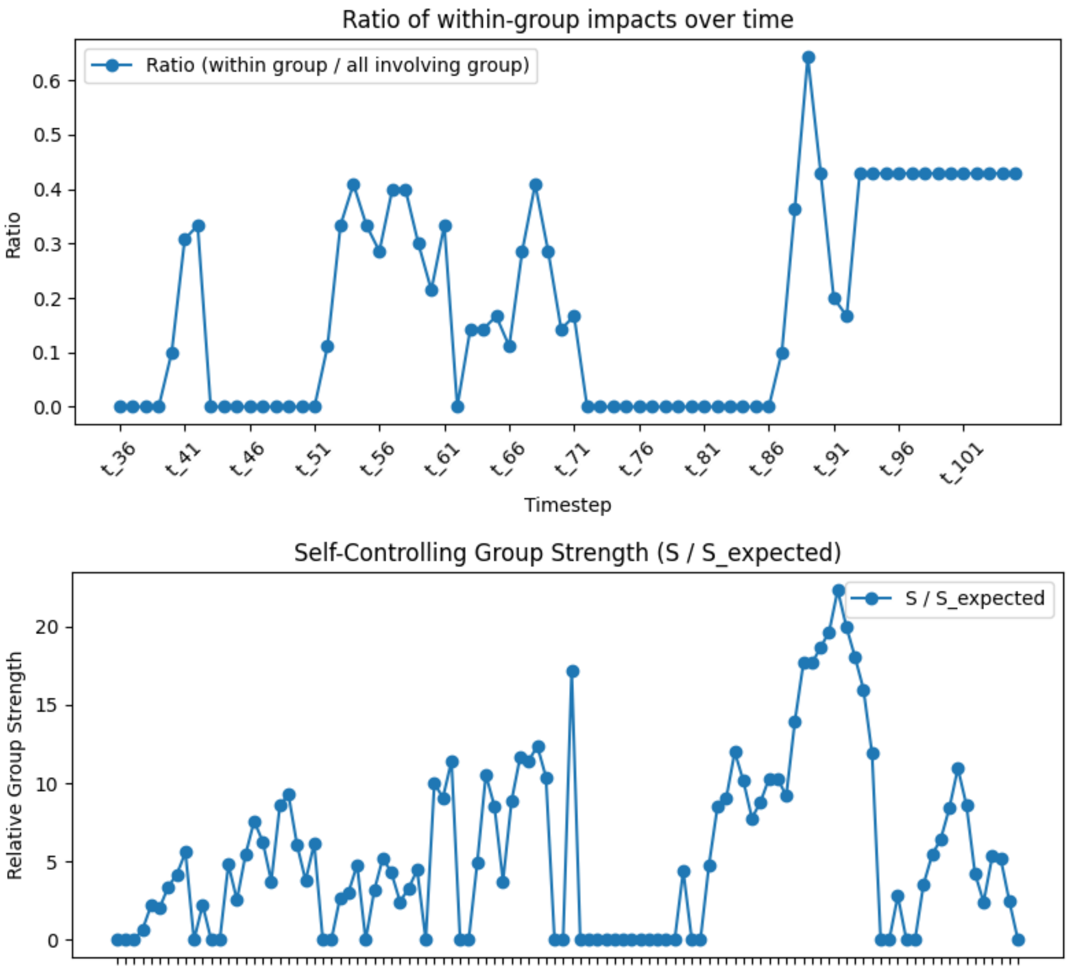

Figure 4.9: Two random groups with (occasionally) high impact. The top one converged into a periodic state, hence why there was no more change in the dynamics.

4.4 Insights using dynamical impact - synergy and dynamical independence

Dynamical impact can provide more complete information for analysis than statistical analysis of observables can not, particularly in not observed patterns, that may have happened based on the mechanisms, had the initial states have been different.

Interesting connections can be drawn between the two frameworks from this perspective. Knowing that dynamical loops are necessary for dynamical independence, and synergistic variables may improve the likelihood of such an event happening. Can it happen, that a group of entities are dynamically independent, and do not have synergistic interactions?
Strictly speaking, yes, and Figure 2.2 is an example of that. The impacts on the next state only have one source, and therefore are not synergistic, yet the system is dynamically independent. In fact, including redundant, nonimpactful interactions, we can easily extend the list of dynamically independent but non-synergistic subsystems.

However, if we make two further conditions, which are practically common:

- Every entity's next state in every timestep is impacted by at least one other entity (otherwise, either the entity is a redundant, unchangeable node at the moment, or its state is constant),
- The impacts on one variable's next state at a timestep are synergistic (which is typical in Boolean logic systems, as functions of a variable that are not synergistic must not have any joint variable terms, which is uncommon in logic expressions, and even the OR function is not linear (as ”sum” would be) in Boolean algebra),

then, by using the fact that there each entity (node) will have at least one impact (in-edge) determining its next state, the only dynamically independent subsystems without synergistic impacts are such cycle dynamics, like the one on Figure 2.2. Any other construction would have a node with two in-edges, therefore there would be two entities synergistically impacting the next state of a third entity.

The framework also can show how an entity may have ”indirect” synergistic impact, that is, not immediate impact, but their impact on other entities eventually lead to causing a difference in the outcome - for this, time-series paths have to be tracked in the dynamical impact.

Chapter 5

Further work

There are two natural next steps building on the current work:

- Formulation of other complexity science concepts using dynamical independence: as an example, complexity is often associated with highly emergent patterns and behaviour, which we have seen can be understood in terms of causal impact. Possibly, a large extension of impacts may lead to high complexity and high description length. This may hint at a possible link between dynamical impact and complexity, which could also serve as a link between complexity and other better understood concepts of the field, such as synergy or dynamical independence. Other concepts of the field may formulated using dynamical impact.
- Improving applicability on real-world systems by tools to attempt to infer dynamical independence. Such work would greatly improve the practicality of this thesis. A possible source for ideas may be causal inference again, as statistical tools to infer causal relationships can also be used in this context.

The framework itself could also be improved. A potential direction: instead of measuring just a difference in the outcome (which is, in a philosophical sense, quite utilitarian), one may measure if something ”happens differently”; having an impact not on the result, but generally on the dynamics. An outcome may exhibit no difference in value due to various cancelling effects, but its state/value may have been built up fundamentally differently.

Chapter 6

Conclusion

To sum up, this work introduced the idea of dynamical impact as a simple way to track how changes spread in a system. I showed how this concept relates to both synergy and redundancy, and how it can be used to rethink some well-known frameworks in complexity science, such as the PID and Dynamical Independence frameworks.Various experiments were explored to check validity, or to look for interesting findings. The experiments showed the novel method to work consistently. While this approach is most useful in systems where we know the rules, like simulations, it still has the potential to help us better understand real complex systems aswekk. The concept of dynamical impact could be used to explore other phenomena and connecting different ideas in complexity science. A next step could be to explore how dynamical impact relates to complexity itself — since it seems likely that systems with higher dynamical impact might also be more complex.

Bibliography

[1] D. C. Krakauer, *The Complex World: An Introduction to the Foundations of Complexity Science*. SFI Press.

[2] D. T. Mihailovic, G. Mimic, E. Nikolic-Djoric, and I. Arsenic, "Novel measures based on the Kolmogorov complexity for use in complex system behavior studies and time series analysis."

[3] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri, "Networks beyond pairwise interactions: Structure and dynamics," vol. 874, pp. 1–92.

[4] F. Battiston, E. Amico, A. Barrat, G. Bianconi, G. Ferraz de Arruda, B. Franceschiello, I. Iacopini, S. Kéfi, V. Latora, and Y. Moreno, "The physics of higher-order interactions in complex systems," vol. 17, no. 10, pp. 1093–1098.

[5] I. Iacopini, G. Petri, A. Barrat, and V. Latora, "Simplicial models of social contagion," vol. 10, no. 1, p. 2485.

[6] T. Robiglio, L. D. Gaetano, A. Altieri, G. Petri, and F. Battiston, "Higher-order Ising model on hypergraphs."

[7] F. E. Rosas, P. A. M. Mediano, A. I. Luppi, T. F. Varley, J. T. Lizier, S. Stramaglia, H. J. Jensen, and D. Marinazzo, "Disentangling high-order mechanisms and high-order behaviours in complex systems," vol. 18, no. 5, pp. 476–477.

[8] D. C. Krakauer, "Symmetry–simplicity, broken symmetry–complexity," vol. 13, no. 3, p. 20220075.

[9] T. Robiglio, M. Neri, D. Coppes, C. Agostinelli, F. Battiston, M. Lucas, and G. Petri, "Synergistic signatures of group mechanisms in higher-order systems."

[10] K. R. Popper, "The Myth of the Framework," in *Rational Changes in Science. Essays on Scientific Reasoning: Boston Studies in the Philosophy of Science, Vol. 98* (J. C. Pitt and M. Pera, eds.), vol. 98, pp. 35–62, 1987.

[11] G. Bianconi, "The topological Dirac equation of networks and simplicial complexes."

[12] P. W. Anderson, "More Is Different," vol. 177, no. 4047, pp. 393–396.

[13] K. Gödel, "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I," vol. 38, no. 1, pp. 173–198.

[14] M. Games, "The fantastic combinations of John Conway's new solitaire game "life" by Martin Gardner," vol. 223, pp. 120–123.

[15] S. Wolfram and M. Gad-el Hak, "A new kind of science," vol. 56, no. 2, pp. B18–B19.

[16] K. Koffka, "Principles of gestalt psychology," vol. 11, no. 44, p. 176.

[17] L. Barnett and A. K. Seth, "Dynamical independence: Discovering emergent macroscopic processes in complex dynamical systems," vol. 108, no. 1, p. 014304.

[18] M. Bedau, "Weak emergence," vol. 11, pp. 375–399.

[19] S. Hempel, R. Pineda, and E. Smith, "Self-reference as a principal indicator of complexity," vol. 6, pp. 22–27.

[20] S. Thurner, P. Klimek, and R. Hanel, *Introduction to the Theory of Complex Systems*. Oxford University Press.

[21] p. u. family=Arcas, given=Blaise Agüera, J. Alakuijala, J. Evans, B. Laurie, A. Mordvintsev, E. Niklasson, E. Randazzo, and L. Versari, "Computational Life: How Well-formed, Self-replicating Programs Emerge from Simple Interaction."

[22] P. L. Williams and R. D. Beer, "Nonnegative Decomposition of Multivariate Information."

[23] F. E. Rosas, P. A. M. Mediano, and M. Gastpar, "Characterising directed and undirected metrics of high-order interdependence."

[24] P. A. M. Mediano, F. E. Rosas, A. I. Luppi, H. J. Jensen, A. K. Seth, A. B. Barrett, R. L. Carhart-Harris, and D. Bor, "Greater than the parts: A review of the information decomposition approach to causal emergence," vol. 380, no. 2227, p. 20210246.

[25] M. Wibral, V. Priesemann, J. W. Kay, J. T. Lizier, and W. A. Phillips, "Partial information decomposition as a unified approach to the specification of neural goal functions," vol. 112, pp. 25–38.

[26] F. E. Rosas, P. A. Mediano, H. J. Jensen, A. K. Seth, A. B. Barrett, R. L. Carhart-Harris, and D. Bor, "Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data," vol. 16, no. 12, p. e1008289.

[27] F. E. Rosas, B. C. Geiger, A. I. Luppi, A. K. Seth, D. Polani, M. Gastpar, and P. A. M. Mediano, "Software in the natural world: A computational approach to hierarchical emergence."

[28] C. Gershenson, "Self-organizing systems: What, how, and why?," vol. 2, no. 1, pp. 1–8.

[29] F. E. Rosas, P. A. M. Mediano, M. Gastpar, and H. J. Jensen, "Quantifying high-order interdependencies via multivariate extensions of the mutual information," vol. 100, no. 3, p. 032305.

[30] T. Scagliarini, D. Nuzzi, Y. Antonacci, L. Faes, F. E. Rosas, D. Marinazzo, and S. Stramaglia, "Gradients of O-information: Low-order descriptors of high-order dependencies," vol. 5, no. 1, p. 013025.

[31] A. B. Barrett, "Exploration of synergistic and redundant information sharing in static and dynamical Gaussian systems," vol. 91, no. 5, p. 052802.

[32] A. I. Luppi, P. A. Mediano, F. E. Rosas, J. Allanson, J. Pickard, R. L. Carhart-Harris, G. B. Williams, M. M. Craig, P. Finoia, A. M. Owen, L. Naci, D. K. Menon, D. Bor, and E. A. Stamatakis, "A synergistic workspace for human consciousness revealed by Integrated Information Decomposition," vol. 12, p. RP88173.

[33] A. M. Proca, F. E. Rosas, A. I. Luppi, D. Bor, M. Crosby, and P. A. M. Mediano, "Synergistic information supports modality integration and flexible learning in neural networks solving multiple tasks," vol. 20, no. 6, p. e1012178.

[34] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.

[35] U. Müller, *Symmetry Relationships between Crystal Structures: Applications of Crystallographic Group Theory in Crystal Chemistry*. Oxford University Press.

[36] S. Weinberg, "A model of leptons," *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, Nov 1967.

[37] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges."