Combining Causal Bayes Nets and Cellular Automata: A Hybrid Modelling Approach to Mechanisms

Alexander Gebharter and Daniel Koch*

Causal Bayes nets (CBNs) can be used to model causal relationships up to whole mechanisms. Though modelling mechanisms with CBNs comes with many advantages, CBNs might fail to adequately represent some biological mechanisms because—as Kaiser ([2016]) pointed out—they have problems with capturing relevant spatial and structural information. In this article we propose a hybrid approach for modelling mechanisms that combines CBNs and cellular automata. Our approach can incorporate spatial and structural information while, at the same time, it comes with all the merits of a CBN representation of mechanisms.

ORCIDs: Koch, https://orcid.org/0000-0002-4893-1774.

1. Introduction

Mechanisms play an important role in many sciences. This is true especially for most of current biomedical research, where to understand a phenomenon often means to identify a mechanism responsible for that phenomenon's occurrence. According to mechanists, certain explanations are best provided by pointing at the phenomenon of interest's underlying mechanism. Once a mechanism is identified, it can be used to make predictions about the system and to control it by means of interventions (for example, by designing a new drug that interferes with parts of that mechanism). Classical approaches to mechanisms—both in philosophical considerations and in biological practice—are typically formulated in qualitative terms and often supported by diagrams that depict how the relevant components interact to produce the phenomenon of interest. Machamer *et al.* ([2000], p. 3) characterize mechanisms

---

as 'entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions'. For other such approaches, see, for example, (Glennan [1996], [2009]; Bechtel and Abrahamsen [2005]; Craver [2007b]); but see also (Jones and Wolkenhauer [2012]).

However, there is also a clear need for quantitative mechanistic explanations and predictions. While quantitative approaches came naturally for many sciences such as physics or chemistry, their application in biomedical research showed a significant delay and still is encountered less frequently than in other sciences. Simple processes such as the oscillation of a pendulum or the concentration of a chemical compound could be readily measured in a quantitative more or less accurate manner rather early in history, allowing the formulation of mathematical theories thereof. Biological processes, in contrast, are often harder to study as they involve a plethora of interwoven entities and interactions on different scales. Many outcomes of biological experiments, such as the phenotype of a cell, are difficult to quantify by simple measures. As a result, biological sciences and the study of mechanisms are still often perceived as something qualitative.

Although the origins of mathematical approaches to biology date back at least to the beginning of the twentieth century, it was long perceived as a niche discipline often with limited recognition by experimentalists.[1] In recent decades, however, formal approaches to biological questions have become increasingly popular and are by now established part within the mainstream of biomedical research. Some of this development is certainly facilitated by technological advancements (for example, by high-resolution and quantitative microscopy approaches such as FRET- or FRAP-microscopy) and the increased production of biological data since the advent of '-omics' techniques. Partly, however, it is also driven by the increasing recognition that quantitative approaches to biology are necessary. Many researchers argue that the reason for this necessity is to capture the complexity of biological systems in a more non-reductionist manner (see, for example, Wolkenhauer and Green [2013]). However, even for a solid understanding of classic mechanistic biology and to provide quantitative mechanistic explanations and predictions, especially where dynamic processes are studied, formal approaches are needed, too (Bechtel and Abrahamsen [2005]; Brigandt *et al.* [2017]).[2]

In scientific practice, this is mirrored by the vast diversity of mathematical and computational modelling approaches that are successfully applied to study biological processes. Examples of popular approaches are ordinary and partial differential equations, agent-based models such as cellular automata, or network models such

---

[1] There are some notable exceptions such as (Lotka [1925]) or the study of enzyme kinetics; see, for example, the famous paper by Michaelis and Menten ([1913]).

[2] For a discussion of examples in which a misunderstanding of quantitative (kinetic and thermodynamic) principles led to misconceptions about the mechanism of how signalling proteins of the small GTPase family are regulated see, for example, (Goody [2014]).

as Boolean and Bayesian networks. The choice of which approach to apply often depends on how suitable the approach is in terms of applicational ease and computational feasibility (and often simply depends on the familiarity with a certain tool). For instance, one of the first comprehensive whole-cell models covering all basic processes of all cellular functions such as cell division, metabolism, DNA replication, protein synthesis, and so on relied on a combination of over twenty different formalisms (Karr *et al.* [2012]).

Philosophy's coverage of formal modelling approaches in general and their application to biological mechanisms in particular is a rather recent development (Winsberg [2009]). While the concept of what a mechanism is has received much attention lately, formal approaches based on a philosophically elaborated concept of mechanisms are only beginning to emerge. Causal Bayes nets (CBNs) are among the most popular philosophical proposals for modelling mechanisms brought forward by authors such as Casini and Baumgartner ([unpublished]), Casini *et al.* ([2011]), Clarke *et al.* ([2014]), and Gebharter and Kaiser ([2014]; Gebharter [2014], [2017b]). CBN approaches to mechanisms come with many advantages: CBNs allow for formulating and testing of causal hypotheses, they can be used for providing probabilistic explanations and predictions, and they can be used to predict what would happen under certain interventions (even if only non-experimental data is available). Unlike other formalisms, CBNs are based on a thorough philosophical characterization of causality that satisfies modern standards for theoretical concepts (Schurz and Gebharter [2016]; Gebharter [2017a]). They have been applied successfully in biology to model signalling pathways, cancer progression mechanisms, and other processes (see, for example, Sachs *et al.* [2005]; Koch *et al.* [2017]). Another advantage of the CBN framework is that it provides a basis for developing powerful algorithms for causal discovery (Spirtes *et al.* [2000]). Murray-Watters and Glymour ([2015]), for example, take up Gebharter's ([2014]) proposal to model mechanisms as bundles of causal arrows in a CBN and develop an algorithmic search procedure for submechanisms whose variables have not yet been measured.

There are, however, also several problems for CBN approaches to mechanisms. Weber ([2016]), for instance, argues that CBN methods cannot handle the dynamics of certain biological mechanisms and are in this regard inferior to differential equations, and Kaiser ([2016]) argues that CBN approaches to mechanisms are problematic because CBNs might fail to capture spatial and structural information relevant for mechanistic explanation in biology. It therefore seems that CBN approaches to mechanisms are forced to remain incomplete if modelling the mechanism of interest requires the incorporation of rich spatial and structural information. This article takes Kaiser's observation as a starting point. We agree that CBN models of mechanisms might lack relevant spatial and structural information. CBNs are, however, quite flexible and we think that one should not be overhasty in rejecting them as useful tools for modelling mechanisms.

In this article we develop a new hybrid approach for modelling mechanisms based on CBNs and cellular automata.[3] We thereby aim to extend the CBN approach while staying within the philosophically elaborated framework of mechanisms. The overall causal structure of mechanisms is, according to Gebharter ([2014]), represented by simple input–output structures (bundles of causal arrows) in CBNs. Cellular automata are used to model submechanisms that feature spatial and structural information that is not directly available in the associated CBNs. Our approach comes with all the powerful merits of a CBN representation of mechanisms while, at the same time, it can avoid problems such as the ones pointed out by Kaiser ([2016]). An additional advantage of implementing cellular automata will be that they can nicely capture the dynamics involved in spatial and structural processes in many mechanisms. This makes our hybrid approach suited for running simulations including lots of spatial and structural information that clearly exceed what could be done on the basis of a pure CBN representation.

The article is structured as follows. In Section 2, we briefly introduce the basics of the CBN and cellular automata frameworks. We then present our hybrid approach for modelling mechanisms in Section 3. As a proof-of-principle we use a simple toy example for illustration of the core ideas how to model mechanisms by supplementing CBNs with cellular automata in this section. In Section 4, we then use a simplified example of protein–receptor binding to illustrate how our approach can avoid problems CBN approaches have with representing rich spatial and structural information in biological mechanisms. In Section 5, we summarize our results and discuss open questions for future research.

2. Causal Bayes Nets and Cellular Automata

In this section we first briefly review the basics of the CBN formalism. We then give a short introduction to cellular automata. We will only present the basics and parts of these formalisms that are relevant for what we will do in the remainder of the article. Both formalisms are illustrated by means of a simple forest fire toy example.

2.1. Causal Bayes nets

The causal interpretation of Bayes nets was mainly developed by Clark Glymour and his students around 1990 (see, for example, Glymour *et al.* [1991]; Spirtes *et al.*

---

[3] We are not the first ones who suggest to combine Bayesian networks and cellular automata. Kocabas and Dragicevic ([2006]), for example, propose to generate transition rules of a CA for modelling the change of land use in an urban environment by applying Bayesian networks to geographic information systems, and Kohler et al. ([2015]) show that dynamic Bayesian networks (Murphy [2002]) and probabilistic cellular automata are intertranslatable under certain conditions. They then use this result to apply Bayesian inference methods in order to tackle problems of parameter estimation in partial differential equations. We are indebted to an anonymous referee for pointing us to these papers. We will come back to the question of how our endeavour in this article differs from these approaches later on.

[1993]) and later by Pearl ([2000]). CBNs are triples $\langle \mathbf{V}, \mathbf{E}, P \rangle$. $\mathbf{V}$ is a set of random variables $X_1, \ldots, X_n$ describing events or event types, $\mathbf{E}$ is a binary relation on $\mathbf{V}$ ($\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$) that is interpreted as direct causal dependence with respect to $\mathbf{V}$, and P is a probability distribution over $\mathbf{V}$ that is intended to provide information about the strengths of the causal influences between the variables in $\mathbf{V}$ propagated over the system of interest's causal structure $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ (which is also called a causal graph).

That X_i is a direct cause of X_j in a causal graph $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ (meaning that $\langle X_i, X_j \rangle \in \mathbf{E}$) is graphically represented by a causal arrow $X_i \rightarrow X_j$. $\mathbf{Par}(X_j)$ is the set of X_j's direct causes (or causal parents) in $\mathbf{G}$. $\mathbf{Des}(X_i)$ is the set of X_i's descendants. All the variables X_j for which there exists a causal path $X_i \rightarrow \ldots \rightarrow X_j$ in $\mathbf{G}$ are descendants of X_i. For technical reasons, X_i is assumed to be a descendant of itself. The descendants of X_i that are not identical with X_i are intended to represent X_i's (direct or indirect) effects in $\mathbf{G}$.

At the very heart of the CBN framework lies the idea that probabilistic dependence is produced by causal structure. How measured probability distributions and causal structures are related is expressed by several core axioms. The most important of these axioms is the one known under the name of the causal Markov condition (CMC) (Spirtes *et al.* [2000], p. 29):

> Causal Markov Condition: $\langle \mathbf{V}, \mathbf{E}, P \rangle$ satisfies the causal Markov condition if and only if every $X_i \in \mathbf{V}$ is probabilistically independent of its non-descendants $\mathbf{V} \backslash \mathbf{Des}(X_i)$ conditional on its causal parents $\mathbf{Par}(X_i)$.[4]

The CMC is assumed to hold for causally sufficient variable sets $\mathbf{V}$.[5] Whenever a model $\langle \mathbf{V}, \mathbf{E}, P \rangle$ satisfies CMC, the model's graph $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ determines the following Markov factorization:

$$P(X_1, \ldots, X_n) = \prod_{i=1}^{n} P(X_i \,|\, \mathbf{Par}(X_i)) \tag{1}$$

The conditional probabilities $P(X_i \,|\, \mathbf{Par}(X_i))$ are called the causal model's parameters. According to Equation 1, specifying a model's parameters fully determines the joint probability distribution over $\mathbf{V}$. There are also further axioms such as the causal minimality and the causal faithfulness condition that will hold in many contexts and play an important role for several purposes, first and foremost for causal discovery (Spirtes *et al.* [2000]). For this article, however, we will only require CMC.

Consider as a simple example a CBN that models how forest fires emerge and might destroy huts (see Figure 1). *Cmp* stands for the number of campfires in a certain region, *Moi* for the number of subregions of a certain size with high moisture, *Fire* for the number of subregions affected by forest fire, and *Hut* for the number of

---

[4] Probabilistic independence of X and Y conditional on Z can be defined as $(P(x|y, z) = P(x|z)) \vee P(y, z) = 0$ for all X-, Y-, and Z-values x, y, and z, respectively, where X, Y, and Z can be variables or sets of variables. Probabilistic dependence can be defined as the negation of probabilistic independence.

[5] A set of variables $\mathbf{V}$ is causally sufficient if and only if 'every common cause of any two or more variables in $\mathbf{V}$ is in $\mathbf{V}$, or has the same value for all units in the population' (Spirtes *et al.* [2000], p. 22).
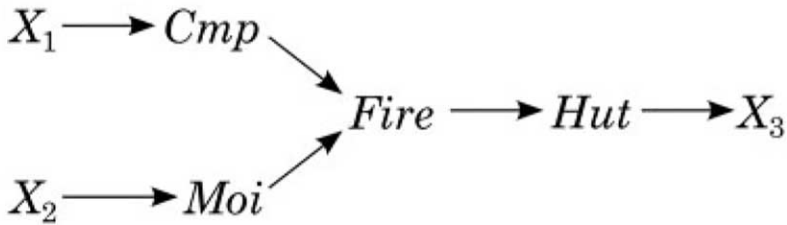
Figure 1. Graph of a simple exemplary CBN.

huts in the region of interest. The causal structures of interest can be imbedded in wider causal contexts. X_1 and X_2, for example, stand for not further specified causally relevant factors for *Cmp* and *Moi*, respectively, while X_3 is a not further specified effect of *Hut*. We assume that the number of camp fires has a positive causal influence on the number of forest fires (camp fires might cause forest fires), while the number of subregions with high moisture has a negative effect on the number of forest fires (moisture might prevent forest fires). In addition, we assume that the number of forest fires has a negative causal influence on the number of huts in the area of interest (forest fires destroy huts). Every variable in the CBN is, according to CMC, independent of its non-descendants conditional on its direct causes. The graph in Figure 1 together with the probability distribution over $\mathbf{V} = \{X_1, X_2, X_3, Cmp, Moi, Fire, Hut\}\}$ could now be used for causal reasoning and for generating predictions about what would happen under possible interventions.

2.2. Cellular automata

Cellular automata were originally developed by von Neumann ([1966]). They can be used to model dynamic processes in many different fields such as biology, physics, complexity science and many others (for an overview and examples, see Nagel and Schreckenberg [1992]; Ermentrout and Edelstein-Keshet [1993]; Schiff [2008]; Bandini *et al.* [2011]). They allow for modelling complex and self-regulatory systems and are suitable for simulating and visualizing spatial and structural changes systems might undergo over time.

Cellular automata describe a system's development over discrete time steps $t_0, t_1, t_2, \ldots$. A cellular automaton (CA) consists of an n-dimensional (finite or infinite) grid of cells (which is called the universe), a set of possible states $s_1, \ldots, s_m$ these cells might be in, a defined neighbourhood $\mathbf{N}_j(x)$ for each cell x (where j is the radius of the neighbourhood), and a set of rules that determine the state of every cell x at the next time step t_{i+1} on the basis of the cells in its neighbourhood at step t_i. The rules are typically the same for all time steps t_i. They can be deterministic or probabilistic. When formulating rules, it is important to consider boundary conditions. These boundary conditions determine how cells close to the edges of the grid behave and how their

neighbourhoods are defined. For running simulations, it is also important to specify the states of every cell in the universe at the first time step t_0. The assignment of states to all cells in the universe at a time step t_i is called a configuration. How exactly parameters such as the neighbourhood or the rules are choosen is a practical question that has to be considered by the modeller in order to make the model of the system of interest empirically adequate.

We use the same forest fire example we used above for illustrating CBNs to give an impression what can be done on the basis of CAs. For the sake of simplicity, we use a two-dimensional 8×8 grid as our universe. Each cell $x_{i,j}$ stands for a subregion of the region of interest and can be in five different states: s_0 (white) for the cell is empty (except for trees), s_1 (black) for hut, s_2 (grey) for forest fire, s_3 (circle) for high moisture, and s_4 (cross) for campfire. We specify a cell's neighbourhood as the Moore neighbourhood $\mathbf{N}_1$ with a radius of one (see Figure 2a). We assume that forest fires can spread through the region of interest, but cannot spread beyond that region's borders. Hence, the neighbourhood $\mathbf{N}_1(x_{i,j})$ of cells $x_{i,j}$ lying at the border of our universe consists of fewer cells (see Figure 2b).

We specify the following simple rules:

R.1. If a cell $x_{i,j}$ is occupied by a hut at t_k and there is a forest fire in its neighbourhood $\mathbf{N}_1(x_{i,j})$ at t_k, then $x_{i,j}$ becomes a forest fire cell at t_{k+1}.

R.2. If a cell $x_{i,j}$ is empty at t_k and there is a forest fire in $\mathbf{N}_1(x_{i,j})$ at t_k, then: (i) if there is a high moisture cell in $\mathbf{N}_1(x_{i,j})$ at t_k, then $x_{i,j}$ becomes a forest fire cell at t_{k+1} with probability 0.05. (ii) if there is no high moisture cell in $\mathbf{N}_1(x_{i,j})$ at t_k, then $x_{i,j}$ becomes a forest fire cell at t_{k+1} with probability 0.75.

R.3. If a cell $x_{i,j}$ is empty at t_k and there is a campfire (but no forest fire) in $\mathbf{N}_1(x_{i,j})$ at t_k, then: (i) if there is a high moisture cell in $\mathbf{N}_1(x_{i,j})$ at t_k, then
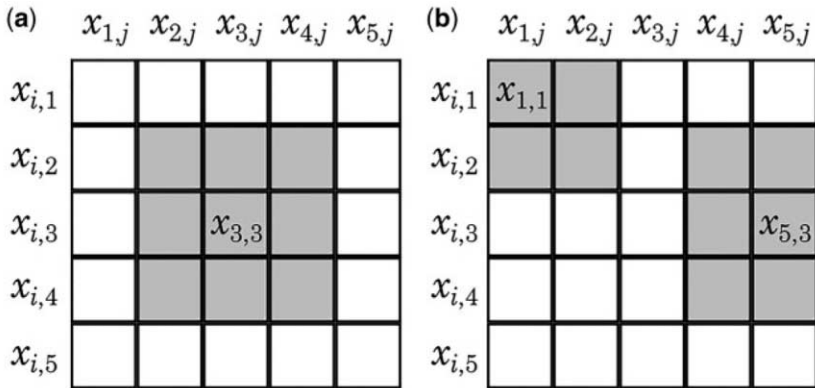

Figure 2. (a) shows the Moore neighbourhood $\mathbf{N}_1$ of cell $x_{3,3}$ in a 5×5 universe; (b) shows the Moore neighbourhood $\mathbf{N}_1$ of the cells $x_{1,1}$ and $x_{5,3}$ lying at the border of that universe.
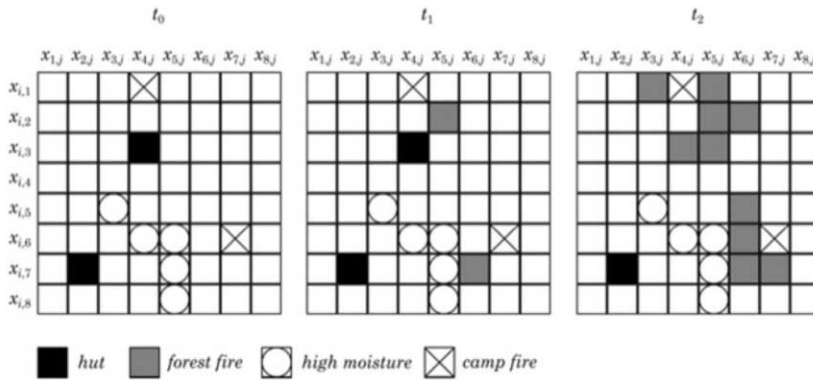
Figure 3. CA simulation of the forest fire example over 3 time steps.

$x_{i,j}$ becomes a forest fire cell at t_{k+1} with probability 0.01. (ii) if there is no high moisture cell in $\mathbf{N}_1(x_{i,j})$ at t_k, then $x_{i,j}$ becomes a forest fire cell at t_{k+1} with probability 0.05.

Rule R.1 is deterministic, while R.2 and R.3 are probabilistic. In each round, R.1 is first applied to each cell, then R.2 is applied to each cell, and then R.3 is applied to each cell. R.1 says that huts burn down if there was a forest fire in their neighbourhood at the time step before. R.2 dictates how forest fires spread through our 8×8 cell universe, and R.3 how forest fires emerge. In particular, R.2 says that forest fires spread with a certain probability and R.3 says that there is a certain probability for forest fires to occur in the vicinity of camp fires. These probabilities are higher if the ground is dry and lower if it is wet. For a simple exemplary simulation over three time steps, see Figure 3.

Note that while CAs can nicely capture the dynamics of systems involving lots of spatial and structural information, they are—contrary to CBNs—not backed up by a thorough philosophical theory of causation. They provide no intrinsic information about the causal structure underlying the systems of interest other than what is provided by the modeller. Often, one might want to interpret the transition rules regulating how cells behave over time steps as descriptions of causal laws or mechanisms. However, these rules can also express non-causal laws of nature, they could be a by-product of certain boundary conditions, or they could simply hold because one or more hidden common causes are at work. Formally, they are therefore blind with respect to the difference between causation and correlation.

3. A Hybrid Approach for Modelling Mechanisms

Our hybrid approach for modelling mechanisms is based on Gebharter's ([2014]) proposal to represent mechanisms as simple input–output structures consisting of one or more causal arrows. In Machamer *et al.*'s ([2000]) terminology, one could

say that the variables at the tails of these arrows describe the start or set-up conditions, while the variables at the arrows' heads describe the finish or termination conditions. What is going on inside the mechanism can then be described by supplementing these input–output structures by more detailed graphs describing the entities and activities (modelled by additional variables) involved in bringing about the phenomenon of interest as well as their organization (modelled by additional arrows). We will use the forest fire toy example introduced in Section 2 to illustrate our hybrid approach for modelling mechanisms.[6] In Figure 4a, for example, the arrows *Cmp → Hut* and *Moi → Hut* inside the dashed box could represent a mechanism that is part of a larger causal structure.[7] *Cmp* and *Moi* would stand for the mechanism's input, and *Hut* for its output. Submechanisms could then be represented by further bundles of arrows that describe in more detail what is going on inside the original mechanism. Two possible submechanisms of the mechanism depicted in Figure 4a could, for example, be represented by the two structures *Cmp → Fire ← Moi* and *Fire → Hut* in the dashed boxes in Figure 4b. For details and explicit conditions that have to be satisfied such that the structures in the dashed boxes in Figure 4b can be submechanisms of the structure in Figure 4a, see (Gebharter [2014], Section 4).

Now the problem with CBN approaches to mechanisms that Kaiser ([2016], p. 921) highlights is 'their failure to include relevant spatio-structural information in a way that does not render the models non-explanatory, unmanageable, or inconsistent with basic assumptions of causal graph theory'. Kaiser's main worry with CBN models of mechanisms is that though they can nicely represent causal dependencies at a quite abstract level, they fail to capture spatio-temporal details relevant for mechanistic explanation in a feasible way. The variables whose causal connections our forest fire CBNs describe, for example, represent quite abstract features of the system of interest: They stand for numbers of specific objects in a certain region of interest. These variables do not provide more specific spatial and structural information about camp fires, subregions with high moisture, forest fires, and huts, how forest fires emerge and spread through the region of interest and how they interact with huts if they come close to them. One could, in principle, replace the CBNs in Figure 4 by much more complex CBNs featuring lots of variables that could then describe the positions, behaviours, and so on of all the possible objects involved

---

[6] This is clearly not the kind of mechanism biologists and other scientists are interested in. However, we think that it is still a nice example to illustrate our approach. Also note that we do not restrict ourselfs to biological mechanisms, but have a much broader notion of mechanism in mind here. Basically, all kinds of causal input–output structures qualify as mechanisms as long as one can tell a more detailed causal story about what is going on inside the causal arrows. In Section 4, we will demonstrate how our hybrid approach can avoid the problems CBNs might have with representing spatial and structural information by providing a more sophisticated hybrid model of a biological mechanism.

[7] Note that information about which bundles of causal arrows represent mechanisms is left open by the formalism. Which arrows should stand for mechanisms (or submechanisms), that is how to draw dashed boxes, is formalism-external information that might to some extent depend on the particular interests of the causal modeller.
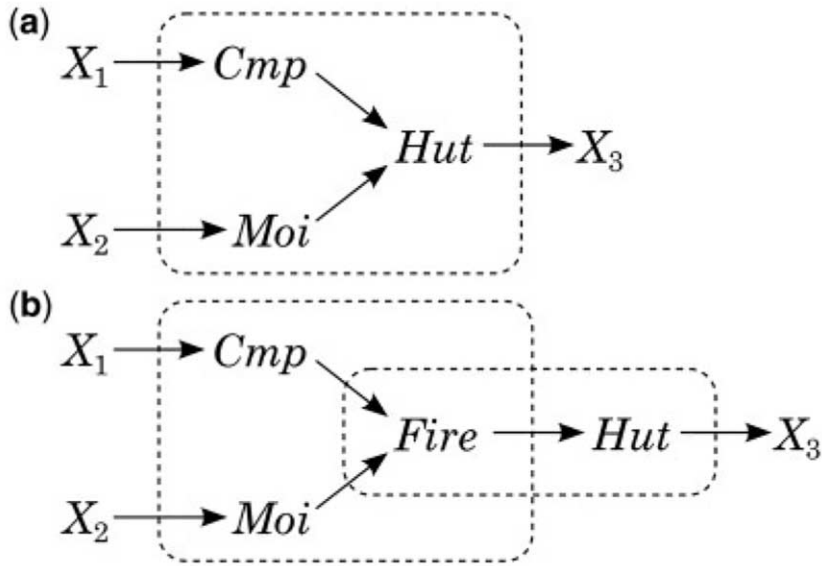
Figure 4. The substructures in the dashed boxes represent mechanisms. The two mechanisms in (b) are submechanisms of the mechanism in (a).

in our toy forest fire mechanism. But such a move would make the nice and simple CBN models in Figure 4 unnecessarily complex, confusing, and probably more or less useless from a practical point of view. It might, in addition, lead to conceptual problems. (For details, see Kaiser [2016].) It seems to be better to use CBN models for what they do best—representing causal structures at a more abstract level—and to leave capturing the relevant spatial and structural properties to the CA formalism, which is much better suited for that purpose.

We now have the CBNs in Figure 4 as well as the CA specified in Section 2 at hand to model the forest fire toy mechanism. The CBNs capture the causal dependencies involved at a more abstract level, and the CA provides detailed spatial and structural information about possible processes within that mechanism. As a last step, we have to combine the CBNs and the CA. In the approach presented in this article, the key idea is to outsource complex spatial and structural features relevant to (sub)mechanisms represented in a CBN into a CA. Recall that the CBNs' variables were intended to represent numbers of certain objects in a region of interest. The variable *Cmp* modelling the number of camp fires could, for example, describe how many of the 8×8 cells of our CA are in state s_4, *Moi* how many of the cells are high moisture subregions (are in state s_3), *Fire* how many cells are occupied by forest fire (are in state s_2), and *Hut* how many of our 8×8 cells feature huts (are in state s_1). These variables are still quite abstract: They do not provide specific information about how states $s_1, \ldots, s_4$ are distributed over the cells in our 8×8 grid. Their values

are compatible with many different configurations in the CA that can be further specified according to the needs of the modeller.

Note that our CBNs' parameters and the independencies implied by our CBNs' causal structures and CMC have to fit the forest fire CA. In general, there are several possibilities how such a fit could be established. If the CA will reach a steady state in the long run, then our CBNs' parameters can be identified with the corresponding conditional probabilities featured by the CA's steady state distribution. Instead of computing these conditional probabilities, one can also calculate them on the basis of simulations. If the system under study will not reach steady state in the long run, then one might try to identify the parameters with the corresponding conditional probabilities of the CA-distribution over t_{i+n} steps. The parameters $P(Fire|Cmp, Moi)$ and $P(Hut|Fire)$ of the CBN whose graph is depicted in Figure 4b, for example, might then be identified with the CA-probabilities of *Fire*'s values at step t_{i+1} given the values of *Cmp* and *Moi* at t_i and with *Hut*'s values at step t_{i+2} given the values of *Fire* at t_{i+1}. This strategy, identifying the parameters at an operating point of interest, seems to work well for our simple forest fire mechanism. Another strategy for specifying parameters in the case of non-steady state processes would be to use dynamic causal models (Murphy [2002]; Gebharter and Schurz [2016]). Every time step in the dynamic causal model would then correspond to a time step in the CA. Since space is limited, we will confine ourselves to applying the first two non-dynamic strategies. The investigation of how exactly dynamic CBNs and cellular automata can be combined to model more complex non-steady state behaviour of mechanisms shall be left for future work.

According to the considerations above, a model of our CBN–CA hybrid approach intended to represent one or more mechanisms can be defined as follows[8]:

> CBN–CA Hybrid Model: $\langle \mathbf{V}, \mathbf{E}, P, \mathbf{U}_1, \dots, \mathbf{U}_n, CA_1, \dots, CA_n \rangle$ is a CBN–CA hybrid model if and only if $\langle \mathbf{V}, \mathbf{E}, P \rangle$ is a causal Bayes net and $CA_1, \dots, CA_n$ are cellular automata such that
>
> > (1) $\mathbf{U}_1, \dots, \mathbf{U}_n$ are (non-empty) subsets of $\mathbf{V}$ of which none contains variables X, Y connected by a directed path $X \to \dots \to Y$ featuring more than one arrow;
> >
> > (2) the values of each variable in a set $\mathbf{U}_i$ describe a number of possible configurations of CA_i;
> >
> > (3) the parameters of each Y in a set $\mathbf{U}_i$ that has causal parents $X_1, \dots, X_m$ in $\mathbf{U}_i$ equal the conditional probabilities $P(Y|X_1, \dots, X_m)$ determined by CA_i (or calculated on the basis of simulations); and
> >
> > (4) whenever $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ together with CMC determines an independence among variables in a set $\mathbf{U}_i$, these independence is also featured by the probability distribution

---

[8] Note that this is only one possibility to combine CBNs and CAs to model mechanisms. We discuss further possibilities in Section 5. A detailed investigation of these alternative possibilities has to be carried out in future work.

over $\mathbf{U}_i$ determined by CA_i (or calculated on the basis of simulations).

The sets $\mathbf{U}_i$ are intended to contain the input and output variables of the mechanisms modelled. Condition (1) says that mechanisms are represented by structures that might have several input and output variables, but do not feature intermediate causes. This reflects the basic idea of representing mechanisms by means of simple input–output structures in CBNs. Condition (2) connects the variables in a set $\mathbf{U}_i$ representing a mechanism to the CA corresponding to $\mathbf{U}_i$. Condition (3) reflects the requirement mentioned above that the parameters of a CBN involving mechanism variables must fit the probabilities determined by $CA_1, \ldots, CA_n$. Condition (4) finally says that the probability distributions determined by $CA_1, \ldots, CA_n$ must feature the screening off relations determined by the CBN's causal structure and CMC. Note that the CBN–CA hybrid model definition can—at least in principle—cover the steady state as well as the t_{i+n} step behaviour of models.

Note that our hybrid approach to modelling mechanisms differs from other approaches to combine Bayesian networks and CAs such as the ones proposed by Kocabas and Dragicevic ([2006]) or Kohler *et al.* ([2015]). Kocabas and Dragicevic use Bayesian networks to generate transition rules for a CA intended to model change of land use on the basis of geographic information systems. Kohler *et al.*, on the other hand, use results about the intertranslatability of dynamic Bayesian networks and CAs to apply Bayesian inference tools for purposes of parameter estimation in partial differential equations. It is irrelevant for these approaches whether the edges of the Bayesian networks are causally interpreted or not. Also note that because the goals these authors have in mind are quite different from ours, the ways how they combine Bayesian networks and CAs differ from ours as well. We are mainly interested in whether causally interpreted Bayesian network models of mechanisms can—contrary to what Kaiser ([2016]) claims—be used as a basis for modelling mechanisms involving complex spatio-structural features. Here the CA formalism comes to help. Whenever such information is required, one can supplement the input–output structures representing mechanisms in CBNs by a CA that involves the missing details about how exactly the mechanism's component parts might spatially and structurally interact to bring about the phenomenon of interest. The minimal conditions that have to be satisfied in order for a CA to play this role in a CBN model of a mechanism are specified in the CBN–CA hybrid model definition. They are intentionally formulated quite loosely and general in order to grant a maximum of flexibility to the modeller.

4. Modelling Spatio-Structural Mechanisms Using CBN–CA Hybrids

In Section 3, we made a proposal how to combine CBNs and CAs to model mechanisms. Before we proceed in showing our approach's potential to include complex spatio-structural information, we shall shortly recapitulate the shortcomings of CBN

approaches to mechanisms highlighted by Kaiser ([2016]). Based on a case study of protein–DNA binding Kaiser argues that CBNs are not able to incorporate complex spatial and structural information relevant in a plethora of biological mechanisms. More precisely, she claims that three distinct kinds of information are necessary to explain how a protein can recognize and bind a specific DNA motif:

> (1) Binding information about the kinds of bonds that are formed between amino acid residues and DNA base-pairs (direct contacts), and about the parts of the DNA and the [transcription factor (TF)] that indirectly interact with each other (indirect contacts);
> (2) Chemical-structural information about the primary sequences of DNA and TF binding sites, about the complementarity of their chemical structures;
> (3) Spatial information about the conformations of DNA and the TF and why they spatially fit, about how the DNA shape contributes to recognition (Kaiser [2016], p. 926).

Clearly, these types of information are relevant for understanding mechanisms of macromolecular interactions (such as interactions between TFs and DNA). Kaiser ([2016]) argues that CBNs are not designed to incorporate such data and that desperately trying to capture them (for example, by means of adding additional variables and values) not only is cumbersome, but can result in contradictions with basic assumptions of graph theory such as the requirement that the variables of a causal model have to be conceptually independent of each other. The problem at heart can be aggravated even more: Although it might be sufficient for many biological applications to incorporate the above information, others also require to clarify how the involved macromolecules acquire their conformation, that is, a certain geometrical shape. The problem with such cases is that the three kinds of information listed above are not independent of each other. If one wants to gain a fundamental understanding how a certain protein acquires its shape one needs to know its primary structure, that is its amino acid sequence, and chemical-structural and physical information about the amino acids such as shape, charge, and reactive groups as well as lots of biophysics to understand how these properties determine secondary and tertiary structure of the protein (its resulting three-dimensional shape). Most likely, a good deal of the required biophysics will contain rather fundamental and mechanistically inexplicable laws.[9] In such cases then, a modeller interested in a mechanism that requires a rather detailed understanding (for example of protein folding) faces the challenge to implement a plethora of spatio-structural, chemical, and physical information as well as mechanically inexplicable laws into a model of a biological mechanism. It is highly questionable that CBNs alone are suitable for such a task.

Our hybrid CBN–CA approach, however, might fare better in cases such as the one described. In order to demonstrate its potential, we choose an example analogue

---

[9] Note that even Glennan ([1996]), who advocates a mechanistic analysis of causal laws, accepts that more fundamental physical laws might not be explicable in terms of mechanisms.

$$S_u \longrightarrow Protein \longrightarrow Receptor \longrightarrow S_d$$

Figure 5. Graph of a simple CBN that models how a certain protein activates a receptor.

to the case of protein–DNA binding used by Kaiser ([2016]). The phenomenon on which we will focus in the remainder of this article is protein–receptor binding, a type of macromolecular interaction in which a ligand protein binds to a receptor protein. Let us assume that we are interested in modelling a biological mechanism that involves a regulatory protein that is expressed upon a specific upstream signal (S_u) and binds to a certain intercellular receptor in order to activate it. Once activated, the receptor engages in other downstream signals (S_d). This biological process could be modelled by a CBN with a simple linear graph as shown in Figure 5.

To keep the example simple, suppose the upstream and downstream signals S_u and S_d are binary variables that can either take the value 'present' or 'not present'. The values of *Protein* represent the concentration of our regulatory protein of interest within the biological cell. Although more fine grained values are possible, we assume that in our example only the values 0, 8, 16, 32 and 64 [proteins/biological cell] are relevant. The values of *Receptor* represent the ratio of activated receptors to total receptor count (ranging from 0 to 1 in steps of 0.05). For the purpose of our study (considering only one type of ligand), we make a rapid equilibrium assumption and assume that once a protein–receptor complex is formed it is instantaneously in its active form.

For some applications such a CBN representation is perfectly suitable and sufficient (see, for example, the model from Sachs *et al.* [2005]). Suppose, however, that we are also interested in understanding the molecular and structural basis of the interaction between the protein and its receptor underlying the phenomenon of receptor binding. Such level of detail in a mechanism is, for example, relevant for explaining how certain mutations interfere with protein–receptor binding on a molecular level by changing the protein's structure. We propose to implement these details by supplementing the causal arrow *Protein → Receptor* with a CA. The CA is intended to provide detailed information about how component parts of the mechanism for protein–receptor binding spatially and structurally interact in order to bring about this phenomenon. As a biologically realistic CA of protein–receptor interaction, due to the immense complexity of the process, is beyond the scope of this article, we will simplify the biology in our example a little bit.[10]

---

[10] Computational structural biology is a very active field of research and predicting the molecular structure of biological macromolecules and their interactions is one of the most challenging problems in computational biology. Although cellular automata have been used for modelling protein folding, we are not aware of an already existing CA model that we could use to model the structure and the interactions between multiple proteins and receptor molecules.

4.1. The biophysics of a tiny peptide world

We assume a two-dimensional universe in which life has evolved. To understand the structural basis of protein folding and receptor binding in this universe, we construct the following CA model of a biological cell that contains a cell membrane, receptors, and proteins. These will be the component parts of the mechanism. We specify a cell[11] x's neighbourhood $\mathbf{N}_j(x)$ as the Moore neighbourhood (with a radius of j) without x itself. Each cell in the CA model can either be empty or occupied by an amino acid or a not further specified particle with a postive ($+$), negative ($–$), or no charge. Although the basic building bricks of life are very similar to the ones in our world, there are some important differences. Proteins, for example, are rather small polypeptides that are composed of only three distinct kinds of amino acids: posimine (positively charged; $+$), negatiline (negatively charged; $–$), and neutramine (no charge). These amino acids share the same squared shape. In peptide sequences we will indicate them with a one-letter code: P for posimine, N for negatiline, O for neutramine. Analogously to our world, protein folding can induce a specific conformation of proteins. Whole molecules either diffuse within the cytosol or move according to repulsive and attractive forces. The following biophysical laws govern this behaviour:

(1) Each charged particle x is subject to repulsive and attractive forces determined by its own charge and by the charges within its neighbourhood $\mathbf{N}_j(x)$ that sum up to the resulting force $\vec{F}_{Res}$:

$$\vec{F}_{Res}(x) = \sum_{i \in \mathbf{N}_j(x)} \vec{F}_{i,x},$$

where $\vec{F}_{i,x}$ is the force a charged cell i exerts on x. $\vec{F}_{i,x}$'s norm depends on the distance $d_{i,x}$ between i and x by $\frac{1}{d_{i,x}}$, its direction depends on whether i and x attract or repel each other. While the same charges ($+/+$ and $–/–$) repel each other, distinct charges ($+/–$ or $–/+$) attract each other.

(2) At each time step, a peptide's charged amino acids x change their position in accordance to $\vec{F}_{Res}(x)$ if possible. Each amino acid x can only change its position with an empty cell within the overlap of the $\mathbf{N}_1$-neighbourhoods of the predecessing and subsequent amino acids $x - 1$ and $x + 1$ in the peptide sequence, that is with an empty cell in $\mathbf{N}_1(x - 1) \cap \mathbf{N}_1(x + 1)$. If the cell in the direction of $\vec{F}_{Res}$ is occupied, but there is an empty cell in the direction of $\vec{F}_{Res}$'s part-dimension, the amino acid changes its position with this cell (and thereby takes the next best option). Although neutramines are not influenced by electrostatic forces, they have a chance of 0.05 to change their position with a random empty cell within $\mathbf{N}_1(x - 1) \cap \mathbf{N}_1(x + 1)$.

---

[11] The term 'cell' denotes a location within the CA's grid. We will use the term 'biological cell' to refer to living cells.

(3) Peptides as whole molecules are also subject to attractive and repulsive forces. The resulting force $\vec{F}_{Res}^{\mathbb{P}}$ on a peptide $\mathbb{P}$ is the sum of the resulting forces $\vec{F}_{Res}(x)$ on each of its amino acids.

(4) In each time-step a peptide will move one cell in the direction of its $\vec{F}_{Res}^{\mathbb{P}}$ provided the movement is not blocked by other structures and $\vec{F}_{Res}^{\mathbb{P}} > 0$. If $\vec{F}_{Res}^{\mathbb{P}} = 0$, there is a chance of 0.2 to move towards a random direction. Due to a crowded cytosol that causes friction, movement is decelerated by factor 0.85 after each step.

Although very primitive, these imaginary laws of biophysics are sufficient to describe a system that shares crucial similarities with the behaviour of real proteins. First of all, proteins can take a specific conformation depending on their primary sequence: Figure 6a–d shows the conformations of different peptides resulting from simulations with a CA model that uses the above laws as rules (with radius $j = 2$ for the Moore-neighbourhood in law 1 and 3). The conformations shown in Figure 6a and b are very stable, hard to change and almost always return to this state after pertubation of the conformation. Figure 6c depicts different semi-stable conformations of the peptide PPNNONNPP in which it remains as long as there are no major pertubations. Due to being charged only negatively, the peptide NNNNN in contrast has no stable conformation but constantly oscillates between different shapes in which the amino acids have a high distance to each other (Figure 6d). Moreover, the laws also allow for interactions between molecules. The peptide PPPOPOONONNN, for example, has a polar charge pattern that enables it to bind molecules of its own kind (see Figure 6e). As we will later see also the kinetics of binding reactions occurring in such systems is similar to those of real chemical systems.

In summary the entities of our tiny peptide world are similar to real biological macromolecules with respect to their structural organization (linear polymers composed of monomeric units with different properties), possible interactions and activities (repulsion, attraction, diffusion) and resulting behaviour (formation of higher-order structures, binding kinetics similar to real binding reactions).

4.2. A hybrid model of receptor activation in the tiny peptide world

The described CA now enables us to construct a CBN–CA hybrid model of the mechanism describing how exactly the component parts of the mechanism bring about the phenomenon of interest, that is how different amounts of regulatory proteins cause different frequencies of activated receptors. Recall that we have already specified which values the variables from the graph in Figure 5 can take. Next, we specify the model's parameters. This could, for example, be done on the basis of empirical data. Suppose the parameters for the variables S_u, *Protein*, and S_d are $P(S_u = present) = 0.1, P(Protein = 0 | S_u = not\ present) = 1,\quad P(Protein = 16 | S_u = present) = P(Protein = 64 | S_u = present) = 0.25,\quad P(Protein = 32 | S_u =$
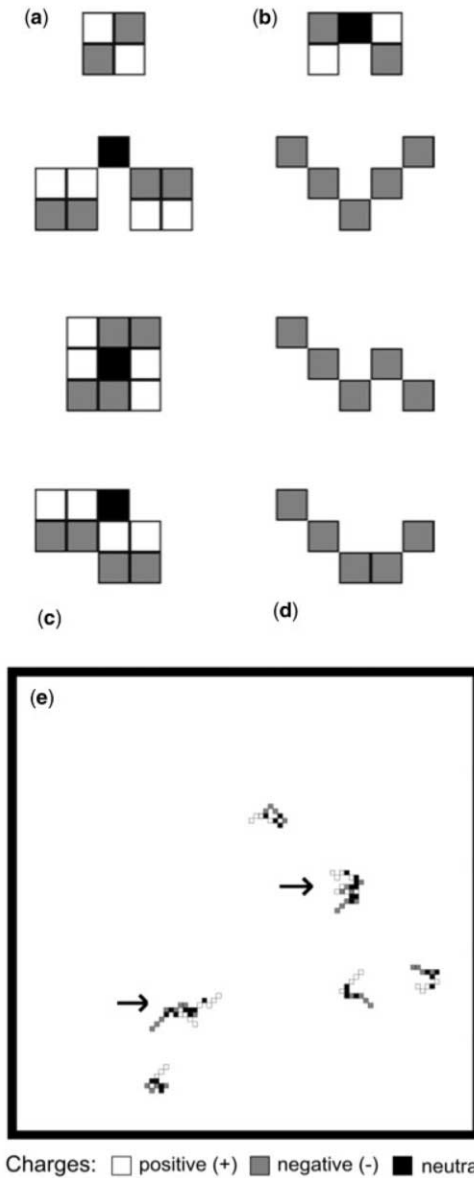
Figure 6. General features of peptides in the tiny peptide world: (a) Stable conformation of a peptide with the sequence NPNP; (b) stable conformation of a peptide with the sequence PNONP; (c) semi-stable conformations of a peptide with the sequence PPNNONNPP; (d) unstable conformations of a peptide with the sequence NNNNN; (e) excerpt from a simulation with eight peptides with sequence PPPOPOONONNN. The arrows mark two interacting peptides.
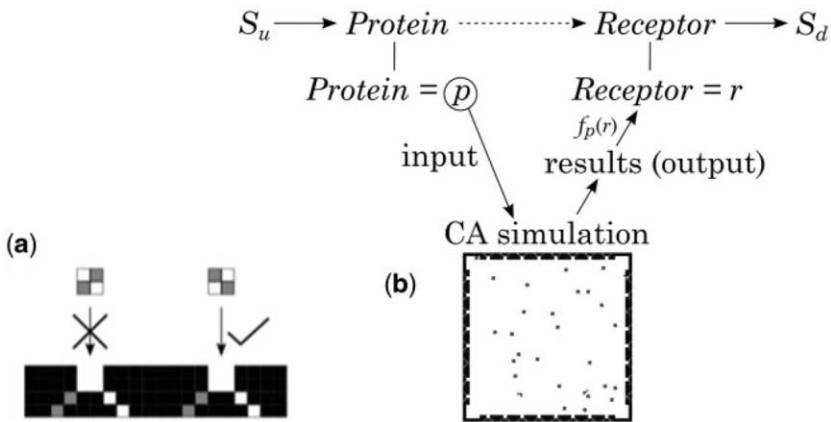
Figure 7. (a) Structure of the receptor. Only proteins with the right shape and orientation can bind into the ligand pocket (which is highlighted in yellow). (b) Sketch of the hybrid modelling approach.

present) = 0.5, $P(S_d = present | Receptor \geq 0.5) = 1$, and $P(S_d = not\ present | Receptor < 0.5) = 1$. Furthermore, we need to connect the CBN with the CA by means of specifying the remaining parameters for the variable *Receptor* from the simulations with the CA. In our example, the binding of the regulatory protein to its receptor will reach an equilibrium that persists as long as the regulatory protein is not degraded. We use this equilibrium to determine the parameters $P(Receptor = r | Protein = p)$ from several runs of a simulation using the CA with input p by the following function:

$$f_p(r) = \frac{\left| \forall\ t : \dfrac{number\ of\ active\ receptors\ at\ t}{total\ number\ of\ receptors\ at\ t} = r \right|}{t}, \tag{2}$$

where t are all generations from all runs in which the CA is at equilibrium. Suppose finally that our protein has the sequence NPNP and its receptor has the structure[12] depicted in Figure 7a. To bind to its receptor, our protein not only has to be in the right conformation, but also in the right orientation. Once the protein occupies the whole ligand pocket, the receptor switches to its active state. Together the CBN and the CA represent a hybrid model in the sense of the CBN–CA hybrid model definition in which the CBN describes the causal structure of the mechanism of interest on a general and more abstract level, while the CA model contains the spatial and structural information relevant to the mechanism. The whole approach is summarized in Figure 7b: $S_u \rightarrow Protein \dashrightarrow Receptor \rightarrow S_d$ is the CBN's graph. The

---

[12] Although in reality receptors are proteins, too, and thus are obliged to the same biophysics as any other protein, we assume that in the tiny protein world receptors feature additional properties that, in contrast to other proteins, ensure they have a fixed conformation.

rest of Figure 7b illustrates how the parameters for the dashed arrow are computed: For each possible value p of *Protein* the conditional probabilities $P(Receptor = r | Protein = p)$ for every possible value r of *Receptor* are computed via Equation 2 on the basis of CA simulations.

With this hybrid model at hand, we can perform simulations. Figure 8a shows the results for our protein at different concentrations and with different ranges for the attractive and repulsive forces, that is with a different radius j for the Moore-neighbourhood N_j. For N_1 and N_2, protein–receptor binding and activation reaches, in an exponentially increasing but saturating manner, an equilibrium (Figure 8a, left and middle graph).[13] For N_3, protein–receptor binding is unstable and results in an equilibrium near base-line (Figure 8a, right graph). From the equilibrium states we can now determine the CBN's missing parameters by the above parameter function that we have done exemplarily for N_2 (see Table 1).

Taking a closer look at the CA simulations (for N_1 and N_2) reveals how the protein–receptor interaction takes place at a molecular level: First, the protein takes a specific square shaped conformation (see Figure 6a) that can occasionally be disturbed by intermolecular interactions, but quickly returns to its original conformation. As it diffuses through the cell, it sometimes encounters a receptor to whose ligand pocket it can bind provided the protein has the right conformation and fits into the pocket. Once the ligand pocket is occupied, the receptor is activated. The attractive forces between the posimine and negatiline of the protein and the opposite charges of the receptor's ligand pocket prevent dissociation.[14]

Although it is a rather phenomenological description of the spatial aspects of our mechanism of interest, it contains all the kinds of information Kaiser ([2016]) claims to be relevant: information about the bonds between amino acids (1), the primary sequence of the protein and the sequence of the surface charges in the ligand pocket (2) as well as the conformation of the protein and the shape of the binding pocket (3). Moreover, the description could be easily explicated further by reference to the underlying laws of biophysics that can also explain how these three kinds of information are at least partially connected.

The outsourcing of spatio-structural information to the CA in CBN–CA hybrid models furthermore allows for predicting new effects on basis of the CA simulations. If, for instance, we are interested in understanding what impact a certain mutation of our protein will have on the activation of receptors, we can simply explore how the protein with the altered sequence behaves in the CA simulations. The mutations N1O and N1P for example (point mutations in which the negatiline at position 1 has been exchanged for a neutramine or posimine, respectively) impair the protein's conformation and binding capacity in most scenarios and thereby lead to a reduced

---

[13] Note that the kinetics exhibited in our simulations are remarkably similar to the chemical model of binding between a receptor R and a ligand L forming a complex RL, R + L $\rightleftarrows$ RL where the kinetics are described by ordinary differential equations.

[14] See the animation available at dclps.phil.hhu.de/ext-data/SupplVideo-1-NonAnon.mp4.
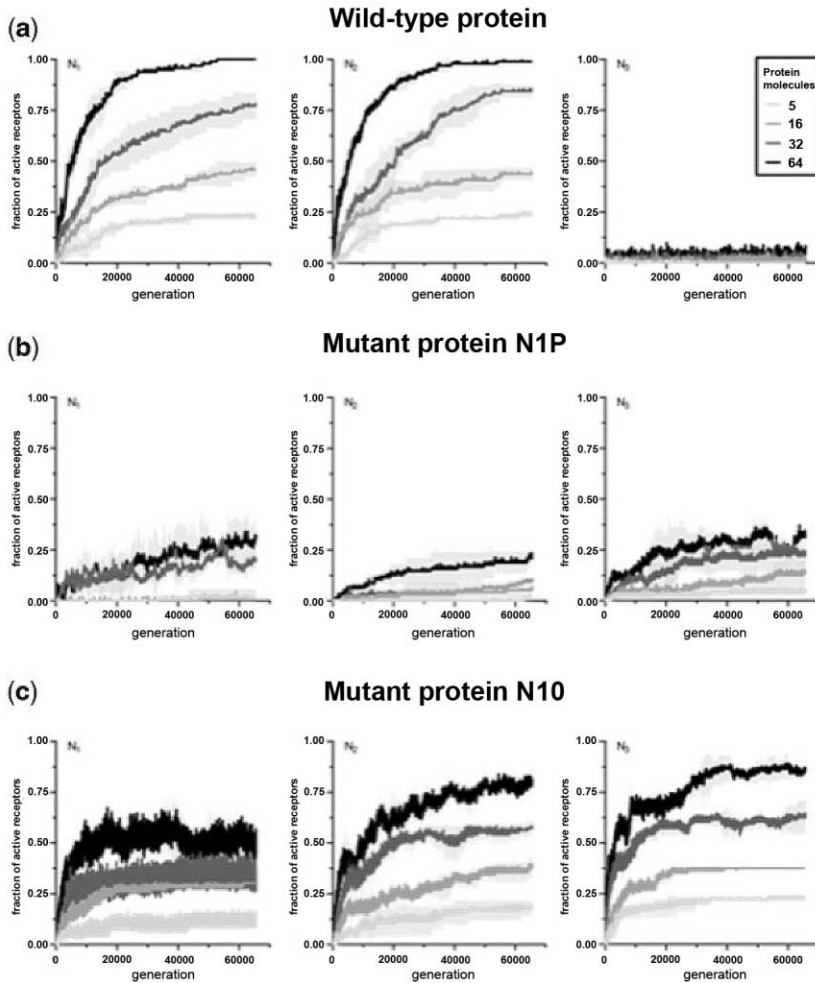
Figure 8. Simulation results for the relative receptor activation under variation of the protein concentration and radius of the Moore-neighbourhood ($N_1 - N_3$). Curves are plotted on the basis of $n = 3$ independent simulation runs with a random distribution of the proteins at g_0. The solid line shows the mean value, the shaded area the standard deviation of the relative receptor activation. (a) Results for the wild type protein with the sequence NPNP. (b) Results for the protein with the point mutation N1O in which the negatiline at position 1 has been exchanged for a neutramine. (c) Results for the protein with the point mutation N1P in which the negatiline at position 1 has been exchanged for a posimine.

receptor activation to varying degrees (see Figures 8b and 8c).[15] Interestingly, the opposite is the case for the radius N_3 as in these scenarios the mutant proteins exhibit a better binding to the receptor. This, too, can be attributed to the lack of square shaped

[15] See also the animation available at dclps.phil.hhu.de/ext-data/SupplVideo-2-NonAnon.mp4.

Table 1. Parameters for *Receptor* determined on the basis of the equilibrium states of three runs of the CA simulations with N_2 and wild type sequence of the protein

	$P(Receptor = r \mid Protein = p)$																				
r	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1
P																					
0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0.1	0.2	0.32	0.38	0	0	0	0	0	0	0	0	0	0
32	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.62	0.38	0	0
64	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.03	0.97

conformation since the increased range of attractive/repulsive forces allows the protein to take shapes and orientations that can still bind to the receptor.

5. Conclusion

We started this article with a problem for CBN representations of mechanisms highlighted by Kaiser ([2016]): Though CBN models of mechanism come with many advantages, they seem to have problems with capturing spatial and structural information relevant for some mechanisms, especially in biology. In this article we proposed to overcome such problems by representing such mechanisms within a hybrid modelling approach. In particular, we propose to supplement CBNs with CAs that can represent complex spatial and structural features of a system. In general, we represent mechanisms by causal input–output structures in CBNs. Whenever spatial and structural information is relevant, such an input–output structure can be supplemented with a CA. The input and the output variables in such a structure then describe global features of the CA, and the probability distribution generated by the CA must fit the distribution over the corresponding input and output variables in the CBN.[16]

Let us finally briefly discuss some possible open problems. We illustrated our hybrid modelling approach and discussed how it can handle spatial and structural information by means of the example of a molecular mechanism underlying protein–receptor binding. Though this mechanism is a very coarse simplification with mostly fictive biophysics, we think that this nicely shows that CBN–CA hybrids are at least in principle capable of dealing with the problem of modelling spatio-structural complex mechanisms. In reality, structural biology is, of course, much more complex and involves plenty kinds of different molecules, intra- and intermolecular bonds and forces. But even if it should turn out that CAs are not suitable for some mechanisms that require a high degree of sophisticated molecular detail, it is clear that the idea of multi-formalism modelling is not limited to CAs. We conjecture that also other methods established in computational structural biology—such as, for example, molecular dynamics simulations (reviewed, for example, in Karplus and McCammon [2002])—might be combined with CBN models of mechanisms.

Although our hybrid approach allows for modelling spatio-structural complex mechanisms, the modeller has to face some additional decisions, for example, how

---

[16] Our approach bears some similarity to the ideas of Chalupka *et al.* ([2017]) who distinguish between macro- and micro-variables and define the former in terms of partitions of the latter in order to automatically generate and test causal hypotheses on the macro-level variables; we are indebted to an anonymous referee for pointing us to this similarity. It thus seems that in the presented model (see Figure 7b) *Receptor* is a macro-variable based on micro-variables such as position and interactions of individual molecules in the CA. However, Chalupka *et al.* explicitly understand macro-variables as task/question specific and independent of the specific instantiation of the micro-variables. In some mechanisms, however, the specific micro-variable instantiation can have a significant influence on the mechanism's output (for example, when hysteresis or stochasticity in biochemical reaction networks play a role) and so these requirements are not necessarily fulfilled.

to specify the parameters. Especially when it comes to modelling non-steady state behaviour of dynamical systems, combining static CBNs with dynamical CAs might violate the Markov condition. Furthermore, our approach leads to a variety of yet unanswered questions, for example: When exactly is spatial information too complex to be represented in a CBN and what are the features that make it complex? How can the transition rules of a CA supplementing a causal input–output structure in a CBN representing a mechanism be learned from data? However, as the aim of this article was to show that hybrid modelling approaches can overcome some of the limitations of conventional CBN models of mechanisms, we leave these and other issues to future research.

A promising future philosophical question would be whether and, if yes, then to which extent CBN–CA hybrids can provide a unified approach to mechanistic explanations and explanations based on fundamental laws analogue to the biophysical laws we have specified for our CA model in Section 4. The simulations based on our fictive tiny peptide world, for example, support the importance and advantages of a unified approach as they have shown that variations of parameters in a law of nature (variation of the size of the Moore-neighbourhood from law 1 and 3) can directly affect the working of the investigated mechanism.

In the hybrid models in the sense of Section 4, the CAs capture what is going on inside of certain arrows or paths, that is the processes underlying these arrows or paths. Another fairly obvious way of combining CBNs and CAs is to use CBNs to describe what is (partly) going on in the CA. In such hybrids, the interface between both formalisms could be defined by transition-rules that feature variables that also occur in the CBN. The CBN could, for example, either describe global processes of the system that influence the CA through at least one of its variables, or the CBN could describe some aspects of an agent or cell within the CA and could be instantiated multiple times for each agent or cell. This class of hybrid models promises a multitude of intriguing possibilities for applications in the sciences, too. Exploiting the rich possibilities of learning and data-mining capabilities of the Bayesian network machinery, CBN models describing certain aspects of the behaviour of different biological cell types can be learned from data. Integration of these models into a supramechanistic CBN–CA hybrid model—maybe similar to Kocabas and Dragicevic's ([2006]) work—could then help to explore the interactions between cells and other components of the cellular environment. Such hybrid approaches could offer valuable insights for example into how different cell types within a certain tissue influence each other and contribute to the overall properties of that tissue or how the microenvironment of a tumour (consisting of different cell types and extracellular components) influences cancer cells.

Summarizing, we believe that CBN–CA hybrids and generally other hybrid modelling approaches to mechanisms might be a promising new field of research that offers new perspectives on philosophical questions (such as the ones mentioned) as well as on possible applications in the sciences.

Acknowledgements

Alexander Gebharter
Department of Theoretical Philosophy
University of Groningen
Groningen, Netherlands
alexander.gebharter@gmail.com

Daniel Koch*
Randall Centre for Cell & Molecular Biophysics
King's College London
London, UK
daniel.koch@kcl.ac.uk

References

Bandini, S., Rubagotti, F., Vizzari, G. and Shimura, K. [2011]: 'A Cellular Automata Based Model for Pedestrian and Group Dynamics: Motivations and First Experiments', in V. Malyshkin (*ed*.), *Parallel Computing Technologies*, Dordrecht: Springer, pp. 125–39.

Bechtel, W. and Abrahamsen, A. [2005]: 'Explanation: A Mechanist Alternative', *Studies in the History and Philosophy of the Biological and Biomedical Sciences*, **36**, pp. 421–41.

Brigandt, I., Green, S. and O'Malley, M. A. [2017]: 'Systems Biology and Mechanistic Explanation', in S. Glennan and P. Illari (*eds*), *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, London: Routledge.

Casini, L. and Baumgartner, M. [unpublished]: 'A Bayesian Theory of Causation'.

Casini, L., Illari, P. M., Russo, F. and Williamson, J. [2011]: 'Models for Prediction, Explanation and Control: Recursive Bayesian Networks', *Theoria*, **26**, pp. 5–33.

Chalupka, K., Eberhardt, F. and Perona, P. [2017]: 'Causal Feature Learning: An Overview', *Behaviormetrika*, **44**, pp. 137–64.

Clarke, B., Leuridan, B. and Williamson, J. [2014]: 'Modelling Mechanisms with Causal Cycles', *Synthese*, **191**, pp. 1651–81.

Craver, C. [2007b]: *Explaining the Brain*, Oxford: Clarendon Press.

Ermentrout, G. B. and Edelstein-Keshet, L. [1993]: 'Cellular Automata Approaches to Biological Modeling', *Journal of Theoretical Biology*, **160**, pp. 97–133.

Gebharter, A. [2014]: 'A Formal Framework for Representing Mechanisms?', *Philosophy of Science*, **81**, pp. 138–53.

Gebharter, A. [2017a]: *Causal Nets, Interventionism, and Mechanisms: Philosophical Foundations and Applications*, Cham: Springer.

Gebharter, A. [2017b]: 'Uncovering Constitutive Relevance Relations in Mechanisms', *Philosophical Studies*, **174**, pp. 2645–66.

Gebharter, A. and Kaiser, M. I. [2014]: 'Causal Graphs and Biological Mechanisms', in M. I. Kaiser, O. R. Scholz, D. Plenge and A. Hüttemann (*eds*), *Explanation in the Special Sciences*, Dordrecht: Springer, pp. 55–85.

Gebharter, A. and Schurz, G. [2016]: 'A Modeling Approach for Mechanisms Featuring Causal Cycles', *Philosophy of Science*, **83**, pp. 934–45.

Glennan, S. [1996]: 'Mechanisms and the Nature of Causation', *Erkenntnis*, **44**, pp. 49–71.

Glennan, S. [2009]: 'Mechanisms', in H. Beebee, C. Hitchcock and P. Menzies (*eds*), *The Oxford Handbook of Causation*, Oxford: Oxford University Press, pp. 315–25.

Glymour, C., Spirtes, P. and Scheines, R. [1991]: 'Causal Inference', *Erkenntnis*, **35**, pp. 151–89.

Goody, R. S. [2014]: 'How Not to Do Kinetics: Examples Involving GTPases and Guanine Nucleotide Exchange Factors', *FEBS Journal*, **281**, pp. 593–600.

Jones, N. and Wolkenhauer, O. [2012]: 'Diagrams as Locality Aids for Explanation and Model Construction in Cell Biology', *Biology and Philosophy*, **27**, pp. 705–21.

Kaiser, M. I. [2016]: 'On the Limits of Causal Modeling: Spatially-Structurally Complex Biological Phenomena', *Philosophy of Science*, **83**, pp. 921–33.

Karplus, M. and McCammon, J. A. [2002]: 'Molecular Dynamics Simulations of Biomolecules', *Nature Structural and Molecular Biology*, **9**, pp. 646–52.

Karr, J. R., Sanghvi, J. C., Macklin, D. N., Gutschow, M. V., Jacobs, J. M., Bolival, B., Assad-Garcia, N., Glass, J. I. and Covert, M. W. [2012]: 'A Whole-Cell Computational Model Predicts Phenotype from Genotype', *Cell*, **150**, pp. 389–401.

Kocabas, V. and Dragicevic, S. [2006]: 'Coupling Bayesian Networks with GIS-Based Cellular Automata for Modeling Land Use Change', in M. Raubal, H. J. Miller, A. U. Frank and M. F. Goodchild (*eds*), *Geographic Information Science*, Berlin: Springer, pp. 217–33.

Koch, D., Eisinger, R. and Gebharter, A. [2017]: 'A Causal Bayesian Network Model of Disease Progression Mechanisms in Chronic Myeloid Leukemia', *Journal of Theoretical Biology*, **433**, pp. 94–105.

Kohler, D., Marzouk, Y. M., Müller, J. and Wever, U. [2015]: 'A New Network Approach to Bayesian Inference in Partial Differential Equations', *International Journal for Numerical Methods in Engineering*, **104**, pp. 313–29.

Lotka, A.J. [1925]: *Elements of Physical Biology*, Philadelphia, PA: Williams and Wilkins.

Machamer, P., Darden, L. and Craver, C. [2000]: 'Thinking about Mechanisms', *Philosophy of Science*, **67**, pp. 1–25.

Michaelis, L. and Menten, M. L. [1913]: 'Die Kinetik der Invertinwirkung', *Biochem Z*, **49**, pp. 333–69.

Murphy, K. P. [2002]: 'Dynamic Bayesian Networks', Ph.D. Thesis, University of California, Berkeley.

Murray-Watters, A. and Glymour, C. [2015]: 'What Is Going on Inside the Arrows? Discovering the Hidden Springs in Causal Models', *Philosophy of Science*, **82**, pp. 556–86.

Nagel, K. and Schreckenberg, M. [1992]: 'A Cellular Automaton Model for Freeway Traffic', *Journal de Physique I*, **2**, pp. 2221–9.

Pearl, J. [2000]: *Causality*, Cambridge: Cambridge University Press.

Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D. A. and Nolan, G. P. [2005]: 'Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data', *Science*, **308**, pp. 523–9.

Schiff, J. L. [2008]: *Cellular Automata: A Discrete View of the World*, Hoboken, NJ: Wiley.

Schurz, G. and Gebharter, A. [2016]: 'Causality as a Theoretical Concept: Explanatory Warrant and Empirical Content of the Theory of Causal Nets', *Synthese*, **193**, pp. 1073–103.

Spirtes, P., Glymour, C. and Scheines, R. [1993]: *Causation, Prediction, and Search*, Dordrecht: Springer.

Spirtes, P., Glymour, C. and Scheines, R. [2000]: *Causation, Prediction, and Search*, Cambridge, MA: MIT Press.

von Neumann, J. [1966]: *Theory of Self-Reproducing Automata*, Urbana, IL: University of Illinois Press.

Weber, M. [2016]: 'On the Incompatibility of Dynamical Biological Mechanisms and Causal Graphs', *Philosophy of Science*, **83**, pp. 959–71.

Winsberg, E. [2009]: 'Computer Simulation and the Philosophy of Science', *Philosophy Compass*, **4**, pp. 835–45.

Wolkenhauer, O. and Green, S. [2013]: 'The Search for Organizing Principles as a Cure against Reductionism in Systems Medicine', *FEBS Journal*, **280**, pp. 5938–48.