

Research Article

Tineke Blom* and Joris M. Mooij

Causality and independence in perfectly adapted dynamical systems

Abstract: Perfect adaptation in a dynamical system is the phenomenon that one or more variables have an initial transient response to a persistent change in an external stimulus but revert to their original value as the system converges to equilibrium. The causal ordering algorithm can be used to construct an *equilibrium causal ordering graph* that represents causal relations and a *Markov ordering graph* that implies conditional independences from a set of equilibrium equations. Based on this, we formulate sufficient graphical conditions to identify perfect adaptation from a set of first-order differential equations. Furthermore, we give sufficient conditions to test for the presence of perfect adaptation in experimental equilibrium data. We apply our ideas to a simple model for a protein signalling pathway and test its predictions both in simulations and on real-world protein expression data. We demonstrate that perfect adaptation in this model can explain why the presence and orientation of edges in the output of causal discovery algorithms does not always appear to agree with the direction of edges in biological consensus networks.

Keywords: dynamical systems, causal modelling, causal discovery, protein signalling networks, Markov property, feedback, equilibrium

1 Introduction

Understanding causal relations is an objective that is central to many scientific endeavours. It is often said that ‘the gold standard’ for causal discovery is a randomized controlled trial, but practical experiments can be too expensive, unethical, or otherwise infeasible. The promise of causal discovery is that we can, under certain assumptions, learn about causal relations by using a combination of data and background knowledge [34, 50, 54]. Roughly speaking, causal discovery algorithms construct a graphical representation that encodes certain aspects of the data, such as conditional independences in the case of constraint-based causal discovery, given some constraints that are imposed by background knowledge. Under additional assumptions on the underlying causal mechanisms (e.g. the causal Markov condition, faithfulness, acyclicity) these graphical representations have a causal interpretation as well [26, 30, 34, 50]. In this work, we specifically consider the equilibrium distribution of *perfectly adapted dynamical systems* that have the property that the class of graphs that encode the conditional independences in the distribution does not have a straightforward causal interpretation in terms of the changes in distribution induced by soft or perfect interventions. Systems for which the causal relations and conditional independences cannot both be unambiguously represented by a single directed graph were recently discussed by Blom et al. [5] and Blom and Mooij [2].

Perfect adaptation in a dynamical system is the phenomenon that one or more variables initially respond to a persistent external stimulus but ultimately revert to their original value. As a consequence, variables in the system change due to an external input, but they become independent of the stimulus change after the system reaches equilibrium again. We study the differences between the causal structure implied by the dynamic equations and the conditional dependence structure of the equilibrium distribution. To do so, we make use of the technique of *causal ordering*, introduced by Simon [47], which can be used to construct a *Markov ordering graph* that encodes conditional independences between variables, as well

*Corresponding Author: Tineke Blom: Informatics Institute, University of Amsterdam; E-mail: t.blom2@uva.nl
 Joris M. Mooij: Korteweg-de Vries Institute, University of Amsterdam; E-mail: j.m.mooij@uva.nl

as a *causal ordering graph* that represents causal relations [5]. We introduce the notion of a *dynamic causal ordering graph* to represent transient causal effects in a dynamical model. We use these graphs to provide a sufficient graphical condition, for dynamical systems to achieve perfect adaptation, which does not require simulations or explicit calculations. Furthermore, we provide sufficient conditions to test for the presence of perfect adaptation in real-world data with the help of the Markov ordering graph and we elucidate the appropriate causal interpretation of the output of causal discovery algorithms when applied to (perfectly adapted) dynamical systems at equilibrium. Finally, we discuss how the notions of the causal Markov condition and the causal faithfulness condition, which are often used to tie graphs that represent conditional independences in a probability distribution to the causal properties of the system that generated the data, become ambiguous in the case of perfectly adapted dynamical systems where the equilibrium and dynamical causal ordering graph are different.

We illustrate our ideas on three simple dynamical systems with feedback: the bathtub model in Dash [11], Iwasaki and Simon [22], the viral infection model in Blom and Mooij [2], De Boer [14], and a chemical reaction network in Ma et al. [28]. We discuss how perfect adaptation may also manifest itself in applications of causal discovery algorithms to a popular protein expression data set [44]. The output of causal discovery algorithms applied to this data sometimes appears to be at odds with the biological consensus presented in Sachs et al. [44], see for example Mooij et al. [34], Ramsey and Andrews [40]. We present a model for the Ras-Raf-Mek-Erk signalling pathway, based on a model in Shin et al. [46], under saturation conditions and test its predictions both in simulations and on real-world data. We demonstrate that perfect adaptation in this model can explain why the presence and orientation of edges in the output of causal discovery algorithms does not always appear to agree with the direction of edges in biological consensus networks that are based on a partial representation of the underlying dynamical mechanisms.

2 Background

In this section we consider the assumptions underpinning popular constraint-based causal discovery algorithms and give a brief description of a simple local causal discovery algorithm, introduced by Cooper [10]. We proceed with a concise introduction to the causal ordering algorithm, which was first introduced by Simon [47] and conclude with a discussion of related work.

2.1 Causal discovery

The main objective in causal discovery is to infer causal relations from experimental and observational data. The most common causal discovery algorithms can be roughly divided into score-based and constraint-based approaches, where the latter are more generally applicable. The idea of constraint-based causal discovery algorithms (e.g PC or FCI and variants thereof, see Colombo et al. [9], Forré and Mooij [18], Spirtes et al. [50], Zhang [54]), which we focus on in the remainder of this section, is that causal relations can be inferred by exploiting conditional independences in the data. These algorithms attempt to construct an equivalence class of graphs that encode a set of conditional independence relations in a probability distribution via a graphical separation criterion. A *d-separation* is a relation between three disjoint sets of vertices in a graph that indicates whether all paths between two sets of vertices are blocked by the vertices in a third, see Pearl [38] or Spirtes et al. [50] for more details. If every d-separation in a graph implies a conditional independence in the probability distribution then we say that it satisfies the *directed global Markov property* w.r.t. that graph. Conversely, if every conditional independence in the probability distribution is due to a d-separation in a graph then we say that it is *d-faithful* to that graph. When a probability distribution satisfies the Markov property w.r.t. a graph and is also faithful to the graph, then this graph is a compact representation of the conditional independences in the probability distribution and we say that it *encodes* its independence relations.

A lot of work has been done to understand the various conditions (e.g. linearity, Gaussianity, discreteness, causal sufficiency, acyclicity) under which a graph that encodes all conditional independences and dependences in a probability distribution has a certain causal interpretation, see Colombo et al. [9], Forré and Mooij [18], Hyttinen et al. [21], Lacerda et al. [24], Mooij and Claassen [30], Mooij et al. [34], Richardson and Spirtes [42], Spirtes et al. [50], Strobl [51], Zhang [54]. Perhaps the simplest assumption is that the data was generated by a causal DAG¹ [50]. In that case, the *causal Markov condition*, which states that variables are independent of their non-effects conditional on all their direct causes, and the *causal faithfulness condition*, which states that there are no other conditional independences than those implied by the causal Markov condition, ensure that there exists a single DAG that represents both conditional independences and causal relations [26, 38]. For the acyclic setting, powerful constraint-based causal discovery algorithms such as PC (under the assumption of causal sufficiency) and FCI (when latent confounders may be present) have been developed [50].

However, many systems of interest in various scientific disciplines (e.g. biology, econometrics, physics) include feedback mechanisms. Cyclic Structural Causal Models (SCMs) [8] can be used to model causal features and conditional independence relations of systems that contain cyclic causal relationships. For linear SCMs with causal cycles, several causal discovery algorithms have been developed [21, 24, 42, 51] that are based on d-separations. The d-separation criterion is applicable to acyclic settings and to cyclic SCMs with either discrete variables or linear relations between continuous variables, but it is too strong in general [49]. Forré and Mooij [17], inspired by the ‘collapsed graph’ in Spirtes [49], developed the alternative σ -separation criterion for graphs that may contain cycles. If every σ -separation in a graph implies a conditional independence in the probability distribution then we say that it satisfies the *generalized directed global Markov property* w.r.t. that graph. Conversely, if every conditional independence in the probability distribution is due to a σ -separation in a graph then we say that it is σ -faithful to that graph. Forré and Mooij [18] propose a sound and complete causal discovery algorithm based on σ -separations and the assumption of σ -faithfulness for data that is generated by a cyclic SCM with non-linear relations between continuous variables. Recently, Mooij and Claassen [30] proved that the PC and FCI algorithms are sound and complete in this setting and showed how to read off causal relations and other features from the output of the algorithm. In earlier work, Richardson [41] proved soundness of a causal discovery algorithm under the generalized directed Markov property and the d-faithfulness assumption, under the additional assumption of causal sufficiency. At the end of this section we will consider the LCD (i.e. Local Causal Discovery) algorithm in Cooper [10], which was proven to be sound in both the σ - and d-separation settings [34].

In this work, we consider equilibrium distributions that are generated by dynamical models. The causal relations in an equilibrium model are defined through the effects of persistent interventions (i.e. interventions that are constant over time) on the equilibrium solution of variables that are endogenous to the model, assuming that the system again converges to equilibrium. It has been shown that directed graphs encoding the conditional independences between endogenous variables in the equilibrium distribution of dynamical systems with feedback do not have a straightforward and intuitive causal interpretation [5, 11, 22]. As a consequence, the output of algorithms such as LCD, PC, or FCI applied to equilibrium data of dynamical systems with feedback at equilibrium, cannot always be interpreted causally in a naïve way. One issue is that the equilibrium distribution of certain (perfectly adapted) dynamical systems can also be generated by a causal DAG (see e.g. the bathtub example in [5, 11, 22] and Section 3.1.1), while the causal mechanisms of the true underlying system are provided by the dynamics of a model that includes feedback. This example illustrates some of the arguments made by Dawid [13] against the use of causal DAGs. The in-depth analysis of causality and independence in perfectly adapted dynamical systems in this paper contributes to this discussion. Representations of dynamical systems at equilibrium as cyclic SCMs may not have a unique solution under perfect interventions [7] and therefore the causal semantics of the

¹ A Directed Acyclic Graph (DAG) is a pair $\langle V, E \rangle$ where V is a set of vertices and E a set of directed edges between vertices such that there are no directed cycles.

system may not be fully captured by the cyclic SCM [4]. Here, we will present methods that supplement existing methods for SCMs to study the properties of perfectly adapted dynamical systems in more detail.

In this paper we will, for the sake of simplicity, limit our attention to one of the simplest causal ordering algorithms, LCD. This algorithm is a straightforward and efficient search method to detect one specific (causal) structure from background knowledge and observation or experimental data [10]. The algorithm looks for triples of variables (C, X, Y) for which (a) C is a context variable that is not caused by any other observed variable and (b) the following (in)dependencies hold: $C \not\perp\!\!\!\perp X$, $X \not\perp\!\!\!\perp Y$, and $C \perp\!\!\!\perp Y \mid X$. Figure 1 shows the graphs that correspond to the LCD triple (C, X, Y) . Note that, in the absence of latent confounders, there are no bi-directed edges, and the graph structure of an LCD triple is a DAG. Under the causal Markov and causal faithfulness assumptions, directed edges in the graph of an LCD triple represent causal relations. For simplicity, we only consider DAGs to encode the conditional independence relations in equilibrium distributions of dynamical models. The ideas in this paper can be extended to a setting with latent variables and more advanced causal discovery algorithms.




Figure 1. Possible graph structures of an LCD triple. In the absence of latent confounders the triple has the structure of the DAG in the figure on the left.

2.2 Causal ordering

The causal ordering algorithm, which was first introduced by Simon [47], applies to sets of equations and returns an ordering of the variables and equations. Here, we give a brief introduction to the causal ordering algorithm of Nayak [37], which is based on the block triangular form of matrices in Pothén and Fan [39]. This algorithm is equivalent but computationally more efficient than the original causal ordering algorithm [5, 20]. It is applicable to sets of equations that can be represented by a bipartite graph with a *perfect matching* (i.e. there exists a subset $M \subseteq E$ of the edges in the bipartite graph $\mathcal{B} = \langle V, F, E \rangle$ so that every vertex in either V or F is adjacent to exactly one edge in M). Although the causal ordering algorithm has been extended to general bipartite graphs by Blom et al. [5], we will, for the most part, assume that a perfect matching exists for the sake of simplicity.

The structure of a set of equations and the variables that appear in them can be represented by a bipartite graph $\mathcal{B} = \langle V, F, E \rangle$, where vertices F correspond to the equations and vertices V correspond to the endogenous variables that appear in these equations. For each endogenous variable $v \in V$ that appears in an equation $f \in F$ there is an edge $(v - f) \in E$. The output of the causal ordering algorithm is a directed cluster graph $\langle \mathcal{V}, \mathcal{E} \rangle$, consisting of a partition \mathcal{V} of the vertices $V \cup F$ into clusters and edges $(v \rightarrow S) \in \mathcal{E}$ that go from vertices $v \in V$ to clusters $S \in \mathcal{V}$.

Application of the causal ordering algorithm to a bipartite graph $\mathcal{B} = \langle V, F, E \rangle$ results in the directed cluster graph $\text{CO}(\mathcal{B}) = \langle \mathcal{V}, \mathcal{E} \rangle$, which we will call the *causal ordering graph* [5]. It is constructed in four steps:

1. Find a perfect matching $M \subseteq E$ and let $M(S)$ denote the vertices in $V \cup F$ that are joined to vertices in $S \subseteq V \cup F$ by an edge in M .
2. For each $(v - f) \in E$ with $v \in V$ and $f \in F$: if $(v - f) \in M$ orient the edge as $(v \leftarrow f)$ and if $(v - f) \notin M$ orient the edge as $(v \rightarrow f)$. Let $\mathcal{G}(\mathcal{B}, M)$ denote the resulting directed graph.
3. Partition vertices $V \cup F$ into strongly connected components \mathcal{V}' of $\mathcal{G}(\mathcal{B}, M)$. Create the cluster set \mathcal{V} consisting of clusters $S \cup M(S)$ for each $S \in \mathcal{V}'$. For each edge $(v \rightarrow f) \in E$ add an edge $(v \rightarrow \text{cl}(f))$ to \mathcal{E} when $v \notin \text{cl}(f)$, where $\text{cl}(f)$ denotes the cluster in \mathcal{V} that contains f .
4. Optionally, exogenous variables appearing in the equations can be added as singleton clusters to \mathcal{V} , with edges towards the clusters of the equations in which they appear in \mathcal{E} .

Example 1. Consider the following set of equations with index set $F = \{f_1, f_2\}$ that contain endogenous variables with index set $V = \{v_1, v_2\}$:

$$f_1 : \quad X_{v_1} - U_{w_1} = 0, \quad (1)$$

$$f_2 : \quad X_{v_2} + X_{v_1} - U_{w_2} = 0, \quad (2)$$

where U_{w_1} and U_{w_2} are exogenous (random) variables indexed by $W = \{w_1, w_2\}$. Figure 2a shows the associated bipartite graph $\mathcal{B} = \langle V, F, E \rangle$. This graph has exactly one perfect matching $M = \{(v_1 - f_1), (v_2 - f_2)\}$, which is used in step 2 of the causal ordering algorithm to construct the directed graph $\mathcal{G}(\mathcal{B}, M)$ in Figure 2b. The causal ordering graph that is obtained after applying steps 3 and 4 of the causal ordering algorithm is given in Figure 2c. \triangle




Figure 2. The bipartite graph \mathcal{B} associated with equations (1) and (2) is given in Figure 2a. The oriented graph $\mathcal{G}(\mathcal{B}, M)$ obtained in step 2 of the causal ordering algorithm, with perfect matching M , in Example 1 is shown in Figure 2b. The causal ordering graph $\text{CO}\mathcal{B}$, with added exogenous variables, is given in Figure 2c. The corresponding Markov ordering graph $\text{MO}(\mathcal{B})$ is displayed in Figure 2d.

Throughout this work, we will assume that sets of equations are *uniquely solvable with respect to the causal ordering graph* [5]. Roughly speaking, this means that the endogenous variables in the model can be solved from the equations in their clusters along a topological ordering of the causal ordering graph. Recently, it was shown that the causal ordering graph represents the effects of soft and certain perfect interventions under the assumption of unique solvability w.r.t. the causal ordering graph [5]. Soft interventions target equations; they do not change which variables appear in the targeted equation and may only alter the parameters or form of the equation. Perfect interventions target clusters in the causal ordering graph and replace the equations in the targeted cluster with equations that set the variables in the cluster equal to constant values. We say that there is a direct path from a vertex x to a vertex y in a directed cluster graph $\langle \mathcal{V}, \mathcal{E} \rangle$ if either $\text{cl}(x) = \text{cl}(y)$ or there is a sequence of clusters $V_1 = \text{cl}(x), V_2, \dots, V_{k-1}, V_k = \text{cl}(y)$ so that for all $i \in \{1, \dots, k-1\}$ there is a vertex $z_i \in V_i$ such that $(z_i \rightarrow V_{i+1}) \in \mathcal{E}$. A soft intervention on an equation or a perfect intervention on a cluster has no effect on a variable in the causal ordering graph whenever there is no directed path to that variable from the intervention target (i.e. the targeted equation or an arbitrary vertex in the targeted cluster, respectively) [5]. Since the equations in Example 1 are uniquely solvable w.r.t. the causal ordering graph in Figure 2c we can use it to read off that, for example, a soft intervention targeting f_1 may have an effect on X_{v_2} and that a perfect intervention targeting the cluster $\{v_2, f_2\}$ has no effect on X_{v_1} .

Given the probability distribution of exogenous random variables, one gets a unique probability distribution on the endogenous variables under the assumption of unique solvability w.r.t. the causal ordering graph. The *Markov ordering graph* is a directed graph $\text{MO}(\mathcal{B})$ that implies conditional independences between the endogenous random variables that solve the system via d-separations [5]. The Markov ordering graph $\langle V, E \rangle$ is obtained from a causal ordering graph $\text{CO}(\mathcal{B}) = \langle \mathcal{V}, \mathcal{E} \rangle$ by putting $V = \bigcup_{S \in \mathcal{V}} S$ and constructing edges $(v \rightarrow w) \in E$ if and only if $(v \rightarrow \text{cl}(w)) \in \mathcal{E}$. The Markov ordering graph for the set of equations in Example 1 is given in Figure 2d. The d-separations in this graph imply conditional independences between the corresponding variables. For instance, since v_1 and w_2 are d-separated we know that X_{v_1} and X_{w_2} are dependent.

Assuming that the probability distribution is d-faithful to the Markov ordering graph and that we have a conditional independence oracle, we know that the output of the PC-algorithm is the Markov equivalence class of the Markov ordering graph. However, Blom et al. [5] and Blom and Mooij [2] already demonstrated that for certain dynamical systems, the directed edges in the Markov ordering graph should not be interpreted as causal relations. Likewise, we will discuss three examples of perfectly adapted systems at equilibrium for which the Markov ordering graph does not have a straightforward causal interpretation in Section 3.3.2. In Section 6.1 we provide a brief discussion about the ambiguity of the causal Markov and faithfulness conditions in these examples.

2.3 Related work

Causal ordering is a technique that can be used to relate the (equilibrium) equations in a dynamical model to causal properties and conditional independence relations [2, 5, 22]. The relationship between dynamical models and causal models has already received much attention over the years. The works of Fisher [16], Mogensen et al. [29], Rubenstein et al. [43], Sokol and Hansen [48], Voortman et al. [53] considered causal relations in dynamical systems that are not at equilibrium, while Blom et al. [4], Hyttinen et al. [21], Lacerda et al. [24], Lauritzen and Richardson [25], Mooij et al. [32, 33] considered graphical and causal models that arise from studying the stationary behaviour of dynamical models. Blom and Mooij [2] study the robustness of model predictions when two subsystems in a dynamical model at equilibrium are combined, and consider opportunities for using causal discovery to detect feedback loops and the presence of variables that are not self-regulating using both models and experimental data for a subsystem. The causal behaviour of dynamical models and their equilibration to an SCM is studied by Bongers and Mooij [7], Dash [11]. In previous work, researchers have noted various subtleties regarding the use of a single graphical model to represent both conditional independence properties and causal properties of certain dynamical systems at equilibrium [4, 11, 13, 24, 25]. Often, restrictive assumptions on the underlying dynamical models are made to avoid these subtleties. In this work we follow Blom and Mooij [2], and directly address these issues by using the causal ordering algorithm to construct separate graphical representations for the causal properties and conditional independence relations implied by these systems. Our approach can be used in the equilibrium setting, but can also be employed to model transient causal effects in non-equilibrium settings, as we will discuss in Section 3.2.1. In this paper we focus on using these ideas to study the properties of perfectly adapted systems and applying this in particular to better understand the causal mechanisms that drive protein signalling networks.

It has been shown that the popular SCM framework [8, 38] is not flexible enough to fully capture the causal semantics (in terms of perfect interventions targeting variables) of certain dynamical systems at equilibrium, and for that purpose Blom et al. [4] proposed to use Causal Constraints Models (CCMs) instead. The drawback of this approach is that the causal constraints do not possess some of the attractive properties of SCMs, although the techniques in Blom et al. [5] can be used to construct graphical representations of causal relations and conditional independences. In the discussion section we consider how the causal ordering technique can be used to obtain graphical presentations and a Markov property for the dynamical model of the basic enzyme reaction that was considered in Blom et al. [4].

The analysis of network topologies that can achieve perfect adaptation is a topic of interest in cell biology, see for example [1, 15, 23, 28, 36]. The present work provides a method that facilitates the analysis of perfectly adapted dynamical systems by providing a principled method to identify perfect adaptation either from model equations or from experimental data and background knowledge. It is our hope that the ideas presented in this paper contribute to increasing the impact of causal inference in cell biology and dynamical modelling.

3 Perfect adaptation

The ability of a system to converge to its original state when a constant and persistent external stimulus is added or changed is referred to as *perfect adaptation*. If the adaptive behaviour does not depend on the precise setting of parameters then we say that the adaptation is *robust*. In the literature, the most interesting of the two is robust perfect adaptation, which is also commonly referred to as perfect adaptation. Henceforth, we will use the term *perfect adaptation* to refer to *robust perfect adaptation*. In this section, we take a look at several examples of simple dynamical systems that can achieve perfect adaptation and consider how we can identify models that are capable of perfect adaptation. Finally, we discuss the correct interpretation of the output of some constraint-based causal discovery algorithms applied to perfectly adapted dynamical systems and possibilities for the identification of perfect adaptation from (equilibrium) data.

3.1 Examples

In this section we present three dynamical systems and show that they are capable of achieving perfect adaptation. The details of simulations that are presented in this section are given in Appendix A.




Figure 3. Simulations of the outflow rate $X_O(t)$ in the bathtub model, the amount of infected cells $X_I(t)$ in the viral infection model, and the concentration $X_C(t)$ in the biochemical reaction network with a negative feedback loop after a change in the input signal. The timing of this change is indicated by a vertical dashed line. The three systems started with input signals $I_K = 1.2$, $I_\sigma = 1.6$, and $I = 1.5$. After a transient response $X_O(t)$, $X_I(t)$, and $X_C(t)$ all converge to their original equilibrium value (i.e. they perfectly adapt to the input signal).

3.1.1 Filling bathtub

We consider the example of a filling bathtub from Iwasaki and Simon [22]. Let $I_K(t)$ be an input signal that represents the size of a drain in the bathtub. The inflow rate $X_I(t)$, water level $X_D(t)$, water pressure $X_P(t)$, and outflow rate $X_O(t)$ are modelled by the following static and dynamic equations:

$$X_I(t) = U_I, \quad (3)$$

$$\dot{X}_D(t) = U_1(X_I(t) - X_O(t)), \quad (4)$$

$$\dot{X}_P(t) = U_2(gU_3X_D(t) - X_P(t)), \quad (5)$$

$$\dot{X}_O(t) = U_4(U_5I_K(t)X_P(t) - X_O(t)), \quad (6)$$

where g is the gravitational constant, and $U_I, U_1, U_2, U_3, U_4, U_5$ are independent exogenous random variables taking value in $\mathbb{R}_{>0}$. Let X_D, X_P , and X_O denote the respective equilibrium solutions for the water

level, water pressure, and outflow rate. The equilibrium equations associated with this model can easily be constructed by setting the time derivatives equal to zero and assuming the input signal $I_K(t)$ to have a constant value I_K :

$$f_I : \quad X_I - U_I = 0, \quad (7)$$

$$f_D : \quad U_1(X_I - X_O) = 0, \quad (8)$$

$$f_P : \quad U_2(g U_3 X_D - X_P) = 0, \quad (9)$$

$$f_O : \quad U_4(U_5 I_K X_P - X_O) = 0, \quad (10)$$

We call the labelling f_D, f_P, f_O that we choose for the equilibrium equations that are constructed from the time-derivatives the *natural labelling* for this dynamical system, which means that the equilibrium equation constructed from $\dot{X}_i(t)$ of variable v_i is labelled as f_i . A solution (X_I, X_D, X_P, X_O) to the system of equilibrium equations satisfies $X_I = U_I$ and $X_O = X_I$ almost surely. From this we conclude that, at equilibrium, the outflow rate is independent of the size of the drain I_K , assuming that U_I is independent of I_K . We recorded the changes in the system after we changed the input signal I_K of the bathtub system in equilibrium. The results in Figure 3a show that the outflow rate has a transient response to changes in the input signal I_K , but it ultimately converges to its original value. The outflow rate X_O in the bathtub model *perfectly adapts* to changes in I_K .

3.1.2 Viral infection model

We consider the example of a simple dynamical model for a viral infection and immune response in Blom and Mooij [2], De Boer [14]. The model describes target cells $X_T(t)$, infected cells $X_I(t)$, and an immune response $X_E(t)$. We will treat $I_\sigma(t)$ as an exogenous input signal that represents the production rate of target cells. The system is defined by the following dynamic equations:

$$\dot{X}_T(t) = I_\sigma(t) - d_T X_T(t) - \beta X_T(t) X_I(t), \quad (11)$$

$$\dot{X}_I(t) = (f\beta X_T(t) - d_I - k X_E(t)) X_I(t), \quad (12)$$

$$\dot{X}_E(t) = (a X_I(t) - d_E) X_E(t). \quad (13)$$

We have that $\beta = \frac{bp}{c}$ where b is the infection rate, p is the number of virus particles produced per infected cell, and c is the clearance rate of viral particles. Furthermore, d_T is the death rate of target cells, a is an activation rate, d_E and d_I are turnover rates and k is a mass-action killing rate. We assume that a, k are constants and that d_T, d_I, d_E , and β are independent exogenous random variables. We use the natural labelling for the equilibrium equations that are constructed from the differential equations:²

$$f_T : \quad I_\sigma - d_T X_T - \beta X_T X_I = 0, \quad (14)$$

$$f_I : \quad f\beta X_T - d_I - k X_E = 0, \quad (15)$$

$$f_E : \quad a X_I - d_E = 0, \quad (16)$$

assuming a constant value I_σ of the input signal. We initialized the model in an equilibrium state and simulated the response of the model after changing the input signal I_σ to three different values. Figure 3b shows that the amount of infected cells $X_I(t)$ has a transient response to a change in the input signal, but then returns to its original value, it perfectly adapts to changes in I_σ .

² Following De Boer [14], we are only interested in strictly positive solutions of this dynamical system. Therefore, we use the equilibrium equation f_I instead of $(f\beta X_T - d_I - k X_E) X_I = 0$ and f_E instead of $(a X_I - d_E) X_E = 0$.

3.1.3 Reaction networks with a negative feedback loop

The phenomenon of perfect adaptation is a common feature in biochemical reaction networks and there exist many network topologies that can achieve (near) perfect adaptation [1, 15]. For networks consisting of only three nodes Ma et al. [28] found by an exhaustive search that there exist two major classes of network topologies that produce robust adaptive behaviour. The reaction diagrams for these networks are given in Figure 4. Here we will only analyse Negative Feedback with a Buffer Node (NFBN), we will examine the other network in the discussion section and in Appendix D. The NFBN system can be described by the following first-order differential equations:

$$\dot{X}_A(t) = I(t)k_{IA} \frac{(1 - X_A(t))}{K_{IA} + (1 - X_A(t))} - F_A k_{FAA} \frac{X_A(t)}{K_{FAA} + X_A(t)}, \quad (17)$$

$$\dot{X}_B(t) = X_C(t)k_{CB} \frac{(1 - X_B(t))}{K_{CB} + (1 - X_B(t))} - F_B k_{FBB} \frac{X_B(t)}{K_{FBB} + X_B(t)}, \quad (18)$$

$$\dot{X}_C(t) = X_A(t)k_{AC} \frac{(1 - X_C(t))}{K_{AC} + (1 - X_C(t))} - X_B(t)k_{BC} \frac{X_C(t)}{K_{BC} + X_C(t)}, \quad (19)$$

where $X_A(t)$, $X_B(t)$, $X_C(t)$ are concentrations of three compounds A , B , and C , while $I(t)$ represents an external input into the system. Assume that k_{IA} , k_{CB} , and k_{AC} are independent exogenous random variables, that we will denote as U_A , U_B , U_C respectively, and that the other parameters are constants. Ma et al. [28] show that perfect adaptation is achieved under saturation conditions, $(1 - X_B(t)) \gg K_{CB}$ and $X_B(t) \gg K_{FBB}$, in which case the following approximation can be made:

$$\dot{X}_B(t) \approx X_C(t)k_{CB} - F_B k_{FBB}. \quad (20)$$

Under the assumption that $I(t)$ has a constant value, the system converges to an equilibrium. We will denote the equilibrium equations that are associated with the time derivatives $\dot{X}_A(t)$ and $\dot{X}_C(t)$ using the natural labelling f_A and f_C . The equilibrium equation f_B is obtained by setting the approximation of the time derivative $\dot{X}_B(t)$ equal to zero. We initialized this model in an equilibrium state and then simulated its response after changing the input signal I to three different values. Figure 3c shows that $X_C(t)$ perfectly adapts to changes in the input signal I .



(a) Negative feedback with a buffer node.

(b) Incoherent feedforward loop with a proportioner node.

Figure 4. The two three-node network topologies that can achieve perfect adaptation in Ma et al. [28]. The motif in Figure 4a shows Negative Feedback with a Buffer Node (NFBN) B , while Figure 4b shows an Incoherent Feedforward Loop with a Proportioner Node (IFFLP) B . Orange edges represent saturated reactions, blue edges represent linear reactions, and black edges are unconstrained reactions. Arrowheads represent a positive influence and edges ending with a circle represent negative influence.

3.2 Identification of perfect adaptation

In this section, we consider graphical representations of dynamical systems that can achieve perfect adaptation. We use these representations to formulate a sufficient graphical criterion to identify perfect adaptation in first-order dynamical models.

3.2.1 Graphical representations

The *functional graph* is a compact representation of the structure of a set of first-order differential equations that are written in canonical form (i.e. derivatives are on the left-hand side of each equation and functions of variables on the right-hand side). Vertices in this graph represent variables or derivatives of variables and directed dashed edges from derivatives to their corresponding variables indicate integration links. Additionally, there is an edge from a variable to a derivative whenever that variable appears in the corresponding differential equation. Figures 5a, 5b, and 5c show the functional graphs for the bathtub model, the viral infection model, and the reaction network respectively.

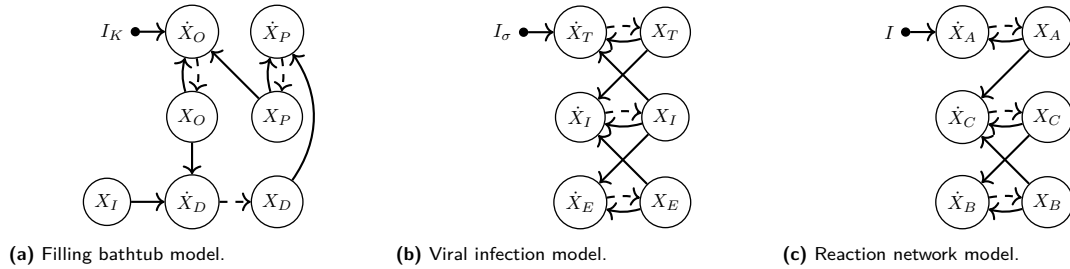


Figure 5. The functional graphs of the dynamics of the bathtub model, the viral infection model, and the reaction network with negative feedback. The input vertices I_K , I_σ , and I are represented by black dots.

Contrary to Iwasaki and Simon [22], we do not interpret the functional graphs in Figure 5 causally, because they may not have an intuitive causal interpretation in terms of regular interventions (e.g. we will not say that X_O causes \dot{X}_O nor that \dot{X}_O causes X_O even though there are directed edges between these vertices in the graph).³ Here, we will consider a graphical representation that represents causal relations in a system of first-order differential equations in canonical form. To do so, we associate both the derivative $\dot{X}_i(t)$ and the corresponding variable $X_i(t)$ with the same vertex v_i . We use the natural labelling for the differential equations, so that a vertex g_i is associated with the differential equation for $\dot{X}_i(t)$. We then construct a *dynamical bipartite graph* $\mathcal{B}_{\text{dyn}} = \langle V, F, E \rangle$ with variable vertices v_i in V and the corresponding dynamical equation vertices $g_i \in F$ and additionally static equation vertices $f_i \in F$. The edge set E has an edge $(v_i - f_j)$ whenever $X_i(t)$ appears in the static equation f_j . Furthermore, there are edges $(v_i - g_j)$ whenever $X_i(t)$ or $\dot{X}_i(t)$ appears in the dynamic equation g_j (which includes the cases $i = j$ due to the natural labelling used).

The dynamical bipartite graphs for the dynamics of the bathtub model, the viral infection, and the reaction network with feedback are given in Figures 6a, 6b, and 6c, respectively. Henceforth, we will assume that the dynamical bipartite graph has a perfect matching that extends the natural labelling of the dynamic equations, i.e. such that all pairs (v_i, g_i) are matched. Application of the causal ordering algorithm to the associated dynamical bipartite graph for the model of a filling bathtub, the viral infection model, and the reaction network results in the *dynamical causal ordering graphs* in Figures 7a, 7b, and 7c, respectively.⁴

The structure of the equilibrium equations can be used to construct an *equilibrium causal ordering graph* that represents the causal structure of dynamical models at equilibrium. The *equilibrium bipartite graphs* for the equilibrium equations of the filling bathtub, the viral infection, and the reaction network with feedback are given in Figures 6d, 6e, and 6f, respectively. Application of the causal ordering algorithm to these equilibrium bipartite graphs results in the equilibrium causal ordering graphs in Figures 7d, 7e,

³ Note that Bongers and Mooij [7] provide an alternative notion of the functional graph that does have an intuitive causal interpretation.

⁴ Our approach here differs from dynamic causal ordering in Iwasaki and Simon [22], who include separate vertices for derivatives and variables that are linked by ‘definitional’ integration links. Their result is similar to the functional graph in Figure 5a.

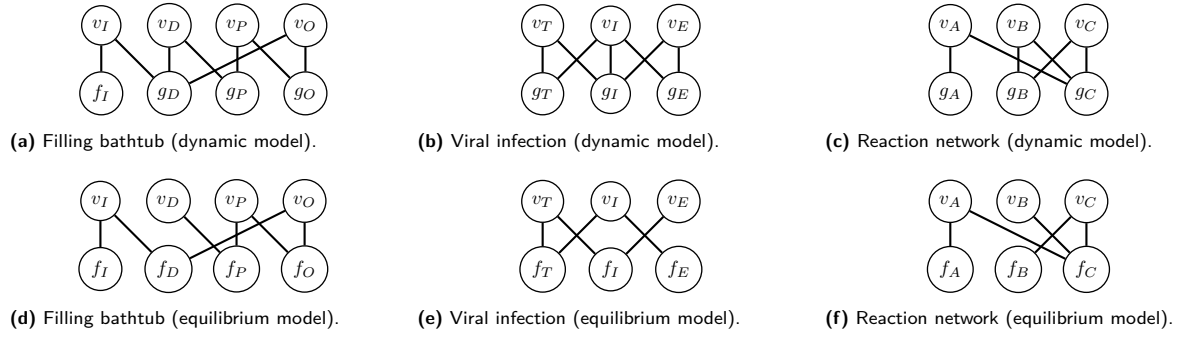


Figure 6. The dynamical bipartite graphs for the bathtub model, the viral infection, and the reaction network with negative feedback are presented in Figures 6a, 6b, and 6c, respectively. The equilibrium bipartite graphs for the bathtub model, the viral infection, and the reaction network with negative feedback are given in Figures 6d, 6e, and 6f, respectively. Comparing the equilibrium bipartite graphs with the dynamic bipartite graphs we note that there is no edge $(v_D - f_D)$ in Figure 6d while $(v_D - g_D)$ is present in Figure 6a, the edges $(v_I - f_I)$ and $(v_E - f_E)$ are not present in Figure 6e whilst the edges $(v_I - g_I)$ and $(v_E - g_E)$ are present in Figure 6b, and there is no edge $(v_B - f_B)$ in Figure 6f while the edge $(v_B - g_B)$ is present in Figure 6c.

and 7f, respectively. Notice that variables v_i do not always end up in the same cluster with the equilibrium equation f_i of the natural labelling. For example, we see in Figure 7d that a soft intervention targeting the equilibrium equation f_O constructed from the time derivative of the outflow rate $X_O(t)$ (e.g. a change in the value of U_5) does *not* affect the value of the outflow rate X_O at equilibrium. Blom et al. [5] showed that, consequently, equations and clusters that may be targeted by interventions should be clearly distinguished from the variables that could be affected by those interventions to preserve an unambiguous causal interpretation.

3.2.2 Identification of perfect adaptation via causal ordering

With the help of the dynamic causal ordering graph and the equilibrium causal ordering graph we can identify perfect adaptation without requiring simulations or explicit calculations. To do so, we require that Assumption 1 below holds.

Assumption 1. If there is a directed path from an input vertex to a variable vertex in the dynamic causal ordering graph of a set of first-order differential equations in canonical form, possibly with static equations as well, then there is a response of that variable to changes in the input signal some (small) time-step later.

We believe that, at least for a large class of dynamical systems, this assumption is satisfied for almost all parameter values w.r.t. the Lebesgue measure on a suitable parameter space (i.e. the property holds *generically*). It seems reasonable to assume that a change in one of the variables or parameters that appear on the right-hand side of a first-order differential equation in canonical form at time t results in a generic change in the value of the variable on the left-hand side of that differential equation at a time $t + \Delta t$. Consider a perfect matching M for the dynamical bipartite graph \mathcal{B}_{dyn} that extends the natural labelling. By construction, directed paths in $\mathcal{G}(\mathcal{B}_{\text{dyn}}, M)$, which coincide with directed paths in the dynamic causal ordering graph $\text{CO}(\mathcal{B}_{\text{dyn}})$, then correspond to transient causal effects (which may persist at equilibrium).

Under Assumption 1, a directed path from the input vertex to a variable vertex in the dynamical causal ordering graph implies a response to a change in the input signal. Lemma 1, which follows directly from Proposition 2 in Blom et al. [5], shows that at equilibrium, a change in the input signal has no effect on the value of a variable if there is no directed path from the input vertex to that variable in the equilibrium causal ordering graph. Theorem 1 then formulates sufficient graphical conditions for the identification of perfect adaptation.

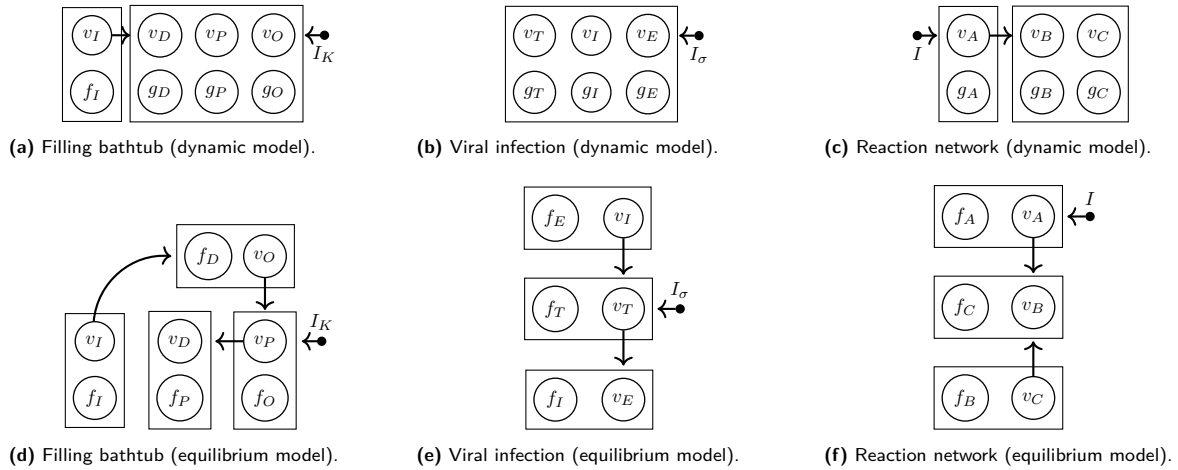


Figure 7. The dynamical causal ordering graphs for the bathtub model, the viral infection, and the reaction network with negative feedback are given in Figures 7a, 7b, and 7c, respectively. The equilibrium causal ordering graphs for the equilibrium equations of the bathtub model, the viral infection and the reaction network with negative feedback are given in Figures 7d, 7e, and 7f, respectively. The input vertices I_K , I_σ , and I are denoted by black dots. The absence or presence of a directed path from a cluster, equation vertex, or input vertex to a variable vertex implies that a causal effect is absent or generically present, respectively.

Lemma 1 (Blom et al. [5]). *Consider a model consisting of static equations, a set of first-order differential equations in canonical form, and an input signal. Assume that the equilibrium bipartite graph has a perfect matching and that the static equations and equilibrium equations derived from the first-order differential equations are uniquely solvable w.r.t. the equilibrium causal ordering graph for all relevant values of the input signal. If there is no directed path from an input vertex to a variable vertex in the equilibrium causal ordering graph then a change in the input signal has no effect on the equilibrium solution of that variable.*

These observations directly lead to our first main result. The apparent simplicity of Theorem 1 is due to it relying on appropriate powerful definitions and concepts such as causal ordering.

Theorem 1. *Consider a model that satisfies the conditions of Lemma 1 and assume that the associated dynamic causal ordering graph has a perfect matching that extends the natural labelling. Under Assumption 1, the presence of a direct path from the input signal I to a variable X_v in the dynamical causal ordering graph and the absence of such a path in the equilibrium causal ordering graph, implies that X_v perfectly adapts to changes in the input signal I .*

We see that there is a directed path from the input signal I_K to v_O in the dynamical causal ordering graph in Figure 7a, while no such path exists in the equilibrium causal ordering graph in Figure 7d. It follows from Theorem 1 that X_O perfectly adapts to changes in the input signal I_K . This is in agreement with the simulation in Figure 3a. Similarly, we can verify that the amount of infected cells X_I in the viral infection model perfectly adapts to changes in the input signal I_σ and that X_C perfectly adapts to I in the reaction network with negative feedback. Clearly, it is easy to verify that perfect adaptation in the bathtub model, the viral infection model, and the reaction network with negative feedback can be identified using the graphical criteria in Theorem 1 using Figure 7.

In Section 6 we construct graphical representations for a dynamical model of a basic enzymatic reaction that achieves perfect adaptation but does not satisfy the conditions in Theorem 1. In Appendix D we will show that the biochemical reaction network in Figure 4b, which Ma et al. [28] identified as being capable of achieving perfect adaptation, does not satisfy the conditions in Theorem 1 either. This shows that these conditions are not necessary for the identification of perfect adaptation in dynamical systems at equilibrium. The further development of methods to analyse perfectly adapted dynamical systems that do not satisfy




Figure 8. The Markov ordering graphs for the bathtub, the viral infection, and the reaction network with a negative feedback loop are given in Figures 8a, 8b, and 8c respectively. Exogenous variables are denoted by dashed circles and input vertices are denoted by black dots.

the conditions of Theorem 1 remains a challenge for future work. We believe that the methods presented in this section are a useful tool for the characterization of a large class of network topologies that are able to achieve perfect adaptation and for the automated analysis of the behaviour of certain perfectly adapted dynamical systems.

3.3 Recognizing perfect adaptation in data

So far we have only considered how perfect adaptation can be identified in mathematical models. In this section we focus on methods for model selection from data that is generated by perfectly adapted dynamical systems. We also discuss how the output of certain constraint-based causal discovery algorithms can be correctly interpreted for such systems.

3.3.1 Conditional independences

The Markov ordering graph can be used to derive conditional independences that are implied by a model at equilibrium and that can be tested in equilibrium data. The Markov ordering graphs for the equilibrium distribution of the dynamical models in the previous sections are constructed after including independent exogenous random variables to the equilibrium causal ordering graph. For the bathtub model, we let vertices $\{v_I, w_1, \dots, w_5\}$ represent independent exogenous random variables U_I, U_1, \dots, U_5 . For the viral infection model we let w_T, w_I, w_E, w_β represent independent exogenous random variables d_T, d_I, d_E , and β in equations (11), (12), and (13). Finally, for the reaction network with negative feedback, we let w_A, w_B , and w_C represent independent exogenous random variables that appear in the differential equations for $X_A(t)$, $X_B(t)$, and $X_C(t)$ respectively.

The Markov ordering graphs for the filling bathtub model, the viral infection model, and the model of a reaction network with a negative feedback loop are given in Figures 8a, 8b, and 8c respectively. Note that the Markov ordering graph for the bathtub model coincides with the result in Dash and Druzdzel [12], who simulated data from the bathtub model until the system reached equilibrium and then applied the PC algorithm to the equilibrium data. Although Dash [11] interprets the learned graphical representation as the ‘causal graph’, this graph does not have a straightforward causal interpretation, see Section 3.3.2 and the discussion in Blom et al. [5] for more details. Instead, the d-separations in these graphs imply conditional independences in the equilibrium distribution between the corresponding variables [5]. For example, since v_I is d-separated from v_D given v_P in the Markov ordering graph of the bathtub model at equilibrium, X_I will be independent of X_D given X_P . The implied conditional independences can for instance be used in the process of model selection. A demonstration of selecting immune responses for a viral infection model using the Markov ordering graph is given by Blom and Mooij [2].

3.3.2 Interpretation of the Markov ordering graph

In this section, we will demonstrate that the Markov ordering graphs in Figures 8a, 8b, and 8c do not have a straightforward causal interpretation in terms of interventions, contrary to what is sometimes claimed [11, 22]. To see this, we first explicitly state what we mean when we talk about ‘causal relations’. In contemporary literature, the common interpretation is that, in the context of a model, an intervention on the cause brings about a change in the effect.

So let us consider an intervention on a dynamical model of a filling bathtub at equilibrium that manipulates the time-derivative $\dot{X}_D(t)$, and consequently the associated equilibrium equation f_D (e.g. by changing one of the parameters that appear in that differential equation). Assuming that the system converges to equilibrium after the intervention, the equilibrium causal ordering graph in Figure 7d tells us that this intervention on f_D generically changes the equilibrium distributions of X_O , X_P , and X_D . Since f_D is not included in Figure 8a, it is not possible to read off the effect of this intervention from the Markov ordering graph of the equilibrium distribution. Clearly, if we would interpret a soft intervention on f_D as an intervention on v_D in the Markov ordering graph, then we would wrongly conclude that the intervention has no effect on X_O and X_P , if we were to interpret the Markov ordering graph causally. Similarly, the equilibrium causal ordering graph in Figure 7d tells us that an intervention targeting f_P only affects X_D , whereas the Markov ordering graph in Figure 8a would incorrectly suggest that an intervention targeting v_P affects both X_P and X_D , if we were to interpret it causally. We conclude that the directed edges in the Markov ordering graph do not represent causal relations in terms of soft interventions.

Analogously, we find that the Markov ordering graph cannot be interpreted in terms of perfect (“surgical”) interventions either. The correct interpretation of a directed edge $(v_i \rightarrow v_j)$ in the Markov ordering graph for the equilibrium distribution of a set of first-order differential equations is that an intervention targeting equations in the cluster of v_i has an effect on the equilibrium distribution of v_j . In many systems, equilibrium equations f_i derived from differential equations for variables $X_i(t)$ end up in the same cluster as the associated variable v_i . In that case, the Markov ordering graph has an unambiguous causal interpretation. In Blom and Mooij [2] it is shown that the Markov ordering graph for the equilibrium distribution of dynamical models in which each variable is self-regulating does have this straightforward causal interpretation. However, a large class of exceptions to this is provided by perfectly adaptive systems.

3.3.3 Detecting perfect adaptation

The most straightforward approach to detect perfect adaptation is to collect time-series data while experimentally changing the input signal to the system. One can then simply observe whether the variables in the system revert to their original values. Unfortunately, this type of data is not always available. Another way to identify feedback loops that achieve perfect adaptation uses a combination of observational equilibrium data, background knowledge, and experimental data. Our second main result, Theorem 2, gives sufficient conditions under which we can identify a system that is capable of perfect adaptation from experimental equilibrium data.

Theorem 2. *Consider a set of first-order dynamical equations in canonical form, satisfying the conditions of Theorem 1, for variables V that has equilibrium equations F with the natural labelling and consider a soft intervention targeting an equation $f_i \in F$. Assume that the system is uniquely solvable w.r.t. the equilibrium causal ordering graph both before and after the intervention and that the intervention alters the equilibrium distribution of all descendants of f_i in the causal ordering graph. If either*

1. *the soft intervention does not change the equilibrium distribution of X_i , or*
2. *the soft intervention alters the equilibrium distribution of a variable corresponding to a non-descendant of v_i in the Markov ordering graph,*

or both, then the system is capable of perfect adaptation.

Proof. If condition 1 holds there is no directed path in the causal ordering graph from f_i to v_i in the equilibrium causal ordering graph, by the assumption that the soft intervention on f_i changes the equilibrium distribution of all its descendants. By definition of the dynamical bipartite graph there is a directed path from g_i to v_i in the dynamical causal ordering graph, because g_i and v_i end up in the same cluster (note that this follows by using the natural labelling as perfect matching and the result that the causal ordering graph does not depend on the chosen perfect matching [5]). It follows from Theorem 1 that X_i perfectly adapts to an input signal I_{f_i} on f_i (i.e. a soft intervention targeting $\dot{X}_i(t)$ and thus the equilibrium equation f_i).

Suppose that 1 does not hold while 2 does hold. By Theorem 4 in [5] (which roughly states that the presence of a causal effect at equilibrium implies the presence of a corresponding directed path in the equilibrium causal ordering graph) we have that f_i is an ancestor of v_i and some v_h in the equilibrium causal ordering graph, while v_i is not an ancestor of v_h in the Markov ordering graph. For a perfect matching M of the equilibrium bipartite graph let $v_j = M(f_i)$. Then v_j is in the same cluster as f_i in the equilibrium causal ordering graph by construction. Note that $j = i$ would give a contradiction, as then v_i would be an ancestor of v_h in the Markov ordering graph. Suppose that the vertex f_j , that is associated with v_j through the natural labelling, is matched to a non-ancestor of v_j in the equilibrium causal ordering graph. Because of the edge $(g_j - v_j)$ in the dynamical bipartite graph, it follows from Theorem 1 that X_j perfectly adapts to an input signal I_{f_j} on f_j . Therefore the system is able to achieve perfect adaptation. Now suppose that f_j is matched to an ancestor v_k of v_j , and consider the vertex f_k . The previous argument can be repeated to show perfect adaptation for X_k is present when f_k is matched to a non-ancestor of v_k in the equilibrium causal ordering graph. Otherwise, f_k must be matched to an ancestor of v_k . Note that the ancestors of v_k are a subset of the ancestors of v_j , which in turn are a subset of the ancestors of v_i . In a finite system of equations, v_i has a finite set of ancestors and therefore we eventually find, by repeating our argument, a vertex f_m that cannot be matched to an ancestor of v_m because v_m has no ancestors that are not matched to one of the vertices f_i, f_j, f_k, \dots that were considered up to that point. Because f_m is matched to a non-ancestor we then find that X_m perfectly adapts to an input signal on I_{f_m} as before. \square

Based on the result in Theorem 2 we can devise the following scheme to detect perfectly adapted dynamical systems from data and background knowledge. We start by collecting observational equilibrium data and use the PC or LCD algorithm to learn a (partial) representation of the Markov ordering graph, assuming the observational distribution to be faithful w.r.t. the Markov ordering graph. We then consider a soft intervention that changes a known equation in the first-order differential equation model (i.e. it targets a known equilibrium equation). If this intervention does not change the distribution of the variable corresponding to this target using the natural labelling, or if it changes the distribution of identifiable non-descendants of the variable corresponding to the target according to the learned Markov equivalence class, we can apply Theorem 2 to identify the perfectly adapted dynamical system. Note that this procedure relies on several assumptions, including faithfulness.

4 Application to a protein signalling model

In cell biology, dynamical systems for protein signalling networks are used to model processes where information is transmitted between and into cells. The underlying dynamics of such models may have unexpected consequences for causal discovery efforts using structure learning methods, see also [45]. Here, we specifically consider the phenomenon of perfect adaptation in a simple model of a well-studied molecular pathway. Using the technique of causal ordering to analyse the conditional independences and causal relations that are implied by the model, we elucidate the causal interpretation of the output of constraint-based causal discovery algorithms like LCD when they are applied to protein expression data.

We do not claim that the model that we analyse here is a realistic model of the protein signalling pathway. Although we will show that the model is able to explain certain observations in real-world data, this is not that surprising for a model with that many parameters.⁵ Instead, our goal is to demonstrate that in systems with perfect adaptation our standard intuitions regarding the output of causal discovery algorithms might fail. Furthermore, we explain the discrepancies between the graphical representations that are produced by causal ordering for equilibrium equations and causal discovery from equilibrium data. In combination, these two techniques help us to better understand causal properties of dynamical systems at equilibrium.

4.1 Dynamical model

We consider the mathematical model for the Ras-Raf-Mek-Erk signalling cascade in Shin et al. [46]. Let $V = \{v_s, v_r, v_m, v_e\}$ be an index set for endogenous variables that represent the equilibrium concentrations X_s , X_r , X_m , and X_e of active Ras, Raf, Mek, and Erk proteins respectively. The dynamics are given by:

$$\dot{X}_s(t) = \frac{I(t)k_{Is}(T_s - X_s(t))}{(K_{Is} + (T_s - X_s(t))) \left(1 + \left(\frac{X_e(t)}{K_e}\right)^{\frac{3}{2}}\right)} - F_s k_{F_s s} \frac{X_s(t)}{K_{F_s s} + X_s(t)} \quad (21)$$

$$\dot{X}_r(t) = \frac{X_s(t)k_{sr}(T_r - X_r(t))}{K_{sr} + (T_r - X_r(t))} - F_r k_{F_r r} \frac{X_r(t)}{K_{F_r r} + X_r(t)} \quad (22)$$

$$\dot{X}_m(t) = \frac{X_r(t)k_{rm}(T_m - X_m(t))}{K_{rm} + (T_m - X_m(t))} - F_m k_{F_m m} \frac{X_m(t)}{K_{F_m m} + X_m(t)} \quad (23)$$

$$\dot{X}_e(t) = \frac{X_m(t)k_{me}(T_e - X_e(t))}{K_{me} + (T_e - X_e(t))} - F_e k_{F_e e} \frac{X_e(t)}{K_{F_e e} + X_e(t)}, \quad (24)$$

where we assume that $I(t)$ is an external stimulus or perturbation.⁶ Roughly speaking, there is a signalling pathway that goes from $I(t)$ to $X_s(t)$ to $X_r(t)$ to $X_m(t)$ to $X_e(t)$ with negative feedback from $X_e(t)$ on $X_s(t)$. As we did for the reaction network with negative feedback in Section 3, we consider the system under saturation conditions. For $(T_e - X_e(t)) \gg K_{me}$ and $X_e(t) \gg K_{F_e e}$ the following approximation holds:

$$\dot{X}_e(t) \approx X_m(t)k_{me} - F_e k_{F_e e}. \quad (25)$$

We let f_s , f_r , f_m , and f_e represent the equilibrium equations corresponding to the dynamical equations in (21), (22), (23), and (24) respectively, where we assume the input signal to have a constant value I . We simulated the model under saturation conditions until it reached equilibrium, and then we recorded the changes in the concentrations $X_s(t)$, $X_r(t)$, and $X_m(t)$ after a change in the input signal I . The results in Figure 3 show that Ras, Raf, and Mek revert to their original values after an initial response. Clearly the equilibrium concentrations X_s , X_r , and X_m perfectly adapt to the input signal I . The details of this simulation can be found in Appendix A. In the next section we will show that the concentration of active Erk does not perfectly adapt to changes in the input signal.

⁵ As mathematician John von Neumann once put it: “With four parameters I can fit an elephant, and with five I can make him wiggle his trunk”.

⁶ For simplicity, we slightly adapted the model so that the feedback mechanism through Raf Kinase Inhibitor Protein (RKIP) is not included. In the differential equation for activated Mek we therefore discarded the dependence on RKIP. The goal here is not to give the most realistic model but to elucidate the phenomenon of perfect interpretation and the causal interpretation of the Markov ordering graph for perfectly adapted dynamical systems.




Figure 9. Perfect adaptation in the Ras-Raf-Mek-Erk signalling pathway. After an initial response to a change of input signal the equilibrium concentrations of active Ras, Raf, and Mek revert to their original values. The concentration of active Erk does not adapt to changes in the input signal. The details of the simulation can be found in Appendix A.

4.2 Graphical representations

We consider graphical representations of the protein signalling pathway. A compact representation of the structure of differential equations (21), (22), (23), and (24) is given in Figure 10a. Using the natural labelling, we construct the dynamical bipartite graph in Figure 10b. The associated dynamical causal ordering graph, with the input signal I included, is given in Figure 10c.

Under saturation conditions, the equilibrium equations f_s , f_r , f_m , and f_e obtained by setting equations (21), (22), (23), and (25) to zero have the bipartite structure in Figure 10d. Note that there is no edge ($f_e - v_e$) in the equilibrium bipartite graph because $X_E(t)$ does not appear in the approximation (25) of (24). The associated equilibrium causal ordering graph is given in Figure 10e, where the cluster $\{I\}$ is added with an edge towards the cluster $\{v_e, f_s\}$ because I appears in equation (21) and in no other equations. So far we have treated all symbols in equations (21), (22), (23), and (24) as deterministic parameters. Let w_s , w_r , w_m , and w_e represent independent exogenous random variables appearing in the equilibrium equations f_s , f_r , f_m , and f_e respectively. After adding them to the causal ordering graph with edges to their respective clusters we construct the Markov ordering graph for the equilibrium distribution in Figure 10f.

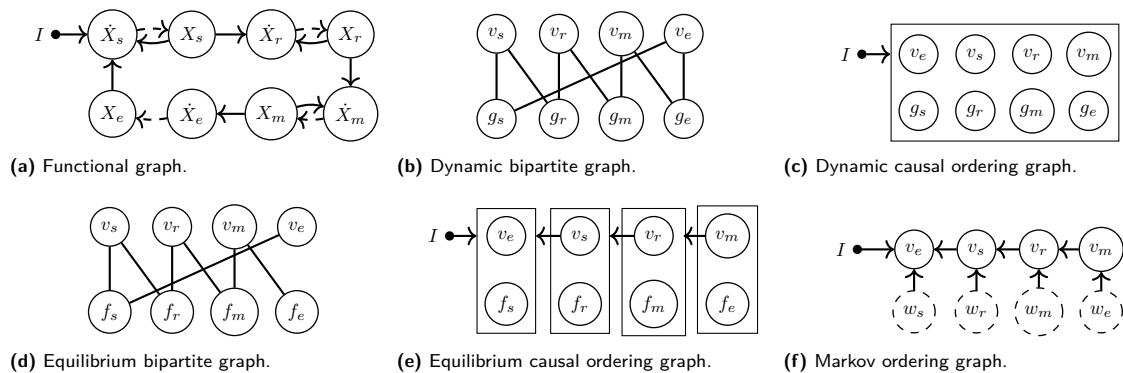


Figure 10. Six graphs associated with the protein signalling pathway model under saturation conditions where indices s, r, m, e correspond to concentrations of active Ras, Raf, Mek, and Erk respectively. The functional graph, dynamic bipartite graph, and equilibrium bipartite graph are compact representations of the model. The dynamic causal ordering graph encodes the presence of transient (generic) causal effects. The equilibrium causal ordering graph represents the effects of manipulations to the equilibrium equations of the model. The Markov ordering graph implies conditional independences in the equilibrium distribution of the variables in the model via d-separations.

4.3 Model predictions and causal discovery

We discuss some of the predictions that can be read off from the equilibrium causal ordering graph and the Markov ordering graph for the equilibrium distribution of the model. In Section 5 we will test these predictions in simulated equilibrium data and real-world protein expression data.

4.3.1 Conditional independences and correlations

The d-separations in the Markov ordering graph imply conditional independences between the corresponding variables [5]. From the graph in Figure 10f we read off the following (implied) conditional independences:

$$I \perp\!\!\!\perp v_s, \quad I \perp\!\!\!\perp v_r, \quad I \perp\!\!\!\perp v_m, \quad v_e \perp\!\!\!\perp v_r \mid v_s, \quad v_e \perp\!\!\!\perp v_m \mid v_s, \quad v_s \perp\!\!\!\perp v_m \mid v_r.$$

A more extensive overview of d-separations and predicted conditional independences can be found in Appendix B. Under the faithfulness assumption, the vertices that are not d-separated in the Markov ordering graph are dependent in the equilibrium distribution. The noise that is introduced into the model by exogenous random variables and the model parameters affect the strength of these dependences. The Markov ordering graph in Figure 10f suggests that the correlation between Mek (i.e. X_m) and Raf (i.e. X_r) should be stronger than the correlation between Mek and Erk (i.e. X_e) because extra noise is introduced along the longer pathway Mek-Raf-Ras-Erk.

4.3.2 Inhibition of MEK activity

A common biological experiment that is used to study protein signalling pathways is the use of an inhibitor that decreases the activity of a protein on the pathway. Such an inhibitor slows down the rate at which the active protein is able to activate another protein. Here, we consider inhibition of Mek activity. Therefore, an experiment where the activity of Mek is inhibited has an effect on parameters in the differential equations in which $X_m(t)$ appears. Since \dot{X}_e is the only child of X_m in the functional graph in Figure 10a, we can interpret this experiment as a soft intervention on g_e in the dynamic model and on f_e in the equilibrium causal ordering graph, where the rate k_{me} at which Erk is activated is decreased. Since there is a directed path from f_e to v_m, v_r, v_s , and v_e in the causal ordering graph in Figure 10e, we expect that a change in an input signal I_e on f_e (e.g. a change in the parameter k_{me}) affects the equilibrium concentrations of active Mek, Raf, Ras, and Erk respectively. Note that Ras, Raf, and Mek are ancestors of Erk in the Markov ordering graph in Figure 10f, so that under the assumptions in Theorem 2 we can use this experiment to detect perfect adaptation in the protein pathway.

4.3.3 Causal discovery

Suppose that we have observational equilibrium data from the protein signalling pathway model and also experimental equilibrium data from a setting where Mek activity is inhibited. The context variable C that indicates from which setting the data was collected (observation or experimental), is not caused by any observed variable, since this variable is set externally by the experimenter at the start of the experiment. This set-up satisfies the conditions of a context variable in the LCD algorithm. In the case of Mek inhibition, this context variable represents a soft intervention on the equation f_e in the causal ordering graph in Figure 10e. The Markov ordering graph that includes the context variable C (but not the independent exogenous random variables) is given in Figure 11. To construct this graph, the context variable C is first added to the equilibrium causal ordering graph in Figure 10e as a singleton cluster with an edge towards the cluster $\{v_m, f_e\}$. The Markov ordering graph is then constructed from the resulting directed cluster graph in the usual way. From this, we can read off (conditional) independences to find the LCD triples that

are implied by the equilibrium equations of the model. We find that (C, v_m, v_r) , (C, v_m, v_s) , (C, v_m, v_e) , (C, v_r, v_s) , (C, v_r, v_e) , and (C, v_s, v_e) are all LCD triples.




Figure 11. Markov ordering graph of the protein signalling pathway with the context variable C included. This context variable indicates whether a cell was treated with a Mek inhibitor or not.

With a conditional independence oracle, and under the faithfulness assumption, the output of complete causal discovery algorithms (like the PC algorithm if causal sufficiency is assumed, or more generally, the FCI algorithm) would be the Markov equivalence class of the Markov ordering graph. Here, it is important to note that the Markov ordering graph does not have a straightforward *causal* interpretation for this perfectly adapted dynamical system. The reasoning is similar to the discussion in Section 3.3.2.

For the protein signalling model, the common biological understanding of the underlying causal mechanism is that Raf activates Mek, Mek activates Erk, and that it is very likely that there is negative feedback from a protein downstream of Erk on Raf [19]. Therefore, even though Raf is a direct cause of Mek (see Figure 10c), in line with the biological consensus, Mek is also an indirect cause of Raf. At equilibrium, Raf is no longer a cause of Mek due to the perfect adaptation. This leads to a situation where there is a directed path from Raf to Mek in biological consensus networks (like the one in Sachs et al. [44] where the feedback loop from Erk to Raf has not been included) while there is a directed path in the *opposite* direction in the Markov ordering graph for the equilibrium equations. If we were to apply the LCD algorithm to the experimental Mek inhibition equilibrium data, we would detect a directed path from Mek to Raf but not from Raf to Mek. The conclusion that ‘Mek is a cause of Raf’ while no causal relation from Raf on Mek can be detected could, at first glance, appear to be at odds with expert knowledge. Similar observations of an apparent “causal reversal” in protein interaction networks have been observed more often, see also Mooij and Heskes [31], Mooij et al. [34], Triantafyllou et al. [52]. The phenomenon of perfect adaptation can help to explain differences between biological consensus networks and the output of causal discovery algorithms. We have shown that a simple model that is capable of perfect adaptation can explain some of the differences between the output of standard constraint-based causal discovery algorithms and biological consensus networks that represent other aspects of the underlying mechanisms. Confusion about causal relations can be avoided by explicitly specifying the interventions that correspond to the causal effects, by distinguishing between statements about the equilibrium distribution and the dynamical model, and analysing models with our approach based on the technique of causal ordering.

In Blom and Mooij [2] it was shown that the causal relations and conditional independences that are implied by the equilibrium equations of a dynamical model may not be preserved when it is combined with another model. They discuss how, for dynamical systems at equilibrium that are only partially modelled and observed, one can reason about the presence of unobserved feedback loops and variables that are not self-regulating in the whole system. In Appendix C, we show that these ideas can also be applied when only $X_s(t)$, $X_r(t)$, and $X_m(t)$ are included as endogenous variables in the perfectly adapted protein signalling model that we presented in this section.

5 Experiments

In this section we present simulations to confirm the qualitative model predictions for the protein signalling model in Section 4. We then consider data from real-world experiments in order to test the validity of the protein signalling model.


5.1 Simulations

We took as input signal $I(t) = i$, with i sampled from a uniform distribution on the interval $(0.5, 1.5)$. We also drew random samples for the parameters k_{Is} , k_{sr} , k_{rm} , and k_{me} from uniform distributions on the intervals $(1.2, 1.5)$, $(2.4, 3.0)$, $(1.7, 2.0)$, and $(0.7, 1.0)$ respectively. We then simulated the dynamical model in equations (21) to (24) with parameter settings: $K_{Is} = 1.0$, $K_e = 1.5$, $F_s = 1.0$, $k_{F_s s} = 1.0$, $K_{F_s s} = 0.9$, $K_{sr} = 1.0$, $F_r = 0.3$, $k_{F_r r} = 1.0$, $K_{F_r r} = 0.8$, $K_{rm} = 0.9$, $F_m = 0.2$, $k_{F_m m} = 1.0$, $K_{F_m m} = 1.2$, $K_{me} = 0.0001$, $F_e = 0.7$, $k_{F_e e} = 1.2$, and $K_{F_e e} = 0.0001$. The parameters were chosen in such a way that the approximation in equation (25) of equation (24) is valid and so that the system converges to an equilibrium where the concentrations of active proteins are strictly between 0 and $T_s = 1.0$, $T_r = 1.0$, $T_m = 1.0$, and $T_e = 5.0$ respectively. We experimented with other parameter values as well, and observed that the analysis of the qualitative behaviour of the model that we present here is valid for many values of the parameters.

5.1.1 Conditional independences and correlation strength

To test whether the conditional independences in Section 4.3.1 hold when the system is at equilibrium, we ran the simulation $n = 500$ times until it reached equilibrium and recorded the equilibrium concentrations X_s , X_r , X_m , and X_e . We tested all (conditional) independences with a maximum of one conditioning variable using Spearman's rank correlation test with a p-value threshold of 0.01. This way, we retrieved all predicted (conditional) independences and all predicted (conditional) dependences. Table 1 in Appendix B provides a list of the estimated correlations and the corresponding p-values.⁷

In Section 4.3 we discussed how the Markov ordering graph for the simple model of a protein signalling pathway suggests that the correlation between Mek and Raf should be stronger than the correlation between Mek and Erk. The scatter plots in Figure 12 below confirm this prediction.



(a) Strong correlation between active Raf and Mek.

(b) Weaker correlation between active Mek and Erk.

Figure 12. Two scatter plots of the Mek-Raf and Mek-Erk concentrations of 100 samples of the simulation experiment of the protein signalling pathway in Section 5.1.1. Note the difference in the signal to noise ratio. The correlation between Mek and Raf is clearly stronger than the correlation between Mek and Erk. The estimate of the rank correlation between Mek and Raf is 0.98 and between Mek and Erk it is -0.51 .

Different parameter regimes correspond to different qualitative behaviour of the protein pathway model. For example, when almost all of the Erk molecules are activated we have that $X_e \approx T_e$. If we repeat the experiment in Section 5.1.1 with $K_e = 100$ and with K_{Is} drawn from a uniform distribution on the interval $(1.9, 2.5)$ then we find that the correlation between X_m and X_e is 0.054 with a p-value of 0.087. The

⁷ Because the LCD algorithm only uses conditional independence test with a maximum of one variable in the conditioning test, we do not consider conditional independence tests with larger conditioning sets in this work. We did experiment with larger conditioning sets but we were not able to retrieve all predicted conditional dependences with our parameter settings and only $n = 500$ samples.

correlation between X_m and X_r is 0.76 with a p-value smaller than $2.2e^{-16}$. The dependence between the concentrations of active Mek and Erk thus disappears under saturation conditions for Erk, while the correlation between Mek and Raf remains strong.

5.1.2 Inhibition of MEK activity and LCD

We assessed the effect of decreasing the activity of Mek on the equilibrium concentrations of Ras, Raf, Mek, and Erk. To that end, we simulated the model with fixed parameters $I = 1.0$, $k_{Is} = 1.0$, $k_{sr} = 1.0$, $k_{rm} = 1.0$, and $k_{me} = 1.1$ until it reached equilibrium. We then decreased the parameter that controls the activity of Mek to $k_{me} = 1.0$. The recorded responses of the concentrations of active Ras, Raf, Mek, and Erk are displayed in Figure 13. From this we confirm our prediction that inhibition of Mek activity affects the equilibrium concentrations of Ras, Raf, Mek, and Erk.




Figure 13. Simulation of the response of the concentrations of active Ras, Raf, Mek, and Erk after inhibition of the activity of Mek. The system starts out in equilibrium with $k_{me} = 1.1$. The concentrations of Ras, Raf, Mek, and Erk are recorded after the parameter controlling Mek activity is decreased to $k_{me} = 1.0$ from $t = 0$ on.

We also simulated a scenario where the inhibition of Mek activity is treated as a context variable, that can be used to apply the LCD algorithm. We ran the simulation $n = 500$ times with $k_{me} = C$, where the context variable C is drawn from a uniform distribution on the interval $(0.98, 1.1)$. To avoid deterministic relations, we drew the parameter $k_{F_{ee}}$ from a uniform distribution on $(0.7, 1.0)$. We ran the simulations until the system reached equilibrium and recorded the equilibrium values of the variables. We then applied the LCD algorithm to search for LCD triples in this equilibrium data with context variable C . For the conditional independence tests we used Spearman's rank correlation with a p-value threshold of 0.01. We found the expected LCD triples (C, v_m, v_r) , (C, v_m, v_s) , (C, v_m, v_e) , (C, v_r, v_s) , (C, v_r, v_e) , (C, v_s, v_e) and no others.

5.2 Protein expression data

In this section we test the predictions of our model on protein signalling data from real-world experiments. For a thorough description of these experiments we refer to Sachs et al. [44] and Lun et al. [27].


5.2.1 Correlation strength

In the simulations of the simple protein signalling pathway model we demonstrated that, as predicted, the correlation between Raf and Mek was much stronger than the correlation between Mek and Erk, and the latter correlation completely disappeared in a setting where Erk was saturated. We test these correlations

in a multivariate single-cell protein expression dataset that was used in Sachs et al. [44]. We considered data that was pooled from different experimental settings, in which cells were exposed to stimulatory and/or inhibitory interventions.⁸ Using Spearman's rank correlation we found a correlation of 0.78 with a p-value smaller than $2.2e^{-16}$ between Raf and Mek. The correlation between Mek and Erk was -0.023 with a p-value of 0.065. The biological consensus according to Sachs et al. [44] is that there is a signalling pathway from Raf to Mek to Erk. Note that the simple model in Section 4.1 provides an explanation as to why we are not able to reject the null hypothesis of zero correlation between Mek and Erk.

5.2.2 Inhibition of MEK activity

The experimental protein expression data used in Sachs et al. [44] includes data from an experiment where cells were perturbed with U0126, which is a known inhibitor of Mek activity. Figures 14a and 14b show the log-transformed concentrations of active Raf, Mek, and Erk proteins after treatment with and without the U0126 perturbation. In both cases the sample was treated with anti-CD3 and anti-CD28, see Sachs et al. [44] for more details on the dataset. These plots clearly show that inhibition of Mek activity results in an increase in the concentrations of active Raf and active Mek and a reduction in the concentration of active Erk. This is in agreement with observations in the simulation study.



(a) Scatter plot for concentrations of active Mek and Raf.

(b) Scatter plot for concentrations of active Mek and Erk.

Figure 14. Two scatter plots of the active Mek, Raf, and Erk concentrations in the Sachs data. The red circles correspond to the sample treated with anti-CD3/CD28 and the Mek-activity inhibitor U0126. The blue circles correspond to the sample treated only with anti-CD3/CD28. The inhibition of Mek results in an increase in the concentration of active Mek and Raf, whereas the concentration of active Erk is reduced.

5.2.3 Detecting LCD Triples

The perturbations with different inhibitors and stimulants in the data that was used in Sachs et al. [44] can be treated as context variables [34]. For the context variable associated with a specific perturbation (e.g. AKT inhibitor) data points collected from the condition with that perturbation were indicated with a 1, while data points collected from other experimental conditions were indicated with a 0. We then searched for LCD triples involving Raf, Mek, and Erk. We detected the following LCD triples using Spearman's rank correlation test with a p-value threshold of 0.01: (AKT inhibitor, Raf, Mek), (LY294002, Raf, Mek), (Psitectorigenin, Raf, Mek), (AKT inhibitor, Raf, Erk), and (β 2cAMP, Raf, Erk). This suggests that, if the system was at equilibrium under saturation conditions, there should be directed paths from Raf to Mek in the Markov ordering graph. Although this observation agrees with the biological consensus networks

⁸ In particular, we used specific perturbation conditions with the following reagents: β 2cAMP, AKT inhibitor, U0126, PMA, G06976, Psitectorigenin, LY294002. In some conditions the general perturbation by the reagent anti-CD3/CD28 also was included. See Sachs et al. [44] for more details.

that include negative feedback from Erk on Raf, there is no directed path from Raf to Mek in the Markov ordering graph for the saturated equilibrium model in Figure 10f. However, the results of LCD seem to strongly depend on the implementation details. For example, both Mooij et al. [34] and Boeken and Mooij [6] report LCD triples in this dataset that imply a directed path from Mek to Raf, as was predicted by the saturated equilibrium model. Furthermore, the assumption that the system was saturated and at equilibrium may have been violated. We also found LCD triples that imply a directed path from Raf to Erk in the Markov ordering graph, if the system was at equilibrium under saturation conditions. Such LCD triples were also reported by Boeken and Mooij [6], but not by Mooij et al. [34]. The detected LCD triples agree with the direction of edges in the Markov ordering graph for the saturated equilibrium model in Figure 10f.

We also searched for LCD triples involving Mek and Erk in the protein signalling data in Lun et al. [27]. This data was collected at different time-points after the abundance of certain proteins was over-expressed in an experiment. We treated the measured expression levels of targeted proteins as context variables in the LCD algorithm, as Blom et al. [3] do. We followed the pre-processing steps in Blom et al. [3], and selected a subset of the perturbations for our analysis. There were three replica's of each experiment and we searched for LCD triples that consistently appeared in all replicas, using Spearman's rank correlation test with a p-value threshold of 0.01. This way, we found the triple (p70RSK, Mek, Erk) at $t = 5$ from the data of experiments where p70RSK was over-expressed and the triple (p38, Erk, Mek) at $t = 60$ from the data of experiments where p38 was over-expressed. The LCD triple (p38, Erk, Mek) suggests that, under saturation conditions and at equilibrium, there is a directed path from Erk to Mek in the Markov ordering graph. The fact that this does not agree with the Markov ordering graph in Figure 10f could be due to a violation of the assumptions of saturation or equilibrium. The LCD triple (p70RSK, Mek, Erk) suggests that there should be a directed path from Mek to Erk in the Markov ordering graph. This is in agreement with the predictions of the protein pathway model at equilibrium and under saturation conditions.

In conclusion, LCD results on real-world data depend on implementation details. In some cases, they agree with the Markov ordering graph in Figure 10f, in other cases they don't.

6 Discussion

In this section, we will discuss that the notions of the causal Markov and faithfulness conditions, which are used to tie causal relations to conditional independences in the setting of causal DAGs, are ambiguous in the context of perfectly adapted systems. We also give an example of a dynamical system for which rewriting of the equilibrium equations reveals a stronger Markov property. We believe these are interesting topics for future work, because understanding the conditions under which the output of constraint-based causal discovery algorithms has a straightforward and intuitive causal interpretation may increase the impact of causal discovery in application domains where perfectly adapted systems frequently occur, if the observed lack of robustness of these methods can be overcome.

6.1 Ambiguity of causal Markov and faithfulness conditions

The causal faithfulness condition and causal Markov condition can be used to relate graphs that represent causal relations between variables to properties of the probability distribution on the space of these variables. In this work, we explicitly differentiate between causal relations in a dynamical model and in the equilibrium model. Furthermore, we also make a clear distinction between the Markov ordering graph (representing conditional independences) and the causal ordering graph (which encodes causal relations). In this context, the commonly used notions of causal faithfulness and the causal Markov condition become ambiguous.

To see this, consider the dynamical causal ordering graph for the viral infection model in Figure 7b. Note that v_T , v_I , v_E share a cluster and that I_σ is a cause of $X_T(t)$, $X_I(t)$, and $X_E(t)$. At the same time, the Markov ordering graph for the equilibrium equations implies that X_I is independent of I_σ , see Figure 7e. If we put a probability distribution on I_σ , we could say that the equilibrium distribution of the variables in the viral infection model is not *faithful* to cause-effect relations implied by the dynamic causal ordering graph.

Additionally, in the equilibrium causal ordering graph for the reaction network with a feedback loop in Figure 7f we see that the only direct cause, in terms of interventions on the equilibrium model, of X_A is I and that X_A is not a cause of X_C . However, the dynamical causal ordering graph in Figure 7c indicates that we cannot expect $X_A(t)$ to be independent of $X_C(t)$ given I , when the system has not yet reached equilibrium. Roughly speaking, the distribution of a system that is initialized with certain initial conditions and that has not yet reached equilibrium at time t is not Markov with respect to the cause-effect relations in the equilibrium model. We consider a study into more generally applicable formulations of these concepts that could be used also for perfectly adaptive systems to be outside the scope of the current paper.

6.2 Rewriting equations may reveal additional structure

Theorem 1 specifies sufficient but not necessary conditions for the presence of perfect adaptation. The equilibrium distribution of some systems is not faithful to the Markov ordering graph associated with the equilibrium equations in the model. Here, we will discuss the dynamical model for a basic enzymatic reaction and we will demonstrate that this model is capable of perfect adaptation, does not satisfy the conditions in Theorem 1, and that the presence of directed paths in the equilibrium causal ordering graph does not imply the presence of a causal effect at equilibrium. The basic enzyme reaction models a substrate S that reacts with an enzyme E to form a complex C , which is converted into a product P and the enzyme E . The dynamical equations for the concentrations $X_S(t)$, $X_E(t)$, $X_C(t)$, and $X_P(t)$ are given by:

$$\dot{X}_S(t) = k_0 - k_1 X_S(t) X_E(t) + k_{-1} X_C(t), \quad (26)$$

$$\dot{X}_C(t) = k_1 X_S(t) X_E(t) - (k_{-1} + k_2) X_C(t), \quad (27)$$

$$\dot{X}_E(t) = -k_1 X_S(t) X_E(t) + (k_{-1} + k_2) X_C(t), \quad (28)$$

$$\dot{X}_P(t) = k_2 X_C(t) - k_3 X_P(t), \quad (29)$$

where $k_{-1}, k_0, k_1, k_2, k_3$ and the initial conditions are independent exogenous random variables S_0 , C_0 , E_0 , and P_0 taking value in $\mathbb{R}_{>0}$ [4, 35]. We included the parameter k_1 into the functional graph of this system in Figure 15a. Since there is a path from k_1 to $X_P(t)$ we would expect that a change in k_1 would generically lead to a transient response of $X_P(t)$. We verified this by simulating this model with $k_{-1} = 1.0$, $k_0 = 1.0$, $k_1 = 1.0$, $k_2 = 0.8$, $k_3 = 2.5$ and with initial conditions $X_S(0) = 1.0$, $X_E(0) = 0.5$, $X_C(0) = 0.5$, and $X_P(0) = 1.0$ until the system reached equilibrium. We then recorded the response after changing the input signal k_1 . Figure 15b shows that X_P perfectly adapts to changes in the input signal k_1 .

The equilibrium equations of the model are given by:

$$f_S : \quad k_0 - k_1 X_S X_E + k_{-1} X_C = 0, \quad (30)$$

$$f_C : \quad k_1 X_S X_E - (k_{-1} + k_2) X_C = 0, \quad (31)$$

$$f_E : \quad -k_1 X_S X_E + (k_{-1} + k_2) X_C = 0, \quad (32)$$

$$f_P : \quad k_2 X_C - k_3 X_P = 0, \quad (33)$$

$$f_{CE} : \quad C + E - (C_0 + E_0) = 0, \quad (34)$$

where the last equation is derived from the constant of motion $C(t) + E(t)$, see Blom et al. [4] for more details. Via the extended causal ordering algorithm [5] the equilibrium causal ordering graph in Figure 16 can be constructed from the equilibrium equations in the model. There is a directed path from k_1 to v_P in the Markov ordering graph. Therefore, even though the basic enzyme reaction does achieve perfect

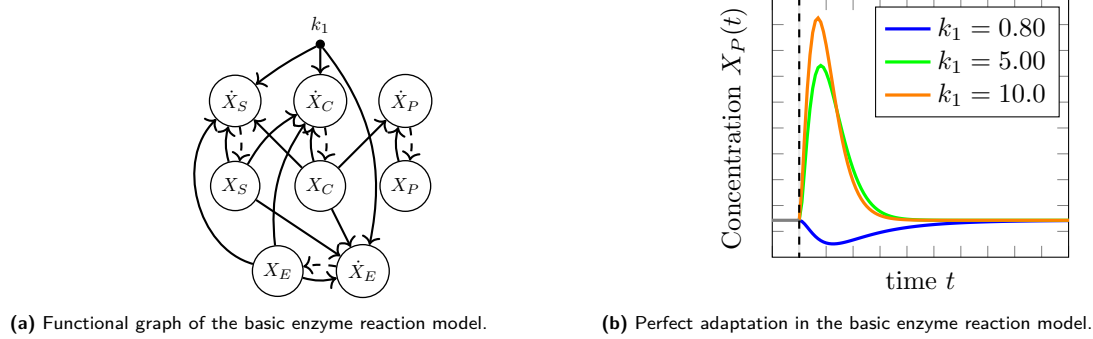


Figure 15. The functional graph of the basic enzyme reaction modelled by equations (26), (27), (28), and (29) in Figure 15a shows that there is a directed path from an input signal that controls the parameter k_1 to all endogenous variables X_S, X_C, X_E, X_P . Figure 15b shows that the concentration X_P perfectly adapts after an initial transient response to a persistent change in the parameter k_1 .




Figure 16. The equilibrium causal ordering graph constructed from the equilibrium equations of basic enzyme reaction modelled by equilibrium equations f_S, f_C, f_E, f_P , and f_{CE} .

adaptation, we see that it does not satisfy the conditions of Theorem 1. The simulation in Figure 15b indicates that there is no causal effect of k_1 on X_P at equilibrium. The basic enzyme reaction is an example of a system for which directed paths in the equilibrium causal ordering graph do not imply generic causal relations between variables.

By combining equilibrium equations we can achieve stronger conclusions for this particular case. For instance, we could consider the equation f'_C , obtained from summing equations f_S and f_C :

$$f'_C : \quad k_0 - k_2 X_C = 0, \quad (35)$$

in combination with f_S, f_P , and f_{CE} . The equilibrium equations f_C and f_E can be dropped because they are linear combinations of the other equations. The equilibrium bipartite graph and equilibrium causal ordering graph associated with f_S, f_{CE}, f'_C , and f_P are given in Figure 17. The equilibrium causal ordering graph in Figure 17b for the rewritten equilibrium equations reveals more structure than the one in Figure 16 for the original equilibrium equations. The two causal ordering graphs do not model the same set of perfect interventions. For example, the (non)effects of an intervention targeting the cluster $\{v_S, f_S\}$ in the causal ordering graph in Figure 17b, where f_S is replaced by an equation $v_S = \xi_S$ setting v_S equal to a constant $\xi_S \in \mathbb{R}_{>0}$, cannot be read off from the equilibrium causal ordering graph in Figure 16.

7 Conclusion

Perfect adaptation is the phenomenon that a dynamical system initially responds to a change of input signal but reverts back to its original value as the system converges to equilibrium. We used the technique of causal ordering to obtain sufficient graphical conditions to identify perfect adaptation in a set of first-order differential equations. The notion of a *dynamical causal ordering graph* was introduced to support our explanation of the differences between the equilibrium and dynamical causal structure. Moreover, we




Figure 17. The equilibrium bipartite graph and equilibrium causal ordering graph associated with the basic enzyme reaction after rewriting the equilibrium equations. The absence of a directed path from f_S to v_E, v_C, v_P indicates that a soft intervention targeting f_S has no effect on those variables at equilibrium.

showed how perfect adaptation can be detected in equilibrium observational and experimental data of soft interventions with known targets.

Constraint-based causal discovery algorithms operate by constructing a graphical representation of the conditional independences in data or a probability distribution, and then reason back about what this implies for the causal relations between variables. Under additional assumptions (such as the causal Markov and faithfulness conditions) the learned graph can be interpreted causally, but these assumptions cannot generally be tested in real-world data. We demonstrated that for perfectly adapted dynamical systems the output of causal discovery algorithms applied to equilibrium data may appear to be at odds with our understanding of the mechanisms that drive the system, suggesting that the standard causal Markov and causal faithfulness conditions are not appropriate for such systems. Therefore, in practical applications of causal discovery to equilibrium data, we should avoid ambiguous terminology that obscures the possible differences between causal relations in the dynamical and equilibrium setting.

We illustrated our ideas on a variety of dynamical models and corresponding equilibrium equations. We applied the technique that we presented in this paper to a model for a well-studied protein signalling pathway and tested our predictions both in simulations and on real-world protein expression data. This turned out to be beneficial for explanation of the differences between the causal interpretation of the results of local causal discovery in real-world data and the biological consensus network that is based on the underlying dynamical equations. We hope that the results presented in this work will bring the world of causal inference closer to application domains that use dynamical models, and vice versa.

Acknowledgements

This work was supported by the ERC under the European Union's Horizon 2020 research and innovation programme (grant agreement 639466).

References

- [1] R. P. Araujo and L. A. Liotta. The topological requirements for robust perfect adaptation in networks of any size. *Nature Communications*, 9:29717141, 2018. 10.1038/s41467-018-04151-6.
- [2] T. Blom and J. M. Mooij. Robustness of model predictions under extension. *arXiv.org preprint*, arXiv:2012.04723v1 [stat.ME], 2020. URL <https://arxiv.org/abs/2012.04723>.
- [3] T. Blom, A. Klimovskaia, S. Magliacane, and J. M. Mooij. An upper bound for random measurement error in causal discovery. In *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- [4] T. Blom, S. Bongers, and J. M. Mooij. Beyond structural causal models: causal constraints models. In *Proceedings of the 35th Annual Conference on Uncertainty in Artificial Intelligence (UAI-19)*, 2019.
- [5] T. Blom, M. M. van Diepen, and J. M. Mooij. Conditional independences and causal relations implied by sets of equations. *arXiv.org preprint*, arXiv:2007.07183v2 [cs.AI], 2020. URL <https://arxiv.org/abs/2007.07183>.

- [6] P. A. Boeken and J. M. Mooij. A Bayesian nonparametric conditional two-sample test with an application to local causal discovery. *arXiv.org preprint*, arXiv:2008.07382v1 [math.ST], 2020. URL <https://arxiv.org/abs/2008.07382>.
- [7] S. Bongers and J. M. Mooij. From random differential equations to structural causal models: the stochastic case. *arXiv.org preprint*, arXiv:1803.08784v2 [cs.AI], 2018. URL <https://arxiv.org/abs/1803.08784>.
- [8] S. Bongers, P. Forré, J. Peters, B. Schölkopf, and J. M. Mooij. Foundations of structural causal models with cycles and latent variables. *arXiv.org preprint*, arXiv:1611.06221v4 [stat.ME], 2020. URL <https://arxiv.org/abs/1611.06221>.
- [9] D. Colombo, M. H. Maathuis, M. Kalisch, and T. S. Richardson. Learning high-dimensional directed acyclic graphs with latent and selection variables. *Annals of Statistics*, 40:294–321, 2012.
- [10] G. F. Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Mining and Knowledge Discovery*, 1:203–224, 1997.
- [11] D. Dash. Restructuring dynamic causal systems in equilibrium. In *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS 2005)*, pages 81–88, 2005.
- [12] D. Dash and M. J. Druzdzel. A note on the correctness of the causal ordering algorithm. *Artificial Intelligence*, 172:1800–1808, 2008.
- [13] A. P. Dawid. Beware of the dag! In *Proceedings of Workshop on Causality: Objectives and Assessment at NIPS 2008*, volume 6 of *Proceedings of Machine Learning Research*, pages 59–86, 2010.
- [14] R. J. De Boer. Which of our modeling predictions are robust? *PLOS Computational Biology*, 8:e1002593, 2012. [10.1371/journal.pcbi.1002593](https://doi.org/10.1371/journal.pcbi.1002593).
- [15] J. E. Ferrell. Perfect and near-perfect adaptation in cell signaling. *Cell Systems*, 2:62–67, 2016.
- [16] F. M. Fisher. A correspondence principle for simultaneous equation models. *Econometrica*, 38(1):73–92, 1970. ISSN 00129682. [10.2307/1909242](https://doi.org/10.2307/1909242).
- [17] P. Forré and J. M. Mooij. Markov properties for graphical models with cycles and latent variables. *arXiv.org preprint*, arXiv:1710.08775v1 [math.ST], 2017. URL <https://arxiv.org/abs/1710.08775>.
- [18] P. Forré and J. M. Mooij. Constraint-based causal discovery for non-linear structural causal models with cycles and latent confounders. In *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- [19] R. Fritsche-Guenther, F. Witzel, A. Sieber, R. Herr, N. Schmidt, S. Braun, T. Brummer, C. Sers, and N. Blüthgen. Strong negative feedback from erk to raf confers robustness to mapk signalling. *Molecular Systems Biology*, 7(1):489, 2011. <https://doi.org/10.1038/msb.2011.27>.
- [20] B. Gonçalves and F. Porto. A note on the complexity of the causal ordering problem. *Artificial Intelligence*, 238:154–165, 2016.
- [21] A. Hyttinen, F. Eberhard, and P. O. Hoyer. Learning linear cyclic causal models with latent variables. *The Journal of Machine Learning Research*, 13(1):3387–3439, 2012.
- [22] Y. Iwasaki and H. A. Simon. Causality and model abstraction. *Artificial intelligence*, 67:143–194, 1994.
- [23] J. Krishnan and I. Floros. Adaptive information processing of network modules to dynamic and spatial stimuli. *BMC Systems Biology*, 13:30866946, 2019. [10.1186/s12918-019-0703-1](https://doi.org/10.1186/s12918-019-0703-1).
- [24] G. Lacerda, P. Spirtes, J. Ramsey, and P. O. Hoyer. Discovering cyclic causal models by independent components analysis. In *Proceedings of the 24th Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 1159–1168, 2008.
- [25] S. L. Lauritzen and T. S. Richardson. Chain graph models and their causal interpretations. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 64:321–361, 2002.
- [26] S. L. Lauritzen, A. P. Dawid, B. N. Larsen, and H. Leimer. Independence properties of directed Markov fields. *Networks*, 20:491–505, 1990.
- [27] X. Lun, V. Zanotelli, J. Wade, D. Schapiro, M. Tognetti, N. Dobberstein, and B. Bodenmiller. Influence of node abundance on signaling network state and dynamics analyzed by mass cytometry. *Nature Biotechnology*, 35:164–172, 2017.
- [28] W. Ma, A. Trusina, H. El-Samad, W. A. Lim, and C. Tang. Defining network topologies that can achieve biochemical adaptation. *Cell*, 138:760–773, 2009.
- [29] S. W. Mogensen, D. Malinsky, and N. R. Hansen. Causal learning for partially observed stochastic dynamical systems. In *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- [30] J. M. Mooij and T. Claassen. Constraint-based causal discovery using partial ancestral graphs in the presence of cycles. In J. Peters and D. Sontag, editors, *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI-20)*, volume 124, pages 1159–1168. PMLR, 8 2020. URL <http://proceedings.mlr.press/v124/m-mooij20a/m-mooij20a-sup.pdf>.
- [31] J. M. Mooij and T. Heskes. Cyclic causal discovery from continuous equilibrium data. In A. Nicholson and P. Smyth, editors, *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 431–439. AUAI Press, 2013. URL <http://auai.org/uai2013/prints/papers/23.pdf>.
- [32] J. M. Mooij, D. Janzing, T. Heskes, and B. Schölkopf. On causal discovery with cyclic additive noise models. *Advances in Neural Information Processing Systems*, pages 639–647, 2011.

- [33] J. M. Mooij, D. Janzing, and B. Schölkopf. From ordinary differential equations to structural causal models: the deterministic case. In *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 440–448, 2013. URL <http://auai.org/uai2013/prints/papers/24.pdf>.
- [34] J. M. Mooij, S. Magliacane, and T. Claassen. Joint causal inference from multiple contexts. *Journal of Machine Learning Research*, 21:1–108, 2020. URL <http://jmlr.org/papers/v21/17-123.html>.
- [35] J. D. Murray. *Mathematical biology I: an introduction*. Springer-Verlag New York, third edition edition, 2002. ISBN 978-0-387-95223-9. 10.1007/b98868.
- [36] D. Muzzey, C. A. Gómez-Urbe, J. T. Mettetal, and A. van Oudenaarden. A systems-level analysis of perfect adaptation in yeast osmoregulation. *Cell*, 138:160–171, 2009.
- [37] P. Nayak. *Automated modeling of physical systems*. Springer-Verlag Berlin Heidelberg, 1995.
- [38] J. Pearl. *Causality : models, reasoning, and inference*. Cambridge University Press, 2000. ISBN 052189560X.
- [39] A. Pothén and C.-J. Fan. Computing the block triangular form of a sparse matrix. *ACM Transactions on Mathematical Software (TOMS)*, 16:303–324, 1990.
- [40] J. Ramsey and B. Andrews. Fask with interventional knowledge recovers edges from the Sachs model. *arXiv.org preprint*, arXiv:1805.03108 [q-bio.MN], 2018. URL <https://arxiv.org/abs/1805.03108>.
- [41] T. S. Richardson. *Models of feedback: interpretation and discovery*. PhD dissertation, Carnegie-Mellon University, 1996.
- [42] T. S. Richardson and P. Spirtes. Automated discovery of linear feedback models. *Computation, Causation, and Discovery*, pages 254–304, 1999.
- [43] P. K. Rubenstein, S. Bongers, B. Schölkopf, and J. M. Mooij. From deterministic odes to dynamic structural causal models. In *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- [44] K. Sachs, O. Perez, D. Pe'er, D. A. Lauffenburger, and G. P. Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308:523–529, 2005.
- [45] K. Sachs, S. Itani, J. Fitzgerald, B. Schoeberl, G. Nolan, and C. Tomlin. Single timepoint models of dynamic systems. *Interface Focus*, 3:24511382, 2013.
- [46] S. Y. Shin, O. Rath, S. M. Choo, F. Fee, B. McFerran, W. Kolch, and K. H. Cho. Positive- and negative-feedback regulations coordinate the dynamic behavior of the ras-raf-mek-erk signal transduction pathway. *Journal of Cell Science*, 122:425–435, 2009.
- [47] H. A. Simon. Causal ordering and identifiability. In *Studies in Econometric Methods*, pages 49–74. John Wiley & Sons, 1953.
- [48] A. Sokol and N. R. Hansen. Causal interpretation of stochastic differential equations. *Electronic Journal of Probability*, 19:1–24, 2014.
- [49] P. Spirtes. Directed cyclic graphical representations of feedback models. In *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI-1995)*, 1995.
- [50] P. Spirtes, C. Glymour, and R. Scheines. *Causation, prediction, and search*. MIT press, 2000.
- [51] E. V. Strobl. A constraint-based algorithm for causal discovery with cycles, latent variables and selection bias. *International Journal of Data Science and Analytics*, 8:33–56, 2019.
- [52] S. Triantafillou, V. Lagani, C. Heinze-Deml, A. Schmidt, J. Tegner, and I. Tsamardinos. Predicting Causal Relationships from Biological Data: Applying Automated Causal Discovery on Mass Cytometry Data of Human Immune Cells. *Scientific Reports*, 7:12724, 2017. 10.1038/s41598-017-08582-x.
- [53] M. Voortman, D. Dash, and M. J. Druzdzal. Learning why things change: the difference-based causality learner. In *Proceedings of the Twenty-Sixth Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010.
- [54] J. Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172:1873–1896, 2008.

A Perfect adaptation simulations

For the simulations in Figures 3 and 9 of the model of a filling bathtub, the viral infection model, the reaction network with a feedback loop, and the protein pathway we used the settings listed below. Since we only simulated a single response, we used constant values for the exogenous random variables as well.

1. Filling bathtub: First we recorded the behaviour of the system for the parameters $I_K = 1.2$, $U_I = 5.0$, $U_1 = 1.1$, $U_2 = 1.0$, $U_3 = 1.2$, $U_4 = 1.0$, $U_5 = 0.8$, $g = 1.0$ until it reached equilibrium. We then changed the input parameter I_K to 0.8, 1.0, and 1.3 and recorded the response until the system reverted to equilibrium.

2. Viral infection: For the parameter settings $I_\sigma = 1.6$, $d_T = 0.9$, $\beta 0.9$, $f = 1.0$, $d_I = 0.3$, $k = 1.5$, $a = 0.1$, $d_E = 0.25$, we simulated the model until it reached equilibrium. We changed the input parameter I_σ to 1.1, 1.3, and 2.0 and recorded the response until equilibrium was reached.
3. Reaction Network: We simulated the model until it reached equilibrium with parameters $I = 1.5$, $k_{IA} = 1.4$, $K_{IA} = 0.8$, $F_A = 1.1$, $k_{FAA} = 0.9$, $K_{FAA} = 1.2$, $k_{CB} = 0.6$, $K_{CB} = 0.0001$, $F_B = 0.7$, $k_{FBB} = 0.7$, $K_{FBB} = 0.0001$, $k_{AC} = 2.1$, $K_{AC} = 1.5$, $k_{BC} = 0.7$, $K_{BC} = 0.6$. The settings were chosen in such a way that the saturation conditions $(1 - X_B(t)) \gg K_{CB}$ and $X_B(t) \gg K_{FBB}$ were satisfied. We then changed the input signal to 0.25, 1.0, and 10.0 and recorded the response.
4. Protein pathway: The parameter settings of the simulation were $I = 1.0$, $k_{Is} = 1.0$, $T_s = 1.0$, $K_{Is} = 1.0$, $K_e = 1.5$, $F_s = 1.0$, $k_{Fss} = 1.0$, $K_{Fss} = 0.9$, $k_{sr} = 1.0$, $K_{sr} = 1.0$, $T_r = 1.0$, $F_r = 0.3$, $k_{Frr} = 1.0$, $K_{Frr} = 0.8$, $k_{rm} = 1.0$, $K_{rm} = 0.9$, $T_m = 1.0$, $F_m = 0.2$, $k_{Fmm} = 1.0$, $K_{Fmm} = 1.2$, $k_{me} = 1.0$, $K_{me} = 0.0001$, $T_e = 1.0$, $F_e = 0.7$, $k_{Fee} = 1.2$, $K_{Fee} = 0.0001$. This ensured that the saturation conditions $(T_e - X_e(t)) \gg K_{me}$ and $X_e(t) \gg K_{Fee}$ were satisfied. After the system reached equilibrium we changed the input signal to 0.9, 1.1, and 1.5 and recorded the response.

The qualitative behaviour that we presented in Figure 3 can be observed for a range of parameter values and does not require exact tuning of the parameters.

B Conditional independences

The Markov ordering graph in Figure 10f was derived from the equilibrium equations of the protein pathway model under saturation conditions. From this we can read off the following d-separations:

$$\begin{aligned}
& I \perp^d v_s, \quad I \perp^d v_s | v_r, \quad I \perp^d v_s | v_m, \quad I \perp^d v_s | \{v_r, v_m\}, \\
& I \perp^d v_r, \quad I \perp^d v_r | v_s, \quad I \perp^d v_r | v_m, \quad I \perp^d v_r | \{v_s, v_m\}, \\
& I \perp^d v_m, \quad I \perp^d v_m | v_s, \quad I \perp^d v_m | v_r, \quad I \perp^d v_m | \{v_s, v_r\}, \\
& v_e \perp^d v_r | v_s, \quad v_e \perp^d v_r | \{v_s, v_m\}, \quad v_e \perp^d v_r | \{v_s, I\}, \quad v_e \perp^d v_r | \{v_s, v_m, I\}, \\
& v_e \perp^d v_m | v_s, \quad v_e \perp^d v_m | v_r, \quad v_e \perp^d v_m | \{v_s, v_r\}, \\
& v_e \perp^d v_m | \{v_s, I\}, \quad v_e \perp^d v_m | \{v_r, I\}, \quad v_e \perp^d v_m | \{v_s, v_r, I\}, \\
& v_s \perp^d v_m | v_r, \quad v_s \perp^d v_m | \{v_e, v_r\}, \quad v_s \perp^d v_m | \{v_r, I\}, \quad v_s \perp^d v_m | \{v_e, v_r, I\}.
\end{aligned}$$

It is easy to check that the equilibrium equations and endogenous variables in this model are uniquely solvable w.r.t. the causal ordering graph (see Blom et al. [5] for details on unique solvability). Therefore, the d-separations above imply conditional independences between the variables in the model [5]. Table 1 shows that the conditional independences with a maximum conditioning set of size one that are implied by the Markov ordering graph are also present in the simulated data.

C Reasoning about feedback loops

Consider the protein signalling model under saturation conditions that is defined by equations (21), (22), (23), and (25). Suppose that the system is only partially modelled and that $X_e(t)$ is treated as a latent

exogenous variable U_e in the submodel for $X_s(t)$, $X_r(t)$, and $X_m(t)$ defined by equations:

$$\dot{X}_s(t) = \frac{I(t)k_{Is}(T_s - X_s(t))}{(K_{Is} + (T_s - X_s(t))) \left(1 + \left(\frac{U_e}{K_e}\right)^{\frac{3}{2}}\right)} - F_s k_{F_{ss}} \frac{X_s(t)}{K_{F_{ss}} + X_s(t)}, \quad (36)$$

$$\dot{X}_r(t) = \frac{X_s(t)k_{sr}(T_r - X_r(t))}{K_{sr} + (T_r - X_r(t))} - F_r k_{F_{rr}} \frac{X_r(t)}{K_{F_{rr}} + X_r(t)}, \quad (37)$$

$$\dot{X}_m(t) = \frac{X_r(t)k_{rm}(T_m - X_m(t))}{K_{rm} + (T_m - X_m(t))} - F_m k_{F_{mm}} \frac{X_m(t)}{K_{F_{mm}} + X_m(t)}. \quad (38)$$

Application of the causal ordering technique to the equilibrium equations associated with these differential equations results in the Markov ordering graph in Figure 18. Assuming faithfulness, the d-connections in this graph indicate that the input signal I is dependent on the equilibrium distribution X_s , X_r , and X_m . However, if X_s , X_r , and X_m were generated by the larger model with the Markov ordering graph in Figure 10f, we know that a statistical test would indicate that they are independent. The discrepancy between the Markov ordering graph for the equilibrium equations of the submodel and these statistical tests would not be due to a faithfulness violation but could be wholly explained by a holistic modelling approach (i.e. by not assuming all unobserved causes to be exogenous to the observed variables). According to Corollary 1 and Proposition 1 in Blom and Mooij [2], the discrepancy between the observed and predicted conditional independences, implies the presence of a non-selfregulating variable and an unobserved dynamical feedback loop involving X_s , X_r , and X_m , if we assume faithfulness. This is in agreement with the fact that the dynamic variable $X_e(t)$ is not self-regulating and that there is a feedback loop in the dynamical causal ordering graph (indicated by the cluster $\{v_s, v_r, v_m, v_e, f_s, f_r, f_m, f_e\}$) in the saturated protein signalling model in Section 4. Interestingly, we can infer the presence of feedback without modelling or observing $X_e(t)$.




Figure 18. Markov ordering graph for the partial model of the protein signalling pathway model given by equations (36), (37), and (38).

D IFFLP Network

The IFFLP topology in Ma et al. [28] that we discussed in Section 3.1.3 could be a graphical representation of the following differential equations:

$$\dot{X}_A(t) = I(t)k_{IA} \frac{(1 - X_A(t))}{K_{IA} + (1 - X_A(t))} - F_A k_{F_{AA}} \frac{X_A(t)}{K_{F_{AA}} + X_A(t)}, \quad (39)$$

$$\dot{X}_B(t) = X_A(t)k_{AB} \frac{(1 - X_B(t))}{K_{AB} + (1 - X_B(t))} - F_B k_{F_{BB}} \frac{X_B(t)}{K_{F_{BB}} + X_B(t)}, \quad (40)$$

$$\dot{X}_C(t) = X_A(t)k_{AC} \frac{(1 - X_C(t))}{K_{AC} + (1 - X_C(t))} - X_B(t)k_{BC} \frac{X_C(t)}{K_{BC} + X_C(t)}, \quad (41)$$

where $I(t)$ represents an external input into the system. This network is capable of perfect adaptation if the first term of $\dot{X}_B(t)$ is in the saturated region $(1 - X_B(t)) \gg K_{AB}$ and the second term is in the linear region $X_B(t) \ll K_{F_{BB}}$, which allows us to make the following approximation:

$$\frac{dX_B(t)}{dt} \approx X_A(t)k_{AB} - \frac{F_B k_{F_{BB}}}{K_{F_{BB}}} X_B(t). \quad (42)$$

Therefore, a steady state solution X_B for B is proportional to the steady state solution X_A for A . Since both terms in equation (41) are proportional to X_A we find that the steady state solution X_C for C is a function of only the parameters k_{AC} , K_{AC} , k_{BC} , and K_{BC} (note that X_A factors out of the equilibrium equation corresponding to (42)), and hence it does not depend on the input parameter I . Since a change in the input signal I changes $\dot{X}_A(t)$ there is a transient effect on $X_A(t)$. Similarly there must also be a transient effect on both $X_B(t)$ and $X_C(t)$. It follows that the system achieves perfect adaptation.

The equilibrium equations associated with equations (39), the approximation (42) to (40), and (41) are given by:

$$f_A : \quad Ik_{IA} \frac{(1 - X_A)}{K_{IA} + (1 - X_A)} - F_A k_{FAA} \frac{X_A}{K_{FAA} + X_A} = 0, \quad (43)$$

$$f_B : \quad X_A k_{AB} - \frac{F_B k_{FBB}}{K_{FBB}} X_B = 0, \quad (44)$$

$$f_C : \quad X_A k_{AC} \frac{(1 - X_C)}{K_{AC} + (1 - X_C)} - X_B k_{BC} \frac{X_C}{K_{BC} + X_C} = 0. \quad (45)$$

The associated equilibrium causal ordering graph in Figure shows that there is a directed path from the input signal I to the cluster $\{v_A, v_B, v_C\}$. Therefore, the conditions of Theorem 1 are not satisfied for the system with input signal I .




Figure 19. The equilibrium causal ordering graph for the IFFLP network modelled by equations (43), (44), and (45).

Table 1. The conditional independences in the simulation of the protein pathway described in Section 5.1 were assessed using Spearman's rank correlations. With a p-value threshold of 0.01, d-separations with a separating set of size 0 or 1 coincide with conditional independences with conditioning sets of size 0 or 1.

Independence test	Correlation	p-value	d-separation
$I \perp\!\!\!\perp X_s$	0.029	0.51	yes
$I \perp\!\!\!\perp X_r$	0.020	0.66	yes
$I \perp\!\!\!\perp X_m$	0.021	0.64	yes
$I \perp\!\!\!\perp X_e$	0.777	$< 2.2e^{-16}$	no
$X_s \perp\!\!\!\perp X_r$	0.957	$< 2.2e^{-16}$	no
$X_s \perp\!\!\!\perp X_m$	0.933	$< 2.2e^{-16}$	no
$X_s \perp\!\!\!\perp X_e$	-0.561	$< 2.2e^{-16}$	no
$X_r \perp\!\!\!\perp X_m$	0.977	$< 2.2e^{-16}$	no
$X_r \perp\!\!\!\perp X_e$	-0.542	$< 2.2e^{-16}$	no
$X_m \perp\!\!\!\perp X_e$	-0.524	$< 2.2e^{-16}$	no
$I \perp\!\!\!\perp X_s \mid X_r$	0.037	0.83	yes
$I \perp\!\!\!\perp X_s \mid X_m$	0.027	0.61	yes
$I \perp\!\!\!\perp X_r \mid X_s$	-0.030	0.51	yes
$I \perp\!\!\!\perp X_r \mid X_m$	-0.005	0.91	yes
$I \perp\!\!\!\perp X_m \mid X_s$	-0.018	0.69	yes
$I \perp\!\!\!\perp X_m \mid X_r$	0.010	0.83	yes
$X_e \perp\!\!\!\perp X_r \mid X_s$	-0.019	0.67	yes
$X_e \perp\!\!\!\perp X_m \mid X_s$	$-2.1 \cdot 10^{-4}$	0.99	yes
$X_e \perp\!\!\!\perp X_m \mid X_r$	0.031	0.49	yes
$X_s \perp\!\!\!\perp X_m \mid X_r$	-0.031	0.48	yes
$I \perp\!\!\!\perp X_e \mid X_s$	0.959	$6.0 \cdot 10^{-275}$	no
$I \perp\!\!\!\perp X_e \mid X_r$	0.937	$1.6 \cdot 10^{-229}$	no
$I \perp\!\!\!\perp X_e \mid X_m$	0.925	$1.2 \cdot 10^{-211}$	no
$I \perp\!\!\!\perp X_s \mid X_e$	0.894	$4.7 \cdot 10^{-175}$	no
$I \perp\!\!\!\perp X_r \mid X_e$	0.832	$1.5 \cdot 10^{-129}$	no
$I \perp\!\!\!\perp X_m \mid X_e$	0.799	$1.1 \cdot 10^{-111}$	no
$X_e \perp\!\!\!\perp X_s \mid X_r$	-0.176	$8.0 \cdot 10^{-5}$	no
$X_e \perp\!\!\!\perp X_s \mid X_m$	-0.236	$9.3 \cdot 10^{-8}$	no
$X_e \perp\!\!\!\perp X_s \mid I$	-0.928	$8.7 \cdot 10^{-216}$	no
$X_e \perp\!\!\!\perp X_r \mid X_m$	-0.164	$2.2 \cdot 10^{-4}$	no
$X_e \perp\!\!\!\perp X_r \mid I$	-0.885	$5.3 \cdot 10^{-167}$	no
$X_e \perp\!\!\!\perp X_m \mid I$	-0.859	$2.8 \cdot 10^{-146}$	no
$X_s \perp\!\!\!\perp X_r \mid I$	0.957	$1.7 \cdot 10^{-269}$	no
$X_s \perp\!\!\!\perp X_r \mid X_e$	0.939	$5.3 \cdot 10^{-232}$	no
$X_s \perp\!\!\!\perp X_r \mid X_m$	0.590	$3.2 \cdot 10^{-48}$	no
$X_s \perp\!\!\!\perp X_m \mid I$	0.933	$3.0 \cdot 10^{-223}$	no
$X_s \perp\!\!\!\perp X_m \mid X_e$	-0.907	$1.2 \cdot 10^{-188}$	no
$X_r \perp\!\!\!\perp X_m \mid I$	-0.977	0	no
$X_r \perp\!\!\!\perp X_m \mid X_e$	0.968	$2.8 \cdot 10^{-302}$	no
$X_r \perp\!\!\!\perp X_m \mid X_s$	0.807	$1.1 \cdot 10^{-115}$	no