

A Statistical Unsupervised Learning Algorithm for Inferring Reaction Networks from Time Series Data

Julien Martinelli, Jeremy Grignard, Sylvain Soliman, François Fages

► **To cite this version:**

Julien Martinelli, Jeremy Grignard, Sylvain Soliman, François Fages. A Statistical Unsupervised Learning Algorithm for Inferring Reaction Networks from Time Series Data. ICML 2019 - Workshop on Computational Biology, Jun 2019, Long Beach, CA, United States. hal-02163862

HAL Id: hal-02163862

<https://hal.inria.fr/hal-02163862>

Submitted on 24 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Statistical Unsupervised Learning Algorithm for Inferring Reaction Networks from Time Series Data

Julien Martinelli^{1,2} Jeremy Grignard^{1,3} Sylvain Soliman¹ François Fages¹

Abstract

With the automation of biological experiments and the increase of quality of single cell data that can now be obtained by phosphoproteomic and time lapse videomicroscopy, automating the building of mechanistic models from these time series data becomes conceivable and a necessity for many new applications. While learning numerical parameters to fit a given model structure to observed data is now a quite well understood subject, learning the structure of the model is a more challenging problem that previous attempts failed to solve without relying quite heavily on prior knowledge about that structure. In this paper, we consider mechanistic models based on chemical reaction networks (CRN) with their continuous dynamics based on ordinary differential equations, and finite time series about the time evolution of concentration of molecular species for a given time horizon and a finite set of perturbed initial conditions. We present a statistical learning algorithm to learn CRNs with a time complexity for inferring one reaction in $\mathcal{O}(t.n^2)$ where n is the number of species and t the number of observed transitions in the traces. We learn both the structure and the reaction rates of the CRN. We evaluate this algorithm and its sensitivity to its statistical threshold parameters, first on simulated data from a hidden CRN, and second on real videomicroscopy single cell time series data over three days about the circadian clock and cell cycle progression of NIH3T3 embryonic fibroblasts. In all cases, our algorithm is able to reconstruct meaningful CRNs. We discuss some limits according to the existence of multiple time

scales and highly variable traces.

1. Introduction

Recent breakthroughs in Machine Learning are paving the way for new kinds of algorithms for analysing data and making diagnosis and predictions in biology and medicine. However, the direct application of machine learning methods, while capable of making accurate predictions, do not provide a biological understanding of the underlying processes nor explanation for the predictions, and may be not accepted in the biomedical domain. For these reasons, a lot of work aims at providing explanations for the predictions made as output of neural networks or other machine learning algorithms trained on data.

Another approach is to try to learn mechanistic models that will make predictions instead of learning directly the predictions from the data. Building mechanistic models of cell processes is however a hard work which necessitates to determine the biochemical mechanisms that are responsible for the high level functions of the cell and its behaviors in normal and perturbed conditions. Automating this process would enable new applications such as automated experiment design or patient-tailored therapeutics.

The main difficulty is to be able to discriminate between causality and correlations in time series data (Pearl, 2009; Preacher & Hayes, 2008). There has been work on this problem for the detection of causal relationship among gene networks. For instance in (Triantafillou et al., 2017), causal relationships are searched between genes and accepted with respect to statistical tests (local causal discovery method). Causal inference can be achieved thanks to randomized experiments, relying on multiple different initial conditions. Supervised learning shows encouraging results as well. Using time series data, the authors of (Ganscha et al., 2018) trained recurrent neural networks and applied them to the DREAM gene regulatory network inference challenge data. Besides good predictions results, their work highlights the need for a simulation of biologically representative time-series data as training set. In the framework of boolean regulatory networks, Answer Set Programming (ASP), a language in which problems and their solutions are respec-

¹Inria Saclay-Île de France, Palaiseau, France ²INSERM U935, Villejuif, France ³Institut de Recherches Servier, Croissy sur Seine, France. Correspondence to: François Fages <Francois.Fages@inria.fr>.

tively expressed in terms of logic programs and models of those programs, has been applied to the discovery of influence relationships between species. In (Ostrowski et al., 2016), considering a prior knowledge graph, a necessary condition for a boolean network to be compatible with the observed time series data is proposed. Then the efficiency of ASP makes it possible to compute an over-approximation of the set of boolean networks best fitting the experimental data. This is illustrated with DREAM challenge data. In (Carcano et al., 2017) the Probably Approximately Correct (PAC) learnability theory of Leslie Valliant’s has been applied to Thomas’s influence models. This algorithm comes with probabilistic guarantees, giving learning bounds depending on the number of samples drawn, underlining the need for diverse enough perturbed condition samples. This shows the existence of a space-time trade-off in the traces between the diversity of initial conditions with zero initial values, and the length of the traces, the former transitions being more informative than the latter. Prior knowledge can also be specified by giving to a species a set of possible species influencing it, resulting in a decreasing number of samples needed. In practice however, these methods need prior knowledge on the structure of the network to obtain its Boolean dynamics with good results.

In (Choi et al., 2018) the problem of inferring chemical reaction networks (CRNs) is defined as the minimization of a fitness criterion based on the compatibility of the learned mechanistic model with the observed traces. This evolutionary algorithm approach manages to recover reactions as well as their mass action law kinetics parameters via a two-step iterative procedure: first a set of reaction is inferred, then its kinetics parameters are estimated.

In this paper, we present a statistical learning algorithm dedicated to the same problem of inferring a CRN from time series data. Our unsupervised learning algorithm does not require prior knowledge nor training. We consider at most binary reactions with mass-action, Michaelian or order 4 Hill kinetics. Based on a pairing of the variations of molecular species in each observed transition, the algorithm repeatedly infers the reaction that minimizes the standard deviation of the inferred rate function among all the observed transitions where the reaction can occur. Once inferred, the contributions of that reaction to state change in the set of observed transitions are subtracted before inferring the next reaction.

Example 1. For instance, on a simple chain of 4 reactions with mass action law kinetics over 5 molecular species, our algorithm is able to reconstruct the CRN from a single simulation trace (Fig. 1).

In the next section, we give some preliminaries on CRNs and time series data. In Sec. 3 we present our statistical learning algorithm with the table of statistical thresholds

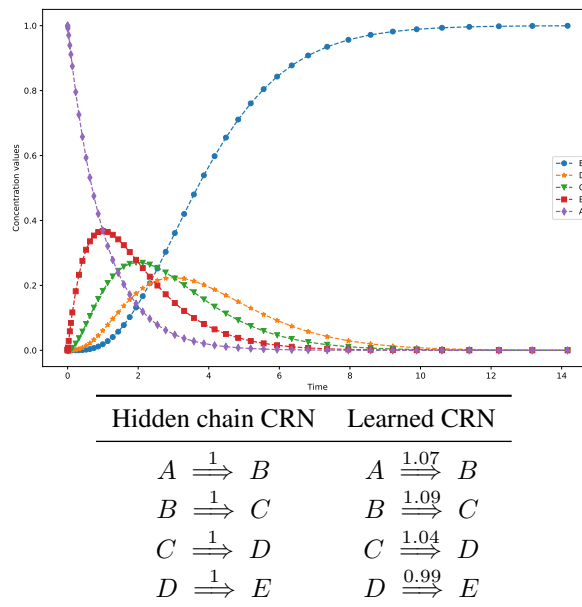


Figure 1. Chain example: simulation trace and learned CRN.

used. We show that its time complexity for inferring one reaction is $\mathcal{O}(t.n^2)$ where n is the number of species and t the number of transitions in the traces. Then we evaluate this algorithm first on simulated data with respectively reactant-parallel, product-parallel, parallel and general hidden CRNs. In each case we analyze the sensitivity of our algorithm to its statistical threshold parameters. Then in Sec. 6 we apply our algorithm to real videomicroscopy time series data on NIH3T3 embryonic mice fibroblasts with FUCCI markers of the cell cycle and Rev-erb- α marker of the circadian clock, over three days, for a total of 91 tracked cell traces and 26000 observed state transitions (Feillet et al., 2014). In all cases, our algorithm is able to reconstruct meaningful yet sometimes incomplete CRNs¹. We conclude on its limits with respect to incompleteness, i.e. false negatives, and erroneous inferences, i.e. false positives, in particular in the cases of highly variable data, and of CRN with multiple time scales.

2. Preliminaries

Unless explicitly noted, we will denote sets and multisets by capital letters (e.g., S or in some cases calligraphic letters like \mathcal{D} , \mathcal{R}), tuples of values by vectors (e.g., \vec{x}) and elements of those sets or vectors (i.e., real numbers, functions, etc.) by small Roman or Greek letters. For a multiset M , $M(x)$ denotes the multiplicity of element x in M . \mathbb{R}_+ denotes the set of non negative real numbers.

¹The code and the examples described in this article are available at <https://lifeware.inria.fr/wiki/Main/Software#CMSB19a>

2.1. Chemical Reaction Networks

We assume given a finite set $S = \{1, \dots, n\}$ of n molecular species. For the sake of readability in the examples, species will also be noted with biological names (e.g., *Cyc* for a Cyclin) or even simple capital letters like A, B, C , instead of by their index $1, \dots, n$.

Definition 1. A reaction over S is a triple (R, P, f) , where R is a multiset of reactants in S , P is a multiset of products in S and $f : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ is a rate function over molecular concentrations.

A catalyst in a reaction is a reactant i such that $R(i) = P(i)$. A strict reactant (resp. strict product, strict participant species) in a reaction is a species such that $R(i) > P(i)$ (resp. $P(i) > R(i)$, $R(i) \neq P(i)$).

A Chemical Reaction Network (CRN) is a finite set of reactions.

For the sake of simplicity, this article only studies reactions with stoichiometry at most 1. In other words, the multisets R and P are actually sets of $\mathcal{P}(S)$. We shall also restrict ourselves to the following rate functions: mass action law kinetics, i.e. product of reactant concentrations by some constant k , Michaelis-Menten kinetics, and Hill of order 4 kinetics.

Example 2. For reactions with mass action law kinetics, the reactions will be written $(\{A, C\}, \{B, C\}, kAC)$ as $((\{A, C\}, \{B, C\}), k)$ or even $A + C \xrightarrow{k} B + C$. For the other kinetics, the rate functions are written explicitly.

Furthermore we introduce the following

Definition 2. A reactant-parallel CRN is a CRN in which any two reactions do not share a same strict reactant. A product-parallel CRN is a CRN in which any two reactions do not share a same strict product. A parallel CRN is a CRN in which any two reactions do not share a same strict participant species.

It is worth noting that Ex. 1 of a linear chain of reactions shows that a both reactant-parallel and product-parallel CRN is not necessarily parallel.

As usual, a CRN can be interpreted in different manners, in a hierarchy of continuous differential, stochastic, discrete and Boolean semantics (Fages & Soliman, 2008). For this article, we consider only the continuous interpretation by ordinary differential equations,

$$\forall s \in S, ds/dt = \sum_{(R,P,f) \in \mathcal{R}} f.(P(s) - R(s)).$$

2.2. Time Series Data

Definition 3. A state vector is a vector $\mathbf{x} \in \mathbb{R}_+^{n+1}$ where x_0 represents the real time, and x_i the concentration of

species i .

A trace, or time series data, is a finite sequence $(\mathbf{x}(1), \dots, \mathbf{x}(d))$ of state vectors at increasing times, i.e. $x_0(1) < \dots < x_0(d)$.

The time step between two time points may be non constant and is given by the first components x_0 of the state vectors. Such traces can be produced either by numerical simulations, e.g. by numerical integration of ordinary differential equations or by stochastic simulation, or by biological experiments, e.g. by time lapse microscopy.

For the sake of evaluating a learning algorithm, synthetic data obtained by simulation of a hidden model must be considered. Fig. 1 shows a simulation trace in dimension 5 for Ex. 1. Furthermore, in order to produce traces close to real data, simulation data can be subsampled at a low frequency, and noise added. It is worth noting however that numerical integration already introduces big time steps and a form of noise on discrete time points due to the discretization and approximation of the derivatives. In the following, we analyze the sensitivity of our learning algorithm to subsampling, trace length, and number of traces with random initial conditions with zero values.

3. Unsupervised Learning Algorithm

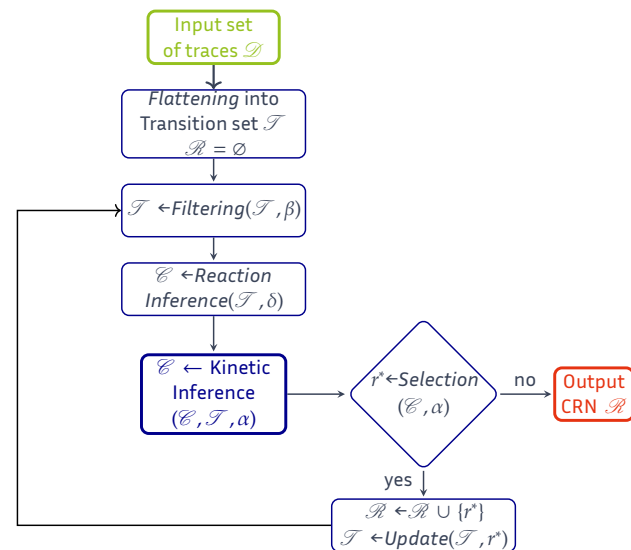


Figure 2. Flowchart of the CRN learning algorithm.

Fig. 2 shows the flowchart of our unsupervised CRN learning algorithm. Table 1 lists the global variables and threshold parameters used with their default values. Alg. 1 describes the main structure of the algorithm. First, the set of traces \mathcal{D} is flattened into a set of transitions \mathcal{T} represented by the triple of predecessor state, difference vector and in-

dex of the trace from where it originates. The transitions are then filtered (Alg. 2) by keeping only the variations that are significant compared to the maximum variation of the species' concentration values along the trace it originates from. Then, from each transition $(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}$, Alg. 3 infers a set of one or two reaction candidates that can explain the most important variation in a transition. Each reaction candidate is returned together with a variable (reactant or product) on which the rate function should be estimated.

Parameter	Value	Threshold description
α	0.3	reaction rate variability in Alg.1
β	0.03	species variation in Alg.2
δ	1	species variation pairing in Alg.3
γ	10^{-4}	zero value in Alg.8
Variable	Description	
\mathbf{x}	state vector	
\mathbf{d}	difference vector	
\mathcal{T}	transition set	
\mathcal{D}	set of traces	
\mathcal{R}	inferred CRN (set of reactions)	
r	(inferred) reaction	
σ	standard deviation of reaction rate	

Table 1. Variables and threshold parameters of the algorithm.

Algorithm 1 Main reaction learning loop

Input: Array of l traces:

$\mathcal{D} = [(\mathbf{x}^1(1), \dots, \mathbf{x}^1(d_1)), \dots, (\mathbf{x}^l(1), \dots, \mathbf{x}^l(d_l))]$

Output: CRN \mathcal{R}

Parameter: α

```

1:  $\mathcal{R} \leftarrow \emptyset, \sigma^* \leftarrow 0$ 
2:  $\mathcal{T} \leftarrow \{(\mathbf{x}^j(t), \mathbf{x}^j(t+1) - \mathbf{x}^j(t), j) \mid$ 
3:    $1 \leq j \leq l, 1 \leq t \leq d_j - 1\}$ 
4: while  $\sigma^* < \alpha$  do
5:    $\sigma^* \leftarrow +\infty, \text{TRANSITION\_FILTERING}(\mathcal{T}, \mathcal{D})$ 
6:   for  $(i, R, P) \in \bigcup_{(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}} \text{REACTION}(\mathbf{d})$  do
7:      $\mathcal{T}_r \leftarrow \{(\mathbf{x}, \mathbf{e}, j) \in \mathcal{T} \mid$ 
8:        $(i, R, P) \in \text{REACTION}(\mathbf{e})\}$ 
9:      $T_R \leftarrow \{(\mathbf{x}, \mathbf{e}, j) \in \mathcal{T}_r \text{ s.t. } \mathbf{x}_m > 0 \forall m \in R\}$ 
10:     $(r, \sigma) \leftarrow \text{KINETIC}(i, R, P, T_R)$ 
11:    if  $\sigma > \alpha$  then  $(r, \sigma) \leftarrow \text{CATALYST}(i, R, P, T_R)$ 
12:     $\sigma \leftarrow \sigma / (1 + \text{SPARSITY}(\mathcal{T}_r, \mathcal{D}))$ 
13:    if  $\sigma < \sigma^*$  then  $\sigma^* \leftarrow \sigma, r^* \leftarrow r$ 
14:    if  $\sigma^* < \alpha$  then  $\mathcal{R} \leftarrow \mathcal{R} \cup \{r^*\},$ 
15:     $\text{TRANSITION\_UPDATE}(\mathcal{T}, r^*)$ 
16: return  $\mathcal{R}$ 
    
```

For each reaction candidate, Alg. 4 returns the rate function with the best statistical score over mass action law kinetics (Alg. 5), Michaelis-Menten and Hill kinetics (Alg. 6). The score is the standard deviation of the ratio of the inferred

kinetics over the observed kinetics in all transitions. If the score is above some threshold α , a catalyst is searched for by the function Alg. 7. Furthermore, as a heuristic to minimize the risk of interference with other not yet discovered reactions, the score of the rate functions estimated on transitions with zero values in the predecessor states are favored by Alg. 8. At the end of the body of the main loop, the reaction with the best score is inferred, and the whole set of transitions is updated by removing its effects on the transitions (Alg. 9), before iterating until no more reactions can be inferred.

Algorithm 2 Slow variation filtering

Parameter: β

```

1: procedure  $\text{TRANSITION\_FILTERING}(\mathcal{T}, \mathcal{D})$ 
2:   for  $(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}$  do
3:     let  $(\mathbf{x}^j(1), \dots, \mathbf{x}^j(d_j)) = \mathcal{D}_j$ 
4:     for  $i \in S$  do
5:       if  $\left| \frac{\mathbf{d}_i}{\mathbf{d}_0} \right| < \beta \cdot \max_{1 \leq t < d_j} \left| \frac{\mathbf{x}_i^j(t+1) - \mathbf{x}_i^j(t)}{\mathbf{x}_0^j(t+1) - \mathbf{x}_0^j(t)} \right|$ 
6:         then  $\mathbf{d}_i \leftarrow 0$ 
    
```

Algorithm 3 Reaction inference

Parameter: δ

```

1: function  $\text{REACTION}(\mathbf{d})$ 
2:    $R \leftarrow \emptyset, P \leftarrow \emptyset$ 
3:    $s^* \leftarrow \underset{s \in S}{\text{argmax}} |\mathbf{d}_s|$ 
4:    $I \leftarrow \{i \in S \text{ s.t. } |\mathbf{d}_{s^*}| \leq (1 + \delta) \cdot |\mathbf{d}_i|\}$ 
5:   if  $|I| > 1$  then
6:     for  $i \in I$  do
7:       if  $\mathbf{d}_i < 0$  then  $R \leftarrow R \cup \{i\}$ 
8:     else  $P \leftarrow P \cup \{i\}$ 
9:      $j \leftarrow i \in P$ 
10:    return  $\{(j, R, P)\}$ 
11:   else
12:      $(i_1, i_2) \leftarrow \underset{s, s' \in S, \mathbf{d}_s \mathbf{d}_{s'} > 0, \mathbf{d}_s \mathbf{d}_{s^*} < 0}{\text{argmin}} \left\{ \left| 1 + \frac{\mathbf{d}_{s^*}}{\mathbf{d}_s + \mathbf{d}_{s'}} \right| \right\}$ 
13:     if  $\left| 1 + \frac{\mathbf{d}_{s^*}}{\mathbf{d}_{i_1} + \mathbf{d}_{i_2}} \right| < \delta$  then
14:       if  $\mathbf{d}_{s^*} < 0$  then
15:         return  $\{(i_1, \{s^*\}, \{i_1\}), (i_2, \{s^*\}, \{i_2\})\}$ 
16:       else
17:         return  $\{(i_1, \{i_1\}, \{s^*\}), (i_2, \{i_2\}, \{s^*\})\}$ 
18:       else
19:         return  $\{(i_1, \{i_1\}, \{s^*\}), (i_2, \{i_2\}, \{s^*\})\}$ 
20:     else
21:       if  $\mathbf{d}_{s^*} < 0$  then return  $\{(s^*, \{s^*\}, \emptyset)\}$ 
22:     else return  $\{(s^*, \emptyset, \{s^*\})\}$ 
    
```


Algorithm 4 Inference of reaction rate function

```

1: function KINETIC( $i, R, P, \mathcal{T}$ )
2:    $(f_1, \sigma_1) \leftarrow \text{MASS\_ACTION}(i, R, P, \mathcal{T})$ 
3:    $(f_2, \sigma_2) \leftarrow \text{HILL}(i, R, P, \mathcal{T}, 1)$ 
4:    $(f_3, \sigma_3) \leftarrow \text{HILL}(i, R, P, \mathcal{T}, 4)$ 
5:    $j = \underset{1 \leq j \leq 3}{\text{argmin}} \sigma_j$ 
6:   return  $(f_j, \sigma_j)$ 
    
```

Algorithm 5 Estimation of mass action law kinetics

```

1: function MASS_ACTION( $i, R, P, \mathcal{T}$ )
2:    $K \leftarrow \emptyset$ 
3:   for  $(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}$  do  $K \leftarrow K \cup \left\{ \frac{\prod_{s \in R} \mathbf{x}_s}{\mathbf{d}_i / \mathbf{d}_0} \right\}$ 
4:    $k \leftarrow |1 / \text{mean}(K)|$ 
5:   return  $((R, P, k \cdot \prod_{s \in R} \mathbf{x}_s), \text{stddev}(K))$ 
    
```

Algorithm 6 Estimation of Hill kinetics of order λ

```

1: function HILL( $i, \{j\}, P, \mathcal{T}, \lambda$ )
2:    $V \leftarrow \emptyset$ 
3:    $(\cdot, \mathbf{d}, \cdot) \leftarrow \underset{(\mathbf{x}, \dots) \in \mathcal{T}}{\text{argmax}} \mathbf{x}_i$ 
4:    $V_{max} \leftarrow |\mathbf{d}_i / \mathbf{d}_0|$ 
5:    $(\mathbf{x}, \cdot, \cdot) \leftarrow \underset{(\cdot, \mathbf{e}, \cdot) \in \mathcal{T}}{\text{argmin}} \left| 1 - \left| \frac{V_{max}/2}{\mathbf{e}_i / \mathbf{e}_0} \right| \right|$ 
6:    $K_m \leftarrow \mathbf{x}_j$ 
7:   for  $(\mathbf{x}, \mathbf{d}, \cdot) \in \mathcal{T}$  do  $V \leftarrow V \cup \left\{ \frac{\mathbf{x}_j^\lambda}{K_m^\lambda + \mathbf{x}_j^\lambda} \cdot \frac{\mathbf{d}_0}{\mathbf{d}_i} \right\}$ 
8:    $V_{max} \leftarrow |1 / \text{mean}(V)|$ 
9:   return  $(\{j\}, P, \frac{V_{max} \mathbf{x}_j^\lambda}{K_m^\lambda + \mathbf{x}_j^\lambda}, \text{stddev}(V))$ 
    
```

Algorithm 7 Search for a catalyst molecule

 Parameter: α

```

1: function CATALYST( $R, P, \mathcal{T}$ )
2:   for  $i \in S \setminus (P \cup R)$  do
3:      $R_i \leftarrow R \cup \{i\}$ 
4:      $P_i \leftarrow P \cup \{i\}$ 
5:      $\mathcal{T}_i \leftarrow \{(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T} \mid \mathbf{x}_i > 0\}$ 
6:      $f_i, \sigma_i \leftarrow \text{KINETIC}((R_i, P_i, f_i), \mathcal{T}_i)$ 
7:      $i^* \leftarrow \underset{i}{\text{argmin}} \sigma_i$ 
8:     if  $\sigma_{i^*} > \alpha$  then return  $(R_{i^*}, P_{i^*}, f_{i^*}), \sigma_{i^*}$  else
       return  $(R, P, f), \sigma$ 
    
```

Proposition 1. *The time complexity of the CRN learning algorithm for inferring one reaction is $\mathcal{O}(t \cdot n^2)$ where t is the number of transitions in the traces and n the number of variables.*

Algorithm 8 Mean number of zero values in supporting transition states.

 Parameter: γ

```

1: function SPARSITY( $\mathcal{T}, \mathcal{D}$ )
2:    $n \leftarrow 0, c \leftarrow 0$ 
3:   for  $(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}$  do
4:      $n \leftarrow n + 1$ 
5:     let  $(\mathbf{x}^j(1), \dots, \mathbf{x}^j(d_j)) = \mathcal{D}_j$ 
6:     for  $i \in S$  do
7:       if  $\mathbf{x}_i \leq \gamma \cdot \max_{1 \leq t < d_j} \mathbf{x}_i^j(t)$  then
8:          $c \leftarrow c + 1$ 
9:   return  $c/n$ 
    
```

Algorithm 9 Remove the effects of an inferred reaction in the transitions.

```

1: procedure TRANSITION_UPDATE( $\mathcal{T}, (R, P, f)$ )
2:   for  $(\mathbf{x}, \mathbf{d}, j) \in \mathcal{T}$  s.t.  $\forall i \in R, \mathbf{x}_i > 0$  do
3:     for  $i \in R$  do
4:        $\mathbf{d}_i \leftarrow \mathbf{d}_i + \mathbf{d}_0 \cdot f(\mathbf{x})$ 
5:     for  $i \in P$  do
6:        $\mathbf{d}_i \leftarrow \mathbf{d}_i - \mathbf{d}_0 \cdot f(\mathbf{x})$ 
    
```

Proof. A simple worst-case analysis of algorithms 2, 7, 9 and 8 shows that they are in $\mathcal{O}(n|\mathcal{T}|)$. Rate inference algorithms 5 and 6 are linear in the size of the set of transitions. The reaction inference of Alg. 3 is quadratic in n since there is a potential lookup for pairs of species. Therefore, noting that in the internal loop over r of Alg. 1 there are at most $2|\mathcal{T}|$ transitions in all the supports \mathcal{T}_r , the complexity of this algorithm is in $\mathcal{O}(n^2|\mathcal{T}|)$. \square

4. Evaluation on CRN Simulation Traces

In the context of evaluating the learning algorithm on simulation traces, the hidden CRN used to generate those traces can be used to compare the learned CRN in terms of correct reactions (true positives), wrong reactions (false positives) and missing reactions (false negatives). The F measure is a classical quality measure combining those criteria: $F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$ where $\text{precision} = \frac{\text{tp}}{\text{tp} + \text{fp}}$ and $\text{recall} = \frac{\text{tp}}{\text{tp} + \text{fn}}$ and tp, fp, fn are the numbers of true positive, false positive and false negative inferences respectively. The F score is equal to 1 for a perfect reconstruction of the hidden CRN and equal to 0 if not a single reaction of the hidden CRN is found. It is worth noting also that the F score penalizes the case where one reaction of the hidden CRN lacks in the learned CRN (creation of a False Negative), but it penalizes twice as much the case where the learned CRN contains a reaction inexistent in the hidden CRN while still missing one.

In the following examples, we shall use traces obtained with two different kinds of initial states: first a so-called “canonical trace” where the initial state contains the CRN input species that are necessary for all reactions to occur, and random traces obtained by generating initial states with probability 0.25 for a species to have a zero value, and an exponential distribution law with mean 100 for non zero values. The random initial states are thus relatively sparse with zero values (somehow similar to knock-outs) that prevent the discovery of some reactions but are useful to disambiguate the effects of concurrent reactions.

4.1. Evaluation on Parallel and Product/Reactant-Parallel CRNs




Figure 3. Parallel CRN example — F score obtained for the learned CRN as a function of the time lapse used for subsampling (from 0% to 10% of the time horizon) the number of random initial conditions with zero values, the time horizon of the trace, and the reaction acceptance threshold α . The statistics (blue curve) have been obtained with 10 runs using traces from perturbed initial states with random zeroes. The results with the canonical trace where A, C, E are initially present, are figured on the red curve.

Fig. 3 shows on a small example of parallel CRN the sensitivity of the F score to, respectively, the level of trace subsampling, the number of traces with random initial conditions with zeroes, the length of the traces, and the reaction acceptance threshold α of the algorithm. In the first plot, we observe not surprisingly that the quality of the learned CRN is negatively impacted by increasing the subsampling steps. This is due to a less accurate estimation of the rate functions. The second curve shows that the number of initial conditions with zeroes in the dataset of random traces increases




Figure 4. Product-parallel CRN example — sensitivity of the F score to the learning parameters (subsampling, number of random initial conditions, length, acceptance).

the score of the learned CRN, while the time horizon of the traces has no impact. This is consistent with the space-time trade-off observed with PAC learning for Boolean regulatory networks (Carcano et al., 2017). The fourth curves draw the score of the learned network as a function of α , the statistical threshold of acceptance of a reaction based on its kinetics standard deviation. A high value of α allows a reaction to be selected even if its score is low.

On a single trace of a non parallel CRNs, it is important that the reaction inference algorithm tries to decompose the variations of products and reactants in several reactions (two in Alg. 3). The example in Fig. 4 is a slight modification of the previous example making it non reactant-parallel with C , but still product-parallel. The algorithm is still able to recover the hidden CRN from the canonical trace, with qualitatively similar sensitivity results.

Fig. 5 shows a reactant-parallel example on which the algorithm is similarly able to recover the structure from the canonical trace, or from random traces.

4.2. Evaluation on the Chain CRN

Example 1 of a chain is a non-parallel CRN for which we are able to recover the full structure of the chain CRN using only one trace, the canonical trace with the input A present. First, reaction $D \Rightarrow E$ is inferred based on the last non zero variations of the trace. Its effect on the variations is then suppressed and the filtering applied. The fact that species D 's variations were updated allows for reaction $C \Rightarrow D$

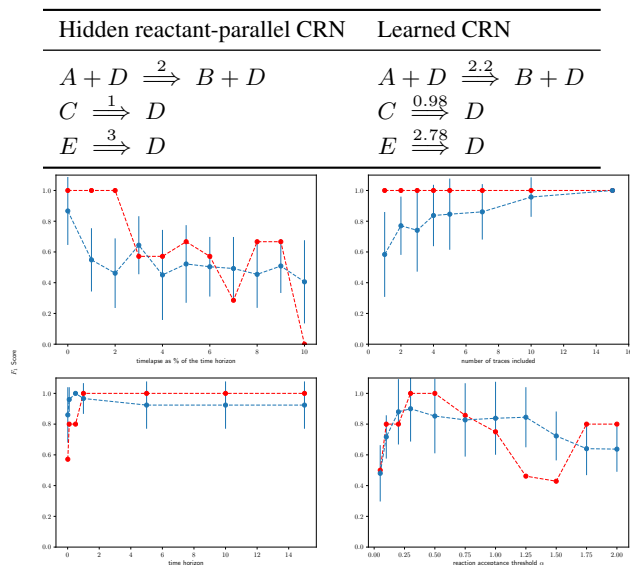


Figure 5. Reactant-parallel CRN — F score versus learning parameters (subsampling, initial conditions, length, acceptance).

to be uncovered, and so on. The sensitivity results, depicted in Fig. 6 are quite the same as for the other models, except when the time length of the traces varies. Unlike for parallel models, some reactions appear in a significant manner (and the state variations they induced are not filtered) only after a certain amount of time. The curve suggests that a time horizon $T \approx 5$ is enough to start having good results, This is due to the chain structure of the CRN, as can be seen on the trace displayed in Fig. 1.




Figure 6. Chain CRN — F score versus learning parameters (subsampling, initial conditions, length, acceptance).

5. Evaluation on a Yeast Cell Cycle Model

On the simulation trace depicted in Fig. 7 of the yeast cell cycle model of Tyson (Tyson, 1991), our algorithm infers the reactions given in Table 2. The discrepancies concerning the synthesis reaction of the cyclin can be very well explained by the existence of multiple time scales in this example. Indeed when it is produced the Cyclin is immediately complexed with Cdc and phosphorylated by very fast reactions. Therefore the free state of the Cyclin cannot be observed and what is inferred is the synthesis of the fast equilibrium state of the complexed Cyclin. On the other hand, the autocatalysis reaction cannot be recovered since we do not consider stoichiometric coefficient.




Figure 7. Numerical simulation of the yeast cell cycle model of (Tyson, 1991).

Hidden CRN	Learned CRN
$\emptyset \xrightarrow{0.015} cy$	$\emptyset \xrightarrow{0.66} cy1 + cdcy2$
$cy + cd1 \xrightarrow{200} cdcy2$	$\emptyset \xrightarrow{0.01} cdcy2$
$cdc2 \xrightarrow{0.018} cdcy1$	$cdc2 \xrightarrow{0.1152} cdcy1$
$cdc2 + 2 * cdcy1 \xrightarrow{180} 3 * cdcy1$	$cdc2 \xrightarrow{0.05} cy1$
$cdc2 + cdcy1 \xrightarrow{1} cy1 + cd$	$cdc2 \xrightarrow{1.62} \emptyset$
$cy1 \xrightarrow{0.6} \emptyset$	$cy1 \xrightarrow{0.4} cdcy1$
$cd1 \xrightarrow{100} cd$	$cd1 \xrightarrow{11259} cd$
$cd \xrightarrow{10000} cd1$	$cd \xrightarrow{5912} cd1$

Table 2. Cell cycle model of Tyson (Tyson, 1991) and learned CRN from the canonical trace where cd is present, $\alpha = 5$.

6. Evaluation on Real Videomicroscopy Data

In this section we consider data-time series obtained by biological experiments on NIH3T3 embryonic fibroblast cells. These data have been used to develop a coupled model of the cell cycle and the circadian clock in this cell

line in (Traynard et al., 2016). The reported experiments have been done using time lapse videomicroscopy and cell tracking using different fluorescent reporters for the cell cycle and the circadian clock observed during 72 hours in proliferating NIH3T3 embryonic mouse fibroblasts (Feillet et al., 2014). This cell line was modified to include three fluorescent markers of the circadian clock and the cell cycle: the RevErb- α :Venus clock gene reporter (Nagoshi et al., 2004) for measuring the expression of the circadian protein RevErb- α , and the Fluorescence Ubiquitination Cell Cycle Indicators (FUCCI), Cdt1 and Geminin, two cell cycle proteins which accumulate during the G1 and S/G2/M phases respectively, for measuring the cell cycle phases (Sakaue-Sawano et al., 2008). The cells were left to proliferate in regular medium supplemented with different concentrations of FBS (20% in this data set). Long-term recording was performed in constant conditions with one image taken every 15 minutes during 72 hours.




Figure 8. Two videomicroscopy traces from the dataset of (Feillet et al., 2014) about cell cycle progression and circadian clocks markers in embryonic NIH3T3 fibroblasts over 3 days.

Fig. 8 shows two traces of this dataset, showing the high variability of the cell behaviors. On the first trace, the algorithm with $\alpha = 1.5$ infers the following reactions with recourse to Michaelis-Menten kinetics:

Learned CRN on the first trace	Rate Functions
$G1 \Rightarrow G2$	$53.66 \frac{G1}{G1 + 11.6}$
$RevErb\alpha \Rightarrow G1$	$33.42 \frac{RevErb\alpha}{RevErb\alpha + 43.7}$
$G1 \Rightarrow \emptyset$	$133.62 \frac{G1}{G1 + 150.4}$

On the second trace, the reactions inferred are

Learned CRN on second trace	Rate Functions
$RevErb\alpha \Rightarrow G1$	$31.9 \frac{RevErb\alpha}{RevErb\alpha + 45.89}$
$G2 \Rightarrow \emptyset$	$21.39 \frac{G2}{G2 + 44.82}$

Interestingly, on the whole dataset of 91 tracked cells and 26000 observed state transitions, the algorithm still manages to infer meaningful reactions:

Learned CRN on all dataset	Rate Functions
$G1 \Rightarrow G2$	$7.1 \frac{G1}{G1 + 3.68}$
$RevErb\alpha \Rightarrow G1$	$22.56 \frac{RevErb\alpha}{RevErb\alpha + 71.45}$
$G1 \Rightarrow \emptyset$	$5.96 \frac{G1}{G1 + 5.0}$
$G2 \Rightarrow \emptyset$	$54.84 \frac{G2}{G2 + 176.23}$

7. Conclusion

The statistical learning algorithm for learning CRNs from time series data presented in this paper provides encouraging results. It is a greedy algorithm in the sense that at each step the reaction candidate with the best score among the observed transitions is inferred and is never reconsidered. The choice of the first reaction inference is thus crucial to the global result of the learning algorithm. This choice is facilitated when the traces contain regimes where some molecular concentrations are quasi-stable, like in the chain example, or in traces obtained from initial states with many zeroes. This is not the case for oscillating trace which are thus more challenging, as illustrated with the simple model of the yeast cell cycle of (Tyson, 1991), and with the real traces obtained by videomicroscopy of biomarkers of the cell cycle and circadian clock of (Feillet et al., 2014).

One crucial feature of our algorithm is the inference of the reaction structure together with a rate function with kinetic parameters estimated on all the transitions where the reaction can occur. The standard deviation of the reaction rates is used to score the reaction candidates and choose to infer the reaction with the best score. This process is iterated after removing the changes due to the inferred reaction in the traces before the computation of the next reaction candidates. One way of possible improvement would be to update the estimation of the kinetic parameters of the previously inferred reactions when a new non parallel reaction is inferred.

The application of our algorithm to the yeast cell cycle model of Tyson which contains multiple time scales necessitated to increase the acceptance parameter of our algorithm, and to content ourselves with the inference of reactions corresponding to the slow timescale, the fast reactions being just not observable. The application to real traces coming from particularly noisy videomicroscopy data on mammalian cells showing moreover very high variability, was very ambitious and necessitated to similarly adapt the threshold parameters to obtain better results. It is remarkable that meaningful reactions coupling the cell cycle progression with the circadian clock marker could be inferred with this dataset. More work is however needed to analyze the behaviour of our algorithm on such highly variable traces and find strategies to automatically adapt the threshold parameters of the algorithm to the quality of the traces.

References

- Carcano, A., Fages, F., and Soliman, S. Probably Approximately Correct Learning of Regulatory Networks from Time-Series Data. In *CMSB'17: Proceedings of the fifteenth international conference on Computational Methods in Systems Biology*, volume 10545, pp. 74–90, May 2017. doi: 10.1007/978-3-319-67471-1_5.
- Choi, K., Hellerstein, J., Wiley, H. S., and Sauro, H. M. Inferring reaction networks using perturbation data. *Bio arXiv*, 2018.
- Fages, F. and Soliman, S. Abstract interpretation and types for systems biology. *Theoretical Computer Science*, 403(1):52–70, 2008. doi: 10.1016/j.tcs.2008.04.024.
- Feillet, C., Krusche, P., Tamanini, F., Janssens, R. C., Downey, M. J., Martin, P., Teboul, M., Saito, S., Lévi, F., Bretschneider, T., van der Horst, G. T. J., Delaunay, F., and Rand, D. A. Phase locking and multiple oscillating attractors for the coupled mammalian clock and cell cycle. *Proceedings of the National Academy of Sciences of the United States of America*, 111(27):9928–9833, 2014. ISSN 1091-6490. doi: 10.1073/pnas.1320474111.
- Ganscha, S., Fortuin, V., Horn, M., Arvaniti, E., and Claassen, M. Supervised learning on synthetic data for reverse engineering gene regulatory networks from experimental time-series. *Bio arXiv*, 2018.
- Nagoshi, E., Saini, C., Bauer, C., Laroche, T., Naef, F., and Schibler, U. Circadian gene expression in individual fibroblasts: cell-autonomous and self-sustained oscillators pass time to daughter cells. *Cell*, 119:693–705, 2004.
- Ostrowski, M., Paulevé, L., Schaub, T., Siegel, A., and Guziolowski, C. Boolean network identification from perturbation time series data combining dynamics abstraction and logic programming. *Biosystems*, 149:139–153, 2016. ISSN 0303-2647. doi: 10.1016/j.biosystems.2016.07.009.
- Pearl, J. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, NY, USA, 2nd edition, 2009.
- Preacher, K. J. and Hayes, A. F. Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, 40(3):879–891, Aug 2008. doi: 10.3758/BRM.40.3.879.
- Sakaue-Sawano, A., Kurokawa, H., Morimura, T., Hanyu, A., Hama, H., Osawa, H., Kashiwagi, S., Fukami, K., Miyata, T., Miyoshi, H., Imamura, T., Ogawa, M., Masai, H., and Miyawaki, A. Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. *Cell*, 132(3):487–498, Feb 2008.
- Traynard, P., Feillet, C., Soliman, S., Delaunay, F., and Fages, F. Model-based investigation of the circadian clock and cell cycle coupling in mouse embryonic fibroblasts: Prediction of reverb-alpha up-regulation during mitosis. *Biosystems*, 149:59–69, November 2016. doi: 10.1016/j.biosystems.2016.07.003.
- Triantafyllou, S., Lagani, V., Heinze-Deml, C., Schmidt, A., Tegner, J., and Tsamardinos, I. Predicting causal relationships from biological data: Applying automated causal discovery on mass cytometry data of human immune cells. *Scientific Reports*, 7:12724, 2017. doi: 10.1038/s41598-017-08582-x.
- Tyson, J. J. Modeling the cell division cycle: cdc2 and cyclin interactions. *Proceedings of the National Academy of Sciences*, 88(16):7328–7332, August 1991.