



HAL
open science

Inferring gene regulatory networks from single-cell data: a mechanistic approach

Ulysse Herbach, Arnaud Bonnaffoux, Thibault Espinasse, Olivier Gandrillon

► To cite this version:

Ulysse Herbach, Arnaud Bonnaffoux, Thibault Espinasse, Olivier Gandrillon. Inferring gene regulatory networks from single-cell data: a mechanistic approach. *BMC Systems Biology*, 2017, 11 (1), pp.105. 10.1186/s12918-017-0487-0 . hal-01646910

HAL Id: hal-01646910

<https://hal.science/hal-01646910>

Submitted on 25 Nov 2017

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.

Inferring gene regulatory networks from single-cell data: a mechanistic approach

Ulysse Herbach^{1,2,3*}, Arnaud Bonnaffoux^{1,2,4}, Thibault Espinasse³, Olivier Gandrillon^{1,2}

¹ Univ Lyon, ENS de Lyon, Univ Claude Bernard, CNRS UMR 5239, INSERM U1210, Laboratory of Biology and Modelling of the Cell, 46 allée d'Italie Site Jacques Monod, F-69007 Lyon, France

² Inria Team Dracula, Inria Center Grenoble Rhône-Alpes, France

³ Univ Lyon, Université Claude Bernard Lyon 1, CNRS UMR 5208, Institut Camille Jordan, 43 blvd. du 11 novembre 1918, F-69622 Villeurbanne cedex, France

⁴ The CoSMo company, 5 passage du Vercors, 69007 Lyon, France

Abstract

The recent development of single-cell transcriptomics has enabled gene expression to be measured in individual cells instead of being population-averaged. Despite this considerable precision improvement, inferring regulatory networks remains challenging because stochasticity now proves to play a fundamental role in gene expression. In particular, mRNA synthesis is now acknowledged to occur in a highly bursty manner. We propose to view the inference problem as a fitting procedure for a mechanistic gene network model that is inherently stochastic and takes not only protein, but also mRNA levels into account. We first explain how to build and simulate this network model based upon the coupling of genes that are described as piecewise-deterministic Markov processes. Our model is modular and can be used to implement various biochemical hypotheses including causal interactions between genes. However, a naive fitting procedure would be intractable. By performing a relevant approximation of the stationary distribution, we derive a tractable procedure that corresponds to a statistical hidden Markov model with interpretable parameters. This approximation turns out to be extremely close to the theoretical distribution in the case of a simple toggle-switch, and we show that it can indeed fit real single-cell data. As a first step toward inference, our approach was applied to a number of simple two-gene networks simulated *in silico* from the mechanistic model and satisfactorily recovered the original networks. Our results demonstrate that functional interactions between genes can be inferred from the distribution of a mechanistic, dynamical stochastic model that is able to describe gene expression in individual cells. This approach seems promising in relation to the current explosion of single-cell expression data.

1 Introduction

Inferring regulatory networks from gene expression data is a longstanding question in systems biology [1], with an active community developing many possible solutions. So far, almost all studies have been based on population-averaged data, which historically used to be the only possible way to observe gene expression. Technologies now allow us to measure mRNA levels in individual cells [2–4], a revolution in terms of precision. However, the network reconstruction task paradoxically remains more challenging than ever.

The main reason is that the variability in gene expression unexpectedly stands at a large distance from a trivial, limited perturbation around the population mean. It is now clear indeed that this variability can have functional significance [5–7] and should therefore not be ignored when

*ulyse.herbach@ens-lyon.fr

dealing with gene network inference. In particular, as the mean is not sufficient to account for a population of cells, a deterministic model – e.g. ordinary differential equation (ODE) systems, often used in inference [8,9] – is unlikely to faithfully inform about an underlying gene regulatory network. Whether such a deterministic approach could still be a valid approximation or not is a difficult question that may require some biological insight into the system under consideration [10]. Another key aspect when considering individual cells is that they generally have to be killed for measurements: from a statistical point of view, temporal single-cell data therefore should not be seen as a set of time series, but rather *snapshots*, i.e. independent samples from a time series of distributions.

On the other hand, single-cell data give the opportunity of moving one step further toward a more accurate physical description of gene expression. Molecular processes of gene expression are overall now well understood, in particular transcription, but precisely how stochasticity emerges is still somewhat of a conundrum. Harnessing variability in single-cell data is expected to allow for the identification of critical parameters and also to provide hints about the basic molecular processes involved [11,12]. Moreover, the variability arising from perturbations in cell populations is often crucial for network reconstruction to succeed [13,14] as the deterministic inference problem suffers from intrinsic limitations [15]. From this point of view, the same information is expected to be contained in the variability between cells in single-cell data. Some of the few existing single-cell inference methods follow this path, for example using asynchronous Boolean network models [16] or generating pseudo time series [9,17]. In this article, we use a mechanistic approach in the sense that every part of our model has an explicit physical interpretation. Importantly, mRNA observations are not used as a proxy for proteins since both are explicitly modeled.

Besides, mechanistic models that are accurate enough to describe gene expression at the single-cell level usually do not consider interactions between genes. For example, the so-called “two-state” (aka random telegraph) model has been successfully used with single-cell RNA-seq data [18], but the joint distribution of a set of genes contains much more information than the marginal kinetics of individual genes: our aim is to exploit this information while keeping the mechanistic point of view.

Namely, we propose to view the inference as a fitting procedure for a mechanistic gene network model. Whereas the goal here is not to achieve global predictability performances (e.g. as in [19]), our framework makes it possible to explicitly implement many biological hypotheses, and to test them by going back and forth between simulations and experiments. The main point of this article is to show that a tractable statistical model for network inference from single-cell data can be derived through successive relevant approximations. Finally, we demonstrate that our approach is capable of extracting enough information out of in silico-simulated noisy single-cell data to correctly infer the structures of various two-gene networks.

2 Methods

In this part, we aim at deriving a tractable statistical model from a mechanistic one. We will use the two-state model for gene expression to build a “network of two-state models” by making the promoter switching rates depend on protein levels. Then, successive relevant simplifications will lead to an explicit approximation of a statistical likelihood.

2.1 A simple mechanistic model for gene regulatory networks

2.1.1 Basic block: stochastic expression of a single gene

Our starting point is the well-known two-state model of gene expression [20–23], a refinement of the model introduced by [24] from pioneering single-cell experiments [25]. In this model, a gene is described by its promoter which can be either active (on) or inactive (off) – possibly representing a transcription complex being “bound” or “unbound” but it may be more complicated [26] – with mRNA being transcribed only during the active periods. Translation is added in a standard way, each mRNA molecule producing proteins at a constant rate. The resulting model (Fig. 1) can be entirely defined by the set of chemical reactions detailed in Table 1, where chemical species G , G^* , M and P respectively denote the inactive promoter, the active promoter, the amount of mRNA and proteins. The mathematical framework generally assumes stochastic mass-action kinetics [27] for all reactions, since they typically involve few molecules compared to Avogadro’s number. In this fully discrete setting, one can use the master equation to compute stationary distributions: for mRNA the exact distribution is a Beta-Poisson mixture [28], and an approximation is available for proteins when they degrade much more slowly than mRNA [29]. In addition, the time-dependent generating function of mRNA is known in closed form [30] and can be inverted in some cases to obtain the transient distribution [28].

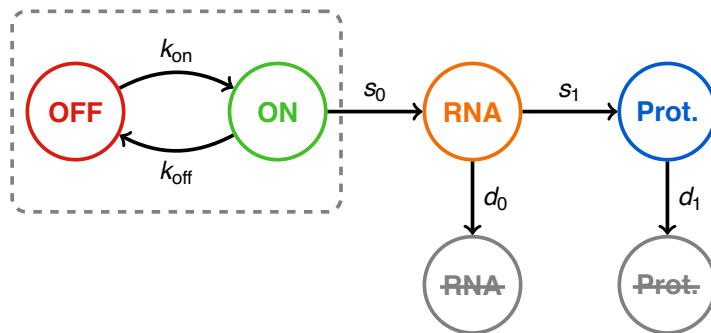


Figure 1: Scheme of the two-state model of gene expression, basic block of our network model.

Reaction	Rate constant	Interpretation
$G \rightarrow G^*$	k_{on}	gene activation
$G^* \rightarrow G$	k_{off}	gene inactivation
$G^* \rightarrow G^* + M$	s_0	transcription
$M \rightarrow M + P$	s_1	translation
$M \rightarrow \emptyset$	d_0	mRNA degradation
$P \rightarrow \emptyset$	d_1	protein degradation

Table 1: Chemical reactions defining the two-state model. The rate constants are usually abbreviated to *rates* as they correspond to actual reactions rates when only one molecule of reactant is present. In the stochastic setting, these rates are in fact propensities, i.e. probabilities per unit of time.

In practice, the formulas involve hypergeometric series that are not straightforward to use in a statistical inference framework. Besides, these series essentially arise from the fact that such a discrete model has to enumerate all potential collisions between molecules (the stochastic

mass-action assumption in the master equation). It is therefore natural to consider keeping only the most important source of noise, that is, keeping a molecular representation for rare species but describing abundant species at a higher level where molecular noise averages out to continuous quantities. A quick look at reactions in Table 1 indicates that the only rare species are G and G^* , with quantities $[G]$ and $[G^*]$ being equal to 0 or 1 molecule and satisfying the conservation relation $[G] + [G^*] = 1$. The other two, M and P , are not conserved quantities in the model and reach a much wider range in biological situations [31], meaning that saturation constants s_0/d_0 and s_1/d_1 are much larger than 1 molecule.

Hence, letting $E(t)$, $M(t)$ and $P(t)$ denote the respective quantities of G^* , M and P at time t , we consider a hybrid version of the previous model, where E has the same stochastic dynamics as before, but with M and P now following usual rate equations:

$$\begin{cases} E(t) : 0 \xrightarrow{k_{\text{on}}} 1, & 1 \xrightarrow{k_{\text{off}}} 0 \\ M'(t) = s_0 E(t) - d_0 M(t) \\ P'(t) = s_1 M(t) - d_1 P(t) \end{cases} \quad (1)$$

This system simply switches between two ordinary differential equations, depending on the value of the two-state continuous-time Markov process $E(t)$, making it a Piecewise-Deterministic Markov Process (PDMP) [32]. From a mathematical perspective, model (1) rigorously approximates the original molecular model when s_0/d_0 and s_1/d_1 are large enough [33, 34] and interestingly, it has already been implicitly considered in the biological literature [22, 23]. Note also that the stationary distribution of mRNA is a scaled Beta distribution that is exactly the one of the Beta-Poisson mixture in the discrete model [28]. Similarly to a recent approach for a two-gene toggle switch [35], we will use (1) as a basic building block for gene networks.

When both $k_{\text{on}} \ll k_{\text{off}}$ and $d_0 \ll k_{\text{off}}$, mRNA is transcribed by *bursts*, i.e. during short periods which make the mRNA quantity stay far from saturation. Hence, the amount transcribed within each burst is approximately proportional to the burst duration, whose mean is $1/k_{\text{off}}$ by definition: this justifies the quantity s/k_{off} often being called “burst size” or “burst amplitude”. Furthermore, promoter active periods are much shorter than inactive ones so they can be seen as instantaneous, justifying the name “burst frequency” for the inverse of the mean inactive time k_{on} . We place ourselves in this situation as it often occurs in experiments [22, 23, 36–38]. Note however that these two notions are not clearly defined when relations $k_{\text{on}} \ll k_{\text{off}}$ and $d_0 \ll k_{\text{off}}$ do not hold.

2.1.2 Adding interactions between genes: the network model

Now considering a given set of n genes, a natural way of building a network is to assume that each gene i produces specific mRNA M_i and protein P_i , and to define a version of model (1) with its own parameters:

$$\begin{cases} E_i(t) : 0 \xrightarrow{k_{\text{on},i}} 1, & 1 \xrightarrow{k_{\text{off},i}} 0 \\ M_i'(t) = s_{0,i} E_i(t) - d_{0,i} M_i(t) \\ P_i'(t) = s_{1,i} M_i(t) - d_{1,i} P_i(t) \end{cases} \quad (2)$$

Still, genes have static parameters and do not interact with each other. To get an actual network, we need to go one step further: reactions $G_i \rightarrow G_i^*$ and $G_i^* \rightarrow G_i$ are not assumed to be elementary anymore, but rather represent complex reactions involving proteins so that promoter parameters $k_{\text{on},i}$ and $k_{\text{off},i}$ now depend on proteins (Fig. 2a), and a fortiori on time. Our network

model will correspond to the explicit definition, for all gene i , of functions $k_{\text{on},i}(P_1, \dots, P_n)$ and $k_{\text{off},i}(P_1, \dots, P_n)$. These functions shall also depend on network-specific parameters quantifying the interactions, thus making the link between “fitting a chemical model” and “inferring a network”. As a toy example, consider replacing $G_i \rightarrow G_i^*$ with two parallel elementary reactions



for which applying the law of mass action directly gives $k_{\text{on},i}(P_1, \dots, P_n) = \theta_{i,0} + \theta_{i,j}P_j$. In a regulatory network (Fig. 2b), it would correspond to adding a directed edge from gene j to gene i , with $\theta_{i,0}$ the basal parameter of gene i , and $\theta_{i,j}$ the strength of activation of gene i by protein P_j . We emphasize that the action of P_j on the promoter G_i is not necessarily direct. For example, P_j can instead indirectly modulate the amount/activity of a transcription factor: we suppose in this article that such hidden reactions are fast enough regarding gene expression dynamics so that protein P_j is a relevant proxy for the transcription factor. Moreover, although we assume here that interactions can only happen at the level of $k_{\text{on},i}$ and $k_{\text{off},i}$, mainly for identifiability purposes, it is also possible to make $d_{1,i}$ and $s_{1,i}$ depend on proteins without fundamentally changing the mathematical approach (e.g. see [39, 40]).

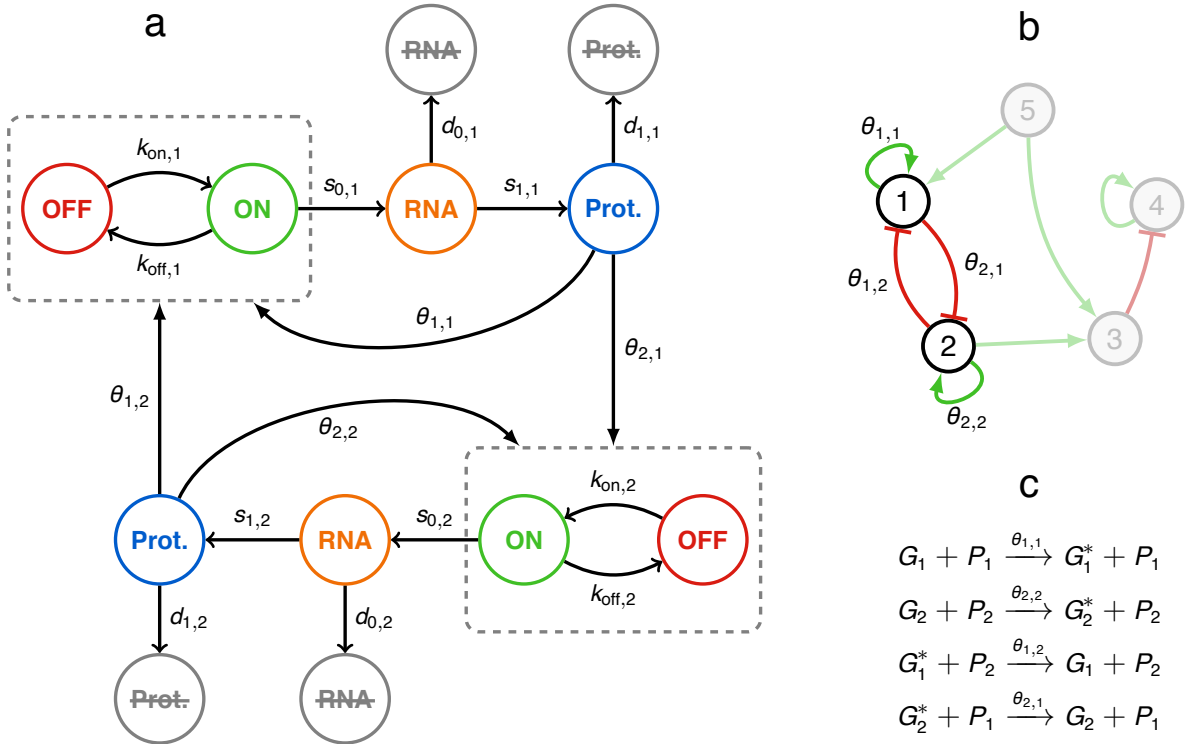


Figure 2: (a) Two genes interacting with each other, forming a network. Interactions are assumed to arise from the dependence of promoter dynamics on protein quantities. (b) A higher level of abstraction leads to the traditional gene regulatory network representation. (c) A toy example of reactions defining the interactions between genes 1 and 2, making the link between representations (a) and (b).

In order to simplify notations, we normalize model (2) into a dimensionless equivalent model: we rewrite it in terms of new variables $\bar{M}_i = \frac{d_{0,i}}{s_{0,i}} M_i$ and $\bar{P}_i = \frac{d_{0,i} d_{1,i}}{s_{0,i} s_{1,i}} P_i$, which have values between 0 and 1, and report this scale change in the definition of $k_{\text{on},i}$ and $k_{\text{off},i}$ (see Appendix A for details). In the remainder of this article, the new variables will still be denoted by M_i and P_i as

no confusion arises. The resulting normalized model is:

$$\begin{cases} E_i(t) : 0 \xrightarrow{k_{\text{on},i}} 1, \quad 1 \xrightarrow{k_{\text{off},i}} 0 \\ M_i'(t) = d_{0,i} (E_i(t) - M_i(t)) \\ P_i'(t) = d_{1,i} (M_i(t) - P_i(t)) \end{cases} \quad (4)$$

still omitting the dependence of $k_{\text{on},i}$ and $k_{\text{off},i}$ on $(P_1(t), \dots, P_n(t))$ for clarity. This form enlightens the fact that $s_{0,i}$ and $s_{1,i}$ are just scaling constants: given a path $(E_i, M_i, P_i)_i$ of system (4), one can go back to the physical path by simply multiplying M_i by $(s_{0,i}/d_{0,i})$ and P_i by $(s_{0,i}/d_{0,i}) \times (s_{1,i}/d_{1,i})$.

Therefore, we get a general network model where each link between two genes is directed and has an explicit biochemical interpretation in terms of molecular interactions. The previous example is very simplistic but one can use virtually any model of chromatin dynamics to derive a form for $k_{\text{on},i}$ and $k_{\text{off},i}$, involving hit-and-run reactions, sequential binding, etc. [41]. Such aspects are still far from being completely understood [42–45] and this simple network model can hopefully be used to assess biological hypotheses. In the next part, we will introduce a more sophisticated interaction form based on an underlying probabilistic model, which is both “statistics-friendly” and interpretable as a non-equilibrium steady state of chromatin environment [43].

2.1.3 Some known mathematical results

Thanks to some recent theoretical results [40, 46], simple sufficient conditions on $k_{\text{on},i}$ and $k_{\text{off},i}$ ensure that the PDMP network model (4) is actually well-defined and that the overall joint distribution of $(E_i, M_i, P_i)_i$ converges as $t \rightarrow +\infty$ to a unique stationary distribution, which will be the basis of our statistical approach. Namely, we assume in this article that $k_{\text{on},i}$ and $k_{\text{off},i}$ are continuous functions of (P_1, \dots, P_n) and that they are greater than some positive constants. These conditions are satisfied in most interesting cases, including the above toy example (3) when $\theta_{i,0} > 0$.

Contrary to creation rates $s_{0,i}$ and $s_{1,i}$, degradation rates $d_{0,i}$ and $d_{1,i}$ play a crucial role in the dynamics of the system. Intuitively, the ratios $(k_{\text{on},i} + k_{\text{off},i})/d_{0,i}$ and $d_{0,i}/d_{1,i}$ respectively control the buffering of promoter noise by mRNA and the buffering of mRNA noise by proteins. A common situation is when promoter and mRNA dynamics are fast compared to proteins, i.e. when $d_{0,i} \gg d_{1,i}$ with $(k_{\text{on},i} + k_{\text{off},i})/d_{0,i}$ fixed. At the limit, the promoter-mRNA noise is fully averaged by proteins and model (4) simplifies into a deterministic system [47]:

$$P_i'(t) = d_{1,i} \left(\frac{k_{\text{on},i}(\mathbf{P}(t))}{k_{\text{on},i}(\mathbf{P}(t)) + k_{\text{off},i}(\mathbf{P}(t))} - P_i(t) \right) \quad (5)$$

where $\mathbf{P}(t) = (P_1(t), \dots, P_n(t))$. The diffusion limit, which keeps a residual noise, can also be rigorously derived [48]. Unsurprisingly, one recovers the traditional way of modelling gene regulatory networks with Hill-type interaction functions. Equation (5) is useful to get an insight into the behaviour of the system (4) for given $k_{\text{on},i}$ and $k_{\text{off},i}$, yet it should be used with caution. Indeed, the $d_{0,i}/d_{1,i}$ ratio has been shown to span a high range, averaging out to the value $d_{0,i}/d_{1,i} \approx 5$ in mammalian cells [31], for which taking the limit $d_{0,i} \gg d_{1,i}$ is not obvious. This is consistent with recent single-cell experiments showing a high variability of both mRNA and protein levels between cells [37]. In that sense, the PDMP model is much more robust than its deterministic/diffusion counterpart while keeping a similar level of mathematical complexity, which motivates our approach.

2.1.4 Simulation

We propose a simple algorithm to compute sample paths of our stochastic network model (4). It consists in a hybrid version of a basic ODE solver, making it efficient enough to perform massive simulations on large scale networks involving arbitrary numbers of molecules, which would be intractable with a classic molecule-based model (Fig. 3). The deterministic part of the algorithm is a standard explicit Euler scheme, while the stochastic part is based on the transient promoter distribution for single genes: this can be justified by the fact that during a small enough time interval, proteins remain almost constant so genes behave as if $k_{\text{on},i}$ and $k_{\text{off},i}$ were constant. We therefore use Bernoulli steps, in a similar way of a diffusion being simulated using gaussian steps.

After discretizing time with step δt , the numerical scheme is as follows. Starting from an initial state $(E_i^0, M_i^0, P_i^0)_i$, the update of the system from t to $t + \delta t$ is given by:

$$\begin{cases} E_i^{t+\delta t} \sim \mathcal{B}(\pi_i^t) \\ M_i^{t+\delta t} = (1 - d_{0,i}\delta t)M_i^t + d_{0,i}\delta t E_i^t \\ P_i^{t+\delta t} = (1 - d_{1,i}\delta t)P_i^t + d_{1,i}\delta t M_i^t \end{cases} \quad (6)$$

where the Bernoulli distribution parameter π_i^t is derived by locally solving the master equation for the promoter [49], i.e.

$$\pi_i^t = \frac{a_i^t}{a_i^t + b_i^t} + \left(E_i^t - \frac{a_i^t}{a_i^t + b_i^t} \right) e^{-(a_i^t + b_i^t)\delta t}$$

with the notation $a_i^t = k_{\text{on},i}(P_1^t, \dots, P_n^t)$ and $b_i^t = k_{\text{off},i}(P_1^t, \dots, P_n^t)$. Intuitively, the algorithm is valid when $\delta t \ll 1/\max_i \{K_{\text{on},i}, K_{\text{off},i}, d_{0,i}, d_{1,i}\}$ where $K_{\text{on},i}$ and $K_{\text{off},i}$ denote the maximum values of functions $k_{\text{on},i}$ and $k_{\text{off},i}$.




Figure 3: Simulations of the two-state model for a single gene. (a) Sample path of the PDMP model using our hybrid numerical scheme (computation time ≈ 0.05 s). (b) Sample path of the classic model using exact stochastic simulation [27] (computation time ≈ 10 s). Parameters values are $k_{\text{on}} = 0.34$, $k_{\text{off}} = 10$, $s_0 = 10^3$, $s_1 = 10$, $d_0 = 0.5$ and $d_1 = 0.1$ (in h^{-1}).

2.2 Deriving a tractable statistical model

We will now adopt a statistical perspective in order to deal with gene network inference, considering a set of observed cells. If they are evolving in the same environment for a long enough time, we can reasonably assume that their mRNA and protein levels follow the stationary distribution of an underlying gene network: this distribution can be used as a statistical likelihood for the cells. Furthermore assuming no cell-cell interactions (which may of course depend on the experimental context), we obtain a standard statistical problem with independent samples. Since the stationary distribution of the stochastic network model (4) is well-defined but a priori not analytically tractable, we will derive an explicit approximation and then reduce our inference problem to a traditional likelihood-based estimation. We will do so in two cases: when there is no self-interaction, and for a specific form of auto-activation.

2.2.1 Separating mRNA and protein timescales

It is for the moment very rare to experimentally obtain the amount of proteins for many genes at the single-cell level. We will therefore assume here that only mRNAs are observed. To deal with this problem, we take the protein timescale as our reference by fixing $d_{1,i}$ and assume that promoter dynamics are faster than proteins, i.e. $(k_{\text{on},i} + k_{\text{off},i}) \gg d_{1,i}$ in a biologically relevant way, say $(k_{\text{on},i} + k_{\text{off},i})/d_{1,i} > 10$ (thus the deterministic limit (5) does not necessarily hold). Furthermore, in line with several recent experiments [37, 50], we assume that $d_{0,i}$ is sufficiently larger than $d_{1,i}$ so that the correlation between mRNAs and proteins produced by the gene is very small: model (4) then can be reduced by removing mRNA and making proteins directly depend on the promoters (see Appendix A). The result is

$$\begin{cases} E_i(t) : 0 \xrightarrow{k_{\text{on},i}} 1, & 1 \xrightarrow{k_{\text{off},i}} 0 \\ P_i'(t) = d_{1,i} (E_i(t) - P_i(t)) \end{cases} \quad (7)$$

which still admits the deterministic limit (5). Since mRNA dynamics are faster than proteins, one can also assume that, given protein levels $\mathbf{P} = (P_1, \dots, P_n)$, each mRNA level M_i follows the quasi-steady state distribution

$$M_i | \mathbf{P} \sim \text{Beta} \left(\frac{k_{\text{on},i}(\mathbf{P})}{d_{0,i}}, \frac{k_{\text{off},i}(\mathbf{P})}{d_{0,i}} \right) \quad (8)$$

corresponding to the single-gene model [28, 39] with constant parameters $k_{\text{on},i}(\mathbf{P})$ and $k_{\text{off},i}(\mathbf{P})$. Numerically, this approximation works well even for moderate values of $d_{0,i}$, such as $d_{0,i} = 5 \times d_{1,i}$ (see the Results section).

Biologically, equations (7) and (8) suggest that correlations between mRNA levels may not directly arise from correlations between promoters *states* (which in fact are weak because of $(k_{\text{on},i} + k_{\text{off},i}) \gg d_{1,i}$), but rather originate from correlations between promoter *parameters* $k_{\text{on},i}$ and $k_{\text{off},i}$, which themselves depend on the protein joint distribution.

Table 2 sums up the successive modelling steps introduced so far. From now on, we will always assume the form (8) for the mRNA distribution, and thus our model is reduced to equation (7) which only involves proteins.

1	Single-gene, discrete [29] ◇ All molecules are discrete ◇ mRNA distribution: Beta-Poisson	↓ Abundant species treated continuously
2	Single-gene, PDMP (1) ◇ Only the promoter is discrete ◇ mRNA distribution: Beta	
3	Network (2), normalized version (4) ◇ Both accurate and fast to simulate ◇ mRNA distribution: unknown	↓ Introduction of interactions via $k_{\text{on}}, k_{\text{off}}$
4	Simplified network (7) ◇ mRNA is removed from the network ◇ Conditional mRNA distribution: Beta (8)	↓ Timescale separation of Protein/mRNA ($d_0 \gg d_1$)

Table 2: Successive dynamical models introduced in this article, recalling for each step the main feature and the form of the mRNA stationary distribution. The full network model (step 3) is used for simulations, while the simplified one (step 4) is used to derive the approximate statistical likelihood.

2.2.2 Hartree approximation

In this section, we present the Hartree approximation principle and provide an explicit formula in the particular case of no self-interaction. The simplified model (7) is still not analytically tractable, but it is now appropriate for employing the *self-consistent proteomic field* approximation introduced in [51, 52] and successfully applied in [53, 54]. More precisely, we will use its natural PDMP counterpart, which will be referred to as “Hartree approximation” since the main idea is similar to the Hartree approximation in physics [51]. It consists in assuming that genes behave as if they were independent from each other, but submitted to a common “proteomic field” created by all other genes. In other words, we transform the original problem of dimension 2^n into n independent problems of dimension 2 that are much easier to solve (see Appendix B for details).

When $k_{\text{on},i}$ and $k_{\text{off},i}$ do not depend on P_i (i.e. no self-interaction), this approach results in approximating the protein stationary distribution of model (7) by the function

$$u(y) = \prod_{i=1}^n \frac{y_i^{a_i(y)-1} (1 - y_i)^{b_i(y)-1}}{B(a_i(y), b_i(y))} \quad (9)$$

where $y = (y_1, \dots, y_n) = (P_1, \dots, P_n) = \mathbf{P}$, $a_i(y) = k_{\text{on},i}(y)/d_{1,i}$, $b_i(y) = k_{\text{off},i}(y)/d_{1,i}$ and B is the standard Beta function. Note that promoter states have been integrated out since they are not required by equation (8).

The function u is a heuristic approximation of a probability density function. It is only valid when interactions are not too strong, that is, when $k_{\text{on},i}$ and $k_{\text{off},i}$ are close enough to constants, and it becomes exact when they are true constants. Besides, it does not integrate to 1 in general. However, this approximation turns out to be very robust in practice and it has the great advantage to be fully explicit (and significantly simpler than in the non-PDMP case), thus providing a promising base for a statistical model.

When $k_{\text{on},i}$ and $k_{\text{off},i}$ depend on P_i , one can still explicitly compute the Hartree approximation in many cases: we will give an example in the next section. Alternatively, it is always possible to use formula (9) even with self-interactions, giving a correct approximation when the feedback is not too strong, as for other proteins.

2.2.3 An explicit form for interactions

We now propose an explicit definition of functions $k_{\text{on},i}$ and $k_{\text{off},i}$. Recent work [36, 55, 56] showed that apparent increased transcription actually reflects an increase in burst frequency rather than amplitude. We therefore decided to model only $k_{\text{on},i}$ as an actual function and to keep $k_{\text{off},i}$ constant. In this view, the activation frequency of a gene can be influenced by ambient proteins, whereas the active periods have a random duration that is dictated only by an intrinsic stability constant of the transcription machinery.

Our approach uses a description of the molecular activity around the promoter in a very similar way as Coulon et al. [42]. Accordingly, we make a quasi-steady state assumption to obtain $k_{\text{on},i}$. This idea based on thermodynamics was also used in the DREAM3 in-Silico Challenge [57] to simulate gene networks. However, only mean transcription rate was described (instead of promoter activity in our work), which is inappropriate to model bursty mRNA dynamics at the single-cell level.

We herein derive $k_{\text{on},i}$ from an underlying stochastic model for chromatin dynamics. We first introduce a set of abstract chromatin states, each state being associated with one of two possible rates of promoter activation, either a low rate $k_{0,i}$ or a high rate $k_{1,i} \gg k_{0,i}$. More specifically, such chromatin states may be envisioned as a coarse-grained description of the chromatin-associated parameters that are critical for transcription of gene i . Second, we assume a separation of timescales between the abstract chromatin model and the promoter activity, so that the promoter activation reaction depends only on the quasi-steady state of chromatin. In other words, the effective $k_{\text{on},i}$ is a combination of $k_{0,i}$ and $k_{1,i}$ which integrates all the chromatin states: its value depends on the probability of each state and a fortiori on the transitions between them. We propose a transition scheme which leads to an explicit form for $k_{\text{on},i}$, based on the idea that proteins can alter chromatin by hit-and-run reactions and potentially introduce a memory component. Some proteins thereby tend to stabilize it either in a “permissive” configuration (with rate $k_{1,i}$) or in a “non-permissive” configuration (with rate $k_{0,i}$), providing notions of *activation* and *inhibition*. A more precise definition and details of the derivation are provided in Appendix C.

The final form is the following. First, we define a function of every protein but P_i ,

$$\Phi_i(y) = \exp(\theta_{i,i}) \prod_{j \neq i} \frac{1 + \exp(\theta_{i,j})(y_j/s_{i,j})^{m_{i,j}}}{1 + (y_j/s_{i,j})^{m_{i,j}}}$$

which may represent the external input of gene i . Then, $k_{\text{on},i}$ is defined by

$$k_{\text{on},i}(y) = \frac{k_{0,i} + k_{1,i}\Phi_i(y)(y_i/s_{i,i})^{m_{i,i}}}{1 + \Phi_i(y)(y_i/s_{i,i})^{m_{i,i}}}. \quad (10)$$

Hence, when the input $\Phi_i(y)$ is fixed, $k_{\text{on},i}$ is a standard Hill function which describes how gene i is self-activating, depending on the Hill coefficient $m_{i,i}$ (Fig. 4). The neutral value is set to $\Phi_i(y) = 1$, so that for this particular value, $s_{i,i}$ is the usual dissociation constant. Moreover, if $\theta_{i,j} = 0$ for all $j \neq i$, then Φ_i becomes the constant function $\Phi_i(y) = \exp(\theta_{i,i})$, and thus $\theta_{i,i}$ may be seen as a “basal” parameter, summing up all potential hidden inputs. On the contrary, if some $\theta_{i,j} > 0$ (resp. $\theta_{i,j} < 0$), then Φ_i becomes itself an increasing (resp. decreasing) Hill-type function of protein P_j , where $m_{i,j}$ and $s_{i,j}$ again play their usual roles.




Figure 4: Different auto-activation types in the network model. Each color corresponds to a fixed value of Φ_i in formula (10), and each curve represents $k_{\text{on},i}$ as a function of y_i for $m_{i,i} = 0$ (no feedback), $m_{i,i} = 1$ (monomer-type feedback) and $m_{i,i} = 2$ (dimer-type feedback). The neutral value $\Phi_i = 1$ is represented by a dashed grey line. Here $k_{0,i} = 0.01$, $k_{1,i} = 2$ and $s_{i,i} = 0.1$.

The $n \times n$ matrix $\theta = (\theta_{i,j})$ therefore plays the same role as the interaction matrix in traditional network inference frameworks [8]. For $i \neq j$, $\theta_{i,j}$ quantifies the regulation of gene i by gene j (activation if $\theta_{i,j} > 0$, inhibition if $\theta_{i,j} < 0$, no influence if $\theta_{i,j} = 0$), and the diagonal term $\theta_{i,i}$ aggregates the “basal input” and the “self-activation strength” of gene i . Note that self-inhibition could be considered instead, but the choice has to be made before the inference since the self-interaction form is notoriously difficult to identify, especially in the stationary regime. In the remainder of this article, we assume that parameters $k_{0,i}$, $k_{1,i}$, $m_{i,j}$ and $s_{i,j}$ are known and we focus on inferring the matrix θ .

A benefit of the interaction form (10) is to allow for a fully explicit Hartree approximation of the protein distribution (see Appendix C for details). In particular, if $m_{i,i} > 0$ and $c_i = (k_{1,i} - k_{0,i})/(d_{1,i}m_{i,i})$ is a positive integer for all i , the approximation is given by

$$u(y) = \prod_{i=1}^n \left(\sum_{r=0}^{c_i} p_{i,r}(y) \frac{y_i^{a_{i,r}-1} (1-y_i)^{b_i-1}}{\text{B}(a_{i,r}, b_i)} \right) \quad (11)$$

with $a_{i,r} = ((c_i - r)k_{0,i} + rk_{1,i})/(d_{1,i}c_i)$, $b_i = k_{\text{off},i}/d_{1,i}$ and

$$p_{i,r}(y) = \frac{\binom{c_i}{r} \text{B}(a_{i,r}, b_i) (\Phi_i(y)/s_{i,i})^r}{\sum_{r'=0}^{c_i} \binom{c_i}{r'} \text{B}(a_{i,r'}, b_i) (\Phi_i(y)/s_{i,i})^{r'}}.$$

In other words, the Hartree approximation (11) is a product of gene-specific distributions which are themselves mixtures of Beta distributions: for gene i , the $a_{i,r}$ correspond to “frequency modes” ranging from $k_{0,i}$ to $k_{1,i}$, weighted by the probabilities $p_{i,r}(y)$. It is straightforward to check that inhibitors tend to select the low burst frequencies of their target ($a_{i,r} \approx k_{0,i}$) while activators select the high frequencies ($a_{i,r} \approx k_{1,i}$). If $m_{i,i} = 0$ for some i , then $k_{\text{on},i}$ does not depend on P_i so one just has to replace the i -th term in the product (11) with the single Beta form as in equation (9), which is equivalent to taking the limit $c_i \rightarrow +\infty$. Finally, when $m_{i,i} > 0$ but c_i is not an integer, using $\lceil c_i \rceil$ instead keeps a satisfying accuracy.

2.2.4 The statistical model in practice

Our statistical framework simply consists in combining the timescale separation (8) and the Hartree approximation (11) into a standard hidden Markov model. Indeed, conditionally to the

proteins, mRNAs are independent and follow well-defined Beta distributions

$$v(x, y) = \prod_{i=1}^n \frac{x_i^{\tilde{a}_i(y)-1} (1-x_i)^{\tilde{b}_i(y)-1}}{B(\tilde{a}_i(y), \tilde{b}_i(y))} \quad (12)$$

where $x = (x_1, \dots, x_n) = (M_1, \dots, M_n) = \mathbf{M}$, $\tilde{a}_i(y) = k_{\text{on},i}(y)/d_{0,i}$ and $\tilde{b}_i(y) = k_{\text{off},i}(y)/d_{0,i}$. Then one can use (11) to approximate the joint distribution of proteins. Hence, recalling the unknown interaction matrix θ , the inference problem for m cells with respective levels $(\mathbf{M}_k, \mathbf{P}_k)_{1 \leq k \leq m}$ is based on the (approximate) complete log-likelihood:

$$\ell = \ell(\mathbf{M}_1, \dots, \mathbf{M}_m, \mathbf{P}_1, \dots, \mathbf{P}_m | \theta) = \sum_{k=1}^m \log(u(\mathbf{P}_k)) + \log(v(\mathbf{M}_k, \mathbf{P}_k)) \quad (13)$$

where we used conditional factorization and independence of the cells.

The basic statistical inference problem would be to maximize the marginal likelihood of mRNA with respect to θ . Since this likelihood has no simple form, a typical way to perform inference is to use an Expectation-Maximization (EM) algorithm on the complete likelihood (13). However, the algorithm may be slow in practice because of the computation of expectations over proteins. A faster procedure consists in simplifying these expectations using the distribution modes: the resulting algorithm is often called “hard EM” or “classification EM” and is used in the Results section. Moreover, it is possible to encode some potential knowledge or constraints on the network by introducing a prior distribution $w(\theta)$. In this case, from Baye’s rule, one can perform *maximum a posteriori* (MAP) estimation of θ by using the same EM algorithm but adding the penalization term $\log(w(\theta))$ to ℓ during the Maximization step (see Appendix D.1 and the Results section). Alternatively, a full bayesian approach, i.e. sampling from the posterior distribution of θ conditionally to $(\mathbf{M}_1, \dots, \mathbf{M}_m)$, may also be considered using standard MCMC methods.

Taking advantage of the latent structure of proteins, we can also deal with missing data in a natural way: if the mRNA measurement of gene i is invalid in a cell k owing to technical problems, it is possible to ignore it by removing the i -th term in the conditional distribution of mRNAs (12). This only modifies the definition of v for cell k in equation (13), ensuring that all valid data is effectively used for each cell.

3 Results

In this part, we first compare the distribution of the mechanistic model (4) to the mRNA quasi-steady state combined with Hartree approximation for proteins, on a simple toggle-switch example. Then, we show that the single-gene model with auto-activation can fit marginal mRNA distributions from real data better than the constant- k_{on} model. Finally, we successfully apply the inference procedure to various two-gene networks simulated from the mechanistic model.

3.1 Relevance of the approximate likelihood

Starting from the normalized mechanistic model (4), two approximations were used to derive the final statistical likelihood (13): the quasi-steady state assumption for mRNAs given protein levels, and the Hartree approximation for the joint distribution of proteins. Crucially, this approximate likelihood has to be close enough to the exact one in order to preserve the equivalence between inferring a network and fitting the mechanistic model. To get an idea of the accuracy, we

considered a basic two-gene toggle switch defined by $k_{\text{on},i}$ following equation (10) with the interaction matrix given by $\theta_{1,1} = \theta_{2,2} = 4$ and $\theta_{1,2} = \theta_{2,1} = -8$ (full parameter list in Appendix F). By computing sample paths (Fig. 5), we estimated the stationary distribution and compared it with our approximation, which appeared to be very satisfying, both for proteins and mRNAs (Fig. 6).




Figure 5: Sample path of a two-gene toggle switch, with gene 1 in red and gene 2 in green. While always staying in a bursty regime regarding mRNAs, genes can switch between high and low frequency modes (here at $t \approx 50\text{h}$). From this example, it is clear that the overall joint distribution can contain correlations even if the bursts themselves are not coordinated.

3.2 Fitting marginal mRNA distributions from real data

A particularity of single-cell data is to often exhibit bursty regimes for mRNA (meaning $k_{\text{on}} \ll k_{\text{off}}$ and $d_0 \ll k_{\text{off}}$) and potentially also for proteins (adding $d_1 \ll k_{\text{off}}$), which are well fitted by Gamma distributions [37]. At this stage, it is worth mentioning that the Gamma distribution can be seen as a limit case of the Beta distribution. Intuitively, when $b \gg 1$ and $b \gg a$ (typically $a = k_{\text{on}}/d_0$ and $b = k_{\text{off}}/d_0$), most of the mass of the distribution $\text{Beta}(a, b)$ is located at $x \ll 1$ so we have the first order approximation

$$\begin{aligned} x^{a-1}(1-x)^{b-1} &= x^{a-1} \exp((b-1) \log(1-x)) \\ &\approx x^{a-1} \exp(-bx) \end{aligned}$$

and thus $\text{Beta}(a, b) \approx \gamma(a, b)$. This way, formulas (11) and (12) can be easily transformed into Gamma-based distributions. Parameters s_0 and k_{off} then aggregate in k_{off}/s_0 because of the scaling property of the Gamma distribution, so only this ratio has to be inferred: from an applied perspective, it simply represents a scale parameter for each gene. This remark leads to a possible preprocessing phase that can be used for estimating the crucial basal parameters of the network, without requiring the knowledge of such scale parameters (see Appendix E).

In addition, our network model is able to generate multiple modes while keeping such bursty regimes (Fig. 5), as noticeable in the stationary distribution (11). Interestingly, this feature has




Figure 6: Exact and approximate stationary distributions for the example of toggle switch. True distributions (left side) were estimated by sample path simulation, while approximations (right side) have explicit formulas. (a) True distribution of proteins. (b) Approximate distribution of proteins, from formula (11). (c) True distribution of mRNAs. (d) Approximate distribution of mRNAs, obtained by integrating the conditional distribution of mRNA (12) against (b).

already been considered in the literature by empirically introducing mixture distributions [58,59]. As a first step toward applications, we compared our model in the simplest case (independent genes with auto-activation) to marginal distributions of single-cell mRNA measurements from [38]. Our model was fitted and compared to the basic two-state model in the bursty regime, i.e. to a simple Gamma distribution: Fig. 7 shows the example of the LDHA gene. Although very close when viewed in raw molecule numbers, the distributions differ after applying the transformation $x \mapsto x^\alpha$ with $\alpha = 1/3$, which tends to compress great values while preserving small values. The data becomes bimodal, suggesting the presence of two bursting regimes, a “normal” one and a very small “inhibited” one: the auto-activation model then performs better than the simple Gamma, which necessarily stays unimodal for $0 < \alpha < 1$. Note that the RTqPCR protocol used in [38] was shown to be far more sensitive than single-cell RNA-seq in the detection of low abundance transcripts [60]. Since the data also contains small nonzero values, this tends to support a true biological origin for the peak in zero. Besides, the case of distributions that are not bimodal until transformed also arises for proteins [61].

3.3 Application of the inference procedure

By construction of the mechanistic model, the interaction matrix θ can describe any oriented graph by explicitly defining causal quantitative links between genes, which is difficult to do within traditional statistical frameworks (e.g. bayesian networks or undirected Markov random fields). The logical downside is that identifiability issues seem inevitable. In a first attempt to assess this

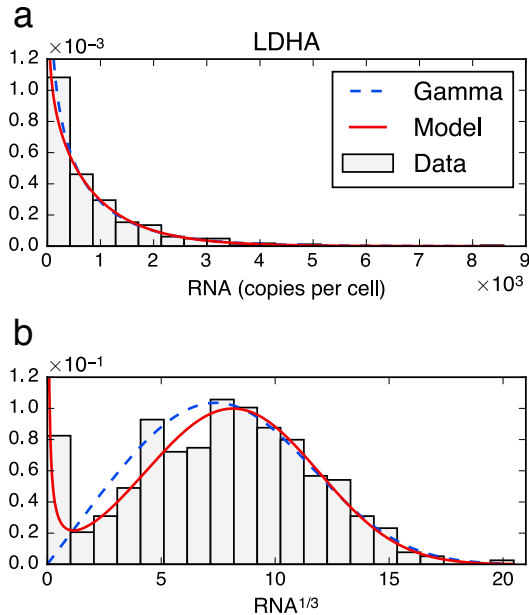


Figure 7: Fitting marginal distributions from real single-cell data: example of the LDHA gene. The red curve is the stationary distribution associated with our interaction form (here a single gene with auto-activation), while the dashed blue curve corresponds to the basic two-state model in the bursty regime (Gamma distribution). (a) The raw data seems to be well fitted by the Gamma distribution, which in this view is close to our model. (b) Same fit viewed after applying the transformation $x \mapsto x^{1/3}$. The data becomes bimodal and the fit appears to be better with the auto-activation model.

aspect, we implemented the inference method presented above and tested it on various two-gene networks, assuming auto-activation for each gene (i.e. $m_{i,i} > 0$) with equation (10) to maximize variability without considering perturbations of the system (parameter list in Appendix F).

We decided to investigate the worst case scenario in terms of cell numbers. We are fully aware of the existence of technologies allowing to interrogate thousands of cells simultaneously, but most of the recent studies still rely upon a much smaller number of cells. For each network, we therefore simulated mRNA snapshot data for 100 cells using the full PDMP model (4). We then inferred the matrix θ using a “hard EM” algorithm based on the likelihood (13), that is, alternatively maximizing the likelihood with respect to θ and with respect to the (unknown) protein levels of each cell. A lasso-like penalization term, corresponding to a prior distribution, was added to the $\theta_{i,j}$ for $i \neq j$ to obtain true zeros – so that the inferred network topology is clear – and to prevent keeping both $\theta_{i,j}$ and $\theta_{j,i}$ when one is significantly weaker (see Appendix D.1 for details of the penalization and the whole procedure).

We obtained highly encouraging results since every structure was inferred with a high probability of success (Fig. 8), meaning that the non-diagonal (i.e. interaction) terms of θ had the right sign and were nonzero at the right places. A list of the inferred values is provided in Table 5 of Appendix F. It is very important at that stage to emphasize that we are not trying to infer θ exactly: we only assess whether it has a zero or nonzero value and its sign. Although the results tend to support the identifiability of the full matrix θ in this simple two-gene case, one has to be aware that the quantity we maximize (an approximate likelihood) is a priori nonconvex and can have several local maxima (i.e. networks that are relevant candidates to explain the data). The result of the inference thus can depend on the starting point: in this first approach

we chose the null matrix to be the starting point for θ , which corresponds to the – biologically relevant – expectation of “balanced” behaviors (e.g. we do not expect $\theta_{1,1} \ll \theta_{2,2}$). Alternatively, one can consider some probabilistic prior knowledge on θ to implement a (possibly rough) idea of parameter values from a Bayesian viewpoint: it is worth mentioning that any knockout information can be implemented this way in our model.

Finally, we assessed the inference behavior in the presence of dropouts, i.e. genes expressed at a low level in a cell that give rise to zeros after measurement [4]. Our first tests tend to indicate that our approach is robust regarding dropouts, in the sense that up to 30% of simulated dropouts does not drastically affect the estimation of θ once the other parameters have been estimated correctly (see Table 6 of Appendix F for an example).

4 Discussion

In this paper, we introduce a general stochastic model for gene regulatory networks, which can describe bursty gene expression as observed in individual cells. Instead of using ordinary differential equations, for which cells would structurally all behave the same way, we adopt a more detailed point of view including stochasticity as a fundamental component through the two-state promoter model. This model is but a simplification of the complexity of the real molecular processes [42]. Modifications have been proposed, from the existence of a refractory period [23] to its attenuation by nuclear buffering [62]. In bacteria, the two states originate from the accumulation of positive supercoiling on DNA which stops transcription [63]. In eukaryotes, although its molecular basis is not quite understood, the two-state model is a remarkable compromise between simplicity and the ability to capture real-life data [18, 22, 36–38]. Our PDMP framework appears to be conceptually very similar to the *random dynamical system* proposed in [64] but it has two major advantages: time does not have to be discretized, and the mathematical analysis is significantly easier. We also note that a similar framework appears in [65, 66] and that a closely related PDMP – which can be seen as the limit of our model for infinitely short bursts – has recently been described in [67].

We then derive an explicit approximation of the stationary distribution and propose to use it as a statistical likelihood to infer networks from single-cell data. The main ingredient is the separation of three physical timescales – chromatin, promoter/RNA, and proteins – and the core idea is to use the self consistent proteomic field approximation from [51, 52] in a slightly simpler mathematical framework, providing fully explicit formulas that make possible the massive computations usually needed for parameter inference. From this viewpoint, it is a rather simple approach and we hope it can be adapted or improved in more specific contexts, for example in the study of lineage commitment [68]. Besides, the main framework does not necessarily has to include an underlying chromatin model and thus it can in principle also be used to describe gene networks in procaryotes.

Mechanistic modelling and statistical inference

An important quality of the PDMP network model is that the simulation algorithm is comparable in speed with classic ODE and diffusion systems, while providing an effective approximation of the “perfect”, fully discrete, molecular counterpart [33, 35]. It is worth noticing that the PDMP – at least the promoter-mRNA system – naturally appears as an example of Poisson representation [28, 69], that is, not a simple approximation but rather the core component of the

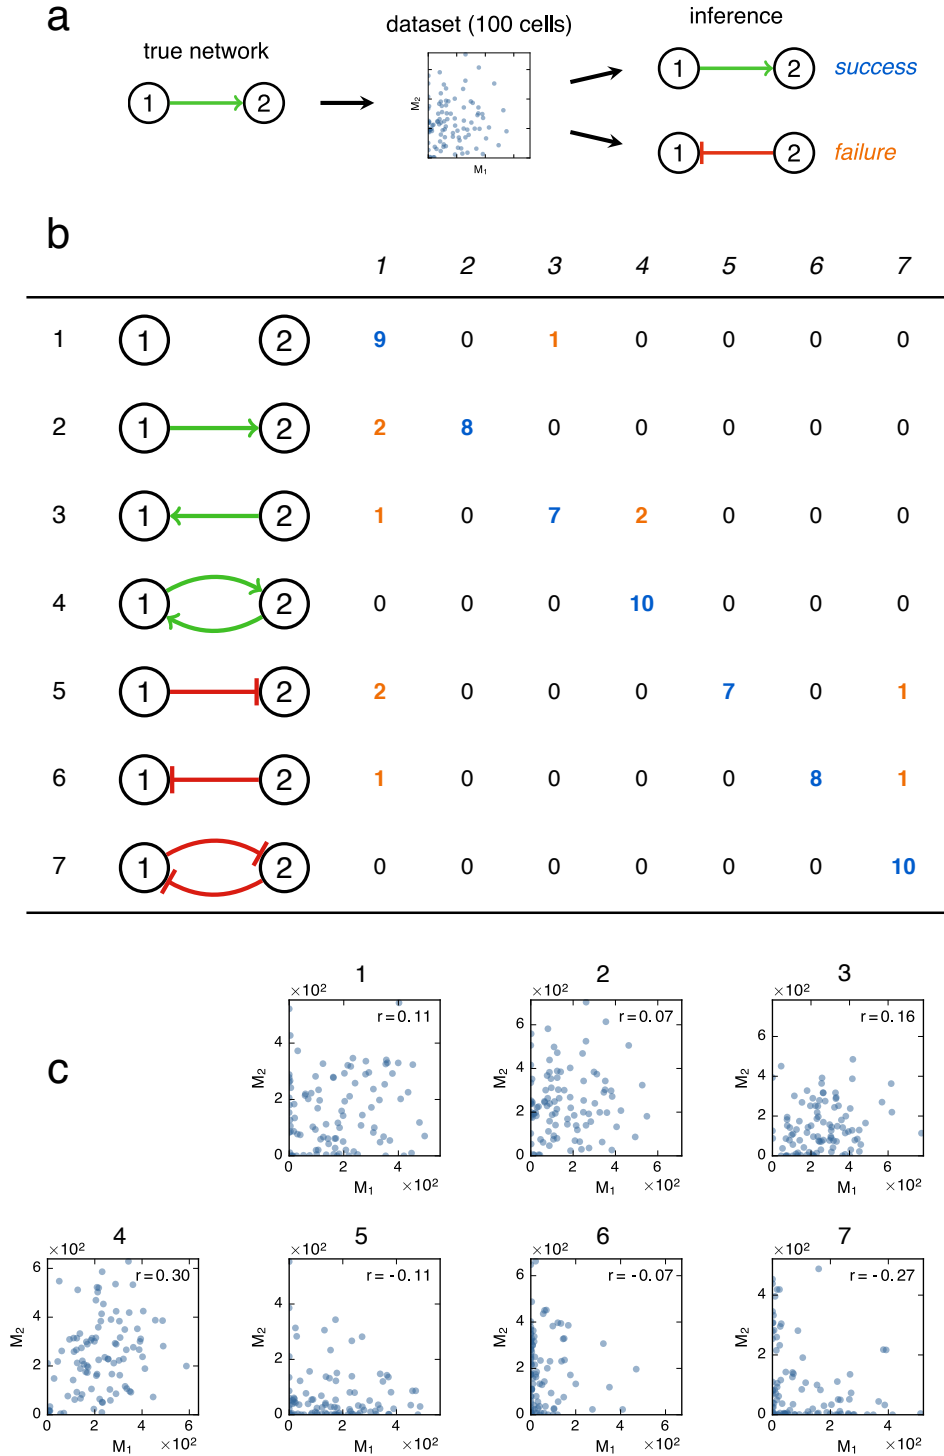


Figure 8: Testing our inference method on simple networks. (a) For each network, numbered from 1 to 7, we simulated 100 cells using the full mechanistic model until the stationary regime was reached. Then we took a snapshot of their mRNA levels and inferred the parameters from this data. The result was called successful when the inferred structure (topology and nature of the links) was the same as the true network. (b) For each network (rows), 10 datasets were simulated and the results were reported by counting the number of inferred θ corresponding to each structure (columns), highlighting successes (blue) and failures (orange). The perfect inference would lead to 10 for all the diagonal terms and 0 everywhere else. (c) Examples of simulated mRNA datasets (one for each network). Although having coherent signs, Pearson's correlation coefficients (top right of each plot) would clearly be insufficient to distinguish between the different networks.

exact distribution of the discrete molecular model. Furthermore, such a simulation speed allowed us to compare our approximate likelihood with the true likelihood for a simple two-gene toggle switch, giving excellent results (Fig. 6). This obviously does not constitute a proof of robustness for every network: a proper quantitative (theoretical or numeric) comparison is beyond the scope of this article but would be extremely valuable. Intuitively, it should work for any number of genes, provided that interactions are not too strong.

Besides, some widely used ODE frameworks [8, 17, 57] can be seen as the fast-promoter limit of the PDMP model: this limit may not always hold in practice, especially in the bursty regime. In particular, Fig. 5 highlights the risk of using mRNA levels as a proxy for protein levels. It also explains why ordering single-cell mRNA measurements by pseudo-time may not always be relevant, as found in [38]. In [70], the authors use a hybrid model of gene expression to infer regulatory networks: it is very close to the diffusion limit of our reduced model (7) with the difference that the discrete component, called “promoter” by the authors, would correspond to the “frequency mode” in the present article, as visible for proteins in Fig. 5. From such a perspective, our approach adds a description of bursty mRNA dynamics that allows for fitting single-cell data such as in Fig. 7.

Finally, our method performed well for simple two-gene networks (Fig. 8), showing that part of the causal information remains present in the stationary distribution: this suggests that it is indeed possible to retrieve network structures with a mechanistic interpretation, even from bursty mRNA data.

Perspectives

We focused here on presenting the key ideas behind the general network model and the inference method: the logical next step is to apply it to real data and with a larger number of genes, which is the subject of work in progress in our group. In particular, we propose a functional preprocessing phase, detailed in Appendix E, that only requires the knowledge of the ratio $d_{0,i}/d_{1,i}$ to estimate all the relevant parameters before inferring θ . The ratio between protein and mRNA degradation rates (or half-lives) hence appears to be the minimum required for such a mechanistic approach to be relevant. Depending upon the species, mRNA and protein half-lives values can be found in the literature (see e.g. [31] for human proteins half-lives), or should be estimated from ad hoc experiments.

From a computational point of view, the main challenge is the algorithmic complexity induced by the fact that proteins are not observed and have to be treated as latent variables. There is a priori no possibility of reducing this without losing too much accuracy, and therefore some finely optimized algorithms may be required to make the method scalable. Furthermore, the identifiability properties of the interaction matrix θ seem difficult to derive theoretically. In this paper we focused on the stationary distribution for simplicity: importantly, several aspects such as time dependence (computing the Hartree approximation in transitory regime) or perturbations (changing the cell’s medium or performing knockouts [71], which can be naturally embedded in our framework) could greatly improve the practical identifiability.

From a biological point of view, our model does not really describe individual cells but rather a concatenation of trajectories obtained by following cells throughout divisions. Experiments suggest that it should be a relevant approximation, providing one considers mRNA and proteins

levels in terms of concentrations instead of molecule numbers [72], which is made possible by the PDMP framework. In this view, the cell cycle results in increasing the apparent degradation rates – because of the increase in cell volume followed by division – and thus plays a crucial role for very stable proteins. However, at such a description level, many aspects of possible compensation mechanisms [73] and chromatin dynamics [74] remain to be elucidated. Regarding the latter aspect, our abstract chromatin states were not modeled from real-life data – chromatin composition for instance – but our approach is relevant in that partitioning into dual-type chromatin states as we did is now known as a pervasive feature of all eukaryotic genomes [75–78].

Acknowledgments

We thank Geneviève Fourel (LBMC/ENSL) for critical reading of the manuscript and Anne-Laure Fougères (ICJ) for her constant support. We also thank all members of the SBDM and Dracula teams for enlightening discussions, and BioSyL Federation and Ecofect Labex for inspiring scientific events. This work was supported by funding from the Institut Rhônealpin des Systèmes Complexes (IXXI) and from French agency ANR (ICEBERG; ANR-IABI-3096).

Appendix A First simplifications

A.1 Normalizing the PDMP network model

In this section we detail the normalization of our network model. Recall that the original model is defined by

$$\begin{cases} E_i(t) : 0 \xrightarrow{k_{\text{on},i}} 1, & 1 \xrightarrow{k_{\text{off},i}} 0 \\ M_i'(t) = s_{0,i}E_i(t) - d_{0,i}M_i(t) \\ P_i'(t) = s_{1,i}M_i(t) - d_{1,i}P_i(t) \end{cases} \quad (14)$$

where $k_{\text{on},i} = k_{\text{on},i}(P_1, \dots, P_n)$ and $k_{\text{off},i} = k_{\text{off},i}(P_1, \dots, P_n)$. First we observe that, given an initial condition

$$(E_1^0, \dots, E_n^0) \in \{0, 1\}^n, \quad (M_1^0, \dots, M_n^0) \in \prod_{i=1}^n \left[0, \frac{s_{0,i}}{d_{0,i}}\right], \quad (P_1^0, \dots, P_n^0) \in \prod_{i=1}^n \left[0, \frac{s_{0,i}s_{1,i}}{d_{0,i}d_{1,i}}\right],$$

the system stays in this set for all $t > 0$, and we introduce the dimensionless variables:

$$\bar{M}_i = \frac{d_{0,i}}{s_{0,i}}M_i \in [0, 1] \quad \text{and} \quad \bar{P}_i = \frac{d_{0,i}d_{1,i}}{s_{0,i}s_{1,i}}P_i \in [0, 1].$$

Then, since $s_{0,i}$, $s_{1,i}$, $d_{0,i}$ and $d_{1,i}$ are constants, we get

$$\bar{M}_i'(t) = \frac{d_{0,i}}{s_{0,i}}M_i'(t) = d_{0,i} \left(E_i(t) - \frac{d_{0,i}}{s_{0,i}}M_i(t) \right) = d_{0,i} \left(E_i(t) - \bar{M}_i(t) \right)$$

and

$$\bar{P}_i'(t) = \frac{d_{0,i}d_{1,i}}{s_{0,i}s_{1,i}}P_i'(t) = d_{1,i} \left(\frac{d_{0,i}}{s_{0,i}}M_i(t) - \frac{d_{0,i}d_{1,i}}{s_{0,i}s_{1,i}}P_i(t) \right) = d_{1,i} \left(\bar{M}_i(t) - \bar{P}_i(t) \right).$$

As a result, we obtain the normalized model:

$$\begin{cases} E_i(t) : 0 \xrightarrow{\bar{k}_{\text{on},i}} 1, & 1 \xrightarrow{\bar{k}_{\text{off},i}} 0 \\ \bar{M}_i'(t) = d_{0,i} \left(E_i(t) - \bar{M}_i(t) \right) \\ \bar{P}_i'(t) = d_{1,i} \left(\bar{M}_i(t) - \bar{P}_i(t) \right) \end{cases} \quad (15)$$

where the rescaled interaction function $\bar{k}_{\text{on},i}$ is defined by

$$\bar{k}_{\text{on},i}(\bar{P}_1, \dots, \bar{P}_n) = k_{\text{on},i} \left(\frac{s_{0,1}s_{1,1}}{d_{0,1}d_{1,1}}\bar{P}_1, \dots, \frac{s_{0,n}s_{1,n}}{d_{0,n}d_{1,n}}\bar{P}_n \right)$$

and $\bar{k}_{\text{off},i}$ is defined analogously. It is straightforward to see that, given a path

$$(E_i(t), \bar{M}_i(t), \bar{P}_i(t))_i$$

of the normalized model (15), the corresponding path of the original model (14) is

$$\left(E_i(t), \frac{s_{0,1}}{d_{0,1}}\bar{M}_i(t), \frac{s_{0,1}s_{1,1}}{d_{0,1}d_{1,1}}\bar{P}_i(t) \right)_i.$$

In this sense, both models are equivalent: in the main text and in the next sections, we always consider model (15) but forget the ‘‘bars’’ to keep the notations simple.

A.2 Separating mRNA and protein timescales

Here we justify the reduced network model involving only promoters and proteins, which is valid when $d_{1,i} \ll d_{0,i}$ for all gene i . A full proof is beyond the scope of this article but we provide a heuristic explanation. We temporarily drop the i index for simplicity. Let $t_1 \geq t_0 \geq 0$ and $E \in \{0, 1\}$, and suppose $E(t) = E$ for all $t \in [t_0, t_1]$. Moreover, let $M_0 = M(t_0) \in [0, 1]$ and $P_0 = P(t_0) \in [0, 1]$. If $d_1 < d_0$, the solution of the linear ODE system

$$\begin{cases} M' = d_0(E - M) \\ P' = d_1(M - P) \end{cases}$$

is given for $t \in [t_0, t_1]$ by

$$\begin{cases} M(t) = E + (M_0 - E)e^{-d_0(t-t_0)} \\ P(t) = E + (P_0 - E)e^{-d_1(t-t_0)} + \frac{d_1}{d_0 - d_1}(M_0 - E)(e^{-d_1(t-t_0)} - e^{-d_0(t-t_0)}) \end{cases}$$

Hence, if $d_1 \ll d_0$, we have

$$P(t) \approx E + (P_0 - E)e^{-d_1(t-t_0)}$$

using the fact that $|M_0 - E| \leq 1$ and $|e^{-d_1(t-t_0)} - e^{-d_0(t-t_0)}| \leq 1$, and thus $P(t)$ approximates the solution of the differential equation $P' = d_1(E - P)$.

Appendix B Hartree approximation

B.1 Hartree approximation for the PDMP model

Before deriving the approximation, we introduce some notation. Let n be the number of genes in the network, $\mathcal{E} = \{0, 1\}^n$ and $\Omega = (0, 1)^n$. At time t , promoter and protein configurations are denoted by $E_t = (e_1, \dots, e_n) = e \in \mathcal{E}$ and $P_t = (y_1, \dots, y_n) = y \in \Omega$, respectively. The distribution of (E_t, P_t) then evolves along time according to its Kolmogorov forward (aka master) equation, which is a linear partial differential equation (PDE) system in our case. This system is high dimensional ($|\mathcal{E}| = 2^n$, the number of possible promoter configurations) but the associated linear operator contains lots of zeros. Using the tensor product notation \otimes , one can write down the equation in a compact form:

$$\frac{\partial u}{\partial t} + \sum_{i=1}^n \frac{\partial (F_i u)}{\partial y_i} = \sum_{i=1}^n K_i u \quad (16)$$

where $u(t, y) = (u_e(t, y))_{e \in \mathcal{E}} \in \mathbb{R}^{2^n} \simeq (\mathbb{R}^2)^{\otimes n}$ represents the probability density function (pdf) of (E_t, P_t) , and matrices $F_i(y_i), K_i(y) \in \mathcal{M}_{2^n}(\mathbb{R}) \simeq \mathcal{M}_2(\mathbb{R})^{\otimes n}$ are defined by

$$F_i(y_i) = I_2 \otimes \dots \otimes \underbrace{F^{(i)}(y_i)}_i \otimes \dots \otimes I_2, \quad K_i(y) = I_2 \otimes \dots \otimes \underbrace{K^{(i)}(y)}_i \otimes \dots \otimes I_2$$

with

$$F^{(i)}(y_i) = \begin{pmatrix} -d_{1,i}y_i & 0 \\ 0 & d_{1,i}(1-y_i) \end{pmatrix} \quad \text{et} \quad K^{(i)}(y) = \begin{pmatrix} -k_{\text{on},i}(y) & k_{\text{off},i}(y) \\ k_{\text{on},i}(y) & -k_{\text{off},i}(y) \end{pmatrix}.$$

The sum in the left side of equation (16) clearly corresponds to a deterministic transport term, while the right side corresponds to the stochastic transitions between promoter configurations.

Furthermore, the PDE system comes with the boundary condition

$$\forall i \in \{1, \dots, n\}, \quad F_i u = 0 \quad \text{on } \partial\Omega \quad (17)$$

and the probability condition

$$u \geq 0 \quad \text{and} \quad \forall t \in \mathbb{R}_+, \quad \sum_{e \in \mathcal{E}} \int_{\Omega} u_e(t, y) \, dy = 1. \quad (18)$$

The self-consistent ‘‘Hartree’’ approximation consists in splitting this 2^n -dimensional problem into n independent 2-dimensional problems by ‘‘freezing’’ the y_j for $j \neq i$ where i is fixed, and then gathering the solutions by taking their tensor product to produce an approximation of the true pdf (see [52] for a heuristic explanation in the discrete protein setting). More precisely, one reduced problem is derived for each gene i from (16)-(17)-(18):

$$\frac{\partial u^i}{\partial t} + \frac{\partial(F^{(i)}u^i)}{\partial y_i} = K^{(i)}u^i \quad (19)$$

where $u^i(t, y) = (u_0^i(t, y), u_1^i(t, y))^{\top} \in \mathbb{R}_+^2$ satisfies the initial condition $u^i(0, y) = u^{i,0}(y)$, the boundary condition $F^{(i)}(y_i)u^i(y) \rightarrow 0$ when $y_i \rightarrow 0$ or 1 , and the probability condition $\int_0^1 [u_0^i(t, y) + u_1^i(t, y)] \, dy_i = 1$ for all $t \geq 0$ and $y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n \in (0, 1)$. Therefore, each u^i is a pdf with respect to $(e_i, y_i) \in \{0, 1\} \times (0, 1)$ but not on $\mathcal{E} \times \Omega$. Finally, the Hartree approximation is given by

$$u(t, y) \approx \bigotimes_{i=1}^n u^i(t, y) \quad (20)$$

where the equality holds if for all i , $k_{\text{on},i}$ and $k_{\text{off},i}$ only depend on y_i .

B.2 Solving the reduced problem

For the moment, the time-dependent closed-form solution of (19) is unavailable, but the unique stationary solution can be easily obtained if one knows a primitive of

$$\lambda_i : y_i \mapsto \frac{k_{\text{on},i}(y)}{d_{1,i}y_i} - \frac{k_{\text{off},i}(y)}{d_{1,i}(1-y_i)}$$

which is the nonzero eigenvalue of the matrix $M^{(i)} = K^{(i)}(F^{(i)})^{-1}$. Indeed, letting $v^i = F^{(i)}u^i$, the stationary equation for v^i from (19) becomes

$$\frac{\partial v^i}{\partial y_i} = M^{(i)}v^i$$

and then, crucially using the fact that $M^{(i)}$ has a constant eigenvector $(-1, 1)^{\top}$ associated with eigenvalue λ_i (the other eigenvalue being 0), one can check that $v^i = e^{\varphi_i}(-1, 1)^{\top}$ is a solution when $\frac{\partial \varphi_i}{\partial y_i} = \lambda_i$. If one has such a φ_i , the stationary solution of (19) is given by

$$u_0^i(y) = Z_i^{-1}y_i^{-1} \exp(\varphi_i(y)) \quad \text{and} \quad u_1^i(y) = Z_i^{-1}(1-y_i)^{-1} \exp(\varphi_i(y)) \quad (21)$$

where Z_i is the normalizing constant (which may still depend on y_j for $j \neq i$). Note that the existence of a positive constant α such that $\min(k_{\text{on},i}, k_{\text{off},i}) \geq \alpha$ imposes the limit 0 for $\exp(\varphi_i(y))$ when $y_i \rightarrow 0$ or 1 , and thus the boundary condition is satisfied. We also obtain the promoter probabilities $p_{0,i} = p(e_i = 0) = Z_{0,i}/Z_i$ and $p_{1,i} = p(e_i = 1) = Z_{1,i}/Z_i$ where $Z_{0,i} = \int_0^1 y_i^{-1} \exp(\varphi_i(y)) \, dy_i$, $Z_{1,i} = \int_0^1 (1-y_i)^{-1} \exp(\varphi_i(y)) \, dy_i$ and $Z_i = Z_{0,i} + Z_{1,i}$.

In particular, when $k_{\text{on},i}$ and $k_{\text{off},i}$ do not depend on y_i (i.e. no self-interaction), we get

$$\varphi_i(y) = \frac{k_{\text{on},i}(y)}{d_{1,i}} \log(y_i) + \frac{k_{\text{off},i}(y)}{d_{1,i}} \log(1 - y_i)$$

which gives the classical solution

$$u_0^i(y) = \frac{b_i}{a_i + b_i} \cdot \frac{y_i^{a_i-1}(1 - y_i)^{b_i}}{\text{B}(a_i, b_i + 1)} \quad \text{and} \quad u_1^i(y) = \frac{a_i}{a_i + b_i} \cdot \frac{y_i^{a_i}(1 - y_i)^{b_i-1}}{\text{B}(a_i + 1, b_i)} \quad (22)$$

with $a_i = k_{\text{on},i}(y)/d_{1,i}$ and $b_i = k_{\text{off},i}(y)/d_{1,i}$. This form makes clear the promoter probabilities $p_{0,i}$ and $p_{1,i}$ and the conditional distributions of protein y_i given the promoter state $e_i = 0$ or 1 , both being Beta distributions. Since the state is usually not observed, one usually considers the marginal pdf of y_i , which is also a Beta:

$$\underline{u}^i(y) = u_0^i(y) + u_1^i(y) = \frac{y_i^{a_i-1}(1 - y_i)^{b_i-1}}{\text{B}(a_i, b_i)}. \quad (23)$$

Note that the conditional distribution of mRNA given proteins also has the form (23) since the PDMP equation is the same, although the argument is not the Hartree approximation but rather the more common quasi-steady state assumption.

B.3 Protein marginal distribution

Given the form of the solution (21), it is in fact always straightforward to integrate over promoters, even for the full (stationary) Hartree approximation (20), and we finally obtain

$$\underline{u}(y) = \sum_{e \in \mathcal{E}} u_e(y) \approx \sum_{e \in \mathcal{E}} \left[\bigotimes_{i=1}^n u^i(y) \right]_e = \sum_{e \in \mathcal{E}} \prod_{i=1}^n \frac{\exp(\varphi_i(y))}{Z_i(y)^{|e_i - y_i|}} = \prod_{i=1}^n \frac{\exp(\varphi_i(y))}{Z_i(y) y_i (1 - y_i)} \quad (24)$$

where we recalled the possible dependence of Z_i on some y_j . Hence, when φ_i and Z_i are known functions, one gets a fully explicit approximation of the joint protein distribution.

Appendix C Explicit interactions

Here we derive an explicit form for the interactions between genes, starting from a coarse-grained biochemical model. That is, for a given gene i , we focus on defining functions $k_{\text{on},i}(y_1, \dots, y_n)$ and $k_{\text{off},i}(y_1, \dots, y_n)$ where y_1, \dots, y_n denote the protein quantities. For simplicity, we drop the i index in this section when there is no ambiguity.

C.1 Simple biochemical model


The basic idea is to slightly refine the two-state model of gene expression: in addition to the usual switching reactions (whose rates are k_{on} and k_{off}), we consider a set of reversible transitions between some chromatin states (e.g. describing enhancer regions). Each chromatin state is then associated with a particular rate for the promoter activation reaction. For simplicity, we consider only two cases: a high rate k_1 (the chromatin will be said *permissive*) and a low rate $k_0 \ll k_1$ (the chromatin will be said *non-permissive*). Once active, the promoter can switch off at a rate that is supposed to be independent from chromatin states. Finally, we assume that the chromatin transitions are due to fast interactions with ambient proteins (binding, hit-and-run, etc.) so that the promoter-switching reactions always see chromatin in its quasi-stationary state. Effective

rates k_{on} and k_{off} can therefore be obtained by averaging over chromatin states: this way, k_{off} is still a constant and k_{on} is now defined by



$$k_{\text{on}} = k_0 p_0 + k_1 p_1$$

where p_0 (resp. p_1) is the probability of the chromatin being non-permissive (resp. permissive).

We now define an explicit model for chromatin dynamics and compute its stationary distribution to derive p_0 and p_1 as functions of y_1, \dots, y_n . We consider 2^n permissive configurations and 2^n non-permissive configurations as follows: for all $I \subset \mathcal{G}$ where $\mathcal{G} = \{1, \dots, n\}$, species C_I (resp. C_I^*) stands for the chromatin being non-permissive (resp. permissive) and in state I . The underlying physics are the following: the chromatin has two “basal” configurations C_\emptyset (non-permissive) and C_\emptyset^* (permissive), which describe dynamics when no protein is present, according to the reactions



Then, each protein P_j is able to modify the chromatin state through a “hit-and-run” reaction, which is kept in memory by encoding the index j in the list I , giving the state C_I or C_I^* . Eventually, this memory can be lost by “emptying” I step by step (going back to the basal configuration). That is, for all $I \subset \mathcal{G}$ and $j \in \mathcal{G} \setminus I$, we consider the reactions



The system then evolves with $[C_I], [C_I^*] \in \{0, 1\}$ and $\sum_I [C_I] + [C_I^*] = 1$, so that only one molecule is present at a time: its species therefore entirely describes the state of the system. Mathematically, we obtain a standard jump Markov process with 2^{n+1} states. For example, the case $n = 2$ leads to the scheme of Figure 9, writing $\bar{a}_j = a_j [P_j]$ and $\bar{c}_j = c_j [P_j]$ for simplicity. The underlying idea is that, depending on a_j, b_j, c_j and d_j , proteins will tend to stabilize the chromatin either in a permissive configuration or in a non-permissive one – providing notions of *activation* and *inhibition*. The basal reactions with rates α and β sum up what we do not observe (i.e. what is likely to happen for the chromatin when none of the P_j are present).




Figure 9: Chromatin states and transitions rates in the case of $n = 2$ proteins.

C.2 Stationary distribution

Letting $\mathcal{S} = \{0, 1\}^{n+1}$, each state can be coded by a vector $s = (s_0, s_1, \dots, s_n) \in \mathcal{S}$ where $s_0 = 1$ if the chromatin is permissive and 0 otherwise, and for $j \geq 1$, $s_j = 1$ if it has been modified by protein P_j and 0 otherwise. If all rates are positive, the system has a unique stationary distribution π which can be exactly computed from the master equation. More precisely, the probability π_s of the chromatin being in state $s \in \mathcal{S}$ is given by

$$\pi_s = \begin{cases} Z^{-1} \alpha \prod_{j=1}^n (\lambda_j [P_j] s_j + 1 - s_j) & \text{if } s_0 = 1 \\ Z^{-1} \beta \prod_{j=1}^n (\mu_j [P_j] s_j + 1 - s_j) & \text{if } s_0 = 0 \end{cases}$$

where $\lambda_j = a_j/b_j$, $\mu_j = c_j/d_j$ and Z is a normalizing constant. Now going back to our initial intention of computing k_{on} , we are only interested in the probability for the chromatin to be permissive,

$$p_1 = \sum_{s_1, \dots, s_n} \pi_{(1, s_1, \dots, s_n)} = Z^{-1} \alpha \sum_{s_1, \dots, s_n} \prod_{j=1}^n (\lambda_j [P_j] s_j + 1 - s_j),$$

and the probability for the chromatin to be non-permissive,

$$p_0 = \sum_{s_1, \dots, s_n} \pi_{(0, s_1, \dots, s_n)} = Z^{-1} \beta \sum_{s_1, \dots, s_n} \prod_{j=1}^n (\mu_j [P_j] s_j + 1 - s_j).$$

Observing that each product term only depends on one s_j , these formulas collapse to

$$p_1 = Z^{-1} \alpha \prod_{j=1}^n (\lambda_j [P_j] + 1), \quad p_0 = Z^{-1} \beta \prod_{j=1}^n (\mu_j [P_j] + 1)$$

and the distribution condition $p_0 + p_1 = 1$ gives $Z = \alpha \prod_{j=1}^n (\lambda_j [P_j] + 1) + \beta \prod_{j=1}^n (\mu_j [P_j] + 1)$. We finally get

$$k_{\text{on}} = \frac{k_0 \beta \prod_{j=1}^n (\mu_j [P_j] + 1) + k_1 \alpha \prod_{j=1}^n (\lambda_j [P_j] + 1)}{\beta \prod_{j=1}^n (\mu_j [P_j] + 1) + \alpha \prod_{j=1}^n (\lambda_j [P_j] + 1)}. \quad (25)$$

From this formula, it is straightforward to see that k_{on} will actually depend on a protein P_j only if $\lambda_j \neq \mu_j$, that is, when reactions involving P_j have unbalanced speeds and tend to favor either permissive configurations ($\lambda_j > \mu_j$) or non-permissive configurations ($\lambda_j < \mu_j$).

C.3 Higher order interactions

So far we only considered that the P_j were interacting as monomers. If they in fact interact after forming dimers or other complexes, and if such complex-forming reactions are even faster than chromatin dynamics, one can take this into account by simply replacing $[P_j]$ in equation (25) with a function of $[P_j]$ corresponding to the quasi-stationary concentration of the complex. This approximation seems to be relevant to capture the overall dependence of k_{on} on the proteins, the main point being to use a continuous description (e.g. rate equations) for proteins, which are abundant, while keeping a discrete (stochastic) description for chromatin. We chose to replace $[P_j]$ with $[P_j]^{m_j}$ where $m_j > 0$, which gives our model a general Hill-type form. Note that $m_j = 2$ (resp. $m_j = 3$) may represent a correct approximation for P_j interacting as a dimer (resp. a trimer) but in general m_j does not necessarily have to be an integer.

C.4 The case of auto-activation

A this stage, it is possible to implement self-interaction for gene i by taking $\lambda_i \neq \mu_i$ in (25) but this leads to obvious identifiability issues: in stationary state, one cannot really distinguish between auto-activation, auto-inhibition and basal level. To cope with these, we restrict ourselves to auto-activation by setting $c_i = d_i = 0$ and keeping only the relevant chromatin states (C_I^* for all I , and C_I for I such that $i \notin I$). The system still has a unique stationary distribution and the formula for k_{on} corresponds to the case $\mu_i = 0$ in (25). Then, starting from the fact that auto-activation is only relevant when the basal level is small enough (for a bistable behaviour to be possible), we take the limit $\alpha \ll 1$ while keeping $\alpha\lambda_i$ fixed: the formula becomes

$$k_{\text{on}} = \frac{k_0\beta \prod_{j \neq i} (\mu_j [P_j]^{m_j} + 1) + k_1\alpha\lambda_i [P_i]^{m_i} \prod_{j \neq i} (\lambda_j [P_j]^{m_j} + 1)}{\beta \prod_{j \neq i} (\mu_j [P_j]^{m_j} + 1) + \alpha\lambda_i [P_i]^{m_i} \prod_{j \neq i} (\lambda_j [P_j]^{m_j} + 1)} \quad (26)$$

where $m_i > 0$ if gene i activates itself and $m_i = 0$ otherwise.

C.5 Parameterization for inference

Parameters of equation (26) are still clearly not identifiable: in order to get a more minimal form, we introduce the following parameterization: $s_j = \mu_j^{-1/m_j}$, $\theta_j = \log(\lambda_j/\mu_j)$ for all $j \neq i$, and $s_i = (\beta/\alpha)^{1/m_i}$, $\theta_i = \log(\lambda_i)$. After simplifying (26), we obtain

$$k_{\text{on}} = \frac{k_0 + k_1\Phi([P_i]/s_i)^{m_i}}{1 + \Phi([P_i]/s_i)^{m_i}}$$

where

$$\Phi = \exp(\theta_i) \prod_{j \neq i} \frac{1 + \exp(\theta_j)([P_j]/s_j)^{m_j}}{1 + ([P_j]/s_j)^{m_j}}.$$

The new parameters have an intuitive meaning: s_j can be seen as a threshold for the influence by protein j , and θ_j characterizes this influence via its sign and absolute value ($\theta_j = 0$ implying that k_{on} does not depend on protein j), with the exception that s_i and θ_i aggregate a basal behaviour and an auto-activation strength.

Finally, we recall the notation $y_j = [P_j]$ and reintroduce the index i of the gene of interest and add it to each parameter. Hence, for every gene i , the function $k_{\text{on},i}$ is defined by:

$$k_{\text{on},i}(y) = \frac{k_{0,i} + k_{1,i}\Phi_i(y)(y_i/s_{i,i})^{m_{i,i}}}{1 + \Phi_i(y)(y_i/s_{i,i})^{m_{i,i}}} \quad (27)$$

with

$$\Phi_i(y) = \exp(\theta_{i,i}) \prod_{j \neq i} \frac{1 + \exp(\theta_{i,j})(y_j/s_{i,j})^{m_{i,j}}}{1 + (y_j/s_{i,j})^{m_{i,j}}}. \quad (28)$$

In our statistical framework, we assume that parameters $k_{0,i}$, $k_{1,i}$, $m_{i,j}$ and $s_{i,j}$ are known and we focus on inferring the matrix $\theta = (\theta_{i,j}) \in \mathcal{M}_n(\mathbb{R})$, which is similar to the interaction matrix in usual gene network inference methods.

C.6 Explicit distribution for an auto-activation model

Here we derive the stationary distribution for a self-activating gene. For simplicity, we drop the i index. In this model, k_{off} is constant and we assume that there are some constants $\Phi \geq 0$, $m \geq 0$, $s > 0$ and $k_1 \gg k_0 > 0$ such that k_{on} has the form

$$k_{\text{on}}(y) = \frac{k_0 + k_1\Phi(y/s)^m}{1 + \Phi(y/s)^m}$$

so the stationary distribution can directly be used in the Hartree approximation of the network model (27), recalling that Φ has to be independent of the gene's own protein but can depend on others. Letting $c = (k_1 - k_0)/(md_1) > 0$, we are in the case of the explicit solution (21) with

$$\varphi(y) = c \log \left(y^{\frac{k_0}{d_1 c}} + \frac{\Phi}{s^m} y^{\frac{k_1}{d_1 c}} \right) + \frac{k_{\text{off}}}{d_1} \log(1 - y)$$

so the protein distribution is

$$\underline{u}(y) = Z^{-1} y^{-1} \left(y^{\frac{k_0}{d_1 c}} + \frac{\Phi}{s^m} y^{\frac{k_1}{d_1 c}} \right)^c (1 - y)^{\frac{k_{\text{off}}}{d_1} - 1}. \quad (29)$$

To get a fully explicit result, i.e. to compute Z , we shall assume that c is a positive integer. If it is not, one can get a satisfying approximation by taking $c = \lceil (k_1 - k_0)/(md_1) \rceil$. Then, expanding (29) using the binomial theorem, we obtain

$$Z = \sum_{r=0}^c \binom{c}{r} B(a_r, b) (\Phi/s^m)^r$$

where $a_r = ((c - r)k_0 + rk_1)/(d_1 c)$ and $b = k_{\text{off}}/d_1$, and a probabilistic representation of \underline{u} in terms of a mixture of Beta distributions:

$$\underline{u}(y) = \sum_{r=0}^c p_r f_r(y) \quad (30)$$

where $f_r(y) = y^{a_r - 1} (1 - y)^{b - 1} / B(a_r, b)$ and $p_r = \binom{c}{r} B(a_r, b) (\Phi/s^m)^r / Z$.

The dissociation constant s is clearly redundant with the input Φ . We fix the particular value

$$s = \left(\frac{B(a_c, b)}{B(a_0, b)} \right)^{\frac{1}{mc}} = \left(\frac{B(k_1/d_1, k_{\text{off}}/d_1)}{B(k_0/d_1, k_{\text{off}}/d_1)} \right)^{\frac{d_1}{k_1 - k_0}} \quad (31)$$

for which the arbitrary neutral case $\Phi = 1$ is ‘‘symmetric’’, i.e. $p_0 = p_c$. Note that s actually only depends on the fundamental parameters k_0 , k_1 , k_{off} and d_1 (and not on c nor m). Figure 10 shows some examples of the resulting distribution, which can be bimodal or not, depending on the value of c (or equivalently, m) when all other parameters are fixed.

Appendix D EM algorithm for network inference

D.1 EM algorithm for MAP estimation

Here we briefly recall the formulation of the Expectation-Maximization (EM) algorithm for *maximum a posteriori* (MAP) estimation. Consider the probabilistic hierarchical model defined by the distribution of proteins $p(y|\theta)$, the distribution of mRNA given proteins $p(x|y, \theta)$, and a prior distribution $p(\theta)$ on the parameters. Assuming we only observe x , we want to infer θ by MAP estimation, that is, find a mode – hopefully the highest – of the posterior distribution $p(\theta|x)$, which satisfies by Baye's rule:

$$p(\theta|x) = \int p(\theta, y|x) dy \quad \text{where} \quad p(\theta, y|x) = p(y|\theta) p(x|y, \theta) \frac{p(\theta)}{p(x)}.$$

As $p(\theta|x)$ has a too complex expression to be efficiently maximized, the EM algorithm rather uses $\ell_\theta(x, y) = \log(p(\theta, y|x))$ by iteratively computing $\theta^{t+1} = \arg \max_\theta \{Q(\theta, \theta^t)\}$ given θ^t , where

$$\theta \mapsto Q(\theta, \theta^t) = \int \ell_\theta(x, y) p(y|x, \theta^t) dy. \quad (32)$$




Figure 10: Protein stationary distributions (red curves) from the auto-activation model for different values of the input Φ , with s set as in (31). The blue curves indicate the underlying weighted Beta distributions in each mixture. (A) $\Phi = \exp(-2)$, (B) $\Phi = \exp(0) = 1$, (C) $\Phi = \exp(2)$. The distribution tends to be strongly bimodal for small c values, while large values make the distribution close to the unimodal no-feedback case (constant k_{on}). Parameters are $k_0 = 0.25$, $k_1 = 1.25$, $k_{\text{off}} = 7.5$, $d_1 = 0.1$ and, only used for scaling, $s_1 = 10$, $d_0 = 0.5$, $s_0 = 10^3$.

A well-known result states that at each step we in fact maximize a lower bound of $p(\theta | x)$, which is the key point of the algorithm and makes it a particular case of “variational method” (see [79] for example). Now, since $p(x)$ (resp. $p(\theta)$) does not depend on θ (resp. y), it turns out that

$$\arg \max_{\theta} \{Q(\theta, \theta^t)\} = \arg \max_{\theta} \{\bar{Q}(\theta, \theta^t) - g(\theta)\}$$

where $g(\theta) = -\log(p(\theta))$ and $\bar{Q}(\theta, \theta^t) = \int [\log p(y | \theta) + \log p(x | y, \theta)] p(y | x, \theta^t) dy$ is the more standard quantity that appears in the “frequentist” EM algorithm for maximum likelihood estimation. Hence, considering a prior on θ simply results in adding a penalization term $g(\theta)$ during the M step in the algorithm.

For example, if we assume that $\theta_{i,j}$ for $i \neq j$ are independent and follow Laplace distributions, i.e. $p(\theta) = \prod_{i \neq j} \frac{\lambda}{2} \exp(-\lambda |\theta_{i,j}|)$, then $g(\theta) = \lambda \sum_{i \neq j} |\theta_{i,j}| + C$ where $C = n(n-1) \log(2/\lambda)$. Since C does not depend on θ , this is equivalent to the standard L^1 (lasso) penalization, which is well known to enforce the sparsity of the network.

D.2 Custom prior on the interactions

Here we consider a custom prior to deal with oriented interactions. Indeed, for every pair of nodes $\{i, j\}$ there are two possible interactions with respective parameters $\theta_{i,j}$ and $\theta_{j,i}$, but it is likely that only one is actually present in the true network. Hence, we want $\theta_{i,j}$ and $\theta_{j,i}$ to “compete” against each other so that only one is nonzero after MAP estimation, unless there is enough evidence in the data that both interactions are present. To this aim, we define the following prior:

$$p(\theta) \propto \exp \left(-\lambda \sum_{i \neq j} |\theta_{i,j}| - \lambda \alpha \sum_{i < j} |\theta_{i,j} \theta_{j,i}| \right) \quad (33)$$

with $\lambda, \alpha \geq 0$. Thus α can be seen as a competition parameter, the case $\alpha = 0$ leading to the standard lasso penalization parametrized by λ .

D.3 The algorithm in practice

As visible in (32), the true EM algorithm involves integration against the distribution $p(y|x, \theta)$, which does not allow for direct numerical integration because of the dimension ($y \in \mathbb{R}^n$). To overcome this problem, a first option is Monte Carlo integration – typically by MCMC – leading to a “stochastic EM” algorithm that is slow but accurate if samples are large enough. A faster option consists in approximating $p(y|x, \theta)$ by its highest mode, i.e. by the Dirac mass $\delta_{\hat{y}}$ where $\hat{y} = \arg \max_y \{p(y|x, \theta)\}$. Then it is worth noticing that since $p(y|x, \theta) \propto p(y|\theta)p(x|y, \theta)$, the whole procedure can be seen as performing a coordinate ascent on the function $(\theta, y) \mapsto p(\theta, y|x)$. We chose this option for the examples: it is sometimes called “hard” or “classification” EM, since a particular case leads to the well-known k -means clustering algorithm [80]. Unfortunately, theoretical foundations of the true EM algorithm are lost by the hard EM (we do not maximize a lower bound of $p(\theta|x)$ anymore), but it often gives satisfying results while requiring much less computational time.

In practice, the procedure is the following. Suppose we observe mRNA levels in m independent cells, and let $\mathbf{x}_k \in \mathbb{R}^n$ (resp. $\mathbf{y}_k \in \mathbb{R}^n$) denote the mRNA (resp. protein) levels of cell k . In line with sections D.1-D.2 and letting $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_m)$ and $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ for simplicity, we define the objective function

$$\mathcal{F}(\mathbf{y}, \theta) = \ell(\mathbf{x}, \mathbf{y}, \theta) - g(\theta) \quad (34)$$

where the complete log-likelihood $\ell(\mathbf{x}, \mathbf{y}, \theta)$ and the penalization $g(\theta)$ are given by

$$\ell(\mathbf{x}, \mathbf{y}, \theta) = \sum_{k=1}^m \log(u(\mathbf{y}_k, \theta)) + \log(v(\mathbf{x}_k, \mathbf{y}_k, \theta)) \quad \text{and} \quad g(\theta) = \lambda \sum_{i \neq j} |\theta_{i,j}| + \lambda \alpha \sum_{i < j} |\theta_{i,j} \theta_{j,i}|,$$

with $u(y, \theta) = p(y|\theta)$ and $v(x, y, \theta) = p(x|y, \theta)$.

The algorithm then simply consists in iterating the following two steps until convergence:

$$\mathbf{y}^{t+1} = \arg \max_{\mathbf{y}} \{\mathcal{F}(\mathbf{y}, \theta^t)\} \quad (35)$$

$$\theta^{t+1} = \arg \max_{\theta} \{\mathcal{F}(\mathbf{y}^{t+1}, \theta)\} \quad (36)$$

The “approximate E step” (35) can be performed using a standard gradient method since u and v are smooth functions of y . The “penalized M step” (36) is a non-smooth maximization problem since g is non-smooth, but it can be performed using a proximal gradient method detailed in the next section. The form of $\ell(\mathbf{x}, \mathbf{y}, \theta)$ is such that we just need to compute $\nabla \log u$ and $\nabla \log v$.

The formulas for u and v derived from the normalized model are given in the main text: they can be applied once the data has been normalized, i.e. after dividing each mRNA i level by $s_{0,i}/d_{0,i}$. In the bursty regime, this scale parameter is neither identifiable nor necessary. Indeed, as explained in the main text, the “Beta-like” distributions collapse to “Gamma-like” ones which we provide below, and for which the scale parameter is identifiable.

D.3.1 Likelihood form in the basic case

In the basic case (no self-interaction), we have:

$$\log(u(y, \theta)) = \sum_{i=1}^n (a_i(y, \theta) - 1) \log(y_i) - b_i(y, \theta) y_i + a_i(y, \theta) \log(b_i(y, \theta)) - \log \Gamma(a_i(y, \theta))$$

and

$$\log(v(x, y, \theta)) = \sum_{i=1}^n (\tilde{a}_i(y, \theta) - 1) \log(x_i) - \tilde{b}_i(y, \theta) x_i + \tilde{a}_i(y, \theta) \log(\tilde{b}_i(y, \theta)) - \log \Gamma(\tilde{a}_i(y, \theta))$$

where $a_i = k_{\text{on},i}/d_{1,i}$, $b_i = (d_{0,i}/s_{1,i}) \times (k_{\text{off},i}/s_{0,i})$, $\tilde{a}_i = k_{\text{on},i}/d_{0,i}$ and $\tilde{b}_i = k_{\text{off},i}/s_{0,i}$.

D.3.2 Likelihood form in the auto-activation case

In the auto-activation case (section C.6), the formula for $\log v$ is the same but $\log u$ is given by

$$\log(u(y, \theta)) = \sum_{i=1}^n \log \left(\sum_{r=0}^{c_i} w_{i,r}(y, \theta) y_i^{a_{i,r}-1} e^{-b_i y_i} \right) - \log \left(\sum_{r=0}^{c_i} w_{i,r}(y, \theta) \Gamma(a_{i,r}) b_i^{-a_{i,r}} \right)$$

where $c_i = \lceil (k_{1,i} - k_{0,i}) / (d_{1,i} m_{i,i}) \rceil$, $a_{i,r} = ((c_i - r)k_{0,i} + r k_{1,i}) / (d_{1,i} c_i)$ and $w_{i,r} = \binom{c_i}{r} (\Phi_i / s_{i,i}^{m_{i,i}})^r$.

D.3.3 Gradients

Explicit computation of the gradients is then straightforward (e.g. with $k_{\text{off},i}$ constant and $k_{\text{on},i}$, Φ_i defined by (27)-(28)) but leads to cumbersome formulas: we implemented them in Scilab and the code is available upon request.

D.4 Proximal gradient method

Here we recall a standard proximal gradient method [81] to solve the M step (36) and provide the proximal operator associated with $g(\theta)$. Note that the method seems to converge in practice, even if g is not convex. It is based on the update

$$\theta^{(k+1)} = \text{prox}_\gamma \left(\theta^{(k)} + \gamma \nabla_\theta \ell(\mathbf{x}, \mathbf{y}, \theta^{(k)}) \right)$$

where $\gamma > 0$ is a step size (learning rate) and prox_γ is the proximal operator associated with $g(\theta)$, defined on $\Theta \simeq \mathbb{R}^{n^2-n}$ by

$$\text{prox}_\gamma(\tau) = \arg \min_{\theta \in \Theta} \left\{ g(\theta) + \frac{1}{2\gamma} \sum_{i \neq j} (\theta_{i,j} - \tau_{i,j})^2 \right\}.$$

In fact, for any $i, j \in \{1, \dots, n\}$ such that $i \neq j$, one can see that $\theta_{i,j}$ and $\theta_{j,i}$ appear in the minimized quantity as independent of all other θ components. Hence, one just has to compute

$$\text{prox}_\gamma(\tau_1, \tau_2) = \arg \min_{(\theta_1, \theta_2) \in \mathbb{R}^2} \left\{ \lambda (|\theta_1| + |\theta_2| + \alpha |\theta_1 \theta_2|) + \frac{1}{2\gamma} \left((\theta_1 - \tau_1)^2 + (\theta_2 - \tau_2)^2 \right) \right\}$$

and use it for any $(\tau_1, \tau_2) = (\tau_{i,j}, \tau_{j,i}) \in \mathbb{R}^2$ to obtain the corresponding components of $\text{prox}_\gamma(\tau)$. Then, letting $\varepsilon = \lambda\gamma$ and assuming γ small enough such that $\alpha\varepsilon < 1$, we obtain

$$\text{prox}_\gamma(\tau_1, \tau_2) = \frac{1}{1 - (\alpha\varepsilon)^2} (h_1, h_2)$$

with 9 cases for the value of (h_1, h_2) depending on (τ_1, τ_2) , given by:

$$\begin{aligned}
1. & \begin{cases} \tau_1 > \varepsilon \\ \tau_1 > \varepsilon(1 + \alpha(\tau_2 - \varepsilon)) \\ \tau_2 > \varepsilon(1 + \alpha(\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 - \varepsilon(1 + \alpha(\tau_2 - \varepsilon)) \\ h_2 = \tau_2 - \varepsilon(1 + \alpha(\tau_1 - \varepsilon)) \end{cases} \\
2. & \begin{cases} \tau_1 > \varepsilon \\ |\tau_2| \leq \varepsilon(1 + \alpha(\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 - \varepsilon \\ h_2 = 0 \end{cases} \\
3. & \begin{cases} \tau_1 > \varepsilon \\ \tau_1 > \varepsilon(1 + \alpha(-\tau_2 - \varepsilon)) \\ \tau_2 < -\varepsilon(1 + \alpha(\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 - \varepsilon(1 + \alpha(-\tau_2 - \varepsilon)) \\ h_2 = \tau_2 + \varepsilon(1 + \alpha(\tau_1 - \varepsilon)) \end{cases} \\
4. & \begin{cases} |\tau_1| \leq \varepsilon(1 + \alpha(-\tau_2 - \varepsilon)) \\ \tau_2 < -\varepsilon \end{cases} \Rightarrow \begin{cases} h_1 = 0 \\ h_2 = \tau_2 + \varepsilon \end{cases} \\
5. & \begin{cases} \tau_1 < -\varepsilon \\ \tau_1 < -\varepsilon(1 + \alpha(-\tau_2 - \varepsilon)) \\ \tau_2 < -\varepsilon(1 + \alpha(-\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 + \varepsilon(1 + \alpha(-\tau_2 - \varepsilon)) \\ h_2 = \tau_2 + \varepsilon(1 + \alpha(-\tau_1 - \varepsilon)) \end{cases} \\
6. & \begin{cases} \tau_1 < -\varepsilon \\ |\tau_2| \leq \varepsilon(1 + \alpha(-\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 + \varepsilon \\ h_2 = 0 \end{cases} \\
7. & \begin{cases} \tau_1 < -\varepsilon \\ \tau_1 < -\varepsilon(1 + \alpha(\tau_2 - \varepsilon)) \\ \tau_2 > \varepsilon(1 + \alpha(-\tau_1 - \varepsilon)) \end{cases} \Rightarrow \begin{cases} h_1 = \tau_1 + \varepsilon(1 + \alpha(\tau_2 - \varepsilon)) \\ h_2 = \tau_2 - \varepsilon(1 + \alpha(-\tau_1 - \varepsilon)) \end{cases} \\
8. & \begin{cases} |\tau_1| \leq \varepsilon(1 + \alpha(\tau_2 - \varepsilon)) \\ \tau_2 > \varepsilon \end{cases} \Rightarrow \begin{cases} h_1 = 0 \\ h_2 = \tau_2 - \varepsilon \end{cases} \\
9. & \begin{cases} |\tau_1| \leq \varepsilon \\ |\tau_2| \leq \varepsilon \end{cases} \Rightarrow \begin{cases} h_1 = 0 \\ h_2 = 0 \end{cases}
\end{aligned}$$

These 9 cases form a partition of \mathbb{R}^2 and are represented in Figure 11. One can check that the case $\alpha = 0$ collapses to the usual proximal operator associated with lasso penalization.

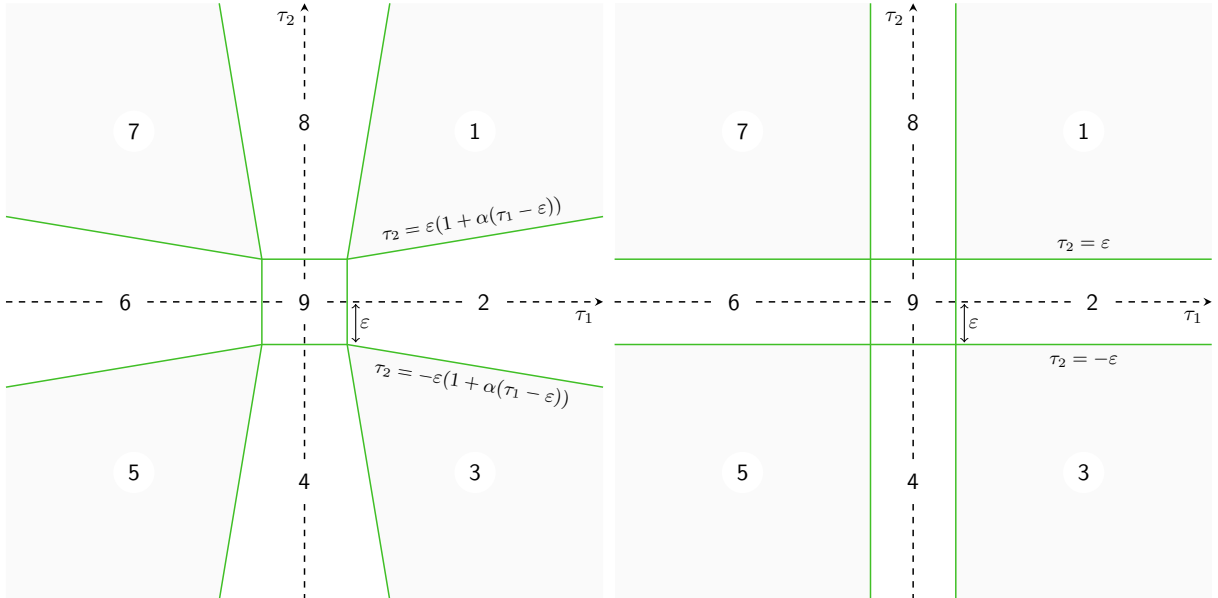


Figure 11: Partition of \mathbb{R}^2 associated with the proximal operator, for $\alpha > 0$ (left) and $\alpha = 0$ (right). Gray areas correspond to a usual gradient and white areas correspond to a threshold.

To obtain the results of Fig 8, we used $\lambda = 10$, $\alpha = 5$ and $\gamma = 10^{-4}$. In a broader context, one may use standard cross-validation to derive appropriate values for λ and α .

Appendix E Dealing with real data

In this section, we propose a pre-processing phase which would be required in order to apply our network inference method to real data. The first step ensures that the approximate likelihood is well-defined, while the second step consists in estimating the basal parameters appearing in functions $k_{\text{on},i}$, $k_{\text{off},i}$. Please note that inferring real networks from real data is beyond the scope of this paper and will be the subject of future papers.

E.1 Spreading zeros

The likelihood does not accept exact zeros (cf. section D.3). This is not a problem with continuous-type data (for instance, based on fluorescence measurements), but it becomes one when dealing with counts (e.g. RNA-Seq). We propose to replace such zeros with relevant positive values. Recall that the PDMP focuses on the promoter and neglects the local molecular noise at the mRNA level. It is therefore natural to consider that, given a value $M > 0$ of mRNA level in the PDMP, the actual number m of molecules in the cell is drawn from the Poisson distribution $\mathcal{P}(M)$. Then, a possible way to replace zeros is to go backwards, i.e. to draw a value M from the PDMP distribution conditioned to $m = 0$. Namely, we propose the following procedure to be applied independently for each gene:

1. Infer a gamma distribution $\gamma(a, b)$ (as a local approximation of the PDMP) from the whole data (possibly at a given time-point) using the standard method of moments;
2. Replace zeros with independent samples from the distribution $\gamma(a, b + 1)$, conditioned to be smaller than the smallest positive value that was measured.

This procedure ensures that zeros are replaced with very small values and that no artificial correlation is introduced. The distribution $\gamma(a, b + 1)$ comes from the fact that, if $\mathcal{L}(M) = \gamma(a, b)$ and $\mathcal{L}(m|M) = \mathcal{P}(M)$, then a simple computation gives $\mathcal{L}(M|m = 0) = \gamma(a, b + 1)$.

E.2 Estimating basal parameters

Here we describe a heuristic method to estimate the model-specific parameters (i.e. everything but the matrix θ) when they cannot be measured through *ad hoc* experiments, in the case of the auto-activation form (27)-(28). Once again we refer to section D.3. Note that for the mechanistic approach to be relevant, one should know at least the ratio $d_{0,i}/d_{1,i}$, which can be obtained by measuring mRNA and protein half-lives. When even this is unavailable, we propose to use the default value $d_{0,i}/d_{1,i} = 5$ (mean value derived from the literature, cf. main text).

The main idea consists in noticing that, when protein i is described by the auto-activation model (27)-(28) (thus following the distribution (30)), mRNA i in quasi-steady state happens to be well described by the same distribution class as (30), with the same $m_{i,i}$ and other parameters being divided by $d_{0,i}/d_{1,i}$. More precisely, we perform the following steps:

1. Estimate $\tilde{a}_{0,i} = k_{0,i}/d_{0,i}$, $\tilde{a}_{1,i} = k_{1,i}/d_{0,i}$, $\tilde{b}_i = k_{\text{off},i}/s_{0,i}$, \tilde{c}_i and Φ from the likelihood

$$f(x_i) \propto \sum_{r=0}^{\tilde{c}_i} \Phi^r x_i^{(1-r/\tilde{c}_i)\tilde{a}_{0,i} + (r/\tilde{c}_i)\tilde{a}_{1,i} - 1} e^{-\tilde{b}_i x_i}.$$

This can be done for instance using an EM algorithm for each value of \tilde{c}_i in some range (e.g. $\tilde{c}_i = 1, 2, \dots, 10$), and then choosing the “arg max” tuple $(\tilde{a}_{0,i}, \tilde{a}_{1,i}, \tilde{b}_i, \tilde{c}_i, \Phi)$. Afterwards, $\tilde{a}_{0,i}$, $\tilde{a}_{1,i}$, \tilde{b}_i and \tilde{c}_i are stored (Φ only serves this step).

2. Consistently with the definition of the model, we set $m_{i,i} = (\tilde{a}_{1,i} - \tilde{a}_{0,i})/\tilde{c}_i$,

$$a_{0,i} = \frac{k_{0,i}}{d_{1,i}} = \frac{d_{0,i}}{d_{1,i}} \cdot \tilde{a}_{0,i}, \quad a_{1,i} = \frac{k_{1,i}}{d_{1,i}} = \frac{d_{0,i}}{d_{1,i}} \cdot \tilde{a}_{1,i}, \quad c_i = \frac{k_{1,i} - k_{0,i}}{d_{1,i}m_{i,i}} = \frac{d_{0,i}}{d_{1,i}} \cdot \tilde{c}_i$$

and we choose $b_i = \frac{d_{0,i}}{d_{1,i}} \cdot \tilde{b}_i$. Then we define, as an approximation of (31) in the bursty regime,

$$s_{i,i} = \frac{1}{b_i} \left(\frac{\Gamma(a_{1,i})}{\Gamma(a_{0,i})} \right)^{1/(a_{1,i} - a_{0,i})}.$$

Note that such b_i is not the “true” value regarding section D.3.2, as we would need to know $\frac{d_{1,i}}{s_{1,i}}$ to apply the formula $b_i = \frac{d_{1,i}}{s_{1,i}} \cdot \frac{d_{0,i}}{d_{1,i}} \cdot \tilde{b}_i$. Fortunately, the network inference does not depend on this scale parameter since the Hill threshold $s_{i,i}$ is proportional to $1/b_i$.

3. Last step consists in extrapolating $m_{i,i}$ and $s_{i,i}$ to the remaining unknown parameters $m_{i,j}$ and $s_{i,j}$ (describing how gene j influences gene i). Since the crucial point is their coherence with respect to the range of protein j , a relevant choice without additional knowledge is, for all $i \neq j$,

$$m_{i,j} = m_{j,j} \quad \text{and} \quad s_{i,j} = s_{j,j}.$$

Appendix F Parameter values

F.1 Models

Table 3: General parameters used in the examples. The $s_{i,j}$ correspond to the normalized model: counterparts in absolute protein numbers are $\bar{s}_{i,j} = s_{i,j} \times (s_0 s_1)/(d_0 d_1) = 2 \times 10^3$ for $i \neq j$ and $\bar{s}_{i,i} = s_{i,i} \times (s_0 s_1)/(d_0 d_1) = 1.9 \times 10^4$.

Parameter	Value	Units
s_0	10^3	mRNA \cdot h ⁻¹
s_1	10	protein \cdot h ⁻¹ \cdot mRNA ⁻¹
d_0	0.5	h ⁻¹
d_1	0.1	h ⁻¹
k_0	0.34	h ⁻¹
k_1	2.15	h ⁻¹
k_{off}	10	h ⁻¹
$m_{i,j}$	2 for $i \neq j$	–
$m_{i,i}$	2 (Fig 5, 6) or 3 (Fig 8)	–
$s_{i,j}$	0.01 for $i \neq j$	proteins (normalized)
$s_{i,i}$	0.095 from eq. (31)	proteins (normalized)

Table 4: Network parameters used in the examples.

Fig 5, 6	$\theta_{1,1}$	$\theta_{1,2}$	$\theta_{2,1}$	$\theta_{2,2}$
	4	-8	-8	4
Fig 8	$\theta_{1,1}$	$\theta_{1,2}$	$\theta_{2,1}$	$\theta_{2,2}$
1	0	0	0	0
2	0	0	1	0
3	0	1	0	0
4	-0.1	1	1	-0.1
5	0	0	-1	0
6	0	-1	0	0
7	0	-1	-1	0

F.2 Results

Table 5: Inferred network parameters used to generate Fig 8b. Each row refers to one of the ten datasets generated for testing. Colors indicate whether the parameters represent the correct topology (blue) or another one (orange) regarding the true networks (cf. Table 4).

Network 1			Network 2			Network 3			Network 4		
	$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$
1	0	0	1	0	0.13	1	0	0	1	0.28	0.22
2	0	0	2	0	0.18	2	0.12	0	2	0.25	0.15
3	0	0	3	0	0.14	3	0.26	0.10	3	0.23	0.21
4	0	0	4	0	0.17	4	0.19	0	4	0.20	0.15
5	0.10	0	5	0	0	5	0.21	0	5	0.18	0.23
6	0	0	6	0	0.20	6	0.36	0.15	6	0.23	0.27
7	0	0	7	0	0.18	7	0.11	0	7	0.17	0.16
8	0	0	8	0	0.11	8	0.11	0	8	0.14	0.13
9	0	0	9	0	0.15	9	0.19	0	9	0.29	0.21
10	0	0	10	0	0	10	0.12	0	10	0.24	0.21

Network 5			Network 6			Network 7		
	$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$
1	0	0	1	-0.14	0	1	-0.20	-0.24
2	0	-0.12	2	-0.11	0	2	-0.34	-0.34
3	0	-0.10	3	-0.13	0	3	-0.17	-0.14
4	0	-0.16	4	-0.15	0	4	-0.23	-0.20
5	-0.10	-0.21	5	-0.21	0	5	-0.17	-0.18
6	0	-0.18	6	-0.16	0	6	-0.19	-0.16
7	0	-0.13	7	-0.29	-0.09	7	-0.21	-0.21
8	0	-0.09	8	0	0	8	-0.30	-0.31
9	0	0	9	-0.13	0	9	-0.11	-0.13
10	0	-0.18	10	-0.21	0	10	-0.11	-0.17

Table 6: Example of inferred networks in the presence of dropouts (30% of the whole dataset) generated by applying a Poisson noise and then a threshold to the “perfect” data. Such zeros were replaced using the procedure described in section E.1 before inferring the networks.

Network 1			Network 2			Network 3			Network 4		
	$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$
1	0	0	1	0	0.27	1	0.31	0	1	0.43	0.41
2	0	0.10	2	0	0.16	2	0.23	0	2	0.44	0.26
3	0	0	3	0	0.31	3	0.32	0.12	3	0.36	0.28
4	0	0	4	0	0.44	4	0.30	0	4	0.44	0.23
5	0.10	0	5	0	0.27	5	0.39	0.11	5	0.38	0.51
6	0	0	6	0	0.55	6	0.40	0.16	6	0.42	0.48
7	0.09	0.16	7	0	0.23	7	0.19	0.11	7	0.48	0.34
8	0	0	8	0	0.21	8	0.14	0	8	0.36	0.32
9	0	0	9	0	0.36	9	0.35	0.16	9	0.59	0.43
10	0	0	10	0	0	10	0.19	0	10	0.41	0.34

Network 5			Network 6			Network 7		
	$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$		$\theta_{1,2}$	$\theta_{2,1}$
1	0	0	1	-0.18	0	1	-0.22	-0.24
2	0	-0.26	2	-0.25	0	2	-0.5	-0.50
3	-0.09	-0.21	3	-0.28	0	3	-0.25	-0.21
4	-0.11	-0.35	4	-0.25	0	4	-0.26	-0.22
5	-0.30	-0.47	5	-0.34	0	5	-0.12	-0.14
6	0	-0.32	6	-0.26	0	6	-0.20	-0.18
7	0	-0.22	7	-0.47	-0.18	7	-0.30	-0.29
8	0	-0.19	8	-0.09	0	8	-0.44	-0.45
9	0	-0.27	9	-0.22	0	9	-0.15	-0.18
10	0	-0.38	10	-0.38	-0.12	10	-0.11	-0.17

References

- [1] M. Hecker, S. Lambeck, S. Toepfer, E. van Someren, and R. Guthke, “Gene regulatory network inference: data integration in dynamic models—a review,” *BioSystems*, vol. 96, no. 1, pp. 86–103, 2009.
- [2] I. Kanter and T. Kalisky, “Single cell transcriptomics: methods and applications,” *Frontiers in Oncology*, vol. 5, p. 53, 2015.
- [3] F. Tang, C. Barbacioru, Y. Wang, E. Nordman, C. Lee, N. Xu, X. Wang, J. Bodeau, B. B. Tuch, A. Siddiqui, K. Lao, and M. A. Surani, “mRNA-Seq whole-transcriptome analysis of a single cell,” *Nature Methods*, vol. 6, no. 5, pp. 377–82, 2009.
- [4] A. Wagner, A. Regev, and N. Yosef, “Revealing the vectors of cellular identity with single-cell genomics,” *Nature Biotechnology*, vol. 34, no. 11, pp. 1145–1160, 2016.
- [5] S. Huang, “Non-genetic heterogeneity of cells in development: more than just noise,” *Development*, vol. 136, no. 23, pp. 3853–3862, 2009.
- [6] A. Eldar and M. B. Elowitz, “Functional roles for noise in genetic circuits,” *Nature*, vol. 467, no. 7312, pp. 167–173, 2010.
- [7] H. Dueck, J. Eberwine, and J. Kim, “Variation is function: Are single cell differences functionally important?,” *Bioessays*, vol. 38, pp. 172–180, 2015.
- [8] A. Mizeranschi, H. Zheng, P. Thompson, and W. Dubitzky, “Evaluating a common semi-mechanistic mathematical model of gene-regulatory networks,” *BMC Systems Biology*, vol. 9, no. 5, pp. 1–12, 2015.
- [9] H. Matsumoto, H. Kiryu, C. Furusawa, M. S. H. Ko, S. B. H. Ko, N. Gouda, T. Hayashi, and I. Nikaido, “Scode: An efficient regulatory network inference algorithm from single-cell rna-seq during differentiation,” *Bioinformatics*, vol. 33, no. 15, pp. 2314–2321, 2017.
- [10] O. Symmons and A. Raj, “What’s luck got to do with it: Single cells, multiple fates, and biological nondeterminism,” *Mol Cell*, vol. 62, no. 5, pp. 788–802, 2016.
- [11] B. Munsky, B. Trinh, and M. Khammash, “Listening to the noise: random fluctuations reveal gene network parameters,” *Molecular Systems Biology*, vol. 5, no. 1, pp. 1–7, 2009.
- [12] C. Zimmer, S. Sahle, and J. Pahle, “Exploiting intrinsic fluctuations to identify model parameters,” *IET Systems Biology*, vol. 9, no. 2, pp. 64–73, 2015.
- [13] X. Cai, J. A. Bazerque, and G. B. Giannakis, “Inference of gene regulatory networks with sparse structural equation models exploiting genetic perturbations,” *PLOS Computational Biology*, vol. 9, no. 5, pp. 1–13, 2013.
- [14] D. Djordjevic, A. Yang, A. Zadorian, K. Rungrueecharoen, and J. W. K. Ho, “How difficult is inference of mammalian causal gene regulatory networks?,” *PLoS One*, vol. 9, no. 11, pp. 1–10, 2014.
- [15] M. T. Angulo, J. A. Moreno, G. Lippner, A.-L. Barabási, and Y.-Y. Liu, “Fundamental limitations of network reconstruction from temporal data,” *J R Soc Interface*, vol. 14, no. 127, pp. 1–6, 2017.
- [16] V. Moignard, S. Woodhouse, L. Haghverdi, A. J. Lilly, Y. Tanaka, A. C. Wilkinson, F. Buettner, I. C. Macaulay, W. Jawaid, E. Diamanti, S.-I. Nishikawa, N. Piterman, V. Kouskoff, F. J. Theis, J. Fisher, and B. Göttgens, “Decoding the regulatory network of early blood development from single-cell gene expression measurements,” *Nat Biotechnol*, vol. 33, no. 3, pp. 1–8, 2015.
- [17] A. Ocone, L. Haghverdi, N. Mueller, and F. Theis, “Reconstructing gene regulatory dynamics from high-dimensional single-cell snapshot data,” *Bioinformatics*, vol. 31, no. 12, pp. i89–i86, 2015.
- [18] J. K. Kim and J. C. Marioni, “Inferring the kinetics of stochastic gene expression from single-cell RNA-sequencing data,” *Genome Biology*, vol. 14, p. R7, 2013.
- [19] D. Marbach, J. C. Costello, R. Küffner, N. M. Vega, R. J. Prill, D. M. Camacho, K. R. Allison, DREAM5 Consortium, M. Kellis, J. J. Collins, and G. Stolovitzky, “Wisdom of crowds for robust gene network inference,” *Nat Methods*, vol. 9, no. 8, pp. 796–804, 2012.
- [20] J. M. Raser and E. K. O’Shea, “Control of stochasticity in eukaryotic gene expression,” *Science*, vol. 304, no. 5678, pp. 1811–1814, 2004.
- [21] A. Becskei, B. B. Kaufmann, and A. van Oudenaarden, “Contributions of low molecule number and chromosomal positioning to stochastic gene expression,” *Nat Genet*, vol. 37, no. 9, pp. 937–944, 2005.
- [22] A. Raj, C. S. Peskin, D. Tranchina, D. Y. Vargas, and S. Tyagi, “Stochastic mRNA Synthesis in Mammalian Cells,” *PLoS Biology*, vol. 4, no. 10, pp. 1707–1719, 2006.

- [23] D. M. Suter, N. Molina, D. Gatfield, K. Schneider, U. Schibler, and F. Naef, “Mammalian genes are transcribed with widely different bursting kinetics,” *Science*, vol. 332, no. 6028, pp. 472–474, 2011.
- [24] M. S. H. Ko, “A stochastic model for gene induction,” *Journal of Theoretical Biology*, vol. 153, pp. 181–194, 1991.
- [25] M. S. H. Ko, H. Nakauchi, and N. Takahashi, “The dose dependence of glucocorticoid-inducible gene expression results from changes in the number of transcriptionally active templates,” *The EMBO Journal*, vol. 9, no. 9, pp. 2835–2842, 1990.
- [26] D. R. Larson, “What do expression dynamics tell us about the mechanism of transcription?,” *Curr Opin Genet Dev*, vol. 21, no. 5, pp. 591–599, 2011.
- [27] D. T. Gillespie, “Exact stochastic simulation of coupled chemical reactions,” *The Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [28] J. Dattani and M. Barahona, “Stochastic models of gene transcription with upstream drives: exact solution and sample path characterization,” *Journal of the Royal Society, Interface*, vol. 14, no. 126, pp. 1–20, 2017.
- [29] V. Shahrezaei and P. S. Swain, “Analytical distributions for stochastic gene expression,” *PNAS*, vol. 105, no. 45, pp. 17256–17261, 2008.
- [30] S. Iyer-Biswas, F. Hayot, and C. Jayaprakash, “Stochasticity of gene products from transcriptional pulsing,” *Phys Rev E Stat Nonlin Soft Matter Phys*, vol. 79, pp. 1–9, 2009.
- [31] B. Schwanhäusser, D. Busse, N. Li, G. Dittmar, J. Schuchhardt, J. Wolf, W. Chen, and M. Selbach, “Global quantification of mammalian gene expression control,” *Nature*, vol. 495, pp. 337–342, 2011.
- [32] M. H. A. Davis, “Piecewise-deterministic markov processes: A general class of non-diffusion stochastic models,” *Journal of the Royal Statistical Society*, vol. 46, no. 3, pp. 353–388, 1984.
- [33] A. Crudu, A. Debussche, and O. Radulescu, “Hybrid stochastic simplifications for multiscale gene networks,” *BMC Systems Biology*, vol. 3, no. 1, p. 89, 2009.
- [34] A. Crudu, A. Debussche, A. Muller, and O. Radulescu, “Convergence of stochastic gene networks to hybrid piecewise deterministic processes,” *The Annals of Applied Probability*, vol. 22, no. 5, pp. 1822–1859, 2012.
- [35] Y. T. Lin and T. Galla, “Bursting noise in gene expression dynamics: linking microscopic and mesoscopic models,” *J. R. Soc. Interface*, vol. 13, pp. 1–11, 2016.
- [36] J. Viñuelas, G. Kaneko, A. Coulon, E. Vallin, V. Morin, C. Mejia-Pous, J.-J. Kupiec, G. Beslon, and O. Gandrillon, “Quantifying the contribution of chromatin dynamics to stochastic gene expression reveals long, locus-dependent periods between transcriptional bursts,” *BMC Biology*, vol. 11, no. 1, p. 15, 2013.
- [37] C. Albayrak, C. A. Jordi, C. Zechner, J. Lin, C. A. Bichsel, M. Khammash, and S. Tay, “Digital quantification of proteins and mrna in single mammalian cells,” *Molecular Cell*, vol. 61, pp. 914–924, 2016.
- [38] A. Richard, L. Boullu, U. Herbach, A. Bonnafoux, V. Morin, E. Vallin, A. Guillemin, N. Papili Gao, R. Gunawan, J. Cosette, O. Arnaud, J.-J. Kupiec, T. Espinasse, S. Gonin-Giraud, and O. Gandrillon, “Single-cell-based analysis highlights a surge in cell-to-cell molecular variability preceding irreversible commitment in a differentiation process,” *PLoS Biology*, vol. 14, no. 12, pp. 1–35, 2016.
- [39] O. Boxma, H. Kaspi, O. Kella, and D. Perry, “On/Off Storage Systems with State-Dependent Input, Output, and Switching Rates,” *Probability in the Engineering and Informational Sciences*, vol. 19, pp. 1–14, 2005.
- [40] M. Benaïm, S. Le Borgne, F. Malrieu, and P.-A. Zitt, “Quantitative ergodicity for some switched dynamical systems,” *Electronic Communications in Probability*, vol. 17, no. 56, pp. 1–14, 2012.
- [41] K. M. Ong, J. A. Blackford, Jr, B. L. Kagan, S. S. Simons, Jr, and C. C. Chow, “A theoretical framework for gene induction and experimental comparisons,” *PNAS*, vol. 107, no. 15, pp. 7107–7112, 2010.
- [42] A. Coulon, O. Gandrillon, and G. Beslon, “On the spontaneous stochastic dynamics of a single gene: complexity of the molecular interplay at the promoter,” *BMC Systems Biology*, vol. 4, p. 2, 2010.
- [43] A. Coulon, C. C. Chow, R. H. Singer, and D. R. Larson, “Eukaryotic transcriptional dynamics: from single molecules to cell populations,” *Nat Rev Genet.*, vol. 14, no. 8, pp. 1–13, 2013.

- [44] N. Friedman and O. J. Rando, “Epigenomics and the structure of the living genome,” *Genome Research*, vol. 25, no. 10, pp. 1482–1490, 2015.
- [45] L. Bintu, J. Yong, Y. E. Antebi, K. McCue, Y. Kazuki, N. Uno, M. Oshimura, and M. B. Elowitz, “Dynamics of epigenetic regulation at the single-cell level,” *Science*, vol. 351, no. 6274, pp. 720–724, 2016.
- [46] M. Benaim, S. Le Borgne, F. Malrieu, and P.-A. Zitt, “Qualitative properties of certain piecewise deterministic Markov processes,” *Annales de l’Institut Henri Poincaré - Probabilités et Statistiques*, vol. 51, no. 3, pp. 1040–1075, 2015.
- [47] A. Faggionato, D. Gabrielli, and M. R. Crivellari, “Non-equilibrium thermodynamics of piecewise deterministic markov processes,” *Journal of Statistical Physics*, vol. 137, pp. 259–304, 2009.
- [48] K. Pakdaman, M. Thieullen, and G. Wainrib, “Asymptotic expansion and central limit theorem for multiscale piecewise-deterministic Markov processes,” *Stochastic Processes and their Applications*, vol. 122, pp. 2292–2318, 2012.
- [49] J. Peccoud and B. Ycart, “Markovian Modelling of Gene Product Synthesis,” *Theoretical Population Biology*, vol. 48, pp. 222–234, 1995.
- [50] G.-W. Li and X. S. Xie, “Central dogma at the single-molecule level in living cells,” *Nature*, vol. 475, no. 7356, pp. 308–315, 2011.
- [51] M. Sasai and P. G. Wolynes, “Stochastic gene expression as a many-body problem,” *PNAS*, vol. 100, no. 5, pp. 2374–2379, 2003.
- [52] A. M. Walczak, M. Sasai, and P. G. Wolynes, “Self-consistent proteomic field theory of stochastic gene switches,” *Biophysical Journal*, vol. 88, pp. 828–850, 2005.
- [53] K.-Y. Kim and J. Wang, “Potential energy landscape and robustness of a gene regulatory network: toggle switch,” *PLoS Computational Biology*, vol. 3, no. 3, pp. 565–577, 2007.
- [54] B. Zhang and P. G. Wolynes, “Stem cell differentiation as a many-body problem,” *PNAS*, vol. 111, no. 28, pp. 10185–10190, 2014.
- [55] A. Senecal, B. Munsky, F. Proux, N. Ly, F. E. Braye, C. Zimmer, F. Mueller, and X. Darzacq, “Transcription Factors Modulate c-Fos Transcriptional Bursts,” *Cell Reports*, vol. 8, pp. 75–83, 2014.
- [56] T. Fukaya, B. Lim, and M. Levine, “Enhancer control of transcriptional bursting,” *Cell*, vol. 166, no. 2, pp. 358–368, 2016.
- [57] D. Marbach, R. J. Prill, T. Schaffter, C. Mattiussi, D. Floreano, and G. Stolovitzky, “Revealing strengths and weaknesses of methods for gene network inference,” *PNAS*, vol. 107, no. 14, pp. 6286–6291, 2010.
- [58] J. Gu, Q. Gu, X. Wang, P. Yu, and W. Lin, “Sphinx: modeling transcriptional heterogeneity in single-cell RNA-Seq,” *bioRxiv preprint*, 2015.
- [59] S. Ghazanfar, A. J. Bisogni, J. T. Ormerod, D. M. Lin, and J. Y. H. Yang, “Integrated single cell data analysis reveals cell specific networks and novel coactivation markers,” *BMC Systems Biology*, vol. 10, p. 127, 2016.
- [60] M. Mojtahedi, A. Skupin, J. Zhou, I. G. Castano, R. Y. Leong-Quong, H. H. Chang, A. Giuliani, and S. Huang, “Cell fate decision as high-dimensional critical state transition,” *PLOS Biology*, vol. 14, no. 12, pp. 1–28, 2016.
- [61] C. Sokolik, Y. Liu, D. Bauer, J. McPherson, M. Broeker, G. Heimberg, L. S. Qi, D. A. Sivak, and M. Thomson, “Transcription factor competition allows embryonic stem cells to distinguish authentic signals from noise,” *Cell Systems*, vol. 1, pp. 117 – 129, 2015.
- [62] N. Battich, T. Stoeger, and L. Pelkmans, “Control of transcript variability in single mammalian cells,” *Cell*, vol. 163, no. 7, pp. 1596–1610, 2015.
- [63] S. Chong, C. Chen, H. Ge, and X. S. Xie, “Mechanism of transcriptional bursting in bacteria,” *Cell*, vol. 158, no. 2, pp. 314–326, 2014.
- [64] F. Antoneli, R. C. Ferreira, and M. R. Briones, “A model of gene expression based on random dynamical systems reveals modularity properties of gene regulatory networks,” *Mathematical Biosciences*, vol. 276, pp. 82–100, 2016.
- [65] D. A. Potoyan and P. G. Wolynes, “Dichotomous noise models of gene switches,” *J Chem Phys*, vol. 143, no. 19, p. 195101, 2015.

- [66] P. G. Hufton, Y. T. Lin, T. Galla, and A. J. McKane, “Intrinsic noise in systems with switching environments,” *Phys Rev E*, vol. 93, no. 5, p. 052119, 2016.
- [67] M. Pájaro, A. A. Alonso, I. Otero-Muras, and C. Vázquez, “Stochastic modeling and numerical simulation of gene regulatory networks with protein bursting,” *Journal of Theoretical Biology*, vol. 421, pp. 51–70, 2017.
- [68] J. Teles, C. Pina, P. Edén, M. Ohlsson, T. Enver, and C. Peterson, “Transcriptional Regulation of Lineage Commitment - A Stochastic Model of Cell Fate Decisions,” *PLoS Computational Biology*, vol. 9, no. 8, pp. 1–13, 2013.
- [69] D. Schnoerr, R. Grima, and G. Sanguinetti, “Cox process representation and inference for stochastic reaction-diffusion processes,” *Nature communications*, vol. 7, pp. 1–11, 2016.
- [70] A. Ocone, A. J. Millar, and G. Sanguinetti, “Hybrid regulatory models: a statistically tractable approach to model regulatory network dynamics,” *Bioinformatics*, vol. 29, no. 7, pp. 910–916, 2013.
- [71] A. Pinna, N. Soranzo, and A. de la Fuente, “From knockouts to networks: establishing direct cause-effect relationships through graph analysis,” *PLoS One*, vol. 5, no. 10, pp. 1–8, 2010.
- [72] G. Corre, D. Stockholm, O. Arnaud, G. Kaneko, J. Viñuelas, Y. Yamagata, T. M. A. Neildes-Nguyen, J.-J. Kupiec, G. Beslon, O. Gandrillon, and A. Paldi, “Stochastic Fluctuations and Distributed Control of Gene Expression Impact Cellular Memory,” *PLoS ONE*, vol. 9, no. 12, p. e115574, 2014.
- [73] O. Padovan-Merhar, G. P. Nair, A. G. Biaesch, A. Mayer, S. Scarfone, S. W. Foley, A. R. Wu, L. S. Churchman, A. Singh, and A. Raj, “Single mammalian cells compensate for differences in cellular volume and dna copy number through independent global transcriptional mechanisms,” *Molecular Cell*, vol. 58, no. 2, pp. 339–352, 2015.
- [74] N. A. Hathaway, O. Bell, C. Hodges, E. L. Miller, D. S. Neel, and G. R. Crabtree, “Dynamics and memory of heterochromatin in living cells,” *Cell*, vol. 149, no. 7, pp. 1447–1460, 2012.
- [75] G. Fourel, F. Magdinier, and E. Gilson, “Insulator dynamics and the setting of chromatin domains,” *BioEssays*, vol. 26, no. 5, pp. 523–532, 2004.
- [76] S. Kueng, M. Oppikofer, and S. M. Gasser, “Sir proteins and the assembly of silent chromatin in budding yeast,” *Annual Review of Genetics*, vol. 47, pp. 275–306, 2013.
- [77] S. S. P. Rao, M. H. Huntley, N. C. Durand, E. K. Stamenova, I. D. Bochkov, J. T. Robinson, A. L. Sanborn, I. Machol, A. D. Omer, E. S. Lander, and E. L. Aiden, “A 3d map of the human genome at kilobase resolution reveals principles of chromatin looping,” *Cell*, vol. 159, no. 7, pp. 1665–1680, 2014.
- [78] M. J. Obersriebnig, E. M. H. Pallesen, K. Sneppen, A. Trusina, and G. Thon, “Nucleation and spreading of a heterochromatic domain in fission yeast,” *Nature Communications*, vol. 7, pp. 1–11, 2016.
- [79] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, “An Introduction to Variational Methods for Graphical Models,” *Machine Learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [80] G. Celeux and G. Govaert, “A classification EM algorithm for clustering and two stochastic versions,” Research Report 1364, INRIA, 1991.
- [81] N. Parikh and S. Boyd, “Proximal Algorithms,” *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 123–231, 2013.