

Sparse learning of Markovian population models in random environments

Conference Paper**Author(s):**

Zechner, Christoph; Wadehn, Federico; Köppl, Heinz

Publication date:

2014-09

Permanent link:

<https://doi.org/10.3929/ethz-a-010362110>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IFAC Proceedings Volumes 47(3), <https://doi.org/10.3182/20140824-6-ZA-1003.01974>

Sparse Learning of Markovian Population Models in Random Environments [★]

Christoph Zechner ^{*} Federico Wadehn ^{*} Heinz Koepl ^{*,**}

^{*} Automatic Control Lab, ETH Zurich, 8092 Zurich, Switzerland

^{**} IBM Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract: Markovian population models are suitable abstractions to describe well-mixed interacting particle systems in situation where stochastic fluctuations are significant due to the involvement of low copy particles. In molecular biology, measurements on the single-cell level attest to this stochasticity and one is tempted to interpret such measurements across an isogenic cell population as different sample paths of one and the same Markov model. Over recent years evidence built up against this interpretation due to the presence of cell-to-cell variability stemming from factors other than intrinsic fluctuations. To account for this extrinsic variability, Markovian models in random environments need to be considered and a key emerging question is how to perform inference for such models. We model extrinsic variability by a random parametrization of all propensity functions. To detect which of those propensities have significant variability, we lay out a sparse learning procedure captured by a hierarchical Bayesian model whose evidence function is iteratively maximized using a variational Bayesian expectation-maximization algorithm.

Keywords: Population models, stochastic chemical kinetics, extrinsic variability, sparse Bayesian learning, variational inference.

1. INTRODUCTION

Markovian population models are ubiquitous in biology to capture the temporal change in abundance for different particle types (i.e. species) caused by interactions or transformations among them. Inferring such models from experimental data is at the core of quantitative biology. Reconstructing models of biochemical cellular processes using the principles of chemical kinetics is an important example. Single-cell technologies provide unprecedented means to perform this task, however novel computational methods are required to deal with the complexity of single-cell data. More specifically, such data represent a heterogeneous aggregate of measurements due to the fact that cells are not exactly identical to start with. Thus, apart from stochastic fluctuations intrinsic to the process under study, *extrinsic* sources of variability contribute to the overall heterogeneity (Elowitz et al. [2002], Colman-Lerner et al. [2005]). The single-cell process gets modulated by its local microenvironment, that can refer to intracellular quantities such as initial copy numbers of participating biomolecules (Koepl et al. [2012]) but also cell-level quantities like the cell's local growth condition (Snijder and Pelkmans [2011]), or its cell-cycle stage (Colman-Lerner et al. [2005]). To capture this variability in a computational model is challenging because the true sources and their strength for a specific cell line are yet to be identified. Hence, recent approaches to address extrinsic noise in the inference procedure have to make an educated guess which quantities of a kinetic model are modulated by extrinsic variability (Zechner et al. [2012, 2014], Ruess

et al. [2013], Hasenauer et al. [2011], Shahrezaei et al. [2008]). For instance, one source of extrinsic noise in gene expression that is believed to be significant are ribosome copy number variations. Taking this as a starting point, we recently developed an inference framework that relies on a hierarchical Bayesian model, where some latent states express the extrinsic variability of an actual model quantity (Zechner et al. [2014]). The structure of this Bayesian model is fixed beforehand and hence can not be changed *a posteriori* when the data is incorporated.

Here we lay out an inference framework where the hierarchical dependency structure among model quantities and extrinsic sources is learned from the data. In order to retrieve results that are interpretable and robust with respect to small sample sizes (i.e. number of cells) we apply a sparse Bayesian learning technique (Neal [1996], Tipping [2001], Bishop [2007]) yielding the named dependency structure with a minimal number of edges. In order to reduce the number of unknown parameters, we make use of the marginalized process introduced in Zechner et al. [2014]. To infer the posterior with respect to the extrinsic variability in an efficient manner we employ a variational Bayesian expectation-maximization (EM) procedure (Beal [2003], Shutin et al. [2012], Dempster et al. [1977]). The outlined method assumes the availability of data in terms of complete and noise free sample paths. The method can be generalized to the more realistic incomplete and noisy data case but the necessary computational machinery for that would sidetrack the exposition and occlude the main idea behind this approach. To this end, the work represents a first step towards a model-based understanding of

^{*} C.Z. and H.K. acknowledge the support from the Swiss National Science Foundation, grant no. PP00P2 128503.

how and which concurrent processes modulate a specific cellular process under study *in vivo*.

The remaining part of the paper is structured as follows. In Section 2 we derive the mathematical models and algorithms. We start with a brief introduction to heterogeneous population models (Section 2.2). Subsequently, in Section 2.3 we develop a suitable hierarchical Bayesian model whose inference is discussed in Section 2.4. In Sections 2.5 and 2.6 we address several practical aspects of the algorithm and briefly discuss how it extends to the incomplete data scenario. The algorithm is analyzed and validated in Section 3 using a few case studies.

2. MATHEMATICAL FRAMEWORK

2.1 Notation

Random quantities and their realizations are denoted by upper- and lowercase symbols, respectively. Symbol p and q are used to indicate the exact and the approximating probability density functions (PDFs), respectively and expectations are denoted by $\mathbb{E}[A] = \int ap(a)da$. For convenience we also introduce expectations of the form $\mathbb{E}_a[f(a, b)] = \int f(a, b)p(a)da$, indicating that the expectation of f is only taken with respect to $p(a)$. We denote the Gamma distribution by $\mathcal{G}(\alpha, \beta)$ with α and β as shape and inverse scale parameters. The exponential distribution is denoted $\text{Exp}(\lambda)$ with inverse scale parameter λ . Furthermore, we express time-dependent quantities on intervals $[0, T]$ by bold symbols, e.g., $\mathbf{x} = \{x(t) \mid 0 \leq t \leq T\}$. The symbol $KL[q(x)||p(x)]$ denotes the Kullback-Leibler divergence between PDFs $q(x)$ and $p(x)$, i.e., $KL[q(x)||p(x)] = \int q(x) \ln \frac{q(x)}{p(x)} dx$. Abbreviations CV and SCV stand for coefficient of variation and squared coefficient of variation, respectively.

2.2 Population Models and Chemical Kinetics

We consider a continuous-time Markov chain (CTMC) X describing the dynamics of a stochastic interaction network comprising d species and ν coupled reaction or transformation channels. The latter are associated with a set of real-valued kinetic parameters $C = \{C_j \mid j = 1, \dots, \nu\}$. The species' abundances at time t defines the random state $X(t) = x$, $x \in \mathbb{Z}_{\geq 0}^d$ of the network. Propensity functions corresponding to each channel are general functions of the state and can often be defined through first principles such as the law of mass-action. Throughout the work we assume propensities to be linear functions of their respective rate parameters, i.e. they take the general form $c_i g_i(x)$, where $g_i(x)$ is an arbitrary nonlinear function. Under knowledge of the parameters C , the dynamics of a single cell follows a conditional CTMC $X \mid C$.

Due to extrinsic cell-to-cell variability, acquired single-cell trajectories can not be thought of as being different realizations of a single CTMC. Although sources of extrinsic variability can be diverse, we throughout the work make the assumption that kinetic parameters are the only source of extrinsic variability that enters the cellular process under study. In contrast to previous approaches where specific parameters were subject to extrinsic variability

(Zechner et al. [2014, 2012]), we assign prior variability to every kinetic parameter in the model. Accordingly, we associate to C a probability distribution, i.e., $C \mid (A = a) \sim p(c \mid a)$, with A a set of hyperparameters. With this, the dynamics of the m -th cell of a population is described by a conditional CTMC $X^m \mid (C^m = c)$. We remark that the parameter dimensionality of the heterogeneous CTMC model increases with every considered cell and hence, scales poorly with the population size M . Fortunately, it was recently shown that a CTMC $X \mid C$ can be integrated over C , yielding a marginalized stochastic process $X \mid A$, which directly depends on the hyperparameters A . While fuller details about the construction and simulation of such a process can be found in Zechner et al. [2013, 2014], Aalen et al. [2008], we only introduce the key quantity needed here, i.e., the marginal path likelihood function. We know from Wilkinson [2006] and Kuechler and Sorensen [1997] that the path likelihood function of an observed sample path $\mathbf{x} = \{x(t) \mid 0 \leq t \leq T\}$ is given by

$$p(\mathbf{x} \mid c) \propto \prod_{i=1}^{\nu} c_i^{r_i(\mathbf{x})} e^{-c_i \int_0^T g_i(x(t)) dt}, \quad (1)$$

with $r_i(\mathbf{x})$ the number of reactions of type i that happened in the path \mathbf{x} . Formally, the marginal path likelihood is obtained via the integral

$$p(\mathbf{x} \mid a) = \int p(\mathbf{x} \mid c) p(c \mid a) dc, \quad (2)$$

whose tractability depends on $p(c \mid a)$. For instance, the Gamma distribution was shown to have convenient analytical properties (Zechner et al. [2014]) and furthermore, appears plausible in the context of gene expression Taniguchi et al. [2010]. Throughout the remaining paper we assume

$$p(c \mid a) = \prod_{i=1}^{\nu} \mathcal{G}(c_i \mid \alpha_i, \beta_i), \quad (3)$$

with $a = \{(\alpha_i, \beta_i) \mid i = 1, \dots, \nu\}$ and $\mathcal{G}(c_i \mid \alpha_i, \beta_i)$ as a Gamma distribution over c_i with shape- and inverse scale parameters α_i and β_i , respectively. A suitable measure of a reaction channel's extrinsic variability is the squared coefficient of variation (or normalized variance), which in the Gamma-case is given by

$$\eta_i = \frac{1}{\alpha_i},$$

indicating that one can detect heterogeneity by merely analyzing the shape parameter α_i . Under assumption (3), the marginal path likelihood function is given by (see e.g., Zechner et al. [2014])

$$p(\mathbf{x} \mid a) = \prod_{i=1}^{\nu} p(\mathbf{x} \mid \alpha_i, \beta_i) \propto \prod_{i=1}^{\nu} \frac{\beta_i^{\alpha_i} \Gamma(\alpha_i + r_i(\mathbf{x}))}{\Gamma(\alpha_i)} \left(\beta_i + \int_0^T g_i(x(t)) dt \right)^{-(\alpha_i + r_i(\mathbf{x}))}. \quad (4)$$

Due to the marginalization, the hidden layer corresponding to the kinetic parameters is entirely removed and hence, does not need to be considered in the following derivations.

2.3 Hierarchical Bayesian Modeling

Assume we have given measurements of M cells of a heterogeneous population, i.e., \mathbf{x}^m for $m = 1, \dots, M$.

The extrinsic variability of each reaction channel i can be quantified by inferring the hyperparameters $\{\alpha_i, \beta_i\}$ from those measurements. According to a Bayesian scenario, this is equivalent to finding the posterior distribution

$$p(a | \mathbf{x}^1, \dots, \mathbf{x}^M) \propto \prod_{m=1}^M p(\mathbf{x}^m | a) p(a) = \prod_{m=1}^M \left(\prod_{i=1}^{\nu} p(\mathbf{x}^m | \alpha_i, \beta_i) p(\alpha_i, \beta_i) \right), \quad (5)$$

which hence, factorizes such that

$$p(a | \mathbf{x}^1, \dots, \mathbf{x}^M) = \prod_{i=1}^{\nu} p(\alpha_i, \beta_i | \mathbf{x}^1, \dots, \mathbf{x}^M). \quad (6)$$

Naively, one could just evaluate the individual terms in (6) and check whether the corresponding values of α_i are below a certain threshold, indicating heterogeneity of the associated reaction. However, since those values are only accessible through the noisy measurements \mathbf{x}^m , it is not clear how to choose such a threshold in order to obtain maximally robust results. For instance, the heterogeneity stemming from the intrinsic molecular fluctuations should be "filtered out" and yield a negative detection result. Positive detections are only desired if there is significant evidence in the data. Technically, this corresponds to solving a *sparse* Bayesian learning problem (Neal [1996], Tipping [2001], Bishop [2007]). The key step to achieve sparsity in empirical Bayesian models is to assign suitable prior- and hyperprior distributions to the model quantities. Since detection of heterogeneity is based on only α_i , we chose

$$p(\alpha_i, \beta_i) = p(\alpha_i | \lambda_i) p(\beta_i), \quad (7)$$

where λ_i controls the shape of $p(\alpha_i | \lambda_i)$ and $p(\beta_i)$ is assumed to be flat over the positive domain, such that $p(\alpha_i, \beta_i) \propto p(\alpha_i | \lambda_i)$. The goal is to define $p(\alpha_i | \lambda_i)$ such that the heterogeneity is forced to zero unless there is significant evidence in the data. Accordingly, suitable distributions will emphasize SCVs around zero while also permitting high values. Here we choose $p(\alpha_i | \lambda_i)$ such that $p(\eta_i) = \text{Exp}(\lambda_i)$. A transformation of random variables yields

$$p(\alpha_i | \lambda_i) = \frac{\lambda_i}{\alpha_i^2} e^{-\frac{\lambda_i}{\alpha_i}}. \quad (8)$$

The resulting prior distributions over α_i are illustrated in Fig. 1 for different values of λ_i .

While standard Bayesian approaches rely on *given* prior knowledge, empirical Bayes techniques aim to infer parameters as well as their hyperparameters from data. In our case, this means that in addition to α_i and β_i , also the hyperparameters λ_i are assumed to be unknown and need to be estimated. In order to obtain a fully Bayesian model, we need to specify hyperprior distributions $p(\lambda_i)$. Again, we assume $p(\lambda_i)$ to be flat but remark that an extension to arbitrary distributions is straight-forward. With the model parameters a and their hyperparameters $b = \{\lambda_i | i = 1, \dots, \nu\}$, we aim to compute the posterior distribution

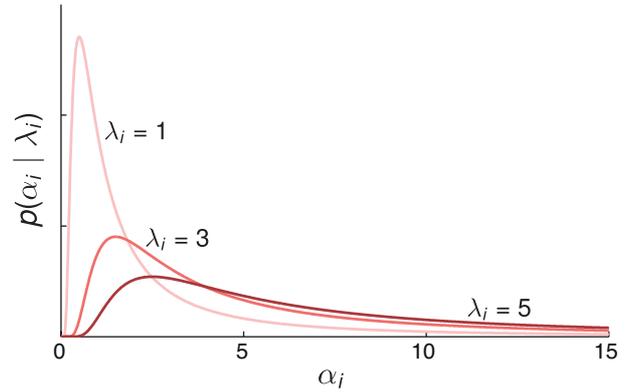


Fig. 1. Prior distributions over α_i for different values of the hyperparameter λ_i . The distributions show a peak for low values of α_i and become more heavy-tailed with increasing λ_i .

$$p(a, b | \mathbf{x}^1, \dots, \mathbf{x}^M) \propto \prod_{i=1}^{\nu} \left(\prod_{m=1}^M p(\mathbf{x}^m | \alpha_i, \beta_i) \right) p(\alpha_i | \lambda_i) p(\lambda_i) = \prod_{i=1}^{\nu} p(\alpha_i, \beta_i, \lambda_i | \mathbf{x}^1, \dots, \mathbf{x}^M) \quad (9)$$

where the r.h.s. of (9) is just the joint distribution over all model quantities. Unfortunately, it turns out that (9) is intractable. In the next section we will develop an variational inference scheme to approximate (9).

2.4 Variational Inference

Variational inference schemes aim to approximate some target posterior $p(z | y)$ by some other distribution $q(z)$. More specifically, one chooses $q(z)$ such as to minimize the Kullback-Leibler divergence (KL) between $q(z)$ and the true distribution. For that sake, note that for every q , the log-evidence function satisfies the decomposition (Beal [2003], Bishop [2007])

$$\ln p(z) = \mathcal{L}[q(z)] + KL[q(z) || p(z | y)], \quad (10)$$

where $\mathcal{L}[q(z)]$ forms a lower bound on $\ln p(z)$ which is given by

$$\mathcal{L}[q(z)] = \int q(z) \ln \frac{p(z, y)}{q(z)} dz. \quad (11)$$

Accordingly, minimizing the KL with respect to q is the same as maximizing its counterpart $\mathcal{L}[q(z)]$, i.e.,

$$q^*(z) = \underset{q(z) \in \mathcal{Q}}{\text{argmax}} \mathcal{L}[q(z)]. \quad (12)$$

It can be seen from (10) and (11) that $\mathcal{L}[q(z)]$ is maximal if and only if $q(z) = p(z | y)$. In order to obtain a tractable $q(z)$, one typically imposes further constraints on its structure. Most commonly, individual components of z are assumed to be independent of each other, i.e.,

$$q(z) = \prod_{l=1}^L q(z_l), \quad (13)$$

also known as the *mean-field* approximation (Beal [2003]). In this case, it can be shown that the optimal variational solution of the individual factors $q(z_i)$ is determined by

$$\ln q^*(z_i) = \mathbb{E}_{j \neq i} [\ln p(z, y)] + \text{const}. \quad (14)$$

where $\mathbb{E}_{j \neq i} [\ln p(z, y)]$ denotes the expectation of the logarithm of the joint distribution, taken with respect to all factors $q(z_j)$ except $q(z_i)$. Since the optimal solution of a particular q -factor depends on all other factors, the mean-field approximation typically induces an iterative inference scheme, where the individual factors are updated in a round-robin fashion. Such schemes stand in close relation with traditional expectation-maximization (EM) algorithms (Dempster et al. [1977]) and accordingly, are often referred to as variational Bayesian EM (VBEM) algorithms (Shutin et al. [2012], Beal [2003]).

In practice, eq. (14) might still be intractable, in which case it is necessary to further restrict the corresponding q -factor. For instance, one could assume $q(z_i)$ to be some parameterized distribution (e.g., a Gaussian with mean and variance) and determine its parameters θ as

$$\theta^* = \operatorname{argmax}_{\theta \in \Theta} \mathbb{E} [\ln p(z, y)], \quad (15)$$

whereas in this case, the expectation is taken with respect to all q -factors. For instance, if one is interested solely in maximum a-posterior (MAP) estimates, $q(z_i)$ can be chosen to be a Dirac-delta function with unknown position.

We will now use the VBEM framework to derive an approximate iterative inference algorithm for the hierarchical Bayesian model from Section 2.3. The goal is to compute an approximate posterior distribution $q(a, b)$ for which we assume that it factorizes as

$$q(a, b) = \prod_{i=1}^{\nu} q(\alpha_i, \beta_i) q(\lambda_i). \quad (16)$$

We remark that in the complete-data scenario considered here, also the true posterior factors over the individual reaction channels $i = 1, \dots, \nu$, however, not over $\{\alpha_i, \beta_i\}$ and λ_i . For analytical simplicity, we further assume $q(\lambda_i) := \delta(\lambda_i - \hat{\lambda}_i)$ with $\hat{\lambda}_i$ as an unknown position parameter. The factor $q(\alpha_i, \beta_i)$ for the i -th reaction channel is determined by

$$\ln q^*(\alpha_i, \beta_i) = \mathbb{E}_{\lambda_i} [\ln p(a, b, \mathbf{x}^1, \dots, \mathbf{x}^M)] + \text{const.}, \quad (17)$$

which becomes

$$\begin{aligned} \ln q^*(\alpha_i, \beta_i) &= \sum_{m=1}^M \ln p(\mathbf{x}^m | \alpha_i, \beta_i) \\ &\quad + \mathbb{E}_{\lambda_i} [\ln p(\alpha_i | \lambda_i)] + \text{const.} \end{aligned} \quad (18)$$

when taking into account the r.h.s. of eq. (9). Together with the marginal path-likelihood function from eq. (4), we further obtain

$$\begin{aligned} \ln q^*(\alpha_i, \beta_i) &= \sum_{m=1}^M \alpha_i \ln \beta_i + \ln \Gamma(\alpha_i + r_i(\mathbf{x}^m)) - \ln \Gamma(\alpha_i) \\ &\quad - (\alpha_i + r_i(\mathbf{x}^m)) \ln \left(\beta_i + \int_0^T g_i(x^m(t)) dt \right) \\ &\quad - \frac{\hat{\lambda}_i}{\alpha_i} - 2 \ln \alpha_i + \text{const.}, \end{aligned} \quad (19)$$

where we have used the fact that

$$\mathbb{E}_{\lambda_i} [\lambda_i] = \int \lambda_i \delta(\lambda_i - \hat{\lambda}_i) d\lambda_i = \hat{\lambda}_i.$$

Although eq. (19) is not of standard form, it can be evaluated analytically or using a suitable sampling algorithm.

The q -factor corresponding to λ_i is found by solving the parametric (instead of variational) optimization

$$\begin{aligned} \lambda_i^* &= \operatorname{argmax}_{\lambda_i \in \mathbb{R}} \mathbb{E} [\ln p(a, b, \mathbf{x}^1, \dots, \mathbf{x}^M)] \\ &= \operatorname{argmax}_{\lambda_i \in \mathbb{R}} \mathbb{E} [\ln p(\alpha_i | \lambda_i)]. \end{aligned} \quad (20)$$

The expectation inside the maximum operator is given by

$$\mathbb{E} [\ln p(\alpha_i | \lambda_i)] = -\hat{\lambda}_i \mathbb{E}_{\alpha_i} [\alpha_i^{-1}] + \ln \hat{\lambda}_i - 2 \mathbb{E}_{\alpha_i} [\ln \alpha_i], \quad (21)$$

whose maximum is found to be

$$\lambda_i^* = \frac{1}{\mathbb{E}_{\alpha_i} [\alpha_i^{-1}]}. \quad (22)$$

2.5 Implementation Aspects

As mentioned earlier, the VBEM scheme leads to an iterative algorithm, where every individual q -factor is estimated given the most recent estimates of all other q -factors (e.g., in a round-robin fashion). For a particular reaction channel i , this means that we first determine $q(\alpha_i, \beta_i)$ given the most recent value of $\hat{\lambda}_i$ and subsequently re-estimate $\hat{\lambda}_i$ given $q(\alpha_i, \beta_i)$ and so forth. Since $q(\alpha_i, \beta_i)$ is not of standard form, we can compute its required statistics either via numerical integration or Monte Carlo sampling. Here we focus on the latter approach and employ a standard Metropolis-Hastings (M-H) sampler with log-normal proposal distributions to draw samples from $q(\alpha_i, \beta_i)$. Those samples are also used for updating the corresponding hyperparameters λ_i , i.e., the expectation in eq. (22) is replaced by a Monte Carlo average. Moreover, we found that replacing $\mathbb{E}_{\alpha_i} [\alpha_i^{-1}]$ by $\mathbb{E}_{\alpha_i} [\alpha_i]^{-1}$ yields a similar estimation performance, while significantly reducing the number of required divisions per iteration.

Note that the parameters corresponding to the homogeneous reaction channels will be driven to infinity, which in theory, causes the algorithm to diverge. Practically – however – one can check whether α_i (or λ_i) is above a critical threshold (e.g., around $10e5$), in which case the i -th reaction is considered homogeneous and excluded from the remaining analysis.

Algorithm 1 summarizes the main structure of the proposed scheme.

Algorithm 1 VBEM algorithm for detecting heterogeneity in stochastic interaction networks.

- 1: Initialize $\hat{\lambda}_i$ for $i = 1, \dots, \nu$
 - 2: **while** not converged **do**
 - 3: **for** $i = 1, \dots, \nu$ **do**
 - 4: Draw samples from $q(\alpha_i, \beta_i)$ using eq. (19) and the current value of $\hat{\lambda}_i$
 - 5: Update $\hat{\lambda}_i$ using eq. (22)
 - 6: **end for**
 - 7: **end while**
-

2.6 Extension to the Incomplete Data Scenario

In principle, the above algorithm can be easily extended for the incomplete data scenario, i.e., if the measurements consist of sparse and noisy readouts Y_n of the Markov chain X at times t_n . Intuitively, this can be understood

as adding another layer on top of the states \mathbf{x}^m in the hierarchical Bayesian model. In this case it turns out that the variational expressions from Section 2.4 also involve expectations with respect to so-called *smoothing* distributions, e.g., $p(\mathbf{x}^m | y_1^m, \dots, y_N^m, a)$ when considering the m -th cell. Computing such distributions (and computing its statistics) is a challenging task on its own and a variety of numerical and analytical approaches have been proposed (Amrein and Künsch [2012], Opper and Sanguinetti [2007], Zechner et al. [2014]). Apart from that, the VBEM framework can be readily applied to the more complicated case of incomplete and noisy measurements.

3. SIMULATIONS

We performed several simulation studies in order to demonstrate and evaluate the proposed method. For each of the case studies, we used the simple reaction network of eukaryotic gene expression illustrated in Fig. 2a. Exemplary trajectories of such a model are shown in Fig. 2b. The model comprises six reaction channels with kinetic parameters c_1, \dots, c_6 , which are either homogenous or heterogeneous – depending on the particular case study.

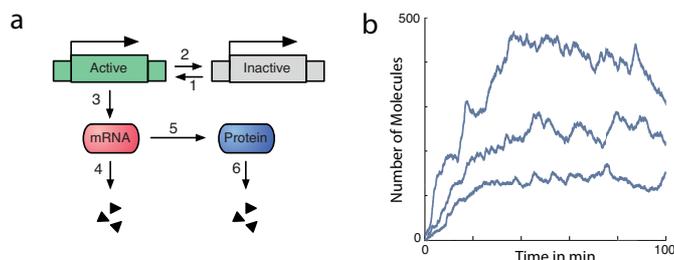


Fig. 2. A simple model of eukaryotic gene expression. (a) Schematic diagram of the reaction network. The numbered arrows indicate chemical events taking place: upon activation of the gene (arrow 1), mRNA can be transcribed (arrow 3) which in turn gets translated (arrow 5) into protein. The remaining arrows indicate gene-deactivation (arrow 2) and degradation events (arrows 4 and 6). (b) Exemplary protein traces of a heterogeneous network. In this case, heterogeneity was simulated by introducing a Gamma-type variability in the translation rate.

Unless otherwise specified, the mean values of the kinetic parameters are chosen according to Table 1.

Table 1. Mean values of the kinetic parameters.

Parameter	c_1	c_2	c_3	c_4	c_5	c_6
Mean (s^{-1})	0.500	0.050	0.100	0.001	0.030	0.008

We first analyzed convergence of the VBEM algorithm using the network from Fig. 2a and assuming a heterogeneity over three out of the six parameters (i.e., c_3 , c_5 and c_6). The results from Fig. 3 indicate that the algorithm is able to correctly identify the extrinsic noise parameters α_i and β_i in presence of heterogeneity. In case of the homogeneous reactions, both α_i and β_i diverge towards infinity, corresponding to a CV of zero and a finite mean of α_i/β_i . Furthermore, we find that in case of the heterogeneous reactions, only very few iterations are necessary until convergence is achieved.

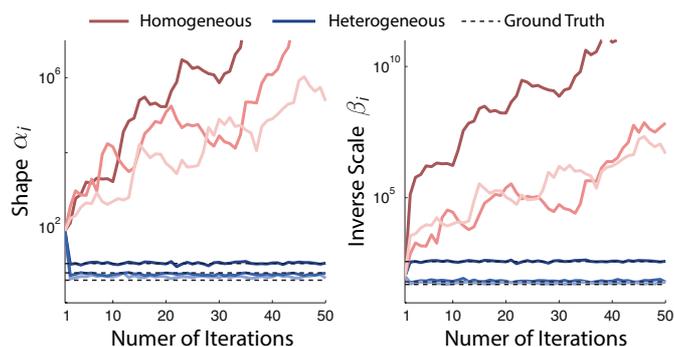


Fig. 3. Convergence of the VBEM algorithm. The algorithm was applied to $M = 30$ cell trajectories between zero and $200min$ with c_3 , c_5 and c_6 being heterogeneous with CVs 0.5, 0.3 and 0.4, respectively. The algorithm was ran for 50 update iterations. The curves correspond to expected values of the respective quantity (i.e., α_i , β_i).

Correct identification of the heterogeneous reactions depends on several parameters such as the population size M or the degree of intrinsic noise. In Fig. 4 we analyze the detection robustness of a single reaction (i.e., the gene-activation event) as a function M . In particular, we computed the ratio of positive detections using 20 independent runs (see figure caption for fuller details). In accordance with our expectations, the results demonstrate that a robust detection of extrinsic variability is possible if enough cells are in place (e.g. around $M > 100$ in this case).

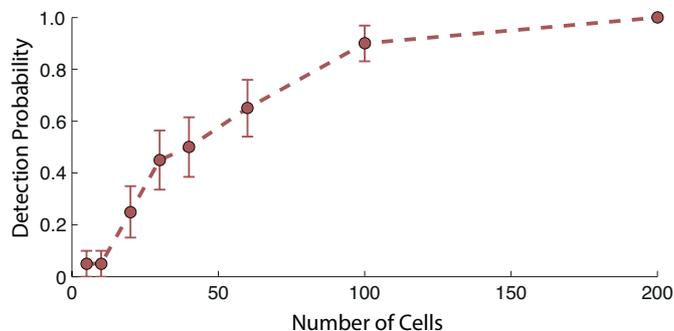


Fig. 4. Detection robustness as a function of the population size. Probabilities for correct detections were computed for different population sizes (i.e., between 5 and 200 cells) using 20 independent runs. Circles denote mean values and whiskers indicate their standard errors (SEM).

Similarly, Fig. 5 shows the probability of successful detection as a function of both *intrinsic* and *extrinsic* variability. Note that intrinsic noise of a reaction firing process scales inversely with its kinetic parameter. Again considering the gene-activation reaction, we computed the detection probabilities for three different values of c_1 (i.e., the intrinsic noise of the expression system) and several degrees of heterogeneity (see figure caption for further details). The parameters c_2 corresponding to the gene-deactivation event was adjusted such as to yield a constant ratio c_1/c_2 .

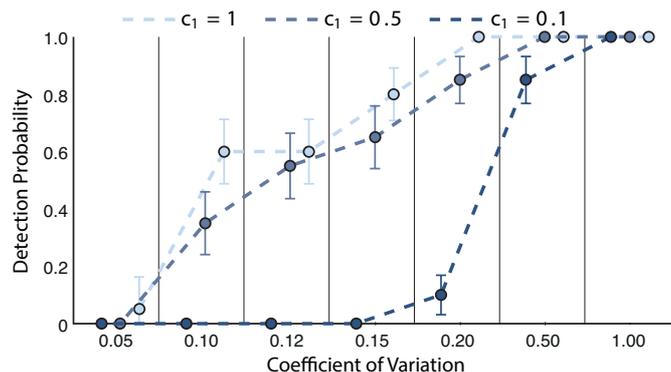


Fig. 5. Detection robustness as a function of intrinsic and extrinsic noise. We computed the rate of positive detection for different values of c_1 yielding different levels of intrinsic noise. For each c_1 we computed the detection robustness for several degrees of extrinsic variability (i.e., CVs between 0.05 and 1) using 20 independent runs. Circles denote mean values and whiskers indicate their standard errors (SEM).

We found that in presence of significant intrinsic noise and only moderate degrees of extrinsic noise, the algorithm facilitates the sparsity constraint and hence, yields negative results. In contrast, when decreasing the level of intrinsic noise, the algorithm is widely able to detect the heterogeneity (see Fig. 5).

4. CONCLUSION

Recent inference approaches that account for extrinsic variability (Zechner et al. [2012, 2014]) are based on *static* model assumptions, which means that one has to anticipate the events that are heterogeneous among individual cells. In this work we lay out a computational framework to automatically detect the events that are characterized by extrinsic variability using time-lapse data. We show that such a scenario can be understood as a sparse learning problem, which we solve using a variational Bayesian inference scheme. We validate the approach under the simplifying assumption of complete data, generated from a model of eukaryotic gene expression. The framework is currently extended for the use with real-world experimental data.

REFERENCES

O. O. Aalen, Borgan, and H. K. Gjessing. *Survival and event history analysis: a process point of view*. Springer Verlag, 2008.

M. Amrein and H. R. Künsch. Rate estimation in partially observed Markov jump processes with measurement errors. *Stat Comput*, 22(2):513–526, 2012.

M. J. Beal. *Variational algorithms for approximate Bayesian inference*. PhD-thesis, University College London, 2003.

C. M. Bishop. *Pattern recognition and machine learning*, volume 4. Springer, New York, 2007.

A. Colman-Lerner, A. Gordon, E. Serra, T. Chin, O. Resnekov, D. Endy, C G. Pesce, and R. Brent. Regulated cell-to-cell variation in a cell-fate decision system. *Nature*, 437(7059):699–706, 2005.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1): 1–38, 1977.

M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–6, August 2002.

J. Hasenauer, S. Waldherr, M. Doszczak, N. Radde, P. Scheurich, and F. Allgower. Identification of models of heterogeneous cell populations from population snapshot data. *BMC Bioinformatics*, 12(1):125, 2011.

H. Koepl, C. Zechner, A. Ganguly, S. Pelet, and M. Peter. Accounting for extrinsic variability in the estimation of stochastic rate constants. *Int J Robust Nonlin*, 22(10): 1103–1119, 2012.

U. Kuechler and M. Sorensen. *Exponential families of stochastic processes*. Springer, New York, 1997.

R. M. Neal. *Bayesian Learning for Neural Networks*. Springer, New York, 1996.

M. Opper and G. Sanguinetti. Variational inference for Markov jump processes. *Adv. in Neural Information Proc. Systems 20*, 2007.

J. Ruess, A. Miliars-Argeitis, and J. Lygeros. Designing experiments to understand the variability in biochemical reaction networks. *Journal of The Royal Society Interface*, 10(88), 2013.

V. Shahrezaei, J. F. Ollivier, and P. S. Swain. Colored extrinsic fluctuations and stochastic gene expression. *Mol Syst Biol*, 4(196):196, 2008.

D. Shutin, C. Zechner, S. R. Kulkarni, and H. V. Poor. Regularized Variational Bayesian Learning of Echo State Networks with Delay & Sum Readout. *Neural Comput*, 24(4):967–995, 2012.

B. Snijder and L. Pelkmans. Origins of regulated cell-to-cell variability. *Nat Rev Mol Cell Biol*, 12(2):119–25, 2011.

Y. Taniguchi, P. J. Choi, G. Li, H. Chen, M. Babu, J. Hearn, A. Emili, and X. S. Xie. Quantifying E. coli proteome and transcriptome with Single-Molecule sensitivity in single cells. *Science*, 329(5991):533–538, 2010.

M. E. Tipping. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, 2001.

D. J. Wilkinson. *Stochastic Modelling for Systems Biology*. Chapman and Hall/CRC, London, April 2006. ISBN 1584885408.

C. Zechner, J. Ruess, P. Krenn, S. Pelet, M. Peter, J. Lygeros, and H. Koepl. Moment-based inference predicts bimodality in transient gene expression. *Proc Natl Acad Sci USA*, 109(21):8340–8345, 2012.

C. Zechner, S. Deb, and H. Koepl. Marginal dynamics of stochastic biochemical networks in random environments. *2013 European Control Conference (ECC)*, pages 4269–4274, 2013.

C. Zechner, M. Unger, S. Pelet, M. Peter, and H. Koepl. Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. *Nat Methods*, 11(2): 197–202, 2014.