

# Switching Regulatory Models of Cellular Stress Response

Guido Sanguinetti<sup>a,b</sup>, Andreas Rutter<sup>c</sup>, Manfred Oppert<sup>c</sup> and Cedric Archambeau<sup>d</sup>

<sup>a</sup> Department of Computer Science, University of Sheffield, Regent Court, 211 Portobello Road, Sheffield, S1 4DP, U.K.,

<sup>b</sup> Biological and Environmental Systems Group, Department of Chemical and Process Engineering, University of Sheffield, Mappin Street, Sheffield, S1 3JD, U.K.

<sup>c</sup> Department of Computer Science, Technische Universität Berlin D-10587 Berlin, Germany

<sup>d</sup> Department of Computer Science University College London, Gower Street, WC1E 6BT, U.K.

Associate Editor: Prof David Rocke

## ABSTRACT

**Motivation** Stress response in cells is often mediated by quick activation of transcription factors. Given the difficulty in experimentally assaying transcription factor activities, several statistical approaches have been proposed to infer them from microarray time-courses. However, these approaches often rely on prior assumptions which rule out the rapid responses observed during stress response.

**Results** We present a novel statistical model to infer how transcription factors mediate stress response in cells. The model is based on the assumption that sensory transcription factors quickly transit between active and inactive states. We therefore model mRNA production using a bi-stable dynamical system whose behaviour is described by a system of differential equations driven by a latent stochastic process. We assume the stochastic process to be a two state continuous time jump process, and devise both an exact solution for the inference problem as well as an efficient approximate algorithm. We evaluate the method on both simulated data and real data describing *E. coli*'s response to sudden oxygen starvation. This highlights both the accuracy of the proposed method and its potential for generating novel hypotheses and testable predictions.

**Availability** MATLAB and C++ code used in the paper can be downloaded from <http://www.dcs.shef.ac.uk/~guido/>.

**Contact** [guido@dcs.shef.ac.uk](mailto:guido@dcs.shef.ac.uk)

## INTRODUCTION

Understanding the molecular basis of stress reaction in cells is one of the most important tasks in Systems Biology. Stress reaction mechanisms are key in a number of biomedical and bio-engineering applications, ranging from drug design to genetic engineering of drought-resistant crops. While most experimental studies have traditionally focused on comparing steady states before and after the imposition of the stress, it is becoming increasingly clear that the dynamics of the immediate reaction to the stress hold important biological information [see e.g. Partridge et al., 2007].

Cells respond to external stimuli in a variety of ways; perhaps the most fundamental is the transcriptional one. The stimulus is mediated by transcription factors (TFs) which transit from inactive to active state and bind to specific genes to activate or inhibit their transcription. Despite its importance, transcriptional regulation is far from being wholly understood. In particular, its experimental exploration is severely hampered by the fact that some of the fundamental key players are very hard to measure: experimental techniques to

measure active TF concentrations or to quantitate their effect on target genes are difficult and time consuming.

In response to these experimental limitations, there has been a significant amount of effort in the modelling community in order to produce statistical models of transcription to infer the activity levels of TFs from time-series measurements of the targets' expression levels. Broadly speaking, there are two categories of models for TF activity inference: coarse models which attempt to capture the simultaneous behaviour of all TFs and all genes in an organism [e.g. Liao et al., 2003, Sanguinetti et al., 2006, Sabatti and James, 2006] and detailed, ordinary differential equations (ODE) based models of small subnetworks involving only a handful of genes and one TF (*Single Input Motifs, SIM*) [e.g. Barenco et al., 2006, Lawrence et al., 2006, Rogers et al., 2007, Khanin et al., 2007, Gao et al., 2008]. While this may seem an overly simple system, it should be pointed out that SIMs are amongst the most over-represented network motifs in bacterial transcriptional networks [Alon, 2006]. Most of these models discretise time assuming TF activity to be constant between observation points; to our knowledge, the only approach to infer a continuous time TF activity profile is Lawrence et al. [2006], where a Gaussian process (GP) prior distribution is placed over TF activity (this was extended further in Gao et al. [2008]).

While these approaches have certainly set important groundwork to understand transcriptional regulation, stress reaction poses new challenges to the statistical modeller. Rapid adaptation to environmental changes is often key to the survival of the cell; in order to cope with this, TFs are often post-transcriptionally regulated via fast processes such as dimerisation and phosphorylation, so that they can be turned on or off as soon as the signal is received [see e.g. Alon, 2006]. This is clearly a problem for discrete time models as the piecewise constant assumption cannot be justified in this case. While the continuous time approach of Lawrence et al. [2006] is more appealing, using a GP prior for TF activity introduces a strong continuity constraint (indeed a smoothness constraint in many cases). Moreover, if a stationary covariance is employed, this will automatically determine a characteristic time-scale over which the latent process can change. This is ill suited to model transcription factors which occasionally vary very quickly, while remaining in a steady state for the majority of time. Of course, this could be

avoided by using a non-stationary covariance (as for example in Archambeau et al. [2007]), but the computational overheads incurred would be very significant.

In this contribution, we present a new approach for continuous time TF activity inference which specifically models quick response to stress signals. We build on the model-based approach initiated in Barengo et al. [2006], where mRNA expression for some genes is controlled by an unobserved TF via a system of ordinary differential equations. We then model the latent process as a Markovian stochastic dynamical process which performs transitions between two states (the *telegraph process*). Thus, our model of transcription is an approximation to the classical Michaelis-Menten dynamics where the time taken for TF concentrations to change from negligible to saturation level is assumed to be extremely short (compared to the time between observations). The Markovian nature of the process means that exact inference is possible for this system. However, the computational overheads are significant, and we devise an efficient approximate inference scheme based on a variational approach. In addition to inferring TF profiles, our method also gives an effective way to estimate a number of parameters, such as mRNA production and decay rates, which often play a critical role in systems biology models but are rarely precisely known.

The rest of the paper is organised as follows. In the model and methods section, we describe our transcriptional regulation model, briefly review inference approaches for continuous time Markov processes, and describe both our exact and variational solutions. In the results section, we test both the exact and approximate inference on a simulated data set, and compare our approach to Lawrence et al. [2006] to underline the importance of appropriate prior assumptions. We then use our approach to study the behaviour of the master regulator FNR in the adaptation of *E. coli* during the sudden change from aerobic to anaerobic conditions. FNR activation coordinates the action of hundreds of genes involved in switching *E. coli* metabolism from aerobic to nitric; our analysis leads to biologically measurable predictions, such as the existence of a finite (measurable) time-lag between stress imposition and FNR activation, and a predicted decrease in activity when the new anaerobic steady state is achieved. In the discussion section, we discuss related approaches and evaluate the relative merits of our approach, as well as highlighting potential extensions.

## 1 MODEL AND METHODS

The starting point for our model is the Michaelis-Menten model of transcription for a single-input motif model with  $i = 1, \dots, m$  gene targets [Alon, 2006, Barengo et al., 2006]

$$\frac{dx_i(t)}{dt} = A_i \frac{c(t)}{\kappa_i + c(t)} + b_i - \lambda_i x_i(t)$$

Here  $x_i(t)$  is mRNA concentration of target gene  $i$  as a function of time,  $b_i$  is its baseline transcription rate,  $\lambda_i$  is the mRNA decay rate and  $c(t)$  is the active TF concentration, itself a function of time. The remaining parameters  $A_i$  and  $\kappa_i$  determine the amplitude and shape of the activation curve;  $A_i$  can be interpreted as the sensitivity of  $x_i$  to the TF, and  $\kappa_i$  represents the concentration at which half the saturation level of activation is achieved. Our aim is to model a situation where a rapid response to a signal makes the TF activity quickly switch between the saturation level and zero. We will therefore simplify the model as

$$\frac{dx_i(t)}{dt} = A_i \mu(t) + b_i - \lambda_i x_i(t) \quad (1)$$

where  $\mu(t) \in \{0, 1\}$ . The model therefore is a bi-stable model with a higher steady state  $x_i = \frac{A_i + b_i}{\lambda_i}$  and a lower steady state for  $x_i = \frac{b_i}{\lambda_i}$ .

Given the time dependence of the driving process  $\mu$ , equation (1) is easy to solve in closed form, and the parameters can be estimated by standard methods (e.g. least squares). However, we are interested in the situation where the process  $\mu$  is not observed. To encode the fact that  $\mu$  can perform an arbitrary number of switches between its two states, we will place a prior on it in the form of a two-states Markov jump process, also known as a *telegraph process*. The telegraph process is characterised by its transition rates  $f_{\pm}(t)$ , which give the rate at which the process switches between the two states. To perform inference, we will be interested in the *single time marginal* probability  $p_1(t)$ , giving the probability that the process is in the on state at time  $t$ . Given transition rates  $f_{\pm}$  for the process, the probability of the system being in a particular state at a given time is given by the following Master equation

$$\frac{dp_1(t)}{dt} = -f_- p_1(t) + f_+ p_0(t)$$

$$\frac{dp_0(t)}{dt} = -f_+ p_0(t) + f_- p_1(t).$$

Using the fact that  $p_0 + p_1 = 1$  at all times, one can reduce the Master equation to a single equation on the probability  $p_1$  as

$$\frac{dp_1(t)}{dt} = -(f_+ + f_-) p_1(t) + f_+. \quad (2)$$

We will assume that we have noise corrupted observations  $\hat{x}_i(t_j)$  of the output  $x_i(t)$  at discrete time points  $t_j$ ,  $j = 0, \dots, N$  with  $x_i(t_0)$  serving as initial conditions for the problem. The observations (conditioned on the true state) will be assumed to be i.i.d. with normal noise model with variance  $\sigma_i^2$ . Thus the probability of making a single observation  $\hat{x}(t)$  given  $x(t)$  at time  $t$  is described by a Gaussian likelihood

$$p(\hat{x}(t)|x(t)) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^m \left(\frac{\hat{x}_i(t) - x_i(t)}{\sigma_i}\right)^2\right). \quad (3)$$

Although this is an incorrect noise model as the quantity  $x$  is clearly positive at all times, the error made will be small for  $b_i/\lambda_i$  much larger than the observation noise  $\sigma_i$ . The remaining parameters of the model  $A_i, b_i$  and  $\lambda_i$  are constrained positive given their physical meaning as production and decay rates and will be given exponential or flat priors.

The inference task consists of two parts: state inference, where we use the noisy observations  $\hat{x}$  to infer the posterior distribution over the true state of the system (both  $x$  and  $\mu$ ), and parameter estimation, where we learn the model parameters  $A_i, b_i, \lambda_i$  and  $\sigma_i$ . In the following subsections, we will present two approaches to performing inference in this model. First, we outline an exact inference approach which exploits the causal structure of the model to derive a forward-backward algorithm for the joint posterior over  $x$  and  $\mu$ . This is closely related to the familiar algorithm for Hidden Markov Models, the main differences being that our state vector is hybrid continuous-discrete, and that time is continuous. The main drawback of this approach is computational: the forward-backward pass requires solving numerically partial differential equations in potentially high dimensions. We then present a more efficient approximate inference algorithm which avoids these problems by directly modelling the posterior distribution over the TF activity  $\mu$ .

### 1.1 Exact inference

Although the process  $x(t)$  with observations  $\hat{x}(t_j)$  looks like a standard Hidden Markov Model, this assumption is not correct. In fact,  $x(t)$  is an integral over the Markov jump process  $\mu(t)$ , as shown by obtaining the general solution of equation (1) using Laplace transform:

$$x_i(t) = e^{-\lambda_i(t-t_0)} \left[ x(t_0) + \int_{t_0}^t e^{\lambda_i(s-t_0)} (A_i \mu(s) + b_i) ds \right]. \quad (4)$$

Therefore  $x(t)$  depends on the whole history of the process  $\mu(t)$  up to time  $t$ . However, the combined process  $(\mu(t), x(t))$  is Markovian, as the dynamics described in (1) and (2) depend only on the current state of the system. Consequently, we can base our exact solution to the state inference problem on the forward-backward algorithm for Markovian stochastic processes,

if we use both  $\mu(t)$  and  $x(t)$  as state variables. There are however still two key differences between our switching model and a standard HMM: the state variable  $(x(t), \mu(t))$  is hybrid continuous-discrete, and time is a continuous variable. Therefore the well-known forward and backward recursion rules for discrete Hidden Markov Models are replaced by partial differential equations (PDEs), the *Chapman-Kolmogorov* equations [e.g. Gardiner, 1996]. Our model is somehow simpler than the general case: jumps only occur in  $\mu(t)$  and there is no diffusion, as  $x(t)$  is a deterministic function if  $\mu(t)$  is known.

We will use the Chapman-Kolmogorov equations to calculate the marginal probability distribution  $q_\mu(x, t)$  of the posterior process of  $\mu$  and  $x = (x_1, x_2, \dots, x_m)^\top$ . Using the Markovian structure of this joint process one can show that

$$q_\mu(x, t) = \frac{1}{Z} p_\mu(x, t) \Psi_\mu(x, t). \quad (5)$$

This decomposition of the posterior marginal is the continuous time version of the well known decomposition in terms of forward and backward messages for Hidden Markov Models [see e.g. Bishop, 2006]. Here  $p_\mu(x, t)$  denotes the marginal probability distribution of the process conditioned on the data before time  $t$ , i.e. the filtered process or forward message, and  $Z$  is a time-independent normalisation constant, which equals the likelihood  $p(\hat{x}_1, \dots, \hat{x}_N | \theta)$  of the data given the model parameters  $\theta = (\lambda, f_0, f_1, A, b)$ . The last part of (5),

$$\Psi_\mu(x, t) = p(\{\hat{x}(t_j) | t_j > t\} | x(t) = x, \mu(t) = \mu), \quad (6)$$

is the likelihood of all observations after time  $t$  under the condition that the process has state  $(x, \mu)$  at time  $t$  (the backward message).

Adapting the general form of the differential Chapman-Kolmogorov equations [Gardiner, 1996] to our case, we obtain the following backward equation satisfied by  $\Psi_\mu(x, t)$ ,

$$\begin{aligned} \frac{\partial \Psi_1}{\partial t} + \sum_{i=1}^m (1A_i + b_i - \lambda_i x_i) \frac{\partial \Psi_1}{\partial x_i} &= f_- \cdot (\Psi_1(x, t) - \Psi_0(x, t)), \\ \frac{\partial \Psi_0}{\partial t} + \sum_{i=1}^m (0A_i + b_i - \lambda_i x_i) \frac{\partial \Psi_0}{\partial x_i} &= f_+ \cdot (\Psi_0(x, t) - \Psi_1(x, t)). \end{aligned} \quad (7)$$

These PDEs must be solved backward in time starting at the last observation  $\hat{x}(t_N)$  using the initial condition

$$\Psi_\mu(x, t_N) = p(\hat{x}(t_N) | x(t_N) = x). \quad (8)$$

The other observations are taken into account by jump conditions

$$\Psi_\mu(x, t_j^-) = \Psi_\mu(x, t_j^+) p(\hat{x}(t_j) | x(t_j) = x) \quad (9)$$

with  $\Psi_\mu(x, t_j^\mp)$  being the values of  $\Psi_\mu(x, t)$  before and after the  $j$ -th observation. Here we use the property of the noise model that the observations  $\hat{x}(t_j)$  are independent conditioned on the process  $(\mu(t), x(t))$ . Therefore the likelihood  $\Psi_\mu(x, t_j^-)$  including  $\hat{x}(t_j)$  is the product of  $\Psi_\mu(x, t_j^+)$  for observations at later time points and the probability  $p(\hat{x}(t_j) | x(t_j) = x)$  given by (3).

In order to calculate  $q_\mu(x, t)$  we need to consider the filtered process described by  $p_\mu(x, t)$ , too. Its time evolution is given by the forward Chapman-Kolmogorov equation

$$\begin{aligned} \frac{\partial p_1}{\partial t} + \sum_{i=1}^m \frac{\partial}{\partial x_i} (1A_i + b_i - \lambda_i x_i) p_1(x, t) &= \\ f_+ p_0(x, t) - f_- p_1(x, t) & \\ \frac{\partial p_0}{\partial t} + \sum_{i=1}^m \frac{\partial}{\partial x_i} (0A_i + b_i - \lambda_i x_i) p_0(x, t) &= \\ f_- p_1(x, t) - f_+ p_0(x, t) & \end{aligned} \quad (10)$$

and the posterior  $q_\mu(x, t)$  fulfils a similar PDE. This can be seen by calculating the time derivative

$$\frac{\partial q_\mu}{\partial t} = \frac{1}{Z} \left( \Psi_\mu(x, t) \frac{\partial p_\mu}{\partial t} + p_\mu(x, t) \frac{\partial \Psi_\mu}{\partial t} \right) \quad (11)$$

of the posterior distribution. Using the PDEs given in (7) and (10) we find

$$\begin{aligned} \frac{\partial q_1}{\partial t} &= -\frac{1}{Z} \sum_{i=1}^m \frac{\partial}{\partial x_i} (1A_i + b_i - \lambda_i x_i) \Psi_1(x, t) p_1(x, t) \\ &+ \frac{1}{Z} (f_+ \Psi_1(x, t) p_0(x, t) - f_- \Psi_0(x, t) p_1(x, t)), \\ \frac{\partial q_0}{\partial t} &= -\frac{1}{Z} \sum_{i=1}^m \frac{\partial}{\partial x_i} (0A_i + b_i - \lambda_i x_i) \Psi_0(x, t) p_0(x, t) \\ &+ \frac{1}{Z} (f_- \Psi_0(x, t) p_1(x, t) - f_+ \Psi_1(x, t) p_0(x, t)). \end{aligned} \quad (12)$$

This equation can be further simplified by introducing time and state dependent posterior jump rates

$$g_+(x, t) = \frac{\Psi_1(x, t)}{\Psi_0(x, t)} f_+ \quad g_-(x, t) = \frac{\Psi_0(x, t)}{\Psi_1(x, t)} f_- \quad (13)$$

and applying (5). We then find

$$\begin{aligned} \frac{\partial q_1}{\partial t} + \sum_{i=1}^m \frac{\partial}{\partial x_i} (1A_i + b_i - \lambda_i x_i) q_1(x, t) &= \\ g_+(x, t) q_0(x, t) - g_-(x, t) q_1(x, t), & \\ \frac{\partial q_0}{\partial t} + \sum_{i=1}^m \frac{\partial}{\partial x_i} (0A_i + b_i - \lambda_i x_i) q_0(x, t) &= \\ g_-(x, t) q_1(x, t) - g_+(x, t) q_0(x, t), & \end{aligned} \quad (14)$$

which is also the forward Chapman-Kolmogorov equation. Consequently, the only differences between prior and posterior process are the jump rates for the telegraph process  $\mu(t)$ .

In the case of a single target gene numerical integration of the PDEs (7) and (14) is computationally feasible. We use the Lax algorithm [Vesely, 1994] for that purpose, because it prevents negative values for  $\Psi_\mu(x, t)$  and  $q_\mu(x, t)$  as long as the step sizes fulfil the condition  $\Delta x > A \Delta t$ . The boundaries are determined by the two steady states  $x_{\text{low}} = b/\lambda$  and  $x_{\text{high}} = (A+b)/\lambda$ . In the forward integration these boundaries are closed, as the process cannot leave the interval between  $x_{\text{low}}$  and  $x_{\text{high}}$ . But it can come from the outside, so that we have to use open boundaries in the backward integration.

Parameter estimation based on the exact solution of the state inference problem is also possible. For that purpose we use the free energy  $F \equiv -\ln p(\hat{x}(t_1), \dots, \hat{x}(t_N) | \theta)$ , i.e. the negative log-likelihood of the data as a function of the model parameters  $\theta = (\lambda, f_0, f_1, A, b)$ . This quantity is given by

$$\begin{aligned} F &= -\ln E_{\text{prior}} [\Psi_\mu(x, t_0)] \\ &= -\ln \int [\Psi_0(x, t_0) p_0(x, t_0) + \Psi_1(x, t_0) p_1(x, t_0)] dx \end{aligned} \quad (15)$$

and only a single backward integration is necessary in order to obtain  $\Psi_\mu(x, t_1)$ . Here  $E_{\text{prior}}$  denotes expectation under the prior distribution for the first observation at  $t_0$ . Minimising the free energy with respect to the model parameters then leads to their type II maximum likelihood estimates  $\theta^* = \arg \min F(\theta)$ .

## 1.2 Variational approximation

As discussed above, the exact solution for the inference problem for the switching model suffers from the curse of dimensionality, so that exact inference in higher dimensions becomes prohibitively expensive. Variational inference [see e.g. Jordan et al., 1999] is a powerful approach to solving

approximately the inference problem. Given an intractable probability distribution  $p$ , the variational approach finds an optimal approximation to  $p$  within a certain family of distributions. This is usually done by minimizing the *Kullback-Leibler (KL) divergence* between the two distribution

$$KL[q||p] = E_q \left[ \log \frac{q}{p} \right] = \int dq \log \frac{q}{p}.$$

By selecting a suitable family of approximating distributions, the inference problem is then turned into an optimisation problem.

We will restrict the discussion to the case of a single target gene, the generalisation to more genes being straightforward. In the following, we will view the stochastic process as a probability measure over the (infinite dimensional) space of possible paths of the TF  $\mu_{0:T}$  (the notation  $\mu_{0:T}$  indicates a specific realisation of the process between 0 and  $T$ ). Given a prior telegraph process  $p(\mu_{0:T}|f_{\pm})$  and a noise model for the observations  $p(\hat{x}|\mu_{0:T})$ , Bayes' theorem allows in principle the computation of a posterior process as

$$p_{post}(\mu_{0:T}|\hat{x}) = \frac{1}{Z} p(\hat{x}|\mu_{0:T}) p_{prior}(\mu_{0:T}|f_{\pm}). \quad (16)$$

As the solution (4) of the model (1) depends on the whole history of the process  $\mu$ , so will the likelihood factor in (16). This means that the posterior process will not be a Markov process. However, it still makes sense to seek a Markov process that approximates optimally the posterior process.

To do this, we will compute the *Kullback-Leibler (KL) divergence* between the posterior process in (16) and an approximating Markov process  $q(\mu|g_{\pm})$ . This is given by

$$KL[q||p_{post}] = \ln Z + KL[q||p_{prior}] - \sum_{j=1}^N E_q [\ln p(\hat{x}_j|x(t_j))]. \quad (17)$$

The KL divergence between two Markov jump processes was computed in the general case in Opper and Sanguinetti [2007] material. A derivation of the KL divergence in the special case of the telegraph process can be found in the supplementary material, the final result being given by

$$KL[q||p_{prior}] = \int_0^T dt q_1(t) \left[ g_-(t) \ln \frac{g_-(t)}{f_-(t)} + f_-(t) - g_-(t) \right] + \int_0^T dt [1 - q_1(t)] \left[ g_+(t) \ln \frac{g_+(t)}{f_+(t)} + f_+(t) - g_+(t) \right].$$

The estimation of the likelihood term in (17) is more challenging; under the assumption of Gaussian noise, it requires the computation of the first two moments of the random variable  $x(t)$  under the approximating process  $q$ . These are given by

$$\begin{aligned} \langle x(t_i) \rangle &= \exp(-\lambda t_i) \left[ x(0) + \frac{b}{\lambda} (\exp(\lambda t_i) - 1) \right. \\ &\quad \left. + A \int_0^{t_i} \exp(\lambda s) q_1(s) ds \right] \\ \langle x^2(t_i) \rangle &= \exp(-2\lambda t_i) \left\{ x(0)^2 + \frac{b^2}{\lambda^2} [\exp(\lambda t_i) - 1]^2 + \right. \\ &\quad \left. + 2x(0) \frac{b}{\lambda} (\exp(\lambda t_i) - 1) + 2x(0) A I_i + 2 \frac{b}{\lambda} (\exp(\lambda t_i) - 1) A I_i + \right. \\ &\quad \left. + A^2 \int_0^{t_i} \int_0^{t_i} \exp[\lambda(t+s)] q_1(t,s) dt ds \right\}. \end{aligned} \quad (18)$$

Here,  $I_i = \int_0^{t_i} \exp(\lambda s) q_1(s) ds$  and we have used the fact that  $\langle \mu(t) \rangle_q = q_1(t)$ .

In general, these integrals are analytically intractable when the rates for the approximating process  $g_{\pm}$  are functions of time. We will therefore solve the optimisation problem on a grid, assuming the approximating process rates to be constant between points in the grid. This allows us to solve explicitly the Master equation on the grid, and therefore allows the calculation

of the integrals needed in (18) (the explicit calculations are given in the supplementary material). By taking the grid to be sufficiently fine, the numerical solution can approximate the true minimum to arbitrary precision. This algorithm scales linearly with the number of genes (since we need to compute a set of moments per gene), making it computationally much more efficient than the exact inference solution.

### 1.3 Parameter estimation

The variational procedure outlined above will obtain an approximation to the posterior distribution of the switching TF  $\mu$  given the observations and the model parameters. This can then be used to estimate the parameters of the model, in an EM-like scheme (often called Variational Bayes EM). In the E-step, an approximate posterior is computed by minimising the KL divergence with respect to the rate parameters  $g_{\pm}^i$ . This can be done by gradient descent or using other search strategies.

In the M-step, we use the approximate posterior  $q(\mu)$  to marginalise the process  $\mu$ , obtaining an approximation to the marginal likelihood of the data as

$$p(\hat{x}|\Theta) \simeq \frac{1}{Z} \exp \{ E_{q(\mu)} [\log p(\hat{x}, \mu|\Theta)] \} \quad (19)$$

where  $Z = \exp(-H[q])$  and  $H[q]$  is the entropy of the approximate posterior. This can be shown to be a lower bound on the true likelihood by invoking Jensen's inequality [see e.g. Jordan et al., 1999]. Equation (19) can be used in Bayes' theorem to compute posterior distributions over the model parameters. The posterior distribution over the parameters  $A$ ,  $b$  and  $\sigma$  is obtained analytically (for suitable choices of priors) in the form of truncated Gaussian and inverse Wishart distributions. Unfortunately, this is not possible for the decay parameter  $\lambda$ ; estimation of its posterior distribution is done by direct evaluation of the un-normalised posterior over a grid in 1 dimension.

## 2 RESULTS


### 2.1 Synthetic data

To benchmark the model and assess the validity of our approximation, we ran the model on a simple synthetic example. We constructed a simple toy data set made up of ten equally spaced observations drawn from the model with input signal (TF activity)

$$\mu(t) = \begin{cases} 1 & t \in [0, 169] \cup [660, 1000] \\ 0 & t \in [170, 659] \end{cases}.$$

The differential equation parameters were chosen as  $A = 3.7 \times 10^{-3}$ ,  $b = 8 \times 10^{-4}$  (production rates) and  $\lambda = 5 \times 10^{-3}$  (decay rate). Gaussian noise with a standard deviation of 0.03 was added to the theoretical values of  $x$  to give the observations.

The results of the inference are shown in Figure 1. Figure 1 (a-b) show the inferred posterior mean (dashed black) compared to the input impulse (thin black) in the approximate and exact case respectively. In order not to clutter the figures, we omitted the confidence intervals for the posterior TF activity; since at each time the TF is a binary variable, these can be obtained from the mean value as  $\sqrt{(q_1 - q_1^2)}$ . Also shown is the posterior first moment for  $x(t)$  (thick black) with confidence intervals and the data (red crosses). The grid used had five grid points for every observation point (for a discussion of how grid size affects model result, see the supplementary material). Both the reconstructions are reasonable although the exact one has much tighter uncertainty. Figure 1 (c) shows the results of applying Lawrence et al. [2006] to the data using a squared exponential covariance function for the GP prior. Although the model produces a good fit to the data, the stationary



**Figure 1.** Results of inference on toy data set. (a) Results from variational approach: inferred posterior mean (dashed black) compared to the input impulse (thin black). Also shown the posterior first moment for  $x(t)$  (thick black) with confidence intervals and the data (red crosses). (b) Results of exact inference. (c) Reconstructed first moment using Lawrence et al. [2006].

covariance used forces the inferred TF profile to have biologically meaningless fast oscillations.

Both the exact and approximate inference slightly underestimate the value of the parameter  $\lambda$  at  $4 \pm 0.3 \times 10^{-3}$  (true value  $5 \times 10^{-3}$ ). The estimates for the model parameters are good both in the approximate and exact case, with  $A = 2.8 \pm 0.3 \times 10^{-3}$  and  $b = 0.8 \pm 0.1 \times 10^{-3}$  (results from approximate inference) and  $A = 3.2 \pm 1.1 \times 10^{-3}$  and  $b = 0.08 \pm 0.6 \times 10^{-3}$  (results from exact inference). The exact inference was carried out with  $\sigma$  fixed to the true value, while the approximate inference obtained  $\sigma = 0.05$ . The whole process took approximately ten minutes on a standard PC for the approximate case and 2 hours for the exact case.

## 2.2 Microaerobic shift in *E.coli*

As a real example on which to test our approach, we considered transcriptomic measurements of the reaction of *E.coli* to sudden oxygen starvation [Partridge et al., 2007]. *Escherichia coli* is a robust organism that can adapt remarkably well to changes in its environment. One of the most dramatic such changes routinely encountered by the bacterium is the change in the availability of oxygen: the bacterium can be expelled from the host's gut and very rapidly moves from an environment with virtually no oxygen to an aerobic environment. This change entails a whole shift in the metabolism of the bacterium from a nitric metabolism to a much more energetically favourable aerobic metabolism. The set of enzymes involved in the two different phases of *E.coli* metabolism is only partly overlapping; in order to perform this shift, a large number of genes must be turned on and off in a coordinated manner. This action is carried out by a few TFs which respond to oxygen signals.

Perhaps one of the most important oxygen sensors in the cell is the iron-sulphur cluster protein FNR. FNR is a master regulator (*i.e.*, one of a dozen TFs which target most of the bacterial genes) which can exist in two states. Its state in the presence of oxygen is an inactive monomer. When oxygen is removed, the protein is dimerized which can in turn bind DNA and activate or repress transcription of a large number of genes. Therefore, one expects that, after a certain time lag, FNR undergoes a fast transition from inactive to active state. Interestingly, total protein concentration of FNR (dimer + monomer) is approximately constant between aerobic and anaerobic conditions [Jervis and Green, 2007].

In the experimental setting considered in Partridge et al. [2007], an aerobically grown culture of *E.coli* K12 was rapidly deprived of oxygen. Microarray measurements were then taken at 5, 10, 15 and 60 minutes following the imposition of the stress. The arrays measure the change in concentration of mRNA *relative* to the initial point. This implies that one of the parameters in our model,  $b$ , becomes unidentifiable as the lower steady state becomes 1. This is easily resolved by setting  $b = \lambda$  in the model. Partridge et al. [2007] also performed genome-wide TF activity inference using the probabilistic model described in [Sanguinetti et al., 2006]. This showed a rapid response of FNR to the signal that appeared to tail off when the system reached steady state.

We considered a subset of five genes which are known to be activated by FNR and with a reasonably simple promoter structure: these are *ompW*, *yjiD*, *hypB*, *moaA* and *aspA*. In practice, all of these genes, with the exception of *yjiD*, are also regulated by other transcription factors. However, these other transcription factors are further downstream than FNR in the stress response cascade, so that one can assume that the initial response to oxygen withdrawal is well modelled as a SIM. Performing the exact inference in this case would require numerically solving a partial differential equation (PDE) with five spatial dimensions, which is infeasible. The results of the approximate inference of the FNR active profile are shown in Figure 2 (a). The system appears to undergo a sharp transition between inactive and active state at around 3 to 6 minutes. This results in an interesting prediction: removing oxygen for a period shorter than 3 minutes will not lead to an FNR-mediated transcriptional response. Therefore, one may view this as an indirect measurement of the time it takes *E.coli* to commit itself to change its metabolic regime between aerobic and nitric. The model also predicts that the activity of FNR will tail off slightly after approximately 20 minutes. While, there is no simple biochemical explanation for this, as FNR will remain dimerized and hence active as long as oxygen is absent, a decrease in activity towards steady state was also predicted by Partridge et al. [2007] using a different computational model. The most plausible explanation is the action of other transcription factors which are downstream targets of FNR and which become active after a reasonable transcriptional delay.

Figure 2 (b) shows the inferred half lives of the five targets (triangles on the right) against their experimentally measured values

Selinger et al. [2003]. In two cases, *yjiD* and *moaA*, the experimental value of the half life of the transcript was not available. In general, there is a good agreement between the inferred values and the experimental measurements, although it should be noticed that the experimental measurements are extremely noisy in some cases. The ability to provide a reasonable indirect estimate of mRNA half lives is potentially precious to biologists: it is known that mRNA decay is a regulated process, implying that mRNA half lives measured in different experimental conditions will in general be different. As it is difficult to measure experimentally decay rates in a dynamic setting like stress reaction, it is essential to be able to identify these parameters in the model.

We can gain further insight into the workings of the model by comparing predicted expression profiles with the observed discrete time points. Figure 3 shows this for three genes, *ompW* (a), *hypB* (b) and *aspA* (c). The solid lines represent the mean of the stochastic process, and the dotted lines are the confidence intervals obtained by adding  $\pm 1$  standard deviation of the time marginals (these error bars include only the variability in the stochastic process, the kinetic parameters were fixed at their maximum likelihood value for the plot). In general, all reconstructed profiles show a saturating behaviour, as implied by the inferred TF activity (Figure 2 (a)). The specific form of the profile though is determined by the kinetic parameters inferred. It is also interesting to notice that the fit of the model to the *hypB* expression profile is not as good as in the other two cases. In particular, *hypB* expression markedly decreases from 15 minutes to 60 minutes, which is incompatible with the other profiles and can hardly be accommodated by the model. This is probably due to the effect of the transcription factor IHF, which also activates *hypB* but is repressed by FNR. It is therefore plausible that, after a certain amount of time, the SIM approximation breaks down in this case, explaining the poor fit to this profile.

### 3 DISCUSSION

In this paper, we presented a novel model-based approach to infer TF activity profiles from microarray time series data. The central assumptions underpinning the model are the SIM assumption (all the genes are targets of a single TF) and the prior model that TFs transit quickly from active to inactive state. This second assumption is likely to be reasonable in many stress reaction experiments, particularly when the stress is applied at the metabolic level, triggering a post-translational modification of the TFs. While this is a fairly broad class of conditions, it is important to point out that there are biological stresses (e.g. heat shock) that do not fit well in this category. It would be interesting to explore stochastic models that can combine fast and slowly reacting components.

The work that is perhaps most closely related to ours is Lawrence et al. [2006], where transcriptional regulation is modelled with a linear system of ODEs with a latent driving factor. However, the choice of a Gaussian process prior for the TF activity has as a natural consequence that the inferred posterior TF profiles vary continuously with time. While this may be a reasonable assumption in certain settings, it clearly is untenable when modelling fast biological processes such as stress reaction. Another important advantage of our model over Lawrence et al. [2006] is the ability to identify the decay parameters. Lawrence *et al.* need to fix at least one of the decay rates to the experimentally measured value (a weakness shared by

Barenco et al. [2006]). Given the very low accuracy of such measurements (see Figure 2 b), the bias introduced by fixing a parameter could potentially lead to serious errors in the estimation of the latent process.

There are also several other papers that attack the problem of inferring TF activity profiles from mRNA time series data. For example, Rogers et al. [2007], Khanin et al. [2007] use the same ODE-based model of transcription, but then restrict themselves to piecewise constant TF level, effectively discretising time. These models are also limited to the SIM case. Many other models address the global case, where hundreds of TFs regulate thousands of genes [e.g. Liao *et al.*, 2003, Sanguinetti *et al.*, 2006, Sabatti and James, 2006]; however, in order to contain complexity, they adopt a simplistic model of transcription where only linear and additive effects of TFs are retained. While genome-wide modelling is still an ambitious target for the model we developed, it should be pointed out that it presents significant computational advantages over other ODE-based approaches. In particular, the complexity of our algorithm scales linearly with the number of target genes and time points, while for example the Gaussian Process-based approach of Lawrence et al. [2006] is cubic in the product of the number of genes and time points. In this light, statistical modelling of a moderate size pathway is certainly within reach.

There is an interesting relationship between our work and the work on re-wiring networks of Guo et al. [2007]. There, links in a biological network were switched on and off according to a discrete time Markov process. In our approach, it is the activity of the regulatory nodes in the network that switches on and off as a (continuous time) Markov process.


Another potentially interesting generalisation of our work is to the case where the stochastic behaviour of gene expression is not solely due to the TFs activity, but also to the intrinsic stochasticity of transcription. A variational approach to inference in bi-stable (or more generally, non-Gaussian) systems of SDEs was recently proposed in [Archambeau *et al.*, 2007].

### ACKNOWLEDGEMENTS


G.S. would like to thank Jeff Green and Robert Poole for many useful discussions on the physiological functions of FNR, and for introducing him to the fascinating world of bacterial stress reaction. C.A. and M.O. gratefully acknowledge support from the EPSRC under the VISDEM project (EP/C005848/1). We would like to thank the anonymous reviewers for their constructive criticism.

### REFERENCES

- U. Alon. *An introduction to systems biology*. Chapman and Hall, London, 2006.
- C. Archambeau, D. Cornford, M. Opper, and J. Shawe-Taylor. Gaussian process approximations of stochastic differential equations. *Journal of Machine Learning Research Workshop and Conference Proceedings*, 1(1):1–16, 2007.
- M. Barenco, D. Tomescu, D. Brewer, R. Callard, J. STark, and M. Hubank. Ranked prediction of p53 targets using hidden variable dynamical modelling. *Genome Biology*, 7(3), 2006.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, Singapore, 2006.



**Figure 2.** Results on *E. coli* data: (a) posterior mean FNR profile; (b) half lives of targets (in minutes) with uncertainty, inferred (triangles on the right) versus experimentally measured. No measurement of the half life of *yjiD* or *moaA* is available.



**Figure 3.** Reconstructed expression profiles versus observed data: (a) *ompW*, (b) *hypB*, (c) *aspA*. Solid lines are mean prediction, dotted black lines confidence intervals ( $\pm 1$  standard deviation), red crosses measured expression levels

- P. Gao, A. Honkela, M. Rattray, and N. D. Lawrence. Gaussian process modelling of latent chemical species: applications to inferring transcription factor activities. *Bioinformatics*, 24(16): i70–i75, 2008.
- C. W. Gardiner. *Handbook of Stochastic Methods*. Springer, Berlin, second edition, 1996.
- F. Guo, S. Hanneke, W. Fu, and E. P. Xing. Recovering temporally rewiring networks: a model-based approach. In *Proceedings of the 24th International Conference on Machine Learning*, 2007.
- A. J. Jervis and J. Green. In vivo demonstration of FNR dimers in response to lower  $O_2$  availability. *Journal of Bacteriology*, 189: 2930–2932, 2007.
- M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- R. Khanin, V. Vinciotti, V. Mersinias, C. Smith, and E. Wit. Statistical reconstruction of transcription factor activity using michaelis-menten kinetics. *Biometrics*, 63:816–823, 2007.
- N. D. Lawrence, G. Sanguinetti, and M. Rattray. Modelling transcriptional regulation using gaussian processes. In *Advances in Neural Information Processing Systems 19*, 2006.
- J. C. Liao, R. Boscolo, Y.-L. Yang, L. M. Tran, C. Sabatti, and V. P. Roychowdhury. Network component analysis: Reconstruction of regulatory signals in biological systems. *Proceedings of the National Academy of Sciences USA*, 100(26):15522–15527, 2003.
- M. Opper and G. Sanguinetti. Variational inference for markov jump processes. In *Advances in Neural Information Processing Systems 20*, 2007.
- J. D. Partridge, G. Sanguinetti, D. Dibden, R. E. Roberts, R. K. Poole, and J. Green. Transition of *escherichia coli* from aerobic to micro-aerobic conditions involves fast and slow reacting regulatory components. *Journal of Biological Chemistry*, 282(15): 11230–11237, 2007.
- S. Rogers, R. Khanin, and M. Girolami. Bayesian model-based inference of transcription factor activity. *BMC Bioinformatics*, 8 (S2), 2007.
- C. Sabatti and G. M. James. Bayesian sparse hidden components analysis for transcription regulation networks. *Bioinformatics*, 22(6):739–746, 2006.
- G. Sanguinetti, N. D. Lawrence, and M. Rattray. Probabilistic inference of transcription factor concentrations and gene-specific regulatory activities. *Bioinformatics*, 22(22):2775–2781, 2006.
- D. W. Selinger, R. M. Saxena, K. J. Cheung, G. M. Church, and C. Rosenow. Global RNA half-life analysis in *Escherichia coli* reveals positional patterns of transcript degradation. *Genome Research*, 13(2):216–223, 2003.
- F. J. Vesely. *Computational Physics: An Introduction*. Plenum Publishing Corporation, New York, 1994.