

# Matching models across abstraction levels with Gaussian Processes<sup>\*</sup>

Giulio Caravagna<sup>1\*\*</sup>, Luca Bortolussi<sup>2,3</sup>, and Guido Sanguinetti<sup>1</sup>

<sup>1</sup> School of Informatics, University of Edinburgh, UK.

<sup>2</sup> DMG, University of Trieste, Trieste, Italy; ISTI-CNR, Pisa, Italy.

<sup>3</sup> MOSI, Dept. of Informatics, Saarland University, Saarbücken, Germany.

**Abstract.** Biological systems are often modelled at different levels of abstraction depending on the particular aims/resources of a study. Such different models often provide qualitatively concordant predictions over specific parametrisations, but it is generally unclear whether model predictions are quantitatively in agreement, and whether such agreement holds for different parametrisations. Here we present a generally applicable statistical machine learning methodology to automatically reconcile the predictions of different models across abstraction levels. Our approach is based on defining a correction map, a random function which modifies the output of a model in order to match the statistics of the output of a different model of the same system. We use two biological examples to give a proof-of-principle demonstration of the methodology, and discuss its advantages and potential further applications.

**Keywords:** Computational abstraction, Emulation, Gaussian Processes, Heteroscedasticity

## 1 Introduction

Computational modelling in the sciences is founded on the notion of abstraction, the process of identifying and representing mathematically the salient features and interactions of a real system. Abstraction is a human led and interdisciplinary activity: many factors influence the decision of which features/ interactions are eventually represented in the abstracted model, including the specialist interests of the scientists formulating the model, as well as computational constraints on the level of detail which can feasibly be implemented on the available hardware. Such factors inevitably vary between different research groups and at different times, leading to a proliferation of different models representing the same underlying phenomena.

---

\* GC and GS gratefully acknowledge support from the European Research Council under grant MLCS306999. LB acknowledges partial support from the EU project QUANTICOL, 600708, and by FRA-UniTTS. We thank Dimitris Milios for useful discussions and for providing us with the MATLAB for heteroscedastic regression.

\*\* Corresponding author, giulio.caravagna@ed.ac.uk.

Systems biology is a prominent field where models with different level of abstraction coexist. As an example, consider the process of gene expression, whereby genetic material stored in the DNA is transcribed into messenger RNA and eventually translated into protein. At the highest level of abstraction, which is frequently employed when studying high-throughput data, the process may be considered as a black box, and one may only be interested in characterising the statistical structures underlying observed levels of gene expression in different conditions [8]. Alternatively, one may want to mechanistically represent the dynamics of some of the steps involved in the process. The choice of which steps to include is again largely driven by extrinsic factors: examples in the literature range from highly detailed models where the synthesis of mRNA/ protein is modelled through the binding and elongation of polymerase/ ribosomes, to models where protein production is modelled as a single reaction with first order kinetics [1, 11].

Representing the same physical process by multiple models at different levels of abstraction immediately engenders the question of how different model outputs can be reconciled. As the models all represent the same underlying physical system, it can be plausibly assumed that such models will agree at least qualitatively for suitable parametrisations. In general, however, models may not agree quantitatively, and their discrepancy may be a function of the parameters. Understanding and quantifying such discrepancies would often be very valuable: first of all, it can shed light on how simplifications within models affect predictions, and secondly it may open the opportunity to construct computationally efficient surrogates of complex models. Such surrogates can be precious when modelling requires intrinsically computationally intensive tasks, like inference.

In this paper, we approach the problem of reconciling models from a statistical machine learning angle. We start by sampling a sparse subset of the parameter space over which we evaluate the models' outputs (generally by simulation). These evaluations are used as a *training set* to learn a *correction map* via a non-parametric regression approach based on Gaussian Processes. We show that our approach yields a consistent stochastic equivalence between models, which provably reconciles the predictions of the two models up to the second moment. We demonstrate the approach on two biological examples, showing that it can lead to non-trivial insights into the structure of the models, and provide an efficient way to simulate a complex model via a simpler model.

The rest of the paper is organised as follows: we start by giving a high level description of the problem and how we attack it. This is followed by a formal definition and a thorough illustration of the proposed solution, discussing its desirable theoretical properties. We then demonstrate the approach on two proof of principle examples, showing the potential for practical application of the method. We conclude the paper by discussing the relationship of our approach to existing ideas in model reduction and statistical emulation, as well as some possible extensions of our results.

## 2 Problem definition

*High level description of the approach.* In this paper we consider the problem of performing analyses which require exhaustive sampling of a model’s outputs,  $M$ , the dynamics of which are expensive to compute. We are not interested in the origin of such a complexity, and *we assume* to be given an *abstraction/surrogate* of this model,  $m$ , which is reasonably less challenging to analyze<sup>4</sup>. For this reason, *we want to investigate a general methodology to use  $m$  as a reliable proxy to get statistics over  $M$ .* Possible applications of this framework, which we discuss in §4, regard *model selection and synthesis, inference, and process equivalence/control*.

In general, we will assume both models to be stochastic processes, e.g. continuous time Markov Chains (CTMCs). Furthermore, we assume that the highly detailed model  $M$  and the reduced model  $m$  share some parameters  $\theta_m$  and some observables/ state variables  $X$ , but the large model will generally have many more state variables and parameters. In general we can compute some statistics of the shared state variables  $X$  (e.g. mean), and that such computation will be considerably more expensive using the detailed model  $M$ .

As both models are devised as abstractions of the same physical system, it is not unreasonable to assume that the expected values of the shared variables will be similar for some parametrisations of the models. However, it is in general not the case that the *distribution* of the shared variables implied by the two models will be the same, and, as we vary the shared parameters  $\theta_m$ , even the expected values may show non-negligible discrepancies. Our aim is to develop a generally applicable machine learning approach to correct the output of the reduced model, in order to match the distribution of the larger model. This has strong practical motivations: one of the primary uses of models in general is to test hypotheses statistically against observed data, and it is therefore essential to capture as accurately as possible the implied variability on observable variables.

The strategy we will follow is simple: we start by sampling values of the shared parameters  $\theta_m$ , and compute the first two statistics of the observed variables as implied by both the large and reduced models (by simulation). In general, one may expect the variance implied by the larger model to be larger, as a more detailed model will imply more stochastic steps. We can therefore correct the first two statistics (mean and variance) of the reduced model output by adding a random function of the shared parameters  $\theta_m$ , which can be learned by rephrasing it as a regression task. We will work with heteroschedastic Gaussian Processes [5].

### 2.1 Formal problem statement

We assume to be given two models of the same underlying physical system:

---

<sup>4</sup>  $M$  could be complex to analyze either because of its structure, e.g., it might have many variables, or its numerical hurdles, e.g., the degree of non-linearity or parameters stiffness. For similar reasons, we do not care whether  $m$  is has been derived by means of independent domain-knowledge or automatic techniques.

- a highly detailed model  $M$ , with state variables  $\mathcal{Y}$  and parameters  $\theta_M$ .
- a reduced model  $m$ , with state variables  $\mathcal{X}$  and parameters  $\theta_m$ .

We will have  $|\mathcal{Y}| \gg |\mathcal{X}|$  and  $|\theta_M| \gg |\theta_m|$ .

*Assumptions.* In general, the problem we want to tackle draws immediate connection to that of using  $m$  as a *statistical emulation* of  $M$ . However, to exploit solutions from *regression analysis* and *machine learning*, we make further assumptions and discuss their limitations thorough the paper.

1. (*Observables*) we assume that it exists a non-empty set of state variables (or, more generally, observables)  $X$ , common to both models, that is sufficient to compute our statistics of interest.
2. (*Parameters*) we assume that model  $M$  is fully parametrized by  $\theta_M$ , a vector of real-valued parameters that breaks down as  $\theta_M = [\theta_m \ \theta_f]^\top$ , with  $\theta_m$  shared between models  $m$  and  $M$ . Here, we assume that  $m$  is fully parametrized<sup>5</sup> by  $\theta_m$ , which ranges in a domain set  $\Theta$ . We term  $\theta_f$  free parameters in  $M$ , given  $\theta_m$ . We further assume to have a probability density  $p(\theta_f)$  for the free parameters, and a probability density  $p(\theta_m)$  for the shared parameters, encoding our knowledge about them.
3. (*Sampling*) we assume that, for every parametrization, each model's dynamics is computable, i.e. it can be simulated.

In this work, we will correction maps conditioned to a reference statistic of interest, in the following sense.

**Definition 1 (Statistic)** *A statistic  $\eta$  is any observable that we can compute from one, or from an ensemble of simulations of a model. We denote its estimator computed from model  $q$  with parameters  $x$  as  $q_{\hat{\eta}}(x)$ , and its true value  $q_\eta(x)$ .*

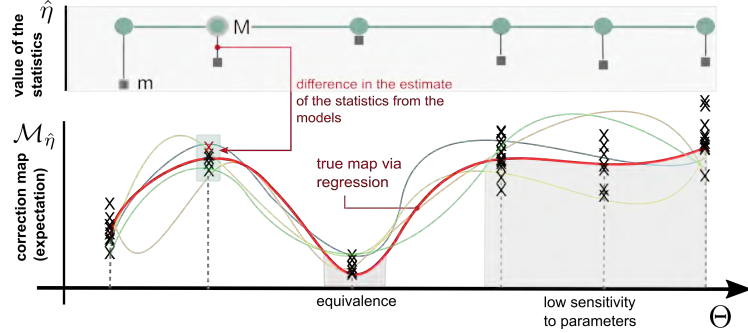
Valid examples of such statistics are, e.g., a single value or the expectation of a variable in  $X$ , the satisfiability of a temporal logical formula depending on variables  $X$  that could be model-checked, etc. The richer the estimator, the higher number of samples are required for the estimator to converge to the true statistics. We will make use of estimators that are *consistent* in the following sense:  $q_{\hat{\eta}}(x) \rightarrow q_\eta(x)$  as the number of samples goes to infinity.

**Definition 2 (Correction map)** *We define an  $\epsilon$ -correction map  $\mathcal{M}_\eta : \Theta \rightarrow \mathbb{R}^w$  for a statistics  $\eta$  to be a function that for any  $\theta_m \in \Theta$ , satisfies*

$$\hat{M}_\eta(\theta_m) \triangleq m_\eta(\theta_m) + \mathcal{M}_\eta(\theta_m) \text{ and } \int_{\Theta} \|\hat{M}_\eta(\theta_m) - \mathbb{E}_{\theta_f}[M_\eta(\theta_m)]\|_2 p(\theta_m) d\theta_m < \epsilon \quad (1)$$

where  $\epsilon > 0$  is the precision, and  $\mathbb{E}_{\theta_f}[M_\eta(\theta_m)] = \int M_\eta(\theta_m; \theta_f) p(\theta_f) d\theta_f$  is the expectation of the statistics computed from  $M$ , with respect to its free parameters  $\theta_f$ .  $\hat{M}$  is our corrected prediction of  $M$ .

<sup>5</sup> In principle, even  $m$  might have a set of free variables, with respect to  $M$ . However, as we have full control over that model, we could assume a parametrization of such variables and all what follows would be equivalent.



**Fig. 1.** Correction maps as regression problems. From sampled estimates of a statistics  $\hat{\eta}$  computed over domain  $\Theta$  by models  $M$  and  $m$ , we can define a correction  $\hat{M}_{\hat{\eta}}(x) \triangleq m_{\hat{\eta}}(x) + \mathcal{M}_{\hat{\eta}}(x)$  from their difference. According to the variables involved, we collect multiple values per parameter value; see also Figure 2. Then, a regression over such values leads to a correction model for  $M$ 's prediction, from  $m$ 's ones; this differs from standard emulation as we retain a mechanistic description of target system, via  $m$ . Maps highlight regions of the parameters where concordance among models and parameters' sensitivity varies, and can be retrieved with regression schema that account for the heteroschedastic variance in the observations.

Thus,  $\mathcal{M}_{\hat{\eta}}$  can correct the outcomes of  $\eta$  assessed over  $m$ ,  $m_{\eta}(\theta_m)$ , to match (with a given precision) those that we would have by computing the statistics over  $M$ . Notice that the corrected outcome has no more dependence from the free parameters, as this is a correction in expectation with respect to  $\theta_f$ .

Notice that the correction is a  $w$ -dimensional vector, and hence  $\|\cdot\|_2$  is used as distance metric, and that the term  $\epsilon$  allows for tolerance in the correction's precision. It is easy to define the optimal, zero-error, correction map:

$$\mathcal{M}_{\eta}^*(\theta_m) \triangleq \mathbb{E}_{\theta_f}[\mathbf{M}_{\eta}(\theta_m)] - \mathbf{m}_{\eta}(\theta_m). \quad (2)$$

However, the correction function  $\mathcal{M}_{\eta}^*(\theta_m)$  is impossible to compute exactly, as we cannot compute neither  $\mathbf{M}_{\eta}$  nor its marginalisation over  $\theta_f$ . Hence, we will learn an approximation of  $\mathcal{M}_{\eta}^*(\theta_m)$  trying to keep its error low so to satisfy Definition 2. We turn this problem into a *regression task*, as graphically explained in Figures 1 and 2, and exploit Gaussian Processes.

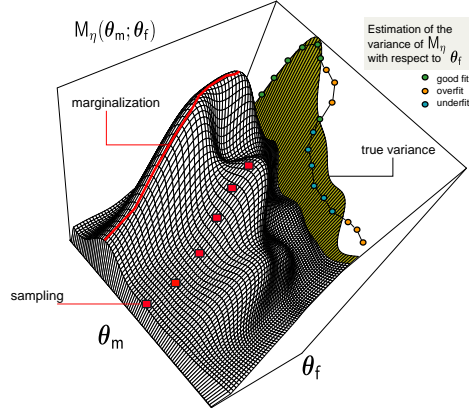
### 3 Learning the Correction Map

In this section we will present our machine learning strategy in more detail. We consider the case of a scalar statistics, as  $w$ -dimensional ones can be treated by solving  $w$  independent learning problems.

#### 3.1 Marginalising $\theta_f$

In order to evaluate (2), we need to be able to compute or approximate the value  $\mathbb{E}_{\theta_f}[\mathbf{M}_{\eta}(\theta_m)]$  with respect to the free parameters of  $M$ , for a any given  $\theta_m$ .

**Fig. 2.** The statistics computed from  $M$ ,  $M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f)$ , has a variance whose contribution depends on the parametrization of all model’s variables, hence on both  $\boldsymbol{\theta}_m$  and  $\boldsymbol{\theta}_f$ . Marginalization is the strategy that, at an exponential cost, ensures that our estimate of such variance is correct; all downsampling strategies, instead, could lead to over or under estimation of such variance. In any case, it is important to account for the variation in such variance as a function of  $\boldsymbol{\theta}_m$ ; this can be achieved by heteroschedastic regression.



As this integral cannot be treated analytically or numerically, due to the large dimensionality involved (the cost is exponential in  $|\boldsymbol{\theta}_f|$ ), we will resort to statistical approximations. Before discussing them, let us comment on the distribution  $p(\boldsymbol{\theta}_f)$ , which is an input for our method. In particular, this distribution encodes our knowledge on the more plausible values of the free parameters. In case we have no information, we can choose a uniform distribution. On the other side of the spectrum, we may know the true value  $\boldsymbol{\theta}_f^*$  of  $\boldsymbol{\theta}_f$ , and choose a delta Dirac distribution, which will dramatically simplify the evaluation of the integral. In this case, we can evaluate (2) as

$$\mathcal{M}_\eta(\boldsymbol{\theta}_m) \triangleq M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f^*) - m_\eta(\boldsymbol{\theta}_m), \quad (3)$$

Moreover, the approximation,  $\int M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f)p(\boldsymbol{\theta}_f)d\boldsymbol{\theta}_f \approx M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f^*)$  is appropriate when the distribution  $p(\boldsymbol{\theta}_f)$  is tightly concentrated around its mode  $\boldsymbol{\theta}_f^*$ .

In general, however, when  $p(\boldsymbol{\theta}_f)$  does not have this special form, we can resort to downsampling  $\mathbb{E}_{\boldsymbol{\theta}_f}[M_\eta(\boldsymbol{\theta}_m)]$ , by generating  $k$  samples  $\boldsymbol{\theta}_f^{(1)}, \dots, \boldsymbol{\theta}_f^{(k)}$  and approximating  $\mathbb{E}_{\boldsymbol{\theta}_f}[M_\eta(\boldsymbol{\theta}_m)] \approx \frac{1}{k} \sum_j M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f^{(j)})$ . In the following, however, we will not necessarily aggregate the values  $M_\eta(\boldsymbol{\theta}_m; \boldsymbol{\theta}_f^{(j)})$ , and treat them individually to better account for the variance in the observed predictions.

### 3.2 Gaussian Processes

We will solve the learning problem resorting to Gaussian Process (GP) regression [13]. GPs are random functions, i.e. probability distributions over a function space, in our case functions  $f : \Theta \rightarrow \mathbb{R}$ , with the property that any finite dimensional projection  $f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_k)$  is a multidimensional Gaussian random variable. It follows that GP are defined by a mean function  $\mu(\boldsymbol{\theta})$ , returning the mean at any point in  $\Theta$ , and by a covariance or kernel function  $k(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ , for giving the covariance between any pair of points in  $\Theta$ . GP can be used to solve regression tasks in a Bayesian setting. The idea is as follows: we put a GP prior on the space of functions  $\{f \mid f : \Theta \rightarrow \mathbb{R}\}$ , typically by assuming constant zero mean

and fixing a kernel function<sup>6</sup>, and then consider the posterior distribution given a set of observations  $Y = \{y^{(i)}\}_I$ ,  $i \in I$ , at input points  $X = \{\theta_m^{(i)}\}$ . If we assume that  $y^{(i)} = f(\theta_m^{(i)}) + \epsilon_i$ , with  $\epsilon_i$  a zero mean Gaussian noise term with variance  $\sigma^2$ , then we obtain that the posterior distribution given the observed data is still a GP, with mean and kernel functions that can be obtained analytically, see [13] for further details. GP regression is essentially the same if the observation noise  $\sigma^2$  is different at different input points, i.e.  $\sigma^2 = \sigma(\theta_m^{(i)})^2$ , in which case we talk about heteroschedastic regression.

### 3.3 The regression task

Let  $\theta_m^{(i)}$  for some index set  $i \in I$  be the *input space*, and  $\{\langle \theta_m^{(i)}, y^{(i)} \rangle\}_I$  our *training points*. In case we use eq. (3) to evaluate the correction map, each  $y^{(i)}$  is a scalar value, and a standard regression schema based on Gaussian Processes can be used. In that case we assume samples of the correction map  $y$  to be observations from a random variable centered at a value given by the *latent function*

$$y^{(i)} \sim \mathcal{N}(\mathcal{M}_\eta(\theta_m^{(i)}), \sigma^2). \quad (4)$$

In this standard Gaussian Processes regression the noise model in the observations is assumed to be a constant  $\sigma^2$  for all sampled points.

In the more general case we work with downsampling solutions that exploit  $k$  samples for the free variable,  $\theta_f^{(1)}, \dots, \theta_f^{(k)}$ . In that case, we have  $k$  correction values per training point,  $\{\langle \theta_m^{(i)}, [y^{(i,1)} \dots y^{(i,k)}]^\top \rangle\}_I$ , that we can use in a straightforward way to reduce to the above schema, or to estimate the variance of  $\mathbf{M}$  conditioned to its free variables. In these cases, the training set is  $\{\langle \theta_m^{(i)}, \bar{y}^{(i)} \rangle\}_I$ , namely we do regression on the point-wise expectation of the correction (i.e.,  $\bar{y}^{(i)} = \frac{1}{k} \sum_{j=1}^k y^{(i,j)}$ ).

**Estimator 1 (Empirical  $\bar{\sigma}$ -estimator.)** Set  $\bar{\sigma}$  to the empirical estimate of the variance across all correction values to exploit the schema in eq. (4) with  $\sigma^2 \triangleq \bar{\sigma}$ .

Besides the simple first case, it is more interesting to account for a *model of the variance in the observations of the predictions from a model*; we discuss how this could be done in two ways.

**Estimator 2 (Point-wise  $\sigma$ -estimator)** Let  $\sigma(\cdot)$  the empirical estimator of the variance of the correction values, per training-point

$$\sigma(\theta_m^{(i)}) \triangleq \text{Var}[y^{(i,1)}, \dots, y^{(i,k)}], \quad (5)$$

then define a model of the variance as a point-wise function of the regression parameter, that is

$$y^{(i)} \sim \mathcal{N}(\mathcal{M}_\eta(\theta_m^{(i)}), \sigma(\theta_m^{(i)})^2). \quad (6)$$

<sup>6</sup> In this work, we use the classic Gaussian kernel fixing hyperparameters by maximising the type-II likelihood; see [13].

In this case, the variance that we expect in each prediction of the latent function is estimated from the data, thus leading to a form of *heteroscedastic* regression.

We can estimate with higher precision a model of the variation in the variance across the input space; to do that we perform *regression of the variance change*, and then inform the outer regression task of that prediction.

**Estimator 3 (Nested  $\sigma$ -estimator)** *Consider the same estimator of the variance as above, and collect the variance estimates as  $\{(\boldsymbol{\theta}_m^{(i)}, w^{(i)})\}_I$ . Learn a latent function model of the true variance  $\sigma_*(\cdot)$ , which we assume to have a fixed variance  $\sigma_*^2$*

$$w^{(i)} \sim \mathcal{N}\left(\sigma_*(\boldsymbol{\theta}_m^{(i)}), \sigma_*^2\right) \quad y^{(i)} \sim \mathcal{N}\left(\mathcal{M}_\eta(\boldsymbol{\theta}_m^{(i)}), \sigma_*(\boldsymbol{\theta}_m^{(i)})^2\right). \quad (7)$$

This is also a form of heteroschedastic regression, but nesting two GP regressions to account in a finer way for the variance's changes.

### 3.4 Properties of the correction map

The correction map that we learn, for all variance schemes, is a statistically sound estimator of  $\mathbb{E}_{\boldsymbol{\theta}_f}[\mathcal{M}_\eta(\boldsymbol{\theta}_m)]$ , in the sense of being consistent.

**Theorem 1 (Correctness)** *Let  $\mathbf{m}_{\hat{\eta}}(\boldsymbol{\theta}_m)$  and  $\mathbb{E}_{\boldsymbol{\theta}_f}[\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)]$  be consistent estimators of  $\mathbf{m}_\eta(\boldsymbol{\theta}_m)$  and  $\mathbb{E}_{\boldsymbol{\theta}_f}[\mathcal{M}_\eta(\boldsymbol{\theta}_m)]$ , then  $\hat{\mathcal{M}}_{\hat{\eta}}(\boldsymbol{\theta}_m) \triangleq \mathbf{m}_{\hat{\eta}}(\boldsymbol{\theta}_m) + \mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)$  is a consistent estimator of  $\mathbb{E}_{\boldsymbol{\theta}_f}[\mathcal{M}_\eta(\boldsymbol{\theta}_m)]$ , for any estimation strategy of  $\mathcal{M}_{\hat{\eta}}$ .*

The result follows because  $\mathcal{M}_{\hat{\eta}}$  converges to  $\mathcal{M}_\eta$  due to properties of GPs [13], and because of the consistency of  $\mathbf{m}_{\hat{\eta}}$  and  $\mathbb{E}_{\boldsymbol{\theta}_f}[\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)]$ . The proof is sketched in Appendix A.2.

The correction map  $\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)$  is estimated from samples of the system, hence it is an approximation of the exact map defined by eq. (2). Thus, it is a correction map in the sense of Definition 2. However, being the result of a GP regression,  $\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)$  is in fact a random function. Therefore, in measuring the error according to eq. (1), we need to consider the average value of the random function  $\mathbb{E}[\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)]$ . The variance  $\text{Var}[\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)]$ , instead, provides a way of computing the error  $\epsilon$ , but in a statistical sense with confidence  $\alpha$ :  $\epsilon$  can be estimated by averaging over  $\Theta$  (with respect to  $p(\boldsymbol{\theta}_m)$ ) the half-width of the region containing  $\alpha\%$  of the probability mass of  $\mathcal{M}_{\hat{\eta}}(\boldsymbol{\theta}_m)$ .

The cost of all our approaches inherently depends on how many samples we pick from  $\Theta$ , the way parameters in  $\boldsymbol{\theta}_f$  are accounted for and the number of parameters in  $\boldsymbol{\theta}_m$ . The sampling cost in general grows exponentially with  $|\boldsymbol{\theta}_m|$ , and each Gaussian regression is cubic in the number of sampled values. Notice that, asymptotically, the cost of the empirical and nested  $\sigma$ -estimators is the same, as the two regressions are executed in series.



## 4 Applications

We discuss now several potential applications of our framework. The advantages of using our approach are mostly computational: the reduced model is simpler to analyze, yet it retains a mechanistic interpretation, compared to statistical emulation.

*Model Building.* Many common problems in the area of dynamical systems can be tackled by resorting to correction maps.

**Problem 1 (Model selection)** *Let  $M$  be a model, and  $m_1, \dots, m_k$  a set of candidate models for correction, each one having a correction map  $\mathcal{M}_{\eta,1}, \dots, \mathcal{M}_{\eta,k}$ . The smallest-correction model  $m^*$  for a statistic  $\eta$  is  $m^* \triangleq \arg \min_{m_i} \int \mathcal{M}_{\eta,i}(\boldsymbol{\theta}) d\boldsymbol{\theta}$ .*

This criterion is certainly alternative to structural Bayesian approaches [3], which can be used to select the structurally smaller model *within* an equivalence class (see below). Also, allows to frame a model synthesis problem.


**Problem 2 (Model synthesis)** *For a model  $M$  with parameters  $\boldsymbol{\theta}_M$  and for a statistic  $\eta$ : (i) partition  $\boldsymbol{\theta}_M$  into sets  $\boldsymbol{\theta}_m$  and  $\boldsymbol{\theta}_f$ , (ii) generate a finite set of plausible reduced models with parameters  $\boldsymbol{\theta}_m$  and (iii) select the best one, according to the above model selection criterion.*

In this case, the partition might aim at identifying in  $\boldsymbol{\theta}_M$  the model's parameters which contribute the most to the variance for the statistics  $\eta$ . Opportunities for *control* are also plausible if one can reduce the problem of controlling  $M$  to “controlling and correcting” a reduced model  $m$ . This should be easier as  $m$  is structurally smaller than  $M$ , and so is lower the complexity of solving a *controller synthesis problem*.

*Parameter Estimation.* Another application of our framework is in Bayesian parameter estimation of the parameters  $\boldsymbol{\theta}_m$  of the big model  $M$ , given observations  $D$  of variables  $X$ , using the small model and the corrected statistics to build approximations of the likelihood function  $p(D | \boldsymbol{\theta}_m)$  to plug into a Bayesian approach (i.e. to compute an approximate posterior). For instance, using the mean of variables  $X$  (and the variance of the correction map as a proxy of the variance  $X$  in  $M$  after marginalisation of  $\boldsymbol{\theta}_f$ ), we can easily compute the distribution of  $X$  under a linear noise approximation and use a Bayesian inference scheme such as [10].

*Model Equivalence.* Correction maps can also be used to compare processes, via weak forms of *bisimilarity* with respect to the observations and a statistics.

**Definition 3 (Model equivalence)** *Models  $M_1$  and  $M_2$ , for a statistic  $\eta$  and parameter sets  $\boldsymbol{\theta}_m$  and  $\boldsymbol{\theta}_f$ , are  $\eta$ -bisimilar conditioned to  $m$ ,  $M_1 \equiv_m^\eta M_2$ , if and only if for all  $\boldsymbol{\theta}_m \in \Theta$ , it holds  $\mathcal{M}_{\eta,1}(\boldsymbol{\theta}_m) = \mathcal{M}_{\eta,2}(\boldsymbol{\theta}_m)$ . A class of equivalence of models with respect to  $m$  and  $\eta$  is the set of all such bisimilar models.*



**Fig. 3.** Example models tested in this paper. Top panel: the Henri-Michaelis-Menten model, where  $\mathbf{m}$  is derived when  $\mathbf{C} : [E]_0 + [ES]_0 \ll [S]_0 + K_{\text{MM}}$ . Bottom panel, a protein translation network where  $\mathbf{m}$  when  $\mathbf{C} : \beta \gg \alpha$ .

This notion of bisimilarity resembles conditional dependence, as we are saying that two models are equivalent if an observer corrects both the same way. In this case,  $\mathbf{m}$  is a *universal corrector* of  $\equiv_{\mathbf{m}}^{\eta}$ , as it can correct for all the models in the class. The class of models that are equivalent to a model  $\mathbf{M}$  is  $\{\mathbf{M}^* \mid \mathbf{M}^* \equiv_{\mathbf{M}}^{\eta} \mathbf{M}\}$  – i.e., the class of models with correction zero; notice that in this case  $\theta_{\mathbf{f}} = \emptyset$ . The previous definition can be relaxed by asking that  $|\mathcal{M}_{\eta,1}(\theta_{\mathbf{m}}) - \mathcal{M}_{\eta,2}(\theta_{\mathbf{m}})| \leq \epsilon$ , for all  $\theta_{\mathbf{m}} \in \Theta$ .

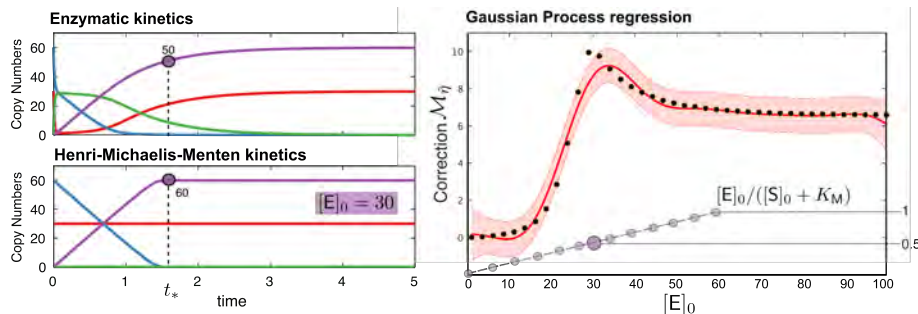
*Remark 1.* Criterion  $\equiv_{\mathbf{m}}^{\eta}$  is a weaker form of *probabilistic bisimilarity*, namely if  $\mathbf{M}_1 \equiv \mathbf{M}_2$  are bisimilar, then  $\mathbf{M}_1 \equiv_{\mathbf{m}}^{\eta} \mathbf{M}_2$  for some  $\mathbf{m}$  and any statistics of interest. For instance, for CTMCs, this follows because  $\equiv$  implies that  $\mathbf{M}_1$  and  $\mathbf{M}_2$  have the same infinitesimal generators for any parameter  $\theta_{\mathbf{m}}$  and  $\theta_{\mathbf{f}}$ , hence the outcomes of  $\mathbf{M}_1$  and  $\mathbf{M}_2$  are indistinguishable, and so are their statistics.

## 5 Examples

We investigate two examples to better illustrate our method.

### 5.1 Model reduction via QSSA

Consider the irreversible canonical enzyme reaction with its exact representation (here, model  $\mathbf{M}$ ), for enzyme  $E$ , complex  $ES$ , substrate  $S$  and product  $PR$  (Figure 3, top left panel). When the concentration of the intermediate complex does not change on the time-scale of product formation, product is produced at rate  $f \triangleq V_{\text{M}}S/(K_{\text{M}} + S)$  where  $V_{\text{M}} = k_2([E]_0 + [ES]_0)$  and  $K_{\text{MM}} = (k_{-1} + k_2)/k_1$ . This is the Henri-Michaelis-Menten kinetics and is technically derived by a quasi-steady-state assumption (QSSA), i.e.,  $\dot{ES} = \dot{E} = 0$ , that is in place when  $\mathbf{C} : [E]_0 + [ES]_0 \ll [S]_0 + K_{\text{MM}}$ , where  $[x]_0$  is the initial amount of species  $x$ .  $\mathbf{m}$  is thus the QSSA reduced model (Figure 3, top right panel).



**Fig. 4.** Correction of the product formation at the transient time  $t_* = 1.5$ , for a mean field model of irreversible canonical enzyme reaction and its simplified Henri-Michaelis-Menten form. Here  $k_1 = 2$ ,  $k_{-1} = 1$  and  $k_2 = 1.5$ ,  $[S]_0 = 60$  and  $[P]_0 = 0$ . Regression over  $[E]_0$  is done with 40 training points from  $(0, 100]$ , and the correction in eq. (3) as M’s free variables are part of the Michaelis-Menten constant.

We interpret these two models as two systems of *ordinary differential equations*, see Appendix A.1, and learn a correction for the following statistics

$$\eta : \mathbb{E}[\text{PR}(t_*)], \text{ with } \text{PR}(t_*) \text{ the number of products at time } t_* \quad (8)$$

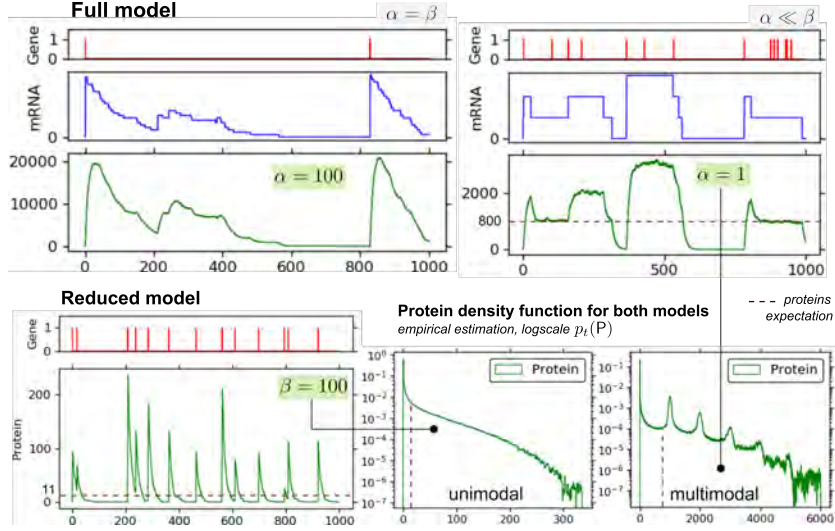
For non-degenerate parameter values both models predict the same equilibrium state, where a total transformation of substrate into product has happened,  $\mathbb{E}[\text{PR}(t)] \rightarrow [S]_0$  for large  $t$ . Thus, we are not interested in correcting the dynamics of  $m$  for long-run times, but rather in the transient (i.e., for small  $t_*$ ).

Also, as the QSSA hypothesis does not hold for certain initial conditions, we set  $\theta_m = \{[E]_0\}$  as the regression variable, and set  $[S]_0 = 60$  and  $[P]_0 = 0$ . The other parameters are  $k_1 = 2$ ,  $k_{-1} = 1$  and  $k_2 = 1.5$  with unit  $(\text{mole per second})^{-1}$ . In terms of regression, we pick 40 samples of the initial enzyme amount from  $(0, 100]$ , and set  $t_* = 1.5$  as it is a time in the transient (manual tuning). This particular example is rather simple, as the free parameters of M are part of the Michaelis-Menten constant and fixed, so we use the simpler correction of eq. (3). Also, knowing when the QSSA holds gives us an interval where we expect the correction to shrink towards zero. The map is shown in Figure 4, which depicts the expected concordance among the correction map and validity of the QSSA.

## 5.2 Model reduction via time-scale separation

Consider a gene switching among active and inactive states of mRNA transcription to be ruled by a *telegraph process* with rates  $k_{\text{off}}/k_{\text{on}}$ . A reaction model of such gene G, protein PR, messenger mRNA with transcription/translation rates  $\alpha/\beta$  as in Figure 3, bottom left panel.

Here the gene switches among active and inactive states, with rates  $k_{\text{on}}$  and  $k_{\text{off}}$ , and PR feeds back positively on inactivation. Proteins and mRNAs are degraded with rates  $\delta_P$  and  $\delta_{\text{RNA}}$ . In the uppermost part of the diagram species



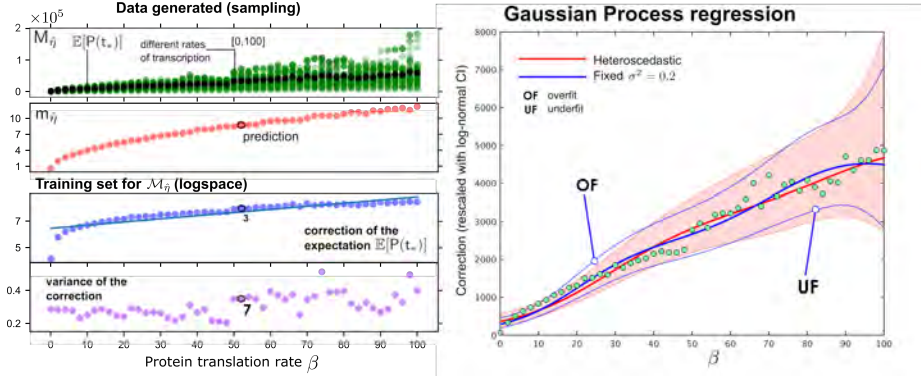
**Fig. 5.** Comparison between the dynamics of the full and the reduced models from §5.2. Parameters are  $k_{\text{off}} = k_{\text{on}} = 10^{-2}$ ,  $\delta_{\text{RNA}} = 10\delta_{\text{P}} = 10^{-2}$ ,  $[\text{G}_{\text{act}}]_0 = 1$ . Values for transcription ( $\alpha$ ) and translation ( $\beta$ ) are reported in the figure. The analysis of empirical proteins' distributions highlights the emergence of multimodal equilibria; the expectation on the number of proteins in the long run spans over different order of magnitudes, according to the relation between  $\alpha$  and  $\beta$ .

marked with a \* symbol are not consumed by a reaction, i.e., mRNA transcription is  $\text{G}_{\text{act}} \rightarrow \text{G}_{\text{act}} + \text{mRNA}$ . This model can be easily encoded in a Markov process, as discussed in A.1.

We can derive an approximation to  $\mathbf{M}$  where the transcription step is omitted. This is valid when  $\mathbf{C} : \beta \gg \alpha$  (*time-scales separation*), namely for every copy of mRNA a protein is immediately translated.

*Correction of protein dynamics.* We build a correction map with  $\theta_{\text{m}} = \{\beta\}$ . In this case the telegraph process is common to both models, but  $\alpha$  and  $\delta_{\text{RNA}}$  are free variables of  $\mathbf{M}$ ; here we assume to have a prior delta distribution over the values of mRNA's degradation, so we set  $\theta_{\text{f}} = \{\alpha\}$ . For some values of the transcription rate condition  $\mathbf{C}$  does not hold; in this case it is appropriate to account for  $\alpha$ 's contribution to the variance in the statistics that we want to correct, when we do a regression over  $\beta$ . Note that also  $\beta$  is part of  $\mathbf{C}$ .

Model  $\mathbf{M}$  leads to stochastic bursts in PR's expression when the baseline gene switching is slower than the characteristic times of translation/transcription. Here we set  $k_{\text{off}} = k_{\text{on}} = 10^{-2}$ , and assume mRNA's lifespan to be longer than protein's ones (also in light of condition  $\mathbf{C}$ ), so  $\delta_{\text{RNA}} = 10\delta_{\text{P}} = 10^{-2}$ . We simulate both models with one active gene,  $[\text{G}_{\text{act}}]_0 = 1$ ; example dynamics are shown in Figure 5, for  $\beta = 100$  and  $\alpha \in \{1, 100\}$ . We observe that, when  $\mathbf{C}$  does not hold



**Fig. 6.** For the regression task we sample 50 values of  $\beta$  from  $[0.1, 100]$ . For  $\alpha$ , instead, we sample 50 random values in the same interval, for each value of  $\beta$ . Correction values are transformed logarithmically (left panel) before doing regression; observe that, on the linear scale, the correction is of the order of  $10^3$  with variance  $10^7$  (midpoint value  $\beta \approx 50$ ). Gaussian Process regression (right panel) is performed with a constant  $\sigma^2 = 0.2$ , eq. (4) and with the  $\sigma$ -estimator, eq. (6). Values are re-scaled linearly, and 95% log-normal confidence intervals are shown; regression highlights that the posterior variances are similar, but the fixed-variance schema tends to underfit or overfit the heteroscedastic variance (*assumed it to be closer to the true one*).

( $\alpha = \beta$ ) the protein bursts increases of one order of magnitude, and the long-run probability density function for the proteins,  $p_t(\text{PR})$ , becomes multimodal.

We define two statistics. One measures the *first moment* of the random variable that models the number of proteins in the long run; the other is a metric interval temporal logic formula [2], expressing the probability of a protein burst within the first 100 time units of simulation.

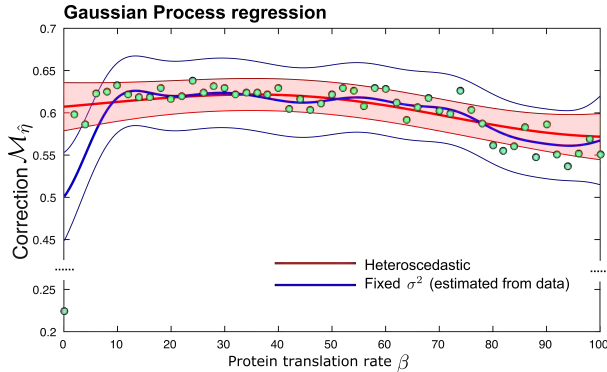
$$\eta_1 : \mathbb{E}[\text{PR}(t_*)], \text{ with } \text{PR}(t_*) \text{ the number of proteins at time } t_* \gg 0 \quad (9)$$

$$\eta_2 : \mathbb{E}[p(\varphi)], \text{ with } \varphi \triangleq \mathbf{F}_{[0,100]} \text{PR}(t) > 200 \quad (10)$$

The former is evaluated by a unique long-run simulation of the model, as its dynamics are ergodic. For the latter we estimate the *satisfaction probability of the formula* via multiple ensembles, as in a parametric *statistical model checking problem* [4].

For the regression task we sample 50 values of  $\beta$ , in the range  $[0.1, 100]$ . For  $\alpha$ , instead, we sample 50 random values in the same interval, for each value of  $\beta$ ; notice that with high probability we pick values where  $\mathbf{C}$  does not hold, so we might expect high correction factors. Data generated and the regression results are shown in Figure 6, for both fixed-variance regression, empirical  $\sigma$ -estimator in eq. (4) and with the  $\sigma$ -estimator, eq. (6). Because variance spans over many orders of magnitude, regression is performed in the logarithmic space. Results highlight a general difference between the posterior variance between the two estimators.

**Fig. 7.** Correction map for the expected satisfaction probability of the linear logic formula  $\eta_2$  in eq. (9). Comparison between the point-wise  $\sigma$ -estimator and the empirical  $\bar{\sigma}$ -estimator. Sampled data is in Appendix Figure 8.



For the second statistics, data is generated from an initial condition where one gene is inactive,  $[G_{in}]_0 = 1$ . Notice that the expected time for the gene to trigger its activation is  $1/k_{on} = 100$  (the time upper-bound of the formula), so for some parametrization there will be non-negligible probability of having no protein spike above threshold 200. The formula satisfaction probability is evaluated with 100 independent model runs, and data generated are shown in Appendix Figure 8. Regression results are shown in Figure 7, where the point-wise  $\sigma$ -estimator and the empirical  $\bar{\sigma}$ -estimator are compared, highlighting the more robustness of the former with respect to the sampled bottom-left outlier.

## 6 Conclusions

Abstraction represents a fundamental tool in the armoury of the computational modeller. It is ubiquitously used in biology as an effective mean to reduce complexity, however systematic analysis tools to quantify discrepancies introduced by abstraction are lacking. Prominent examples, beyond the ones already mentioned, include models with delays, generally introduced to approximately lump multiple biochemical steps [6], or physiological models of organ function which adopt multi-scale abstractions to model phenomena at different organisational levels. These include some of the most famous success stories of systems biology, such as the heart model of [12], which also constitutes perhaps the most notable example of a physical systems which has been modelled multiple times at different levels of abstraction. Employing our techniques to clarify the quantitative relationship between models of cardiac electrophysiology would be a natural and very interesting next step.

Our approach has deep roots in the statistical literature. In this paper we have focussed on the scenario where we try to reconcile the predictions of two separate models, however the complex model was simply used as a sample generator black box. As such, nothing would change if instead of a complex model we used a real system which can be interrogated as we vary some control parameters. Our approach would then reduce to fitting the simple model with a structured, parameter dependent error term. This is closely related with the use

of Gaussian Processes in the geostatistics literature [7], where simple (generally linear) models are used to explain spatially distributed data, with the residual variance being accounted for by a spatially varying random function. Another important connection with the classical statistics literature is with the notion of *emulation* [9]. Emulation constructs a statistical model of the output of a complex computational model by interpolating sparse model results with a Gaussian Process. Our approach can be viewed as a *partial emulation*, whereby we are interested in retaining mechanistic detail for some aspects of the process, and emulate statistically the residual variability. In this light, our work represents a novel approach to bridge the divide between the model-reduction techniques of formal computational modelling and the statistical approximations to complex models.

## References

1. S.Aitken, R.D.Alexander, and J.D.Beggs. A rule-based kinetic model of rna polymerase ii c-terminal domain phosphorylation. *J R Soc Int*, 10(86):20130438, 2013.
2. R. Alur, T. Feder, and T.A. Henzinger. The Benefits of Relaxing Punctuality. *J. ACM*, 43(1):116–146, 1996.
3. D. Barber. *Bayesian reasoning and machine learning*. Cambridge Un Press, 2012.
4. L. Bortolussi, D. Milios, and G. Sanguinetti. Smoothed model checking for uncertain Continuous-Time Markov Chains. *Inf and Comp*, 247:235–253, 2016.
5. L. Bortolussi and G. Sanguinetti. Learning and Designing Stochastic Processes from Logical Constraints. In *Proc. of the 10th Int Conf on Quantitative Evaluation of Systems*, volume 8054 of *LNCS*, pages 89–105, 2013.
6. G. Caravagna. *Formal modeling and simulation of biological systems with delays*. PhD thesis, University of Pisa, 2011.
7. N.Cressie and C.K.Wikle. *Statistics for spatio-temporal data*. Wiley & Sons, 2015.
8. D.C. Hoyle, M Rattray, R. Jupp, and A. Brass. Making sense of microarray data distributions. *Bioinformatics*, 18(4):576–584, 2002.
9. M.C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *J Ro Stat Soc: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
10. M. Komorowski, B. Finkenstadt, C.V. Harper, and D.A. Rand. Bayesian inference of biochemical kinetic parameters using the linear noise approximation. *BMC Bioinf*, 10:343, 2009.
11. N.D Lawrence, G. Sanguinetti, and M. Rattray. Modelling transcriptional regulation using gaussian processes. In *Adv Neural Inf Process Syst*, 785–792, 2006.
12. D. Noble. Modeling the heart—from genes to cells to the whole organ. *Science*, 295(5560):1678–1682, 2002.
13. C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*. MIT Press, Cambridge, Mass., 2006.

## A Appendix

All the code that replicate these analysis is available at the corresponding author's webpage, and hosted on Github (repository GP-correction-maps).

### A.1 Further details on the examples

The two models from §5.1 correspond to these systems of differential equations

$$M := \begin{cases} \dot{E} = -k_1 E \cdot S + k_{-1} ES + k_2 ES \\ \dot{S} = -k_1 E \cdot S + k_{-1} ES \\ \dot{ES} = k_1 E \cdot S - k_{-1} ES - k_2 ES \\ \dot{PR} = +k_2 ES \end{cases} \quad m := \begin{cases} \dot{E} = 0 \\ \dot{S} = -V_M S / (K_M + S) \\ \dot{ES} = 0 \\ \dot{PR} = +V_M S / (K_M + S) \end{cases}$$

which we solved in MATLAB with the ode45 routine with all parameters (InitialStep, MaxStep, RelTol and AbsTol) set to 0.01.

Concerning the Protein Translation Network (PTN) in §5.2, the set of reactions and their propensity functions that we can use to derive a Continuous Time Markov Chain model of the network are the following. Here  $\mathbf{x}$  denotes a generic state of the system and, for instance,  $\mathbf{x}_{\text{mRNA}}$  the number of mRNA copies in  $\mathbf{x}$ .

<i>event</i>	<i>reaction</i>	<i>propensity</i>
activation	$G_{\text{in}} \rightarrow G_{\text{act}}$	$a_1(\mathbf{x}) = k_{\text{on}} \cdot \mathbf{x}_{G_{\text{in}}}$
inactivation	$G_{\text{act}} \rightarrow G_{\text{in}}$	$a_2(\mathbf{x}) = k_{\text{off}} \cdot \mathbf{x}_{PR}$
transcription	$G_{\text{act}} \rightarrow G_{\text{act}} + \text{mRNA}$	$a_3(\mathbf{x}) = \alpha \cdot \mathbf{x}_{G_{\text{act}}}$
mRNA decay	$\text{mRNA} \rightarrow \emptyset$	$a_4(\mathbf{x}) = \delta_{\text{RNA}} \cdot \mathbf{x}_{\text{mRNA}}$
translation	$\text{mRNA} \rightarrow \text{mRNA} + PR$	$a_5(\mathbf{x}) = \beta \cdot \mathbf{x}_{\text{mRNA}}$
PR decay	$PR \rightarrow \emptyset$	$a_6(\mathbf{x}) = \delta_P \cdot \mathbf{x}_{PR}$

The reduced PTN model is a special of this reactions set where transcription and mRNA decay are omitted. In this case we used StochPy to simulate the models and generate the input data per regression – see <http://stochpy.sourceforge.net/>; data sampling exploits python parallelism to reduce execution times.

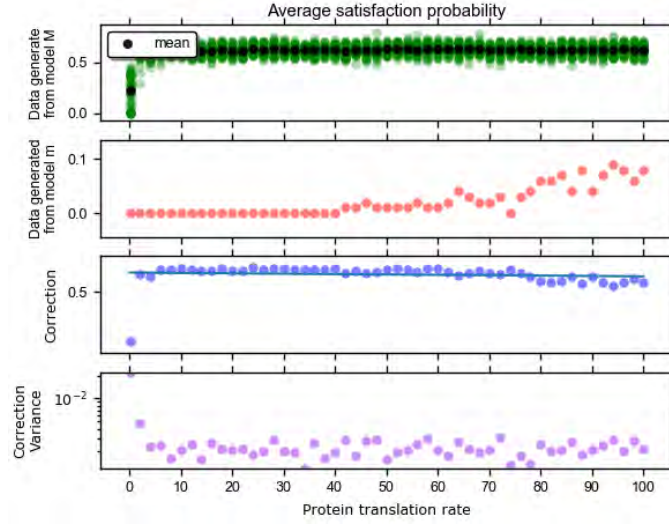
For regression, we used the Gaussian Processes for Machine Learning toolbox for fixed-variance regression, see <http://www.gaussianprocess.org/gpml/code/matlab/doc/> and a custom implementation of the other forms of regression.

### A.2 Proofs

Proof of Theorem 1.

*Proof.* Both the empiricals and nested estimator rely on an *unbiased* estimator of the mean/variance, which means that if  $k \rightarrow \infty$ , i.e., we sample all possible values for the free variables, we would have a true model of  $\bar{y}/\sigma$ . This means that, for each sampled value from  $\Theta$ , even the simplest  $\bar{\sigma}$ -estimator would be equivalent, in expectation, to the marginalization of the free variables. This is





**Fig. 8.** Data generated to compute the satisfaction probability of the linear logic formula  $\eta_2$  in eq. (9). For each model 100 independent simulations are used to estimate the expectation of the probability. The regression input space is the same used to compute  $\eta_1$ , but the models are simulated with just one inactive gene in the initial state. The heteroscedastic variance in the regression is computed as the variance of the correction of the expected satisfaction probability (point-wise  $\sigma$ -estimator); the fixed-variance regression is computed by estimating the variance from the data (empirical  $\bar{\sigma}$ -estimator).

enough, combined with properties of Gaussian Processes regression (i.e., the convergence to the true model with infinite training points), to state that the overall approach leads to an unbiased estimator of the correction map.