

Experimental Design for Dynamics Identification of Cellular Processes

Vu Dinh · Ann E. Rundell · Gregory T. Buzzard

Received: 8 April 2013 / Accepted: 29 January 2014 / Published online: 13 February 2014
© Society for Mathematical Biology 2014

Abstract We address the problem of using nonlinear models to design experiments to characterize the dynamics of cellular processes by using the approach of the Maximally Informative Next Experiment (MINE), which was introduced in W. Dong et al. (PLoS ONE 3(8):e3105, 2008) and independently in M.M. Donahue et al. (IET Syst. Biol. 4:249–262, 2010). In this approach, existing data is used to define a probability distribution on the parameters; the next measurement point is the one that yields the largest model output variance with this distribution. Building upon this approach, we introduce the Expected Dynamics Estimator (EDE), which is the expected value using this distribution of the output as a function of time. We prove the consistency of this estimator (uniform convergence to true dynamics) even when the chosen experiments cluster in a finite set of points. We extend this proof of consistency to various practical assumptions on noisy data and moderate levels of model mismatch. Through the derivation and proof, we develop a relaxed version of MINE that is more computationally tractable and robust than the original formulation. The results are illustrated with numerical examples on two nonlinear ordinary differential equation models of biomolecular and cellular processes.

Keywords Experimental design · System identification · Response surface modeling · Cellular dynamics

V. Dinh (✉) · G.T. Buzzard

Department of Mathematics, Purdue University, 150 N. University Street, West Lafayette, IN 47907, USA

e-mail: vdinh@purdue.edu

A.E. Rundell

Weldon School of Biomedical Engineering, Purdue University, 206 S. Martin Jischke Drive, West Lafayette, IN 47907, USA

1 Introduction

The development and simulation of mathematical models of cellular processes can enhance our understanding of the underlying biological mechanisms (Kitano 2002). Two important components of model development are the collection of data and the tuning of parameters for a given model structure to approximate the data. In many settings, the collection of data is difficult and/or expensive, while tuning model parameters to data often involves a difficult nonlinear optimization, with potentially many local optima. Moreover, the choice of data may make this tuning more or less difficult; thus we aim to design experiments to collect the most informative data for a given model structure.

The review (Kreutz and Timmer 2009) provides a broad overview of model-based experimental design methodologies for systems biology, including methods for various optimality conditions governing unique, structural and practical parameter identification. Of course, many books and articles have been written about experimental design, both from the frequentist and Bayesian points of view. We make no attempt to review them all here; a classic mathematical reference is Pukelsheim (1993). Many methods for experimental design focus on identifying the parameters—designing experiments to minimize some measure of uncertainty in the parameter values given a model structure.

In contrast, we are more concerned with developing a method to explore and elucidate the observable response (which we refer to as the output dynamics) of a cellular process rather than identifying the model parameters themselves. One motivation for this is that for a systems biology model with N parameters, the set of possible output dynamics is often contained in a space of lower dimension $l \ll N$ (or perhaps in a small neighborhood of such a space). This feature, which is an obstacle for the problem of unique parameter identification, is an advantage for designing experiments to identify dynamics; that is, we can choose a good design with very few experiments (approximately $O(l)$) but still obtain enough information to identify the dynamics.

Methods related to approximating the observable response as a function of independent input variables fall under the broad heading of regression or response surface methodology. Once again, there are many books and articles on the topic of experiment design for fitting response surfaces, and there are many approaches for representing a response surface. Most such approaches (e.g., Kriging and generalized polynomial chaos) seek to approximate the response surface with a linear combination of a fixed set of basis functions, such as polynomials or trigonometric functions. See Myers and Montgomery (2002) for an overview.

In this paper we focus on experiment design for accurate approximation of the response surface using a given biologically based model. However, beyond this, as explained in Dong et al. (2008), the method we apply acts as a kind of imaging method for understanding the behavior of a cell. Based on an initial understanding of cell behavior (which is likened to an image from a microscope), we choose the next experiment to provide as much resolving power as possible in our next measurement (which is likened to focusing the microscope to enhance a particular feature). Instead of a linear combination of basis functions, as is often assumed in the experiment design literature, we assume that the model structure encodes the dominant interactions

and mechanisms using a nonlinear system of differential equations. To fix ideas, suppose that our model output has the form $y = f(\omega, t)$, where $\omega \in \Omega \subset \mathbb{R}^N$ is a fixed vector of unknown parameters, and where our quantity of interest is the dynamics output, which is a function of time, $t \in [0, T]$ (more generally t could be a vector of inputs to represent any independent variables such as time, voltage, etc.). Measurements of y at a given time t can be modeled as a random variable. An estimate of ω based on these random variables is also a random variable. This estimate can be used to estimate the output $y(t)$ for any t , again giving a random variable. Classical experiment design typically seeks to minimize the variance in the estimate of ω or y . One approach to designing experiments for accurate response surface modeling is to use the condition of G-optimality. Roughly, this condition chooses an experiment design to minimize the variance in the output. In the case of a model that is linear in parameters, the Kiefer–Wolfowitz Equivalence Theorem states that this is equivalent to D-optimality, in which the design is chosen to minimize the determinant of the covariance matrix for the estimate of ω (the inverse of the Fisher Information Matrix) (Pukelsheim 1993, Chap. 9). There is an extension of this result to nonlinear models (White 1973); however, this result depends upon knowing the true parameters in the model, which are not known in general. In Fedorov and Hackl (1997, Sect. 5.6) this problem is addressed by either (i) using a minimax approach, in which the design is chosen to minimize over experiment designs the maximum over parameter space of the determinant; or (ii) using a Bayesian approach, in which an optimality criterion for a design (such as the determinant of the dispersion matrix or the maximum output variance) is averaged using a prior distribution on parameter space, and then the design is chosen to maximize this expected criterion. A computational difficulty with both of these approaches is the need to evaluate a complex optimality criterion at many points in parameter space for each candidate design.

Alternatively, the Maximally Informative Next Experiment (MINE) algorithm proposed in Dong et al. (2008) and later in Donahue et al. (2010) uses a sequential approach to experiment design in which existing data is used to construct a probability distribution on the parameter space; this distribution is then used to calculate the variance in the output as a function of time (perhaps normalized by expected experimental variance). Based on this calculation, the next measurement will be taken at the time point with highest current (normalized) variance. That is, the next sampled time point will be chosen at the time point that has highest current uncertainty in the output. This method of design is modified to produce a parallel (nonsequential) design in Bazil et al. (2011). Intuitively, each new point in such a design should provide the maximum possible information about the dynamics of the system and hence lead to convergence to the true system dynamics. This approach is theoretically appealing in that it does not depend on an estimate of the true parameter values, and it is computationally appealing in that it requires a relatively simple sampling over the parameter space according to a specified distribution; this may be achieved reasonably efficiently with Monte Carlo methods.

However, little is known about the convergence properties of this method: Is this method sufficient to characterize the response surface (in the limit as the number of experimental points tends to ∞)? In fact, in general, this scheme will not sample densely over the interval $[0, T]$, so it is not at all clear that it is sufficient to completely

characterize the dynamics over this interval. Moreover, it is not entirely clear how to use these nondense samples to estimate the dynamics.

Motivated by the MINE algorithm, we address the following two problems for the identification of systems dynamics:

- (A) Specify the Dynamics Estimator: Given a set of data $(t_1, d_1), (t_2, d_2), \dots, (t_m, d_m)$ and a model $y = f(\omega, t)$, how should we estimate the system dynamics?
- (B) Prove convergence of Dynamics Estimator: For a given sequential approach to choosing measurement points t_j and given the dynamics estimator in (A), do the estimated dynamics converge to the true dynamics?

In the derivation of the solutions to these problems, we developed variations of the MINE algorithm that are more computationally efficient than the original. We describe these variations and solutions to Problems (A) and (B) with various assumptions in the body of the paper.

Most approaches to problem (A) use the data to estimate parameter values and then use these parameters to obtain the corresponding dynamics. For a complex, nonlinear model this is a difficult optimization problem with possibly many local optima and perhaps even multiple global optima. In place of using an estimated vector of parameters to estimate the dynamics, we propose what we call the expected dynamics estimator (EDE). This uses the available data to induce a probability distribution on parameter space and then averages the dynamic output using this distribution. There are a number of advantages of this method of dynamics identification over parameter identification. First, since the dynamics for a deterministic system are unique, we do not need to worry about multiple global solutions. Second, by using the EDE, we look for the average behavior of the system (with respect to a carefully constructed probability distribution). This task is typically much simpler than solving a nonlinear optimization problem. Furthermore, the Markov Chain Monte Carlo method can be employed to reconstruct the system's true dynamics.

Another important advantage of the probabilistic framework over parameter estimation (via optimization) is that it is a feasible approach in cases of unidentifiability. A crucial problem in parameter estimation is the calculation of confidence intervals for the estimated parameters. In the simplest scenario when the number of data points is smaller than the number of parameters, any parameter estimation (via optimization) method will fail to provide a reliable estimate of the confidence region. Such methods (that return a single parameters estimate) will never be able to predict unknown output with high confidence (or any confidence at all). In order to do so, it needs to compute all possible parameter values that are consistent with available data, which is very unlikely in practice. This also extends to the case when the model parameters are unidentifiable, which is a common phenomenon in systems biology. Our probabilistic framework provides a feasible way to address this issue: a given set of measurements gives a probability distribution on parameters, which can be used to construct confidence interval for output dynamics in addition to the EDE.

Problem (B) is a question about the consistency of the estimator (the ability to recover the true dynamics) as a function of a particular choice of measurement points. This question highlights the fact that the ability of the EDE to recover the true dynamics (consistency) depends heavily on the experimental design algorithm. We note

here that although the MINE method shares some features of a Bayesian approach, in that a probability distribution on parameters is updated based on new data samples, it does not fall into the class of Bayesian experimental design since the design points are not chosen to maximize an expected utility function.

This paper is organized to prove and illustrate the consistency of the EDE in various situations that progressively increase in complexity toward practical applicability. In Sect. 2, we introduce the mathematical framework and the main assumptions about the behavior of the investigated model that we use throughout the paper. We also define the EDE to address problem (A). Section 3 addresses problem (B) for the ideal case where the investigated model is a correct model (can reproduce the true dynamics exactly) and data are noiseless. Theorem 3.1 deals with the case in which the experiments are made at random time points; this result is provided primarily to illustrate the ideas to be used in later results but in a setting that avoids some technical assumptions that are needed later. Theorem 3.2 and Theorem 3.4 provide results in the case where the experiments are designed sequentially as in Donahue et al. (2010) and Dong et al. (2008) in two different settings: when the parameter space is discretized and when the sets of possible measured time points and output values are discretized. We then extend the consistency result to a larger class of designs by relaxing the choice of a point with maximal variance to a point with variance within a fixed constant multiple of maximal variance (Theorem 3.3). Our results imply that for these designs, we can always recover the true dynamics, even if all the measurements are made in a small portion of the time interval $[0, T]$. Section 4 extends the result about EDE's consistency to the case where the experimental data are subject to random noise (Theorem 4.2). In this section we require that the set of possible measurement points is finite in order to guarantee convergence even in the face of noisy data. This assumption is reasonable for the practical implementation of any experiment design. Section 5 relaxes the requirement of a correct model by allowing for a bounded error between the true dynamics and the closest approximation of the model (Theorem 5.1). From this result we also justify the use of approximation methods in the algorithm to design experiments. In Sect. 7, we illustrate our theoretical findings and demonstrate the efficacy of our method to design sequential experiments for dynamics identification with various biological models. We also give an example to show that the choice of a design point within a fixed constant of maximal variance can lead to a faster rate of convergence of the EDE relative to the original MINE algorithm. It is worth noting that although the framework we use in this paper is sequential, one can extend the result to the parallel case following the approach suggested in Bazil et al. (2011).

2 Mathematical Framework

2.1 Model Formulation

We assume a mathematical model of a cellular process in the form

$$\begin{aligned}\dot{x} &= \alpha(\omega, x) && \text{(System of ODEs),} \\ x(0) &= x_0(\omega) && \text{(Initial conditions),} \\ y(t) &= f(\omega, t) = \beta(\omega, x(t)) && \text{(Output),}\end{aligned}$$

where $x = (x_1, x_2, \dots, x_{n_x}) \in M \subset \mathbb{R}^{n_x}$ is the state variable, with M a subset of \mathbb{R}^{n_x} containing the initial state, and $f(\omega, t) \in \mathbb{R}^L$ is the observable response (output dynamics) that corresponds to L different experimentally observable quantities. Throughout this paper, for the sake of simplicity, we will assume that $L = 1$. However, all of the arguments can be extended to the case of multidimensional output without any difficulty.

It is worth noting that the set of possible outputs is not necessarily the same as the number of dynamic variables occurring in the system. An output could be any kind of prediction, e.g., also a sum or ratio or even integral of dynamic variables. However, in the case $L = 1$, there is only one observable output y . Identification of y will lead to identification of all possible outputs as well as a characterization of the uncertainty in unidentifiable outputs.

The purpose of our experimental design framework is to determine as accurately as possible the output dynamics based on measurements. This is a kind of interpolation problem. We do not address the extrapolation problem, in which measurements of one output are used to make inference about an unobservable quantity.

The vector of unknown parameters is denoted by $\omega = (\omega_1, \dots, \omega_N) \in \mathbb{R}^N$ and is assumed to belong to a subset Ω of \mathbb{R}^N . In most parts of the paper, the parameter space Ω will be assumed to be an open set along with a probability measure on Ω , or a discrete subset of \mathbb{R}^N along with a probability measure. The components of α and β are assumed to be C^1 functions of their arguments. These functions and initial conditions may depend on the parameter vector $\omega \in \Omega$.

The system will therefore be associated with the mapping $F : \Omega \rightarrow C^1([0, T], \mathbb{R})$ defined by $F(\omega) = f(\omega, \cdot)$, where $f(\omega, \cdot)$ is the observable response of the system as a function of $t \in [0, T]$ for a given ω . The image of Ω under f , $Y = f(\Omega, \cdot) \subset C^1([0, T], \mathbb{R})$ will be referred to as the dynamics space in this paper.

Throughout this paper, the true dynamics and the data values at a given time, t , will be denoted by $g(t)$ and $d(t)$, respectively. We assume that $d(t) = g(t) + \epsilon$, where ϵ is a random variable describing the noise in measurements. In Sect. 3, we assume that $\epsilon = 0$, so that the data are completely noise-free. In later sections we address the case of noisy data. In Sects. 3 and 4 we assume that the model is correct; that is, there is some $\omega_0 \in \Omega$ so that $f(\omega_0, t) = g(t)$ for all $t \in [0, T]$. We relax this assumption in Sect. 5.

2.2 Expected Dynamics Estimator (EDE)

A given data set $(t_1, d_1), \dots, (t_n, d_n)$ will be used to induce a probability distribution on the parameter space. We do this through the normalized likelihood function,

$$p_n(\omega) = c_n \exp\left(-\sum_{i=1}^n (d_i - f(\omega, t_i))^2\right)$$

(or a variant of this expression), where c_n is a constant such that p_n is a probability distribution on Ω . (Note that if no data has been observed, the distribution p_0 is just the uniform distribution in Ω .)

The expected dynamics estimator (EDE) with respect to this probability distribution is then

$$\hat{D}_n(t) = E_{p_n(\omega)}[f(\omega, t)],$$

which we use as an estimator of the system's true dynamics. Thus, instead of trying to maximize the likelihood function in order to estimate dynamics, we average the output dynamics, weighted by the probability as determined by the likelihood function. It is also worth noting that the EDE is the natural estimator that is used frequently as a part of the ensemble method and is usually computed by Monte Carlo Markov Chain methods.

3 EDE Consistency for Noise-Free Data

In this section, we establish results about the consistency of the expected dynamics estimator, that is, the ability to recover the true dynamics under a specified experimental design. The proof will be provided for two different cases:

- (1) When the sampled time points $\{t_n\}$ are chosen at random from an absolutely continuous probability distribution μ on $[0, T]$.
- (2) When the sampled time points are chosen sequentially as in Donahue et al. (2010) and Dong et al. (2008), where the next sampled time point will be the point with highest current uncertainty (output variance).

Before moving forward to analyze the convergence of the EDE in these two cases, it is worth mentioning the distinction between two different sources of uncertainties (in both parameters and output dynamics): noise in data (aleatoric uncertainty), and structural uncertainty (epistemic uncertainty) in the model. Given a set of noise-free data, the corresponding set of parameter values that fit the data perfectly well can still be an infinite set (usually, is a union of low-dimensional manifolds). The simplest example for this phenomenon is when the number of data is less than the number of model parameters.

In unidentifiable nonlinear systems, this set of “fitted” parameters may not collapse to a point mass even if all measurable outputs are known completely. This uncertainty in parameters may never be eliminated. The forward propagation of this uncertainty to the output space is the target in this noise-free framework.

The likelihood proposed in the noise-free setting, therefore, is not associated with noise in data, but with the structural uncertainty in model parameters from available data (how well a parameter set fits the data). Instead of focusing on a low-dimensional set of “fitted” parameters, we use an everywhere positive likelihood function to constrain the parameter space. From a methodological point of view, the idea here is similar to those behind simulated annealing methods for optimization and multiple Monte Carlo Markov Chains method for statistical inference: since the objects of interest is difficult to identify, we relax it by heated objects that are easier to study and use our experimental design algorithm to sequentially reduce the temperature in an optimal way to identify the true output dynamics.

3.1 Randomly Chosen Experimental Design Points

To illustrate the ideas used in later results, we consider the case where the sampled time points $\{t_n\}$ are chosen independently at random from an absolutely continuous probability distribution, with the assumption that the data are noise free (i.e., $d(t_i) = g(t_i)$ for all i). In this setting, we have the following theorem, which says that in the limit as $n \rightarrow \infty$, the expected dynamics estimator converges to the system's true dynamics.

Theorem 3.1 *Suppose there exists $\omega_0 \in \Omega$ such that $f(\omega_0, t) = g(t)$ for all $t \in [0, T]$. Suppose also that $\{t_n\}$ are chosen independently at random from an absolutely continuous probability distribution μ on $[0, T]$ and that $1 \leq r < \infty$. Let*

$$p_n(\omega) = c_n \exp\left(-\sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r\right),$$

where c_n is the normalizing constant to ensure that p_n is a probability distribution on Ω . Then for all $t \in [0, T]$,

$$\lim_{n \rightarrow \infty} E_{p_n}[f(\omega, t)] = g(t).$$

Moreover, the convergence is uniform in t .

Before proving the theorem, we provide some intuition. Every time a new time point is sampled, the likelihood function is multiplied by a new term of the form $\exp(-|f(\omega, t_{n+1}) - g(t_{n+1})|^r)$. If ω does not correspond to the true dynamics, there must be a region of $[0, T]$ where $f(\omega, t)$ differs from $g(t)$. Since the $\{t_n\}$ are chosen independently at random from an absolutely continuous probability distribution, eventually multiple time points will be sampled in this region, causing the value of the likelihood at ω go to 0. Therefore, in the limit as $n \rightarrow \infty$, the distribution $p_n(\omega)$ will concentrate more and more on the set of ω that corresponds to the true dynamics. Hence the expected dynamics will also converge to the system's true dynamics.

We use the following two lemmas, whose proofs will be provided in Sect. 6. The first is a result on the convergence of Monte Carlo integration. The second is a result on the convergence of the EDE.

Lemma 3.1 *Let points t_i be chosen as in Theorem 3.1, and let $1 \leq r < \infty$. Define*

$$h_n(\omega) = \exp\left(-\frac{1}{n} \sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r\right)$$

and

$$h(\omega) = \exp\left(-\int_0^T |f(\omega, t) - g(t)|^r d\mu(t)\right).$$

Then

$$\frac{h_n(\omega)}{h(\omega)} \rightarrow 1 \quad \text{uniformly in } \omega \in \Omega$$

and

$$\lim_{n \rightarrow \infty} \|h_n\|_n = \|h\|_\infty.$$

Lemma 3.2 *Let a and b be continuous functions on $\Omega \times [0, T]$ and $[0, T]$, respectively, and let $\{p_n\}$ be a sequence of probability distributions on Ω .*

(a) *Define*

$$h(\omega) = \exp\left(-\int_0^T |a(\omega, t) - b(t)|^r d\mu(t)\right),$$

and suppose that

- (i) *for any $\alpha < 1$, there exists $\delta < 1$ and $C > 0$ such that if $\omega \in \Omega$ with $h(\omega) \leq \alpha \|h\|_\infty$, then $p_n(\omega) < C\delta^n \forall n$;*
- (ii) *there exists $\omega_0 \in \Omega$ such that $a(\omega_0, t) = b(t)$ for all $t \in [0, T]$.*

Then

$$\lim_{n \rightarrow \infty} E_{p_n}[a(\omega, t)] = b(t) \quad \forall t \in [0, T]$$

and

$$\lim_{n \rightarrow \infty} \text{Var}_{p_n}[a(\omega, t)] = 0 \quad \forall t \in [0, T].$$

Moreover, for both limits, the convergence is uniform in t .

(b) *Assume that Ω is finite and that there exists a set $S \subset [0, T]$ such that*

$$\{\omega \in \Omega : p_n(\omega) \not\rightarrow 0\} \subset \{\omega \in \Omega : a(\omega, t) = b(t) \forall t \in S\}.$$

Then for all t in S ,

$$\lim_{n \rightarrow \infty} E_{p_n}[a(\omega, t)] = b(t)$$

and

$$\lim_{n \rightarrow \infty} \text{Var}_{p_n}[a(\omega, t)] = 0.$$

Moreover, for both limits, the convergence is uniform in t .

Proof of Theorem 3.1 Let

$$q_n(\omega) = \exp\left(-\sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r\right).$$

Then $p_n = c_n q_n$, and $q_n = (h_n)^n$, where h_n is defined as in Lemma 3.1, so

$$p_n(\omega) = \frac{q_n(\omega)}{\int_\Omega q_n(\omega) d\omega} = \frac{h_n^n(\omega)}{\int_\Omega h_n^n(\omega) d\omega} = \left(\frac{h_n(\omega)}{\|h_n\|_n}\right)^n.$$

Let $0 < \alpha < 1$, and suppose $\omega \in \Omega$ with $h(\omega) \leq \alpha \|h\|_\infty$. By Lemma 3.1 we have $\lim_{n \rightarrow \infty} h_n(\omega) = h(\omega)$ and $\lim_{n \rightarrow \infty} \|h_n\|_n = \|h\|_\infty$. Let $\epsilon > 0$ and $\delta = \alpha(1 + \epsilon)^2$. For ϵ small, we have $\delta < 1$, and for this ϵ , there exists N (independent of ω) large enough such that if $n > N$, then

$$h_n(\omega) \leq (1 + \epsilon)h(\omega) \leq \alpha(1 + \epsilon)\|h\|_\infty \leq \alpha(1 + \epsilon)^2\|h_n\|_n.$$

Hence, for all $n > N$,

$$p_n(\omega) = \left(\frac{h_n(\omega)}{\|h_n\|_n} \right)^n \leq \delta^n$$

with $\delta < 1$. Since there exists $\omega_0 \in \Omega$ such that $f(\omega_0, t) = g(t)$ for all $t \in [0, T]$, we can apply Lemma 3.2(a) with $a = f$ and $b = g$ to obtain the uniform convergence

$$\lim_{n \rightarrow \infty} \int_{\Omega} p_n(\omega) f(\omega, t) d\omega = g(t) \quad \text{for all } t \in [0, T].$$

The integral on the left is $E_{p_n}[f(\omega, t)]$, so this gives the desired equality. \square

Note that the proof depends on the sequence $\{t_i\}$ only through Lemma 3.1, so the result holds for any sequence that yields the conclusion in that lemma. A quasi-random sequence satisfying a low-discrepancy condition (Morokoff and Caflisch 1994) is one such sequence, so we make the following remark.

Remark 3.1 The conclusion of Theorem 3.1 is still valid if $\{t_i\}$ is a low-discrepancy sequence, i.e.,

$$D_N(\{t_1, \dots, t_N\}) := \sup_{B \subset \Omega} \left| \frac{\#\{1 \leq i \leq N : t_i \in B\}}{N} - \text{Vol}(B) \right| \rightarrow 0$$

as N approaches infinity.

The results in this section imply that if data is collected uniformly at random, we can recover the true dynamics from the sampled data. This is one example of a so-called space-filling design (Cox and Reid 2000). However, in practice, randomly chosen points do not produce an efficient experimental design since many of the measurements will not give much information about the system; the convergence, although guaranteed, may be slow.

3.2 Design Points Using the Maximally Informative Next Experiment

Intuitively, we expect the sequential designs of Donahue et al. (2010) and Dong et al. (2008), for which the next sampled time point is the one that has the highest current uncertainty (variance), to increase the convergence rate relative to randomly selected design points. On the other hand, the measured points may no longer be dense in $[0, T]$, so it is not clear that the dynamics may be recovered on the entire interval.

In the following theorem, we extend the consistency result in the previous subsection to this type of sequential design, with the additional assumption that Ω is finite.

This assumption was also used in the context of parameter identification in Pronzato (2009, 2010). We conclude that we can recover the entire true dynamics, even if all the measurements are made in a small subset of $[0, T]$ (in the extreme case, at one point). As in the previous subsection, we still assume that data are subject to no error.

Theorem 3.2 *Let ω_0, r, p_n be as in Theorem 3.1 and assume that Ω has finite cardinality. Suppose that each t_{n+1} is chosen so that*

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \leq \text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \quad \forall t \in [0, T]. \quad (1)$$

Then

$$\lim_{n \rightarrow \infty} E_{p_n}[f(\omega, t)] = g(t) \quad \forall t \in [0, T].$$

That is, the EDE converges to the true dynamics of the system. Moreover, the convergence is uniform in t .

By choosing the next time point to be the point with the highest variance, we put a constraint on the variance of the whole dynamics: the variance at other points must be smaller than the variance at the measured points, which in turn converges to 0. In this case we deduce that the expected dynamics on the whole interval converges to some limit dynamics. If we can prove further that a “true” parameter vector ω_0 is still in the support of the limit distribution, then obviously this limit dynamics is equal to the true system dynamics.

As above, this is straightforward when the t_i are chosen at random from an absolutely continuous distribution. However, the case where Ω is an open set and t_i are chosen according to (1) is a bit different. In a continuous framework, a parameter vector has measure zero, and good performance of the true parameter vector does not guarantee that it will stay in the support of the limit distribution. Such a situation can happen in the case where the model is not robust around the true parameter and at the chosen time points, the neighborhood around true parameters in the parameter space fit the data worse than some other regions. This may cause the expected dynamics to converge to incorrect dynamics. Though this situation is perhaps unlikely to happen in practice, we cannot exclude such a possibility for a convergence result.

To resolve this issue, we assume in Theorem 3.2 that Ω is a finite set. This may be achieved, for example, by subdividing each coordinate axis using a fixed step size and taking the set of points in Ω that lie on the resulting grid. An alternative approach in which the outputs and the set of possible measured time points are discretized instead of Ω is also suggested in Theorem 3.4. Both assumptions are natural and do not hinder the applicability of the method in practice.

Proof of Theorem 3.2 As in the proof of Theorem 3.1, let

$$q_n(\omega) = \exp\left(-\sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r\right),$$

and recall that $p_n = c_n q_n$. Also, let A be the set of cluster points of $\{t_n\}$: points $t \in [0, T]$ such that there exists a subsequence $\{t_{n_k}\}$ of $\{t_n\}$ with $t_{n_k} \rightarrow t$.

We claim first that if $p_n(\omega)$ does not tend to 0 with n (so that ω has probability above some fixed $\rho > 0$ for infinitely many n), then $f(\omega, t) = g(t)$ for all $t \in A$. Indeed, consider any $\omega \in \Omega$, $t \in A$ such that $|f(\omega, t) - g(t)| = c > 0$. Since A is the set of limit points of $\{t_n\}$, there exists a subsequence $\{t_{n_k}\}$ of $\{t_n\}$ such that $t_{n_k} \rightarrow t$. Since f and g are continuous, for k large enough, we have $|f(\omega, t_{n_k}) - g(t_{n_k})| \geq c/2$. Hence,

$$\sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r \rightarrow \infty$$

as $n \rightarrow \infty$, and so $q_n(\omega) \rightarrow 0$.

On the other hand, the assumption that some ω_0 gives the true dynamics implies that $f(\omega_0, t) - g(t) = 0$ for all t , hence $q_n(\omega_0) = 1$. Therefore, $p_n(\omega_0)/p_n(\omega) \rightarrow \infty$. Since Ω is a finite space, $p_n(\omega_0) \leq 1$, and hence $p_n(\omega) \rightarrow 0$. Hence, $p_n(\omega) \not\rightarrow 0$ implies $f(\omega, t) = g(t)$ for all $t \in A$.

Using Lemma 3.2(b) (for finite Ω) with $a = f$ and $b = g$, we deduce that

$$\lim_{n \rightarrow \infty} E_{p_n} [f(\omega, t)] = g(t) \quad \forall t \in A$$

and

$$\text{Var}_{p_n(\omega)} [f(\omega, t)] \rightarrow 0 \quad \forall t \in A.$$

On the other hand, the choice of t_{n+1} gives

$$\text{Var}_{p_n(\omega)} [f(\omega, t)] \leq \text{Var}_{p_n(\omega)} [f(\omega, t_{n+1})] \quad \forall t \in [0, T]. \quad (2)$$

Now we claim that

$$\text{Var}_{p_n(\omega)} [f(\omega, t_{n+1})] \rightarrow 0. \quad (3)$$

Indeed, by contradiction, assume that there exists a subsequence $\{t_{n_k}\}$ and a positive constant C such that

$$\text{Var}_{p_{n_k}(\omega)} [f(\omega, t_{n_k+1})] \geq C$$

for all k . Since $[0, T]$ is compact, we can drop to a subsequence to assume that t_{n_k+1} converges to some $t_0 \in A$. By the continuity of f and its derivatives on the compact set $\Omega \times [0, T]$, there is $C_0 > 0$ such that for all $k > 0$ and $\omega \in \Omega$,

$$|f(\omega, t_{n_k}) - f(\omega, t_0)| \leq C_0 |t_{n_k} - t_0|. \quad (4)$$

Hence, by using this inequality, we have

$$\limsup_{k \rightarrow \infty} E_{p_{n_k}(\omega)} |f(\omega, t_{n_k+1}) - f(\omega, t_0)| \leq \lim_{k \rightarrow \infty} C_0 |t_{n_k+1} - t_0| = 0,$$

which implies that

$$\lim_{k \rightarrow \infty} E_{p_{n_k}(\omega)} [f(\omega, t_{n_k+1})] = \lim_{k \rightarrow \infty} E_{p_{n_k}(\omega)} [f(\omega, t_0)].$$

By a similar argument, we also have

$$\lim_{k \rightarrow \infty} E_{p_{n_k}(\omega)}[f^2(\omega, t_{n_k+1})] = \lim_{k \rightarrow \infty} E_{p_{n_k}(\omega)}[f^2(\omega, t_0)].$$

Therefore,

$$\lim_{k \rightarrow \infty} \text{Var}_{p_{n_k}(\omega)}[f(\omega, t_{n_k+1})] = \lim_{k \rightarrow \infty} \text{Var}_{p_{n_k}(\omega)}[f(\omega, t_0)] = 0,$$

which contradicts the choice of C .

From (2) and (3) we obtain

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \leq \text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \rightarrow 0 \quad \forall t \in [0, T].$$

In other words, for all t in $[0, T]$,

$$\lim_{n \rightarrow \infty} \sum_{\omega \in \Omega} p_n(\omega) (f(\omega, t) - E_{p_n}[f(\omega, t)])^2 = 0. \quad (5)$$

The fact that ω_0 gives the true dynamics implies that ω_0 is a maximum for q_n , hence for p_n . Hence, $p_n(\omega_0) \geq p_n(\omega)$ for all $\omega \in \Omega$, and from the fact that Ω is finite we deduce that $p_n(\omega_0) \geq 1/|\Omega|$. Using this with (5) gives

$$\begin{aligned} (f(\omega_0, t) - E_{p_n}[f(\omega, t)])^2 &\leq |\Omega| p_n(\omega_0) (f(\omega_0, t) - E_{p_n}[f(\omega, t)])^2 \\ &\leq |\Omega| \sum_{\omega \in \Omega} p_n(\omega) (f(\omega, t) - E_{p_n}[f(\omega, t)])^2 \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. Hence,

$$E_{p_n(\omega)}[f(\omega, t)] \rightarrow f(\omega_0, t) = g(t) \quad \forall t \in [0, T]. \quad \square$$

As in the discussion before the proof, the only reason we use the variance criterion is to bound the variance of dynamics at unmeasured points by the variance at measured points. This criterion can be relaxed as follows.

Theorem 3.3 *The result from Theorem 3.2 is still valid if Condition (1) (that the next time point has maximum variance) is replaced by the condition that the variance at the next time point is within a fixed constant of the maximum variance. That is, there exists $C > 1$ such that for all $t \in [0, T]$,*

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \leq C \text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})]. \quad (6)$$

There are several motivations for this relaxation of criterion (1). First, in practice, the optimization of the variance function (which is usually done by Markov Chain Monte Carlo methods) is subject to random effects arising in the sampling process. By using criterion (6), we look for a near-optimal solution of the optimization problem; this condition is stable with respect to an MCMC scheme. Second, in real experiments, some sets of measurements may be more expensive and technically difficult

than the others. By looking for a near-optimal solution, we make it possible for experimenters to find an alternative measurement when the optimization problem gives rise to an optimum that is experimentally difficult to implement. Finally, as we will see in Sect. 4, the new criterion allows us to apply additional criteria for point selection in order to facilitate resampling, which will be used to establish the convergence rate of the EDE in the face of noisy data.

As noted above, instead of discretizing Ω , we may discretize the output space and the measurement space. This gives the following result.

Theorem 3.4 *The result from Theorem 3.2 is still valid if we assume that Ω is open and bounded, but that the possible outputs of the system and the set of possible measured time points are both finite. In this case, we get convergence of the EDE on the full (finite) set of possible measured time points.*

Proof We denote by $\mathcal{T} = \{\tau_i\}_{i=1}^K$ the set of all possible measured time points and assume that at each time point, the output function $f(\omega, t)$, as well as the true dynamics $g(t)$, are discretized by a finite grid. That is, the continuous function $g(t)$ is approximated by values $(Rg(\tau_1), \dots, Rg(\tau_K))$, where $Rg(t)$ is obtained from $g(t)$ by rounding to the nearest of some finite set of allowable output values.

By making arbitrarily small perturbations to the possible output values if necessary, we can assume without loss of generality that $\forall \tau \in \mathcal{T}$, the true output value $g(\tau)$ does not lie midway between two allowable output values. Therefore, there exists an open neighborhood U_{ω_0} of ω_0 such that if $\omega \in U_{\omega_0}$, then $Rf(\omega, \tau) = Rf(\omega_0, \tau) = Rg(\tau) \forall \tau \in \mathcal{T}$. For the remainder of the proof, we use f and g to mean Rf and Rg .

As in the proof of Theorem 3.2, we now consider any $\omega \in \Omega$ such that $p_n(\omega)$ does not tend to 0 with n . Assume that $f(\omega, t) \neq f(\omega_0, t)$ for some t in the cluster set, A , of $\{t_i\}$ (since \mathcal{T} is finite here, this is the set of points that are measured infinitely many times). Using the same argument as in the proof of Theorem 3.2, we deduce that $p_n(\omega) \leq p_n(\omega_0)$ and $p_n(\omega_0)/p_n(\omega) \rightarrow \infty$. Note that in this case, although Ω is not finite, the argument is still valid since p_n is constant on the open set U_{ω_0} , so that $p_n(\omega_0)$ is bounded above by $1/\text{Vol}(U_{\omega_0})$, where $\text{Vol}(U)$ denotes the volume of a set U .

Therefore, $p_n(\omega) \not\rightarrow 0$ implies $f(\omega, t) = g(t)$ for all $t \in A$. Since Ω may be infinite, Lemma 3.2 cannot be applied directly in this case. However, by denoting $U_A = \{\omega \in \Omega : f(\omega, t) = g(t) \forall t \in A\}$, we have for all $t \in A$

$$\left| \int_{\Omega} p_n(\omega) f(\omega, t) d\omega - g(t) \right| \leq \int_{\Omega \setminus U_A} p_n(\omega) |f(\omega, t) - g(t)| d\omega.$$

Since f and g are bounded and $p_n(\omega_0) < 1/\text{Volume}(U_{\omega_0})$, we have $|p_n(\omega)| |f(\omega, t) - g(t)| \leq Cp_n(\omega_0) \leq C/\text{Vol}(U_{\omega_0})$. Also, $p_n(\omega) \rightarrow 0$ on $\Omega \setminus U_A$, so by the Dominated Convergence Theorem, the right-hand side converges to 0 as n tends to ∞ . We deduce that $E_{p_n(\omega)}[f(\omega, t)] \rightarrow g(t) \forall t \in A$. By a similar argument, we also have $\text{Var}_{p_n(\omega)}[f(\omega, t)] \rightarrow 0 \forall t \in A$.

We next use the fact that the set $\mathcal{T} = \{\tau_i\}_{i=1}^K$ of possible measured time points is finite to deduce that $\text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \rightarrow 0$. Indeed, assume that

$$\text{Var}_{p_{n_k}(\omega)}[f(\omega, t_{n_k+1})] \geq C$$

for some subsequence $\{n_k\}$ and positive constant C . Since \mathcal{T} is finite, there exists $t_0 \in A$ that appears in the subsequence $\{t_{n_k}\}$ infinitely many times; this implies that $\text{Var}_{p_{n_k}(\omega)}[f(\omega, t_0)] \not\rightarrow 0$, which is a contradiction.

Hence, $\text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \rightarrow 0$. Combining this with (1), we see that in fact $\text{Var}_{p_n(\omega)}[f(\omega, t)] \rightarrow 0$ for all $t \in \mathcal{T}$. This proves that the EDE converges to the true system dynamics on \mathcal{T} . \square

Note that the condition of discrete measured time points in the previous result allows us to avoid the need for the regularity condition (4), which does not hold for the piecewise constant functions obtained by discretizing the system outputs. In the next section we provide further justification for a finite set of measurement points.

4 EDE Consistency with Noisy Data

In practice, of course, data from experiments are subject to noise. Hence, in this section we extend the results from previous sections to the case of additive Gaussian noise. As is common in many settings, we assume that

$$d(t_i) = g(t_i) + \epsilon_i,$$

where $g(t)$ is the true dynamics (which is unknown), $d(t_i)$ is the measured data at the sampled time point t_i , and ϵ_i are i.i.d. Gaussian random variables (see Dong et al. (2008) for empirical support of this noise model).

The analysis in the case of noisy data is a bit different from that used in the previous section. Intuitively, if “close”, but not exactly at the same points in time, t_1 and t_2 , are measured, and if there is a functional relation between the output at t_1 and at t_2 , then the measurement at t_1 will also help refine the information about data at t_2 . However, theoretically, this assertion is difficult to prove and may even be incorrect, due to nonlinearity: if the relations between outputs at t_1 and t_2 are nonlinear, using data at t_2 to constrain t_1 may create a bias in the fitted output.

For example, if $f_2 = f_1^2$, then

$$(f_1 + e)^2 = f_1^2 + 2ef_1 + e^2 = f_2 + ef_1 + e^2.$$

When using averaging, the linear error term will go away by the strong law of large number, but the quadratic term will have positive expectation, which results in a bias in estimation of f_2 . The stronger the nonlinearity, the larger the bias, and that makes it hard to clarify the convergence.

In order to obtain a convergence result using noisy data, we need to be able to average over multiple trials, which makes sense only if we measure repeatedly at a given time. In the theorem below, as in Theorem 3.4, we discretize the time interval and allow measurements to be taken at only finitely many specified points and use a slightly different form of probability distribution. This guarantees that experiments will be replicated many times at “important” points. When data are collected multiple times, the average value is used to constrain the dynamics: the larger the number of times we make the measurement at a time point t , the more confidence we put on the

average data at that point. With this framework, we again have the convergence of the EDE to the true system dynamics.

The idea of using a finite grid to replace the whole time interval to facilitate re-sampling is a common technique in the problem of parameter identification (Pronzato 2009, 2010). In studies of ODEs, under the assumption that f is analytic, this is further supported by the following theorem from Sontag (2002), which guarantees that if we can identify the system dynamics on a finite grid, we can identify the dynamics on the whole interval.

Theorem 4.1 (Sontag 2002) *Assume $f(w, t)$ depends analytically on ω and t , and let N be the dimension of the parameter space. Then, for Lebesgue almost every randomly chosen set of $2N + 1$ experiments, the following property holds: For any two parameters that have distinct dynamics, one of the experiments in this set will distinguish them.*

We further note that the assumption of analyticity may be replaced by the assumption that Ω is a finite set. That is, with this assumption we can find a finite grid $\mathcal{T} \subset [0, T]$ that satisfies the above property: For any two parameters that have distinct dynamics on $[0, T]$, one of the experiments $t \in \mathcal{T}$ in this grid will distinguish them.

Finally, even in the case where the parameter space Ω is an open set, by choosing the discretized time points to be the nodes for an efficient interpolation scheme, then by interpolating on this finite set, convergence on the finite set of times \mathcal{T} converts to uniform approximation on the entire interval $[0, T]$.

Hence throughout this section, we discretize $[0, T]$ to a finite grid $\mathcal{T} = \{\tau_i\}_{i=1}^K$ and assume that the experiments can be made only at the nodes of this grid. We also continue to assume that Ω has finite cardinality. With these assumptions, we have the following theorem.

Theorem 4.2 *Let $C > 1$. Assume that Ω is finite and at step n , $t_{n+1} \in \mathcal{T}$ is chosen so that*

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \leq C \text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \quad \forall t \in \mathcal{T}.$$

For $1 \leq k \leq K$, let $k_n(\tau_i)$ be the number of experiments made at time τ_i up through step n , and $\{d_j(\tau_i)\}_{j=1}^{k_n(\tau_i)}$ be the data values from those experiments, with $d_j(\tau_i) = g(\tau_i) + \epsilon_j$; the ϵ_j are iid $N(0, \sigma^2)$. Define $B_n = \{\tau_i : k_n(\tau_i) > 0\}$ and

$$p_n(\omega) = c_n \exp\left(-\sum_{\tau_i \in B_n} k_n(\tau_i) \left(f(\omega, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i)\right)^r\right),$$

where c_n is the normalizing constant and $r > 2$. Then

$$\lim_{n \rightarrow \infty} E_{p_n}[f(\omega, t)] = g(t) \quad \forall t \in \mathcal{T}.$$

Moreover, the convergence is uniform in $t \in \mathcal{T}$.

The same result with $r = 2$ is also valid if the following condition is satisfied:

$$\lim_{n \rightarrow \infty} \frac{\log \log k_n(\tau_1)}{k_n(\tau_2)} = 0 \quad \forall \tau_1, \tau_2 \in A, \quad (7)$$

where A is the set of all cluster points of $\{t_n\}$.

Note that if \mathcal{T} satisfies the conditions in the discussion following Theorem 4.1, then determining the dynamics in \mathcal{T} is sufficient to determine the dynamics on the whole $[0, T]$.

Proof Let $A = \{\tau_i : \lim_{n \rightarrow \infty} k_n(\tau_i) = \infty\}$ and

$$q_n(\omega) = \exp \left(- \sum_{\tau_i \in B_n} k_n(\tau_i) \left(f(\omega, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right)^r \right).$$

We claim that

$$\{\omega : q_n(\omega) \not\rightarrow 0\} \subset \{\omega : f(\tau_i, \omega) = g(\tau_i) \quad \forall \tau_i \in A\}.$$

Indeed, consider any $\omega \in \Omega$ and $\tau_{i_0} \in A$ such that $|f(\omega, \tau_{i_0}) - g(\tau_{i_0})| = c > 0$. Let $X_j = d_j(\tau_{i_0}) = g(\tau_{i_0}) + \epsilon$ and note that $\{X_j\}$ is a Gaussian sequence of iid random variables with $E[X_j] = g(\tau_{i_0}) = f(\omega, \tau_{i_0})$. By the law of large numbers, we have with probability 1

$$g(\tau_{i_0}) = \lim_{n \rightarrow \infty} \frac{1}{k_n(\tau_{i_0})} \sum_{j=1}^{k_n(\tau_{i_0})} d_j(\tau_{i_0}).$$

Hence, there exists N such that for all $n > N$,

$$\left| g(\tau_{i_0}) - \lim_{n \rightarrow \infty} \frac{1}{k_n(\tau_{i_0})} \sum_{j=1}^{k_n(\tau_{i_0})} d_j(\tau_{i_0}) \right| \leq c/2,$$

which implies (by the triangle inequality)

$$\left| f(\tau_{i_0}, \omega) - \lim_{n \rightarrow \infty} \frac{1}{k_n(\tau_{i_0})} \sum_{j=1}^{k_n(\tau_{i_0})} d_j(\tau_{i_0}) \right| \geq c/2.$$

Therefore,

$$\sum_{\tau_i \in B_n} k_n(\tau_i) \left(f(\omega, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right)^r \geq (c/2)^r k_n(\tau_{i_0}) \rightarrow \infty \quad (8)$$

as $n \rightarrow \infty$.

Now consider any $\tau_i \in A$. By the law of the iterated logarithm, we have with probability 1

$$\limsup_{n \rightarrow \infty} \sqrt{\frac{k_n(\tau_i)}{\log \log k_n(\tau_i)}} \left(f(\omega_0, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right) = \sqrt{2}.$$

If $r > 2$, there exists a constant C such that

$$k_n(\tau_i) \left(f(\omega_0, \tau_i) - \frac{1}{k} \sum_{j=1}^k d_j(\tau_i) \right)^r \leq C \frac{\log \log k_n(\tau_i)}{(k_n(\tau_i))^{(r/2-1)}} \rightarrow 0 \quad (9)$$

as $n \rightarrow \infty$.

Since this is true for any τ_i in the finite set A , we have

$$\lim_{n \rightarrow \infty} \sum_{\tau_i \in B_n} k_n(\tau_i) \left(f(\omega_0, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right)^r < \infty.$$

Therefore, $q_n(\omega_0)$ is bounded below, so $\frac{p_n(\omega_0)}{p_n(\omega)} = \frac{q_n(\omega_0)}{q_n(\omega)} \rightarrow \infty$. Since Ω is a finite space, $p_n(\omega_0) \leq 1$. This makes $p_n(\omega) \rightarrow 0$.

In the case where $r = 2$, we have

$$k_n(\tau_i) \left(f(\omega_0, \tau_i) - \frac{1}{k} \sum_{j=1}^k d_j(\tau_i) \right)^2 \leq C \log \log k_n(\tau_i),$$

which implies

$$\limsup_{n \rightarrow \infty} \sum_{\tau_i \in B_n} k_n(\tau_i) \left(f(\omega_0, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right)^2 \leq C_1 + \sum_{\tau_i \in A} C \log \log k_n(\tau_i).$$

Equation (7) implies that there exists N such that for all $n \geq N$,

$$\log \log k_n(\tau_i) \leq \frac{(c/2)^2}{2C(\#A)} k_n(\tau_{i_0}) \quad \forall \tau_i \in A,$$

where $\#A$ denotes the cardinality of A , and c is the constant in (8).

Combine this inequality and (8) (with $r = 2$), we have

$$\frac{p_n(\omega)}{p_n(\omega_0)} = \frac{q_n(\omega)}{q_n(\omega_0)} \leq \exp \left(C_1 - \frac{c^2}{8} k_n(\tau_{i_0}) \right) \rightarrow 0$$

as $n \rightarrow \infty$. Hence, $p_n(\omega) \rightarrow 0$.

At this point, we have proved that

$$\{\omega : p_n(\omega) \not\rightarrow 0\} \subset \{\omega : f(\omega, \tau_i) = g(\tau_i) \quad \forall \tau_i \in A\}.$$

Hence, by Lemma 3.2(b),

$$\lim_{n \rightarrow \infty} E_{p_n}[f(\omega, \tau_i)] = g(\tau_i) \quad \forall \tau_i \in A$$

and

$$\lim_{n \rightarrow \infty} \text{Var}_{p_n}[f(\omega, \tau_i)] = 0 \quad \forall \tau_i \in A.$$

On the other hand, we have

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \leq C \text{Var}_{p_n(\omega)}[f(\omega, t_{n+1})] \quad \forall t \in \mathcal{T}.$$

Using the same argument as in the proof of Theorem 3.2, we have

$$\text{Var}_{p_n(\omega)}[f(\omega, t)] \rightarrow 0 \quad \forall t \in \mathcal{T}$$

and

$$E_{p_n(\omega)}[f(\omega, t)] \rightarrow f(\omega_0, t) = g(t) \quad \forall t \in \mathcal{T}. \quad \square$$

By an argument similar to that used in the proof of Theorem 3.4, we obtain the following result.

Theorem 4.3 *The result of Theorem 4.2 is still valid if we replace the condition of finite cardinality of Ω with the condition of a finite set of output values and possible measurement time points.*

5 EDE Consistency with Model Mismatch

So far we have investigated various schemes to design experiments for dynamics identification, under the assumption that the investigated model is a correct model, i.e., there exists $\omega_0 \in \Omega$ such that $f(\omega_0, t) = g(t)$ for all $t \in [0, T]$. Here we relax this condition using the concept of ϵ -equivalence.

Definition 5.1 Let $\epsilon > 0$ and suppose that g and h are continuous on $[0, T]$. Then g and h are ϵ -equivalent means that $\|g - h\|_\infty < \epsilon$.

To obtain the main result in this section, we also need to assume that the function outputs are discretized by a finite grid of resolution ϵ , similar to the discretization in Theorem 3.4. However, here we use an adaptive discretization in that it changes based on the measurements obtained so far and based on the time point.

Definition 5.2 Let h be continuous on $[0, T]$, let \mathcal{T} and $d_j(\tau_i)$ be as in Theorem 4.2, and let $\epsilon > 0$. Define

$$d_n^*(\tau_i) = \begin{cases} \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) & \text{if } k_n(\tau_i) > 0, \\ 0 & \text{if } k_n(\tau_i) = 0, \end{cases}$$

and

$$R_n^\epsilon h(\tau_i) = d_n^*(\tau_i) + \operatorname{sgn}(h(\tau_i) - d_n^*(\tau_i)) \left\lfloor \frac{|h(\tau_i) - d_n^*(\tau_i)|}{\epsilon} \right\rfloor \epsilon.$$

This choice of discretization is needed to guarantee the convergence of the estimated dynamics. With this setting, we have the following theorem, in which we use the framework of Theorem 4.2 but with the output discretization just given. The proof of this theorem is a combination of the techniques employed in Theorem 4.2 and Theorem 3.4. Here $\lfloor x \rfloor$ denotes the largest integer less than or equal to x .

Theorem 5.1 *Let Ω , C , \mathcal{T} , B_n , k_n , $d_j(\tau_i)$ be as in Theorem 4.2, $\epsilon_0 > 0$, and assume that there is $\omega_0 \in \Omega$ such that $f(\omega_0, t)$ and $g(t)$ are ϵ_0 -equivalent. For $\epsilon > \epsilon_0$, define*

$$p_n(\omega) = c_n \exp \left(- \sum_{\tau_i \in B_n} k_n(\tau_i) \left(R_n^\epsilon f(\omega, \tau_i) - \frac{1}{k_n(\tau_i)} \sum_{j=1}^{k_n(\tau_i)} d_j(\tau_i) \right)^2 \right),$$

where c_n is the normalizing constant, and assume that at each step, the next measurement is chosen so that

$$\operatorname{Var}_{p_n(\omega)} [R_n^\epsilon f(\omega, t)] \leq C \operatorname{Var}_{p_n(\omega)} [R_n^\epsilon f(\omega, t_{n+1})] \quad \forall t \in \mathcal{T}.$$

Then, for almost every $\epsilon > \epsilon_0$, the expected dynamics converges (uniformly in $t \in \mathcal{T}$) to limit dynamics that are ϵ -equivalent to $g(t)$.

Proof Denote by A the set of all $t \in \mathcal{T}$ that are measured infinitely many times. By the strong law of large numbers, we have $d_n^*(\tau) \rightarrow g(\tau)$ for all $\tau \in A$. Since Ω and \mathcal{T} are finite, there is a full measure set of $\epsilon > \epsilon_0$ such that for all $\tau \in \mathcal{T}$ and $\omega \in \Omega$, the distance between $g(t)$ and $f(\omega, \tau)$ is not a multiple of ϵ . This implies that $\lim_{n \rightarrow \infty} |f(\omega, \tau) - d_n^*(\tau)|/\epsilon$ is not an integer for any $\omega \in \Omega$ and $\tau \in A$, hence that $\lim_{n \rightarrow \infty} R_n^\epsilon f(\omega, \tau)$ exists for all such ω and τ . Also, if $\tau \notin A$, then $R_n^\epsilon f(\omega, \tau)$ is constant for n large enough. Hence, $\lim_{n \rightarrow \infty} R_n f(\omega, \tau)$ exists for all $\omega \in \Omega$, $\tau \in \mathcal{T}$.

On the other hand, the assumption on ω_0 implies that for all $\tau \in \mathcal{T}$,

$$|d_n^*(\tau_i) - f(\omega_0, \tau)| \leq |d_n^*(\tau_i) - g(\tau)| + \epsilon_0.$$

For n sufficiently large, the right-hand side is less than ϵ for all $\tau \in \mathcal{T}$. For such n , we have $R_n^\epsilon f(\omega_0, \tau) = R_n^\epsilon g(\tau)$ for all τ in \mathcal{T} .

Now consider $\omega \in \Omega$ such that $\lim_{n \rightarrow \infty} R_n^\epsilon f(\omega, \tau) \neq \lim_{n \rightarrow \infty} R_n^\epsilon f(\omega_0, \tau)$ for some τ in A . Using the same argument as in the proof of Theorem 4.2, we deduce that $p_n(\omega_0)/p_n(\omega) \rightarrow \infty$. Since Ω is a finite space, $p_n(\omega_0) \leq 1$. This makes $p_n(\omega) \rightarrow 0$. We have proved that

$$\{\omega : p_n(\omega) \not\rightarrow 0\} \subset \left\{ \omega : \lim_{n \rightarrow \infty} R_n f(\omega, \tau) = \lim_{n \rightarrow \infty} R_n g(\tau) \quad \forall \tau \in A \right\}.$$

Then, by Lemma 3.2(b),

$$\lim_{n \rightarrow \infty} E_{p_n} [R_n^\epsilon f(\omega, \tau)] = \lim_{n \rightarrow \infty} E_{p_n} [R_n^\epsilon g(\tau)] = g(\tau) \quad \forall \tau \in A$$

and

$$\lim_{n \rightarrow \infty} \text{Var}_{p_n} [R_n^\epsilon f(\omega, \tau)] = 0 \quad \forall \tau \in A.$$

On the other hand, we have

$$\text{Var}_{p_n(\omega)} [R_n^\epsilon f(\omega, t)] \leq C \text{Var}_{p_n(\omega)} [R_n^\epsilon f(\omega, t_{n+1})] \quad \forall t \in \mathcal{T}.$$

Using the same argument as in the proof of Theorem 3.2, we have

$$\text{Var}_{p_n(\omega)} [R_n^\epsilon f(\omega, t)] \rightarrow 0 \quad \forall t \in \mathcal{T}$$

and

$$\lim_{n \rightarrow \infty} E_{p_n(\omega)} [R_n^\epsilon f(\omega, t)] = \lim_{n \rightarrow \infty} R_n g(t) \quad \forall t \in \mathcal{T}.$$

This proves that the EDE converges to limit dynamics that are ϵ -equivalent to the true system dynamics on \mathcal{T} . \square

6 Proofs of Supporting Lemmas

In this section, we provide the proofs of the two lemmas that have been used throughout this paper.

6.1 Lemma 3.1 (Convergence of Monte Carlo Integration)

Proof First, we note that

$$\frac{h_n(\omega)}{h(\omega)} = \exp \left(\int_0^T |f(\omega, t) - g(t)|^r d\mu(t) - \frac{1}{n} \sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r \right).$$

Using the Koksma–Hlawka inequality for convergence of quasi-Monte Carlo integration (Niederreiter 1978), we have

$$\begin{aligned} & \left| \int_0^T |f(\omega, t) - g(t)|^r d\mu(t) - \frac{1}{n} \sum_{i=1}^n |f(\omega, t_i) - g(t_i)|^r \right| \\ & \leq r D_n^* \int_0^T |f(\omega, t) - g(t)|^{r-1} \left| \frac{\partial f}{\partial t}(\omega, t) - g'(t) \right| d\mu(t), \end{aligned} \quad (10)$$

where D_n^* is the discrepancy of the finite sequence $\{t_1, t_2, \dots, t_n\}$ (see Morokoff and Caflisch 1994 for more information about the discrepancy).

Since f is a C^1 function on the compact set $\Omega \times [0, T]$ and g is C^1 on $[0, T]$, there exists M independent of ω and t such that

$$\int_0^T |f(\omega, t) - g(t)|^{r-1} \left| \frac{\partial f}{\partial t}(\omega, t) - g'(t) \right| d\mu(t) \leq M. \quad (11)$$

Since μ is absolutely continuous with respect to Lebesgue measure, we have $D_n^* \rightarrow 0$ as $n \rightarrow \infty$.

From (10) and (11) we have for any $\omega \in \Omega$ that

$$|\log(h_n(\omega)/h(\omega))| \leq rMD_n^* \rightarrow 0.$$

Hence, $h_n(\omega)/h(\omega) \rightarrow 1$ uniformly in $\omega \in \Omega$.

Also, since $h_n(\omega) \leq 1$ for all ω and $h(\omega_0) = 1$, we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|h_n\|_n &= \limsup_{n \rightarrow \infty} \left(\int_{\Omega} h_n^n d\omega \right)^{1/n} \\ &\leq \limsup_{n \rightarrow \infty} \text{Vol}(\Omega)^{1/n} \\ &= 1 = \|h\|_{\infty}. \end{aligned} \quad (12)$$

To get a lower bound, let $\epsilon > 0$. Note that if $h(\omega) \geq 1 - \epsilon/2$ and $|h_n(\omega) - h(\omega)| \leq \epsilon/2$, then by the triangle inequality we have $h_n(\omega) \geq 1 - \epsilon$. Also, since h is continuous and $\|h\|_{\infty} = 1$, we have

$$C_{\epsilon} := \text{Vol}(\{\omega : h(\omega) \geq 1 - \epsilon/2\}) > 0.$$

Since h_n/h converges uniformly on Ω to 1 and $|h| \leq 1$, there exists $N(\delta, \epsilon)$ large enough such that for all $n \geq N$ and all $\omega \in \Omega$, $|h_n(\omega) - h(\omega)| \leq \epsilon/2$. Hence,

$$\text{Vol}(\{\omega : h_n(\omega) \geq 1 - \epsilon\}) \geq C_{\epsilon}.$$

So

$$\int_{\Omega} h_n^n d\omega \geq \int_{\{h_n \geq 1-\epsilon\}} h_n^n d\omega \geq C_{\epsilon}(1-\epsilon)^n$$

and

$$\left(\int_{\Omega} h_n^n dx \right)^{1/n} \geq C_{\epsilon}^{1/n}(1-\epsilon).$$

Taking $n \rightarrow \infty$, we deduce

$$\|h\|_{\infty} \geq \limsup_{n \rightarrow \infty} \|h_n\|_n \geq \liminf_{n \rightarrow \infty} \|h_n\|_n \geq \|h\|_{\infty} - \epsilon.$$

Since ϵ was arbitrary, $\lim_{n \rightarrow \infty} \|h_n\|_n = \|h\|_{\infty}$. □

6.2 Lemma 3.2 (Convergence of the Expected Dynamics Estimator)

Proof We provide the proof for part (a). The proof for part (b) uses a similar argument.

Let $\epsilon > 0$ and define

$$U = \{\omega \in \Omega : |a(\omega, t) - b(t)| < \epsilon \forall t \in [0, T]\}.$$

Since a, b are continuous and $a(\omega_0, t) = b(t)$ for all t , we see that U is a neighborhood of ω_0 . Then for $t \in [0, T]$, we have

$$\begin{aligned} & \left| \int_{\Omega} p_n(\omega) a(\omega, t) d\omega - b(t) \right| \\ & \leq \int_{\Omega} p_n(\omega) |a(\omega, t) - b(t)| d\omega \\ & = \int_{\Omega \setminus U} p_n(\omega) |a(\omega, t) - b(t)| d\omega + \int_U p_n(\omega) |a(\omega, t) - b(t)| d\omega \\ & \leq \int_{\Omega \setminus U} p_n(\omega) |a(\omega, t) - b(t)| d\omega + \epsilon. \end{aligned}$$

Now we claim that there exists $\alpha < 1$ such that $\forall \omega \in \Omega \setminus U, h(\omega) \leq \alpha$. Indeed, assume that $\exists \omega_n \in \Omega \setminus U$ with $h(\omega_n) \rightarrow 1$. Then for each n there is $t_n \in [0, T]$ such that $|a(\omega_n, t_n) - b(t_n)| \geq \epsilon$. Since $\Omega \times [0, T]$ is compact, without loss of generality, we can assume that $\omega_n \rightarrow \omega^* \in \Omega, t_n \rightarrow t^* \in [0, T]$. Since a and b are continuous, we deduce that $|a(\omega^*, t^*) - b(t^*)| \geq \epsilon$ and $h(\omega^*) = 1$.

However, $h(\omega^*) = 1$ implies that $\int_0^T |a(\omega^*, t) - b(t)|^r d\mu(t) = 0$. Since μ is absolutely continuous and a and b are continuous, this implies that $a(\omega^*, t) = b(t)$ for all $t \in [0, T]$, which contradicts $|a(\omega^*, t^*) - b(t^*)| \geq \epsilon$.

Therefore, there exists $\alpha < 1$ such that $\forall \omega \in \Omega \setminus U, h(\omega) \leq \alpha$. Hence, by using hypothesis (i) we have

$$\int_{\Omega \setminus U} p_n(\omega) |a(\omega, t) - b(t)| d\omega \leq \text{Vol}(\Omega) \delta^n \sup_{(\omega, t)} |a(\omega, t) - b(t)|,$$

and hence,

$$\left| \int_{\Omega} p_n(\omega) a(\omega, t) d\omega - b(t) \right| \leq \epsilon + C_1 \delta^n,$$

where C_1 is a constant that does not depend on t and ω . Since ϵ is arbitrary, we deduce that

$$\lim_{n \rightarrow \infty} E_{p_n}[a(\omega, t)] = \lim_{n \rightarrow \infty} \int_{\Omega} p_n(\omega) a(\omega, t) d\omega = b(t)$$

uniformly in $t \in [0, T]$. Note that this argument actually shows the somewhat stronger statement that $E_{p_n}[|a(\omega, t) - b(t)|] \rightarrow 0$ uniformly in t . The same argument shows that $E_{p_n}[|a(\omega, t) - b(t)|^2] \rightarrow 0$ uniformly in $t \in [0, T]$. Hence, taking $\bar{a}(t) = E_{p_n}[a(\omega, t)]$, we have

$$\text{Var}_{p_n}[a(\omega, t)] = E_{p_n}[|a(\omega, t) - b(t)|^2] - |b - \bar{a}(t)|^2,$$


which converges to 0 uniformly in $t \in [0, T]$. \square

7 Numerical Examples

In this section we provide numerical examples to illustrate our theoretical findings and demonstrate the efficacy of our variations on the MINE method to design experiments for dynamics identification.

7.1 A Simple ODE Model

We consider a simple biochemical system that contains three chemicals:



where k_1 and k_2 are the (unknown) degradation rates of A and B , respectively. We also assume that at the beginning, the system only contains A .

We model this system using

$$\begin{aligned} \frac{dA}{dt} &= -k_1 A, & \frac{dB}{dt} &= k_1 A - k_2 B, & \frac{dC}{dt} &= k_2 B, \\ (A(0), B(0), C(0)) &= (1, 0, 0). \end{aligned}$$

In this particular example, we are interested in the dynamics of B . The parameter space is $[0.1, 10] \times [0.1, 10]$, the time interval is $[0, 180]$ (seconds), and the “true” dynamics of the system will correspond to a fixed value ω_0 that is chosen randomly from the uniform distribution on the parameter space.

The experiments are designed sequentially using criteria (1) or (6), depending on the assumptions for a given example. In all cases, at step $n + 1$, the expected dynamics and the corresponding variance function are calculated using the Markov Chain Monte Carlo method. A Markov chain of length 10000 with respect to the invariant measure p_n is sampled on the parameter space using Griddy–Gibbs sampling (Ritter and Tanner 1992). To speed up the sampling process, a sparse grid interpolant (Klimke and Wohlmuth 2005) is used to approximate the model output. At each point of the chain, the corresponding dynamics is evaluated using the polynomial interpolant. The average of these sampled dynamics is computed to approximate the EDE, and the variance is approximated in a similar manner. The interpolant we used in this example has an estimated L^∞ error of order 10^{-4} , which is small in comparison to the experiment error and therefore is negligible. The error of the interpolant is estimated by the difference between the interpolated dynamics and the exact dynamics evaluated using the MatLab solver `ode15s` at 1000 parameter vectors chosen at random from the uniform distribution on the parameter space.

First, we use the framework of Theorem 3.2, in which the data is collected with no noise, the time interval is not discretized, and the experiments are designed using condition (1). The selected sampling times are shown in Fig. 1 (left panel). We see that even without the discretization of the possible sampled time points, the algorithm focuses on two regions in time that are sufficient to capture the system dynamics. This is consistent with the fact that the system is controlled by two parameters. Figure 1 (right panel (i), solid curve) shows how rapidly the EDE approximates the actual




Fig. 1 (Two-dimensional parameter space) *Left*: Measured time points designed by MINE criteria. The algorithm focuses on two regions in time that capture the system dynamics. *Right*: The L^∞ errors of EDE on log-scale in three different cases: (i) Data collected with no noise, using the original MINE criteria (1), (ii) Data with Gaussian noise, using the original MINE criteria (1), (iii) Data with noise, using criterion (6) on a finite set of output values and possible measurement time points

response. After five experiments, the EDE has converged to the true system dynamics within a negligible error.

Next, we consider the case where the data are subject to Gaussian noise with $\sigma^2 = 0.01$. The dashed curve in Fig. 1 (right panel, (ii)) represents the error of the EDE as described in the original algorithm. The dash-dot curve (iii) corresponds to the assumptions of Theorem 4.2. In this case, the experiments are designed sequentially using criteria (6) with $C = 2$, and the time interval is discretized by a uniform grid whose distance between neighbor points is equal to 20. In either case, the algorithm provides a good approximation of the true dynamics after just a few sequential experiments.

7.2 An ODE Model of the T-Cell Signaling Pathway

In this example, we consider a mathematical model of the T-cell signaling pathway proposed by Lipniacki et al. (2008). This is a system of ODEs with 37 state variables, 19 parameters, and fixed initial conditions. We seek to design experiments to identify the dynamics of pZAP, one of the state variables of the system.

In this example, the parameter space is defined relative to a nominal parameter vector. That is, for each component of the nominal vector, we define a range of five times smaller to five times larger than this component. The whole parameter space is the 19-dimensional set formed by the product of these 19 intervals. The time interval is $[0, 201]$ (seconds). The true dynamics of the system are given by a fixed choice of ω_0 that is chosen randomly from the uniform distribution on the parameter space. The expected dynamics and the corresponding variance function are calculated as described in the previous example. To reduce the computational cost, we also construct a sparse grid interpolant to approximate the output of the ODE system. We use a sparse grid with 50,000 points to construct the interpolant. Even so the interpolant has an L^∞ error of order 10^{-2} , so that there is some mismatch in the model.

Fig. 2 (19-dimensional parameter space) Design points in three different cases: (i) Data collected with no noise, using the original MINE criteria (1), (ii) Data with Gaussian noise, using the original MINE criteria (1), (iii) Data with noise, using criterion (6) on a finite set of output values and possible measurement time points




Figure 2 shows the sequence of design points created by the algorithm in three different cases: (i) Data collected with no noise, using the original MINE criteria (1); (ii) Data with Gaussian noise, using the original MINE criteria (1); (iii) Data with noise, using criterion (6) on a finite set of output values and possible measurement time points. In all three cases, the design algorithm focuses on two distinct regions, one of which is precisely defined, the other of which is somewhat nebulous and may perhaps be considered as two regions, particularly in case (iii). This result suggests that although the ODE system is controlled by 19 different parameters, the set of possible system dynamics is contained (at least approximately) in a space of dimension 3. It is worth noting that in Dong et al. (2008), the authors also predicted a pile-up of data points under MINE criteria in the open-loop setting (where multiple measurements are chosen in one step). Our result confirms the same property in the closed-loop case, where the measurements are chosen sequentially with updated probability distributions.

Figure 3 shows the approximation error of the EDE in these three cases. As in the previous example, the error in case (i) decreases quickly to a level consistent with the error in the interpolant and the error in MCMC sampling. This supports the assertion that if we know the exact values of the dynamics at three important points, we are able to recover the whole time course of the dynamics.

Since our algorithms in cases (ii) and (iii) are data-dependent, we present two different realizations of the performance.

In the first case (left panel of Fig. 3), the original algorithm using the MINE criteria with noisy data does not do very well in recovering the dynamics after the first 15 experiments. The problem here is that a measurement with significant noise, especially in the first few steps, can cause the estimator to shift toward a region of the parameter space in which the dynamics do not agree with the true dynamics. Moreover, if the output function at this point of measurement is relatively insensitive to parameter




Fig. 3 (19-dimensional parameter space). The L^∞ errors of EDE on log-scale in three different cases: (i) Data collected with no noise, using the original MINE criteria (1), (ii) Data with Gaussian noise, using the original MINE criteria (1), (iii) Data with noise, using criterion (6) on a finite set of output values and possible measurement time points

changes in this region, it may take many additional measurements to overcome this initial misestimation.

In our example, we encounter this issue: the second measurement made at $t = 201$ gives a data value of nearly 1 when the true value is approximately 0.75. This measurement shifts the probability distribution toward a broad region in the parameter space where the corresponding dynamics saturate to the maximum value 1. This reduces the system variance at time 201 to a relatively small value in comparison to that of other time points. A direct consequence is that in the next eight experiments, no measurement is made around time 200, and the EDE's error does not improve. However, during this process, the parameters that correspond to the true dynamics gain weight, causing the variance around time 200 to increase. Finally, a measurement at time 201 is made in step 11, which significantly decreases the error of the expected dynamics estimator.

This example illustrates the fact that although the convergence of the original algorithm is guaranteed, the convergence may be slow. Some drawbacks of the original algorithm are removed in case (iii) by replacing criteria (1) by criteria (6) and by restricting the set of possible time points to be finite. By making the set of possible measurement points finite, we collapse the important regions in the time interval to single points and facilitate resampling to get more accurate data at these important points. Also, by using criteria (6), we obtain the freedom to select the next measurement point subject to multiple criteria, as described next.

For Fig. 3, as in the previous example, the set of all possible measurement time points is restricted to a uniform grid of resolution 20, starting from 1. To design experiments, we used the following ranking: among time points with relatively high variance (specifically, that have variance larger than half of the maximum), the time points that have already been measured are given more priority (to promote resampling); among time points that have been measured, the points with fewer measurements have more priority; among time points that have the same number of measurements, the ones with higher variance have more priority.




Fig. 4 (19-dimensional parameter space). *Left*: EDEs using different sparse grid interpolators to approximate the dynamics. The EDEs are evaluated after 10 steps. *Right*: Expected dynamics estimator and predicted confidence intervals of the output dynamics with $\epsilon = 0.05$

The advantages of this variation of the algorithm are illustrated in Fig. 3 (left panel). After two measurements that coincide with those of the previous case (including the point with large measurement error), the algorithm then selects a different measurement point that leads the error to drop quickly. After six experiments, the expected dynamics estimator has converged to the true dynamics within an acceptable error.

On the right panel of Fig. 3, we consider a different realization of data on the first measurement. In this case, the random data is obtained with small error and leads to quick convergence of the EDEs corresponding to both criteria.

This example also illustrates the fact that the probabilistic framework in experimental design works well in the case where the number of data is less than the number of parameters, or when the model is unidentifiable: our examples could not have been done using a method of parameter estimation via optimization. Assume that in example 2, we can make measurements with high accuracy at three time points 50, 100, and 200 and want to know the value at time 150. The number of data points in this case is less than the number of parameters, and any method that returns a single parameter estimate will never be able to predict with high confidence (or any confidence at all) the output value at 150. In order to do so, it needs to compute every possible parameter values that fit the data, which is very unlikely in practice. Our probabilistic framework provides a feasible way to address the issue: we considered such an example in Fig. 4, in which we quantify the uncertainty of the dynamics with only 10 noisy measurements ($\sigma = 0.1$) that accumulate at three time points (see also Fig. 2), which is much less than the number of model parameters (19).

Model Mismatch Finally, we illustrate the effect of model mismatch on the convergence of the EDE by using different sparse grid interpolants to approximate the system output. In this particular example, we run the algorithm with the relaxed MINE criteria on a finite set of measured time points and output values with three different sparse grids of 1000, 2000, 9000 grid points, respectively. The estimated L^∞ -errors of the three interpolants are 0.2, 0.1, and 0.05. As in the previous example, the errors

of the interpolants are estimated by the difference between the interpolated dynamics and the exact dynamics evaluated using the MatLab solver `ode15s` at 1000 parameter vectors chosen uniformly at random from the parameter space. These interpolants are considered as approximate models with varying degrees of mismatch. In each case, the EDE is evaluated after 10 points selected according to criteria (6) with $C = 2$.

The results of this example are given in Fig. 3. It is not surprising that all three cases give good estimates of the true dynamics: since we are not concerned with the identification of parameters, as long as the dynamics space of the approximate models are close to the dynamics space of the true model, the algorithm will work well. Although the sparse grid interpolant with 1000 grid points is not a good approximation of the system output, it has enough degrees of freedom to capture the behavior of the system so that a weighted average over parameter space gives a good estimation of the true dynamics.

8 Conclusion

Building upon the Maximally Informative Next Experiment algorithm, we have developed several variants of a model-based experiment design algorithm. This algorithm uses existing data to produce a probability distribution on parameter space and then identifies possible measurement points whose output values have large variance under this distribution. We have also proven the convergence of the associated EDE (expected dynamics estimator) to the true system dynamics under a variety of assumptions on the model and data, even when the chosen experiments cluster in a small finite set of points. This approach provides an effective way to incorporate the knowledge arising from nonlinear models into the experiment design process. We illustrated our results with numerical examples on various models of cellular processes.

There are several avenues for future work. First, in Dong et al. (2008), the authors proposed several MINE criteria for experimental design. In this work, we establish the theoretical foundations for one of them. The next step would be validating other MINE criteria within a more general model setting.

Second, in this paper, we focused on the identification of observable outputs and did not attempt to address the extrapolation problem, in which measurements of one output are used to make inference about an unobservable output. However, it is worth noting that our framework can be naturally extended to identify unobservable outputs that are theoretically identifiable given that all information about possible observable outputs is known. The problem of determining which unobservable outputs are identifiable in a given experimental setting will be addressed in one of our independent but related works.

Acknowledgements All authors partially supported by NSF grant DMS-0900277.

References

- Bazil, J. N., Buzzard, G. T., & Rundell, A. E. (2011). A global parallel model based design of experiments method to minimize model output uncertainty. *Bull. Math. Biol.*, 74, 688–716.

- Cox, D., & Reid, N. (2000). *The theory of the design of experiments. Monographs on statistics and applied probability*. London/Boca Raton: Chapman & Hall/CRC Press.
- Donahue, M. M., Buzzard, G. T., & Rundell, A. E. (2010). Experiment design through dynamical characterisation of non-linear systems biology models utilising sparse grids. *IET Syst. Biol.*, 4, 249–262.
- Dong, W., Tang, X., Yu, Y., Nilsen, R., Kim, R., Griffith, J., Arnold, J., & Schuttler, H. (2008). Systems biology of the clock in *Neurospora crassa*. *PLoS ONE*, 3(8), e3105.
- Fedorov, V. V., & Hackl, P. (1997). *Model-oriented design of experiments*. New York: Springer.
- Kitano, H. (2002). Systems biology: a brief overview. *Science*, 295(5560), 1662–1664.
- Klimke, A., & Wohlmuth, B. (2005). Algorithm 847: Spinterp: piecewise multilinear hierarchical sparse grid interpolation in MATLAB. *ACM Trans. Math. Softw.*, 31(4), 561–579.
- Kreutz, C., & Timmer, J. (2009). Systems biology: experimental design. *FEBS J.*, 276(4), 923–942.
- Lipniacki, T., Hat, B., Faeder, J. R., & Hlavacek, W. S. (2008). Stochastic effects and bistability in T-cell receptor signaling. *J. Theor. Biol.*, 254(1), 110–122.
- Morokoff, W. J., & Caflisch, R. E. (1994). Quasi-random sequences and their discrepancies. *SIAM J. Sci. Comput.*, 15, 1251–1279.
- Myers, R. H., & Montgomery, D. C. (2002). *Response surface methodology: process and product optimization using designed experiments*. New York: Wiley.
- Niederreiter, H. (1978). Quasi-Monte Carlo methods and pseudo-random numbers. *Bull. Am. Math. Soc.*, 84(6), 957–1041.
- Pronzato, L. (2009). Asymptotic properties of nonlinear estimates in stochastic models with finite design space. *Stat. Probab. Lett.*, 79(21), 2307–2313.
- Pronzato, L. (2010). One-step ahead adaptive D-optimal design on a finite design space is asymptotically optimal. *Metrika*, 71(2), 219–238.
- Pukelsheim, F. (1993). *Optimal design of experiments*. New York: Wiley.
- Ritter, C., & Tanner, M. A. (1992). Facilitating the Gibbs sampler: the Gibbs stopper and the Griddy-Gibbs sampler. *J. Am. Stat. Assoc.*, 87(419), 861–868.
- Sontag, E. D. (2002). For differential equations with r parameters, $2r + 1$ experiments are enough for identification. *J. Nonlinear Sci.*, 12, 553–583.
- White, L. V. (1973). An extension of the general equivalence theorem to nonlinear models. *Biometrika*, 60(2), 345–348.