



Prediction uncertainty and optimal experimental design for learning dynamical systems

Benjamin Letham, Portia A. Letham, Cynthia Rudin, and Edward P. Browne

Citation: *Chaos* **26**, 063110 (2016); doi: 10.1063/1.4953795

View online: <http://dx.doi.org/10.1063/1.4953795>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/chaos/26/6?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Identifiability and observability analysis for experimental design in nonlinear dynamical models](#)

Chaos **20**, 045105 (2010); 10.1063/1.3528102

[A Dynamical Approach for Modelling and Control of Production Systems](#)

AIP Conf. Proc. **622**, 199 (2002); 10.1063/1.1487535

[Anticipatory computing in the attitude control of satellites](#)

AIP Conf. Proc. **573**, 257 (2001); 10.1063/1.1388694

[Estimation theory and model parameter selection for therapeutic treatment plan optimization](#)

Med. Phys. **26**, 2348 (1999); 10.1118/1.598749

[Nonlinear prediction, filtering, and control of chemical systems from time series](#)

Chaos **7**, 614 (1997); 10.1063/1.166261



Prediction uncertainty and optimal experimental design for learning dynamical systems

Benjamin Letham,^{1,a)} Portia A. Letham,^{2,b)} Cynthia Rudin,^{3,c)} and Edward P. Browne^{4,d)}

¹Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

²Department of Chemical Engineering, Arizona State University, Tempe, Arizona 85281, USA

³Department of Computer Science and Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina 27708, USA

⁴Koch Institute for Integrative Cancer Research, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

(Received 8 February 2016; accepted 31 May 2016; published online 14 June 2016)

Dynamical systems are frequently used to model biological systems. When these models are fit to data, it is necessary to ascertain the uncertainty in the model fit. Here, we present prediction deviation, a metric of uncertainty that determines the extent to which observed data have constrained the model's predictions. This is accomplished by solving an optimization problem that searches for a pair of models that each provides a good fit for the observed data, yet has maximally different predictions. We develop a method for estimating *a priori* the impact that additional experiments would have on the prediction deviation, allowing the experimenter to design a set of experiments that would most reduce uncertainty. We use prediction deviation to assess uncertainty in a model of interferon-alpha inhibition of viral infection, and to select a sequence of experiments that reduces this uncertainty. Finally, we prove a theoretical result which shows that prediction deviation provides bounds on the trajectories of the underlying true model. These results show that prediction deviation is a meaningful metric of uncertainty that can be used for optimal experimental design. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4953795>]

Nonlinear dynamical systems are used throughout systems biology to describe the dynamics of biomolecular interactions. These models typically have a number of unknown parameters, such as infection rates and decay rates, which are estimated by fitting the model to measurements from the physical system. Two important questions then arise: What is the uncertainty in the model predictions, and how can that uncertainty be reduced? We describe here a new approach for measuring uncertainty in model predictions, by searching for a pair of model parameters that both provide a good fit for the observed data, but make maximally different predictions. We further show how to estimate the impact on the uncertainty of a candidate experiment that has not yet been done, allowing the experimenter to determine beforehand if an experiment will be valuable. We use prediction deviation to analyze a model of HIV infection which can only be partially observed. With prediction deviation, and with appropriately selected experiments, we are able to provide bounds on the behavior of the unobserved quantities and gain insights into inhibition that are otherwise unavailable.

I. INTRODUCTION

Systems of nonlinear differential equations are used throughout biology to model the behavior of complex, dynamical systems. These models have proven particularly useful in systems biology for describing networks of biomolecular interactions.¹ Often the utility of the model depends on being able to estimate a set of unknown parameters, which is typically done by collecting data from the physical system and finding the best-fit parameters. When inferring a dynamical system from data, there are two important questions that arise:

- (1) *Uncertainty quantification*: Is the model sufficiently constrained by the data?
- (2) *Optimal experimental design*: If not, what additional experiments would most reduce the remaining uncertainty?

Uncertainty is often measured by constructing a confidence interval for each parameter estimate. We propose a different approach to the problem of uncertainty quantification and then show that this approach leads naturally to an optimal experimental design strategy. Our fundamental hypothesis is that the purpose of fitting a model is to be able to use it to make predictions. In many situations, the parameter values *per se* are not of interest, rather the goal is to gain insight into the system's behavior. In these situations, the purpose of assessing model uncertainty is to determine if the model's predictions can be trusted.

This paper begins by developing *prediction deviation*, a new measure of uncertainty on predicted behaviors. We then use prediction deviation to measure uncertainty in a partially observed model of human immunodeficiency virus

^{a)}Electronic mail: bletham@alum.mit.edu. Present address: Facebook, Menlo Park, CA 94025, USA.

^{b)}Electronic mail: pletham@asu.edu

^{c)}Electronic mail: cynthia@cs.duke.edu

^{d)}Electronic mail: ebrowne@broadinstitute.org

(HIV) infection, where we found that after one experiment there remained substantial uncertainty in the behavior of the unobserved component. We then show that prediction deviation leads naturally to a way to measure *experiment impact*, which is a maximum uncertainty on predicted behaviors if an additional experiment were to be conducted. This approach is used to determine a sequence of experiments that reduces uncertainty in the HIV infection model and ultimately bounds the behavior of the unobserved component. Finally, we provide a theoretical foundation for prediction deviation by showing that, under reasonable assumptions, it bounds the trajectory of the underlying true model.

A. Confidence intervals do not measure predictive power

Parameter confidence intervals are a poor way of determining if a nonlinear dynamical system’s predictions are constrained by the observed data. Sensitive dependence means that tight confidence intervals do not imply constrained predictions. The classic Lorenz system provides an illustration of this phenomenon

$$\frac{dx}{dt} = \theta_1(y - x), \quad \frac{dy}{dt} = x(\theta_2 - z) - y, \quad \frac{dz}{dt} = xy - \theta_3 z.$$

Fig. 1(a) shows $x(t)$ data generated from the Lorenz system with parameters $\theta^{\text{true}} = [7, 38, 5]$, initial conditions $x(0) = 10$, $y(0) = 20$, and $z(0) = 3$, and a small amount of normally distributed noise. The best-fit estimates for the parameters, θ^* , are very close to the true values θ^{true} , and have seemingly tight confidence intervals: $\theta_1^* = 7.00$ (6.51 – 7.49), $\theta_2^* = 38.03$ (36.08 – 40.28), and $\theta_3^* = 5.00$ (4.82 – 5.17), with 95% simultaneous likelihood-based intervals in parentheses.² Suppose we wished to use these observed data with $y(0) = 20$ to predict the behavior of the system when $y(0) = 7$, with all other factors staying constant. Do the tight confidence intervals allow for confidence in the model’s predictions at this different initial condition? Fig. 1(b) shows that they do

not. This figure compares the predictions made by the best-fit model θ^* to those made by the model with parameters $\bar{\theta} = [6.98, 38.12, 4.99]$. $\bar{\theta}$ is well within the confidence intervals of θ^* , moreover, the fit error of $\bar{\theta}$ is within the 95% confidence interval for the fit error of θ^* , meaning $\bar{\theta}$ is also a good fit for the observed data. However, $\bar{\theta}$ and θ^* make entirely different predictions for the condition we wish to predict. The phase portraits in Figs. 1(c) and 1(d) show that this small change in the parameters is enough to send the trajectory to a different side of the attractor. The Lorenz system is a canonical example of sensitivity, but chaotic dynamics are not required to have tight confidence intervals with unconstrained predictions. For instance, this same result can be had any time a basin boundary lies within the confidence interval.

Tight confidence intervals do not imply constrained predictions, and likewise wide confidence intervals do not imply unconstrained predictions. Parameters in nonlinear dynamical systems may be interrelated such that they individually have large confidence intervals, yet the predictions of interest are actually constrained. The following parameterization of the Lotka-Volterra predator-prey model illustrates this fact:

$$\frac{dx}{dt} = \theta_1 \theta_3 x - \theta_2 \theta_3 xy, \quad \frac{dy}{dt} = \theta_2 \theta_4 xy - \theta_1 \theta_4 y.$$

A symmetry in the parameters renders them all unidentifiable—they have infinite confidence intervals. Suppose we were able to observe $x(t)$ data and wished to use these data to predict the state $y(t)$. Fig. 2 shows that despite the infinite confidence intervals, $x(t)$ data constrain predictions of $y(t)$. The data in Fig. 2 were generated using $\theta^{\text{true}} = [1, 0.05, 1, 1]$ and initial conditions $x(0) = y(0) = 10$, with standard normal noise. $\bar{\theta}$ is the worst-case of how bad the prediction in $y(t)$ could be. Specifically, of all parameters that have fit error within the 95% confidence interval of the fit error of the best-fit (that is, all parameters that provide a good fit of the data), $\bar{\theta}$ is the one that maximized the squared difference between its prediction of $y(t)$

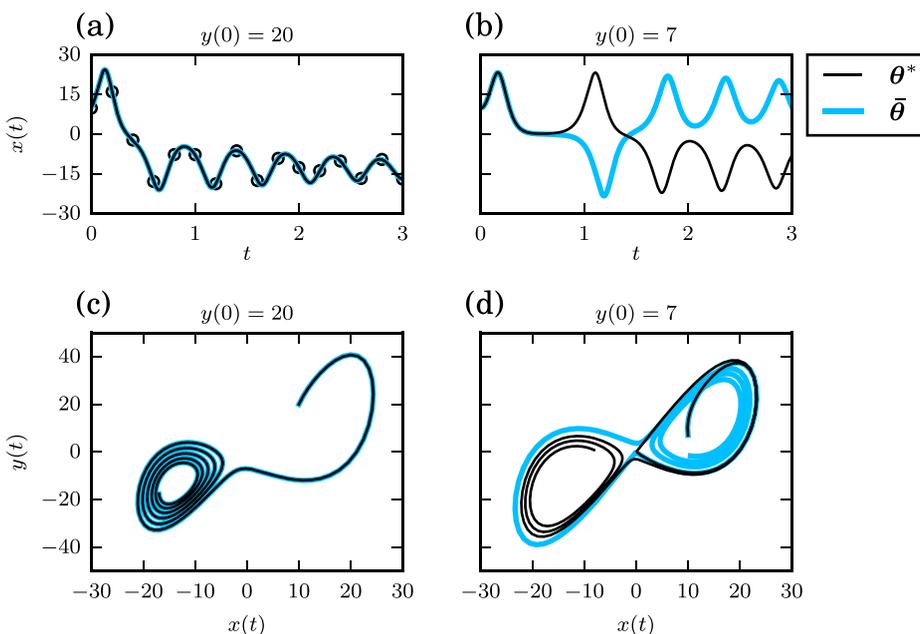


FIG. 1. (a) Circles indicate simulated data points from the Lorenz system, with the best-fit in black and the alternative model in blue. (b) Despite both models providing a good fit to the data in panel (a), they produce very different predictions on a different initial condition. (c) and (d) Phase portraits for trajectories in panels (a) and (b).

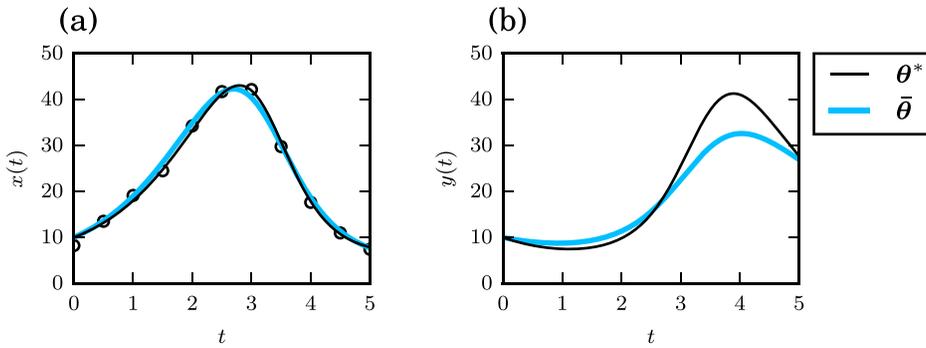


FIG. 2. (a) Circles indicate simulated data points from the Lotka–Volterra model, with the best-fit in black and in blue the model that maximized the difference in panel (b), subject to providing a good fit to these simulated data points. (b) Predictions from the best-fit and alternative models of the state $y(t)$ are constrained, despite unidentifiable parameters.

and that of the best-fit θ^* . Thus any model that fits the $x(t)$ observations will make a prediction on $y(t)$ that differs from the best-fit by no more than the difference seen in $\bar{\theta}$.

This model contains a structural unidentifiability, which could be identified and corrected by a reparameterization.^{3–5} It is also possible to have model parameters that are structurally identifiable but not *practically* identifiable, given the noise in the collected data.² Gutenkunst *et al.*⁶ show that models with parameters that cannot be well constrained by data are ubiquitous in systems biology and conclude that “modelers should focus on predictions rather than on parameters.”

If the purpose of fitting the model to data is to ascertain the values of the parameters, then confidence intervals provide a useful quantification of uncertainty. However, if the purpose is to use the fitted model to make predictions about unobserved variables or unobserved conditions, then confidence intervals serve no purpose for dynamical systems. We propose putting aside the issue of measuring confidence intervals and instead directly measure the uncertainty in the quantity of interest: the predictions.

II. PREDICTION DEVIATION AS A MEASURE OF UNCERTAINTY

Figs. 1(b) and 2(b) provide the motivation for our approach to measuring uncertainty. We consider a scenario, or set of scenarios, for which we are interested in predicting the system behavior. In Fig. 1(b) this was a different initial condition and in Fig. 2(b) it was an unobserved variable. We then pose the following question: Of all parameters that are a good fit to the observed data, what is the largest deviation in predicted behaviors for any pair? We call this deviation the *prediction deviation*. A low prediction deviation, such as that in Fig. 2(b), means that the observed data have constrained the prediction of interest. A high prediction deviation, such as that in Fig. 1(b), means that the observed data have not constrained the prediction of interest.

A. Parameter estimation

We must first introduce notation to make the idea of prediction deviation precise. We suppose that we are learning a system of ordinary differential equations with state variables $\mathbf{x}(t)$, unknown parameters θ , and known external factors \mathbf{v}

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t; \theta, \mathbf{v}). \tag{1}$$

If the initial conditions are known then they are included in \mathbf{v} , and if unknown in θ .

We now provide notation for the observed data and the data fitting problem. Let $\mathcal{P}_j = (I_j, T_j, \mathbf{v}^j)$ represent a particular *experiment*, with I_j being the set of state variables that are observed, $T_j = \{T_{i,j} : i \in I_j\}$ the sets of time points at which these observations are made for each state variable, and \mathbf{v}^j the external factors. We suppose that a total of J experiments have been performed, resulting in observed data $\tilde{x}_i^j(t)$, for $j = 1, \dots, J, i \in I_j$, and $t \in T_{i,j}$. We denote the complete set of observed experiments as $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_J\}$ and the complete set of observed data as $\tilde{\mathbf{x}}$.

The unknown parameters θ are typically estimated from the observed data $\tilde{\mathbf{x}}$ by minimizing the weighted squared error

$$z_{\text{fit}}(\theta; \mathcal{P}, \tilde{\mathbf{x}}) := \sum_{j=1}^J \sum_{i \in I_j} \sum_{t \in T_{i,j}} \left(\frac{x_i(t; \theta, \mathbf{v}^j) - \tilde{x}_i^j(t)}{\sigma_{ijt}} \right)^2, \tag{2}$$

where $x_i(t; \theta, \mathbf{v}^j)$ is obtained by integrating (1) and σ_{ijt}^2 is the noise variance. The best-fit parameters θ^* are the solution to the least squares problem

$$\underset{\theta}{\text{minimize}} \quad z_{\text{fit}}(\theta; \mathcal{P}, \tilde{\mathbf{x}}). \tag{3}$$

B. Prediction deviation

To measure prediction deviation, we wish to search over the set of all models that provide a good fit to the observed data. We consider a model θ to be a “good fit” to the observed data if its fit error $z_{\text{fit}}(\theta; \mathcal{P}, \tilde{\mathbf{x}})$ is not too much worse than that of the best fit, θ^* . Specifically, we measure a 95% confidence interval for $z_{\text{fit}}(\theta^*; \mathcal{P}, \tilde{\mathbf{x}})$, which we denote as $[z_l^*, z_u^*]$. In the event of normally distributed observation noise, a parametric estimate for the interval can be obtained using the χ^2 distribution, or, as we do here, a nonparametric confidence interval can be obtained with the bootstrap.⁷ The prediction deviation is defined as the maximum difference on a prediction problem between any pair of models that both have fit error within the 95% confidence interval of the best-fit error.

The prediction problems for which the prediction deviation is to be measured are defined in the same way as the experiment according to which data were collected. We call $\mathcal{Y}_\ell = (I_\ell, T_\ell, \mathbf{v}^\ell)$ a *prediction problem*, where as before I_ℓ is the set of state variables to be predicted in problem ℓ , $T_\ell = \{T_{i,\ell} : i \in I_\ell\}$ are the sets of time points at which

these predictions are made for each state variable, and \mathbf{v}^ℓ are the external factors. Let $\mathcal{Y} = \{\mathcal{Y}_1, \dots, \mathcal{Y}_L\}$ be the full collection of variables and experiments of interest for prediction. The squared difference between θ^1 and θ^2 on the prediction problems is

$$z_{\text{dev}}(\theta^1, \theta^2; \mathcal{Y}) := \sum_{\ell=1}^L \sum_{i \in I_\ell} \sum_{t \in T_{i,\ell}} \left(\frac{x_i(t; \theta^1, \mathbf{v}^\ell) - x_i(t; \theta^2, \mathbf{v}^\ell)}{\sigma_{i\ell t}} \right)^2. \quad (4)$$

As before, $\sigma_{i\ell t}^2$ is an estimate of the noise level for that measurement, which is important primarily for combining multiple measurements of possibly different scales into one metric. Prediction deviation can now be framed as an optimization problem

$$\text{maximize}_{\theta^1, \theta^2} z_{\text{dev}}(\theta^1, \theta^2; \mathcal{Y}) \quad (5a)$$

$$\text{subject to } z_{\text{fit}}(\theta^1; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_{it}^* \quad (5b)$$

$$z_{\text{fit}}(\theta^2; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_{it}^* \quad (5c)$$

The objective (5a) searches for a pair of models that maximize the difference in predictions on the prediction problems, while the constraints (5b) and (5c) limit the search to those models that provide a good explanation for the observed data $\tilde{\mathbf{x}}$. Let $\bar{\theta}^1$ and $\bar{\theta}^2$ be the maximizers of problem (5). Then, the optimal objective value $z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y})$ is the prediction deviation and can be obtained by solving this single, constrained maximization problem. We show in the supplementary material⁸ results for the Lorenz system from Fig. 1 and now discuss how prediction deviation can be used to understand and reduce uncertainty in a viral infection model.

III. UNCERTAINTY IN A MODEL OF HIV INFECTION

A. The model and data

We now use prediction deviation to assess how well observed data constrain the model predictions of an unobserved component in a model of the innate immune response to HIV infection.⁹ The model describes the dynamics of how interferon-alpha (IFN α) protects CD4 T cells from infection by HIV. IFN α is a signaling protein that endows CD4 T cells with protection from HIV by upregulating genes that disrupt viral replication. In the model, CD4 T cells (C) are infected by HIV (H) and become infected cells (CH) that produce additional viruses. Exposure to IFN α (I) induces a refractory state in both uninfected and infected cells (CI and CHI, respectively) which if uninfected are protected from infection, and if infected no longer produce additional viruses. The refractory state is reversible and CI and CHI cells eventually revert to their original state, C and CH, respectively. The dynamical system that describes the interactions of these quantities is

$$\frac{dC(t)}{dt} = \theta_1 C(t) + \theta_3 CI(t) - \theta_2 C(t) \frac{I(t)}{\theta_8 + I(t)} - \theta_5 C(t) H(t),$$

$$\frac{dCI(t)}{dt} = (\theta_1 - \theta_3) CI(t) + \theta_2 C(t) \frac{I(t)}{\theta_8 + I(t)},$$

$$\frac{dCH(t)}{dt} = (\theta_1 - \theta_4) CH(t) + \theta_5 C(t) H(t) - \theta_2 CH(t) \frac{I(t)}{\theta_8 + I(t)} + \theta_3 CHI(t),$$

$$\frac{dCHI(t)}{dt} = (\theta_1 - \theta_3 - \theta_4) CHI(t) + \theta_2 CH(t) \frac{I(t)}{\theta_8 + I(t)},$$

$$\frac{dH(t)}{dt} = \theta_6 CH(t) - \theta_7 H(t).$$

We use here tissue culture data collected by Browne *et al.*,⁹ who provide full details of the experimental methodology. In short, varying levels of IFN α were added to tissue cultures with CD4 T cells. After allowing the cells to incubate with the IFN α for 6 h, HIV was added to the culture for 1 h and then washed out. The total number of uninfected (C + CI) and infected (CH + CHI) cells, along with the viral count (H) were measured with four replicates every 24 h, for 3 days. This experiment was done separately for a total of 7 initial IFN α levels: 0, 0.002, 0.02, 0.2, 2, 20, and 200 ng/ml. In this tissue culture, the IFN α activity remained constant and so $I(t)$ was a known, external factor. Details of model fitting and prediction deviation implementation are given in Appendix A.

B. Prediction deviation after one experiment

To illustrate how prediction deviation changes with additional experiments and how it can be used for experiment selection, we label each combination of variables and IFN α level as a separate experiment. For example, C + CI measured at $I = 0.002$ ng/ml defines one experiment, and H measured at $I = 2.0$ ng/ml is another. There are a total of 21 such experiments for which data were collected. We begin by using data from only one of these experiments, and then consider the problems of determining uncertainty in model fit, and deciding which additional experiments to add in order to reduce prediction uncertainty.

The purpose of defining the model and collecting experimental data is to understand the dynamics of how IFN α provides protection to CD4 T cells during HIV infection. The experimental data themselves do not explicitly show the interaction of IFN α and CD4 T cells inasmuch as only the sum C + CI can be observed. The natural prediction problem is to then try to predict the CI timecourse, at the same observation times as the C + CI data. We begin with just one experiment, and let \mathcal{P} be the experiment corresponding to C + CI measured at $I = 0.002$ ng/ml. Let \mathcal{Y} be the corresponding prediction problem, CI at $I = 0.002$ ng/ml.

Using prediction deviation, we can determine if the observations of C + CI at $I = 0.002$ ng/ml constrain the predictions of CI at $I = 0.002$ ng/ml, and Fig. 3 shows that they do not. In particular, Fig. 3(a) shows that the two models that maximize prediction deviation both provide a good fit for the observed data, while Fig. 3(b) shows that they provide widely diverging predictions about the CI timecourse: One of the models suggests that nearly all of the CD4 T cells are refractory, while the other one suggests that nearly none of them are. These observed data do not in any way increase our understanding of the IFN α dynamics.

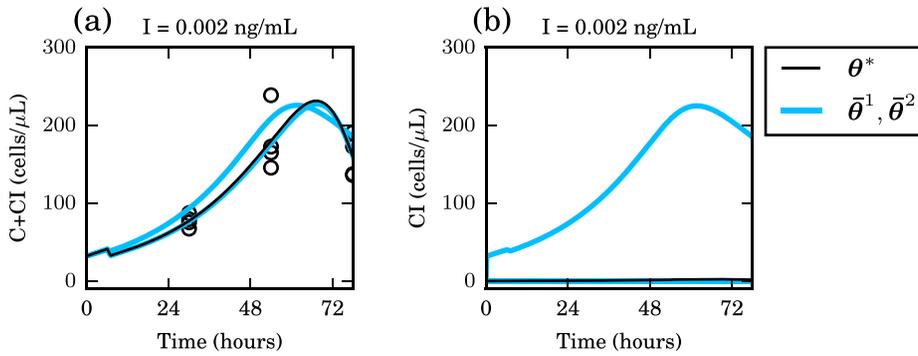


FIG. 3. (a) Circles indicate observed data for total uninfected CD4 T cells. In black is the best-fit model, and in blue are the two prediction deviation models, which also provide a good fit to the data. (b) The prediction deviation models provide widely differing predictions about the number of uninfected cells that are refractory, ranging from nearly none to nearly all.

IV. OPTIMAL EXPERIMENT DESIGN

Knowing that the predictions of CI are entirely unconstrained, the question that naturally follows is to determine which of the remaining 20 experiments should be done next in order to maximally reduce the prediction deviation. More generally, we wish to predict the impact that a particular candidate experiment or set of experiments \mathcal{P}' will have on the prediction deviation, given that we have already completed experiments \mathcal{P} , with $\mathcal{P}' \cap \mathcal{P} = \emptyset$.

Fig. 4 provides some insight into this problem. This figure shows the predictions of the best-fit and prediction deviation models from Fig. 3 on two of the candidate experiments, C + CI at I levels of 0.0 and 200.0 ng/ml. On the candidate experiment in Fig. 4(a), the prediction deviation models are very different. Suppose observations were collected for this experiment and then prediction deviation were recomputed using both these observations and the original set. After collecting data, at least one of the prediction deviation models in Fig. 4(a) would no longer be a good fit. We cannot know *a priori* if the observations will lie close to one of the models and thereby disqualify the other, or if they will lie in the middle, disqualifying both, but at least one model will not be a good fit for the new observations.

Fig. 4(b) shows the alternative situation where the prediction deviation models do not disagree on the candidate experiment. Were this experiment to be done, it is possible that the observations would disqualify both prediction deviation models and there would be a reduction in uncertainty. However, it is also possible for the observations to be such that both models remain feasible, meaning there is no reduction in uncertainty.

The experiment in Fig. 4(a) seems like a good choice for reducing uncertainty, however having a large deviation on the candidate experiment \mathcal{P}' does not necessarily mean that

the deviation on the prediction problem of interest, in this case CI at $I = 0.002$ ng/ml, will actually be reduced. Certainly that pair of prediction deviation models will no longer satisfy both constraints (5b) and (5c), however, there may exist yet another pair of models that do not disagree on \mathcal{P}' but produce the same prediction deviation on \mathcal{Y} . A powerful property of prediction deviation as a measure of uncertainty is that we actually can determine if this is the case.

A. Estimating experiment impact

Collecting observations from experiment \mathcal{P}' would change the prediction deviation by requiring the prediction deviation models to be a good fit for the new observations. In essence, there would be two new constraints that must be satisfied

$$z_{\text{fit}}(\bar{\theta}^1; \mathcal{P}', \tilde{\mathbf{x}}') \leq \eta \quad \text{and} \\ z_{\text{fit}}(\bar{\theta}^2; \mathcal{P}', \tilde{\mathbf{x}}') \leq \eta,$$

for some η , where $\tilde{\mathbf{x}}'$ are the new observations. In Appendix A, we show that these constraints imply

$$z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{P}') \leq 2\eta, \tag{6}$$

which allows us to get some idea of the impact these constraints would have even without knowledge of $\tilde{\mathbf{x}}'$. The essence of this result is that if the prediction deviation models are a good fit for the new data, they must have close trajectories on the new data. We are unable to restrict the prediction deviation models to be a good fit for the new data until we have collected the new data. We can, however, restrict the prediction deviation models to have close trajectories on the candidate experiment, thus estimating the

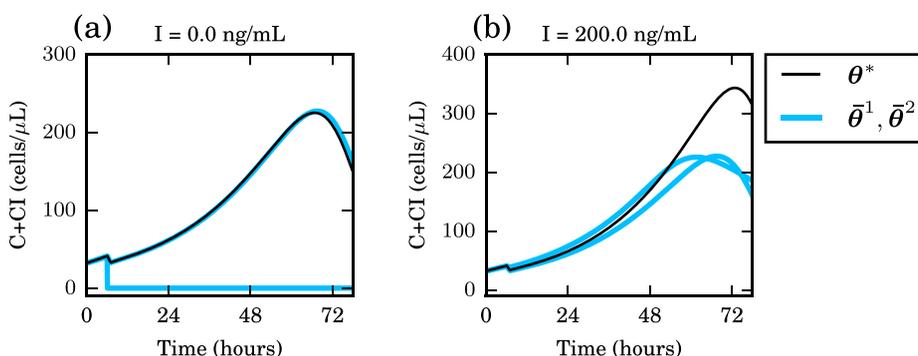


FIG. 4. Trajectories from the same best-fit (black) and prediction deviation (blue) models as Fig. 3, for two candidate experiments. (a) Observations from this experiment would disqualify at least one of the prediction deviation models. (b) Both prediction deviation models might remain feasible after this experiment.

impact that the candidate experiment would have on the prediction deviation. This is done by solving the prediction deviation problem with the added constraint (6), which we call the *experiment impact problem*

$$\text{maximize}_{\theta^1, \theta^2} z_{\text{dev}}(\theta^1, \theta^2; \mathcal{Y}), \tag{7a}$$

$$\text{subject to } z_{\text{fit}}(\theta^1; \mathcal{P}, \bar{\mathbf{x}}) \leq z_u^*, \tag{7b}$$

$$z_{\text{fit}}(\theta^2; \mathcal{P}, \bar{\mathbf{x}}) \leq z_u^*, \tag{7c}$$

$$z_{\text{dev}}(\theta^1, \theta^2; \mathcal{P}') \leq \eta. \tag{7d}$$

This problem is identical to problem (5) used to find the prediction deviation, with the added constraint (7d). Model pairs like that in Fig. 3(a) will not be feasible solutions to this problem, inasmuch as they violate (7d). If there does exist a different pair that produces close trajectories on \mathcal{P}' but still has a large deviation on \mathcal{Y} , this optimization problem will find that pair. We denote the solutions to this optimization problem as $\hat{\theta}^1$ and $\hat{\theta}^2$, and call the optimal objective value $z_{\text{dev}}(\hat{\theta}^1, \hat{\theta}^2; \mathcal{Y})$ the *estimated experiment impact*.

Of all possible outcomes of \mathcal{P}' , the outcome that reduces uncertainty in \mathcal{Y} the least is if the observations follow the trajectories of $\hat{\theta}^1$ and $\hat{\theta}^2$. In this sense, the predicted experiment impact is a worst-case reduction of uncertainty, and we can expect that \mathcal{P}' will reduce the prediction deviation at least as much as $z_{\text{dev}}(\hat{\theta}^1, \hat{\theta}^2; \mathcal{Y})$, subject to the closeness requirement η being appropriate. Appendix A describes how η can be chosen.

V. REDUCING UNCERTAINTY OF IFN α DYNAMICS

We now continue the results on uncertainty in IFN α dynamics and use the predicted experiment impact to find additional experiments that reduce the uncertainty shown in

Fig. 3(b). There are 20 candidate experiments consisting of different component measurements and varying IFN α levels. The estimated experiment impact optimization problem, (7), was solved for each of these candidate experiments, and results for the experiment that predicted the largest reduction of uncertainty, C + CI at I = 0.0 ng/ml, are shown in Fig. 5. Fig. 5(a) shows the predicted experiment impact models on the candidate experiment, which are forced to have close trajectories. Fig. 5(b) shows that for the prediction problem there is a substantial reduction in uncertainty by requiring the models to produce close trajectories on the candidate experiment—this is the estimated experiment impact. The actual experiment impact is shown in Figs. 5(c) and 5(d): in Fig. 5(c) the actual observations, and in Fig. 5(d) the prediction deviation after including those observations. The actual reduction in prediction deviation was very close to that predicted by the estimated experiment impact in Fig. 5(b).

The estimated experiment impact problem predicted a significant reduction in uncertainty from only one of the 20 candidate experiments. Fig. 6 shows the results of separately adding each of the 20 candidates, and the C + CI at I = 0.0 ng/ml experiment of Fig. 5 provided by far the largest reduction of uncertainty. The other two experiments at I = 0.0 ng/ml provided a moderate reduction in uncertainty, while the remaining 17 experiments provided no reduction in uncertainty, by increasing the amount of noise in the fitting.

We denote the prediction deviation models after including data from experiment \mathcal{P}' as $\bar{\theta}^{1'}$ and $\bar{\theta}^{2'}$. Fig. 7(a) compares the estimated experiment impact, $z_{\text{dev}}(\hat{\theta}^1, \hat{\theta}^2; \mathcal{Y})$, to the actual impact of each candidate experiment, $z_{\text{dev}}(\bar{\theta}^{1'}, \bar{\theta}^{2'}; \mathcal{Y})$. As already seen in Fig. 5, for the one candidate that in actuality significantly reduced uncertainty, the predicted impact was very close to the actual impact. There were two

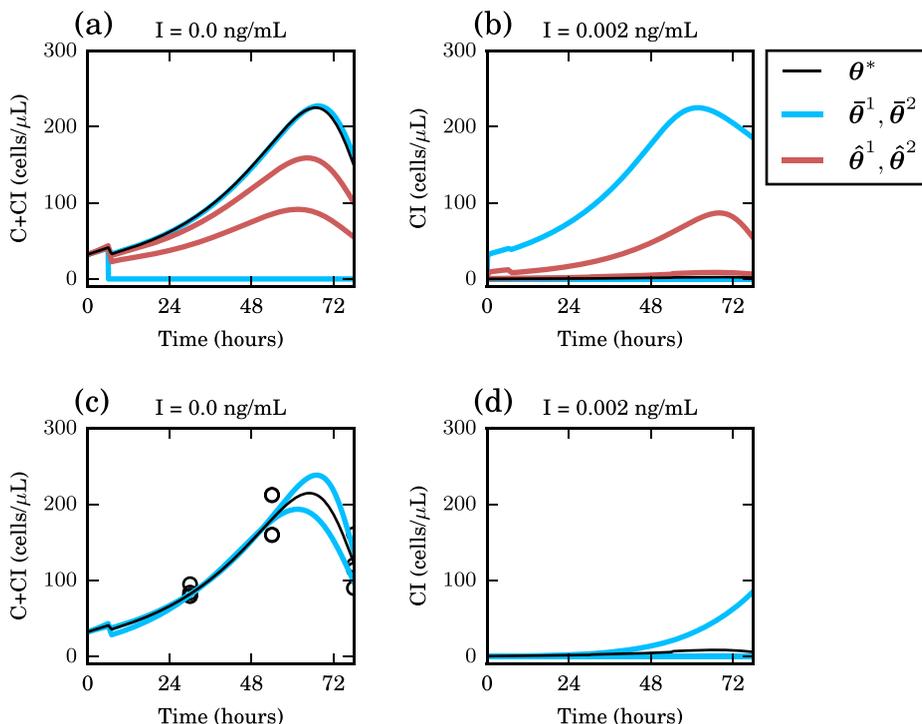


FIG. 5. (a) and (b) The expected experiment impact for a candidate experiment (a) on the prediction problem (b). In black and blue are the best-fit and prediction deviation models, respectively, as in Figs. 4(a) and 3(b). In red are the expected experiment impact models, which predict a substantial reduction in uncertainty from this experiment. (c) and (d) The corresponding figures after adding observations from the candidate experiment in (a). In (c), the updated best-fit and prediction deviation models after adding the data (circles, with overlapping data shown side-by-side). In (d), the updated prediction deviation, reduced from (b) by the new observations.

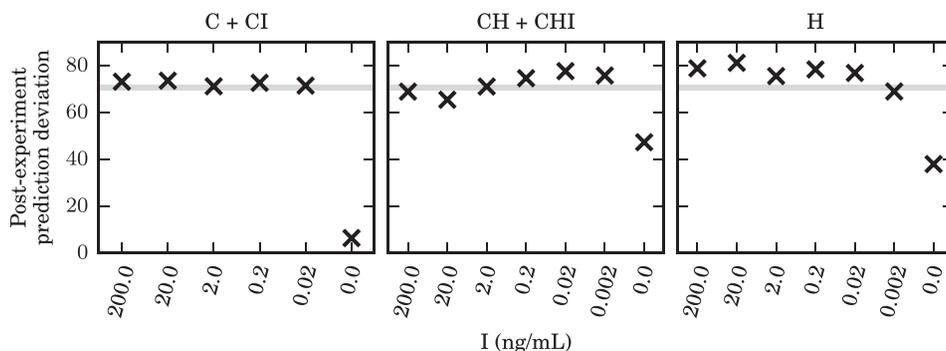


FIG. 6. Markers show the prediction deviation measured after including observations from each of the 20 candidate experiments. The horizontal line shows the prediction deviation prior to incorporating those observations, from Fig. 3. The experiment that most reduced uncertainty was that from Fig. 5.

experiments that provided a moderate reduction in uncertainty which was not matched by the estimated experiment impact. For the remaining 17 experiments, solving the estimated experiment impact problem correctly predicted that these experiments would not reduce uncertainty. Because it comes from adding a constraint to the prediction deviation problem, the estimated experiment impact problem cannot predict an increase in uncertainty, rather it can only predict that uncertainty will not decrease. Thus in Fig. 7(a) the estimated experiment impacts for the 17 ineffectual candidates are very close to the previously measured prediction deviation.

The two experiments with a moderate reduction in uncertainty that was not predicted give insight into how the estimated experiment impact problem works. Fig. 8 shows the pre-experiment and post-experiment prediction deviation models, along with the expected experiment impact models, for one of these two experiments. Estimated experiment impact is a worst-case analysis, and for these two experiments the worst-case models did not reduce uncertainty while the post-experiment models did. Each of these experiments had a potential outcome, consistent with the observed data, which would not have reduced uncertainty. Fig. 8(a) shows this worst-case potential outcome for one of the experiments. The actual data did not follow these worst-case trajectories, and in fact were able to moderately reduce uncertainty. Importantly, there were no experiments for which the estimated experiment impact indicated a reduction of uncertainty where in reality there was none. Because estimated experiment impact is a worst-case analysis, if the model is correct, this type of error will not occur and we

will not do experiments that end up not reducing uncertainty.

The fact that 17 of the 20 experiments produced no reduction of uncertainty could not have been known without solving the estimated experiment impact problem. In particular, measuring the uncertainty in the candidate experiments themselves, as in Fig. 4, could not predict that all of these experiments would have no impact. Fig. 7(b) compares the deviation on the candidate experiments, $z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{P}')$, to the actual experiment impact $z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y})$. The two candidates with the highest pre-experiment deviation did not actually reduce uncertainty at all. Fig. 7(b) shows that the uncertainty in the candidate experiment does not at all predict the impact that the candidate will have in the uncertainty of the prediction problem.

Fig. 9 shows the outcome of using the expected experiment impact in a sequential experimentation setting. Starting from the data in Fig. 3(a), we sequentially added in the data from the candidate experiment whose estimated experiment impact predicted the largest reduction in uncertainty. Each time after adding data from a candidate, we recomputed the prediction deviation with the new set of observations and recomputed the estimated experiment impact of the remaining candidates. Fig. 9(a) shows that adding in the second set of observations (those in Fig. 5(c)) produced a large drop in prediction deviation, shown in Figs. 5(b) and 5(d). Additional experiments continued to reduce uncertainty, but in much smaller amounts, consistent with the findings of Fig. 6. After adding data from just 3 of the 20 candidate experiments, the uncertainty was at nearly the level that was obtained by

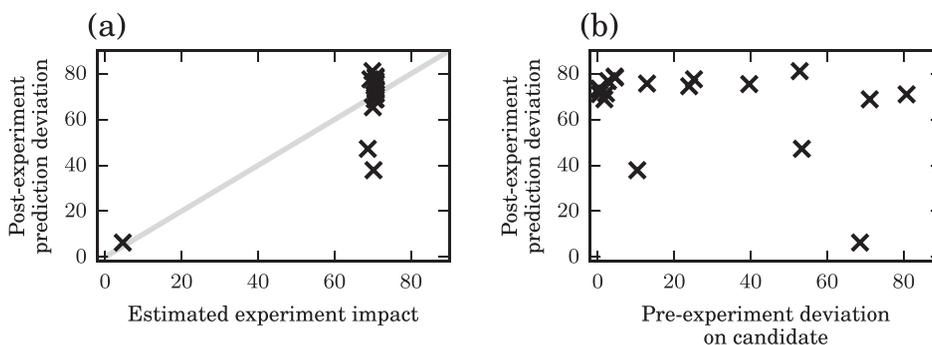


FIG. 7. (a) Markers show for each candidate experiment the estimated experiment impact compared to the actual prediction deviation measured after including the observations from that candidate. The gray line indicates where the estimate matches the actual outcome. (b) For each candidate experiment, the deviation of the prediction deviation models on that candidate experiment (see Fig. 4) compared to the prediction deviation measured after including the observations from that candidate. Candidate deviation does not provide a good prediction of experiment impact.

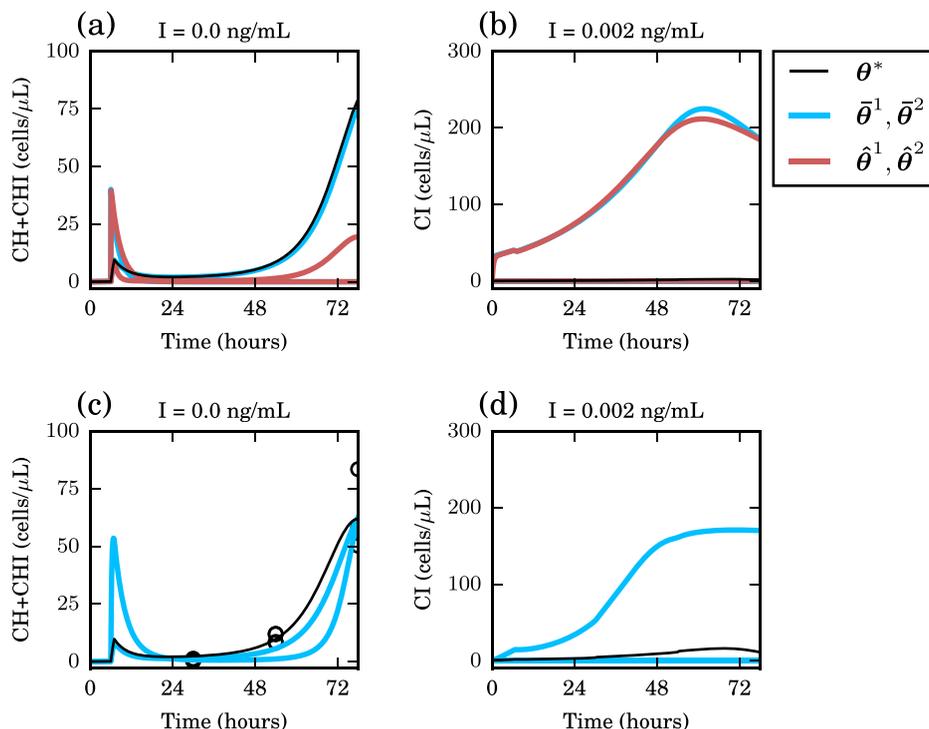


FIG. 8. (a) and (b) The expected experiment impact for a candidate experiment (a) on the prediction problem (b). In black and blue are the best-fit and prediction deviation models, respectively. The expected experiment impact models (red) show a possible outcome of the experiment that does not reduce uncertainty. (c) and (d) The corresponding figures after adding observations from the candidate experiment in (a). The actual data from the experiment (c) were not the worst-case outcome found by the expected experiment impact problem in (a), and actually produced a moderate reduction in uncertainty (d).

including all 20 of the candidate experiments. Fig. 9(b) compares the prediction deviation with only the initial experiment to that obtained after including the first three experiments selected using the estimated experiment impact. Initially, the data supported both the hypothesis that nearly none of the CD4 T cells were refractory, and the hypothesis that nearly all of the CD4 T cells were refractory. With the additional observations, the prediction deviation shows that only a small minority of CD4 T cells are refractory.

VI. THEORETICAL ANALYSIS

Prediction deviation has a strong theoretical guarantee that further motivates its use as a metric of uncertainty. For the purposes of the theoretical analysis, we assume that there exists a true model θ^{true} , and the observed data equal the output of this true model, plus random noise

$$\tilde{x}_i^j(t) = x_i(t; \theta^{true}, \mathbf{v}^j) + \epsilon_{ijt},$$

where ϵ_{ijt} are independent but not necessarily identically distributed random variables. Let α be such that z_u^* used to measure prediction deviation is the upper-bound on a $1 - \alpha$

confidence interval. Under reasonable assumptions on ϵ_{ijt} which are given in Appendix B, the following theorem holds:

Theorem 1. *With probability at least $1 - \alpha$*

$$z_{dev}(\theta^{true}, \bar{\theta}^1; \mathcal{Y}) \leq z_{dev}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y}) \quad \text{and}$$

$$z_{dev}(\theta^{true}, \bar{\theta}^2; \mathcal{Y}) \leq z_{dev}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y}).$$

This theorem means that the trajectory of the true model is in a particular sense bounded by that of the prediction deviation models: With high probability, it does not differ from either of the prediction deviation models by an amount larger than the difference in the prediction deviation models themselves. Thus if the prediction deviation is small and the trajectories of the prediction deviation models are close, then the trajectory of the true model can be specified within a narrow window, with high probability. This guarantee shows that prediction deviation corresponds to bounds on the underlying true model and provides additional support for the validity of prediction deviation as a metric of uncertainty. The proof is given in Appendix B.

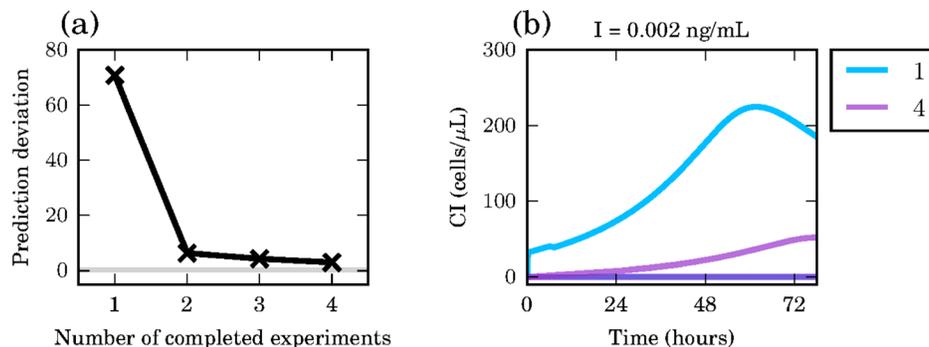


FIG. 9. (a) Observations were added sequentially from the candidate experiment with the best estimated experiment impact, and prediction deviation recomputed after each addition. The horizontal gray line shows the prediction deviation obtained after adding observations from all 20 candidate experiments. (b) The prediction deviation models corresponding to the 1 (blue) and 4 (purple) completed experiment markers from (a).

VII. RELATED WORKS

There are several related lines of work assessing predictive power in dynamical systems. Kreutz *et al.*¹⁰ use an optimization approach to measure prediction confidence intervals. Prediction intervals are measured by solving a sequence of minimization problems, separately for each time point in each prediction problem. Prediction intervals differ from the prediction deviation in that there might be different models that provide the upper and lower bounds at each time interval, whereas prediction deviation produces a single pair of models that maximizes the total deviation across all time points. The main strength of using prediction deviation as a measure of uncertainty is the ability to directly predict the impact of an additional experiment on the prediction deviation, via the estimated experiment impact problem. Kreutz *et al.*¹⁰ propose using the prediction intervals of the candidate experiments to decide which experiment would have the highest impact on the prediction problem. For nonlinear dynamical systems, reducing uncertainty of the model under one condition (the candidate experiment) does not necessarily reduce uncertainty under a different condition (the prediction problem). This is shown clearly in Fig. 7(b), where many candidate experiments had large uncertainty themselves, yet their observations did not reduce the uncertainty in the prediction problem. For prediction deviation, on the other hand, solving the optimization problem in (7) provides a direct estimate of how much reducing uncertainty in the proposed experiment will reduce uncertainty in the prediction problem. Because it is a worst-case analysis, the estimate from solving (7) also will not make the sort of error shown in Fig. 7(b) where the recommended experiments end up not reducing uncertainty. Other approaches to measuring prediction intervals include bootstrapping¹¹ and MCMC sampling in a Bayesian framework.¹² Vanlier *et al.*¹³ provide a review of recent approaches to measuring uncertainty both in parameters and in predictions.

Optimal experimental design has typically been studied in the context of parameter estimation^{14–16} or, more recently, model selection and discrimination.^{17–21} Kreutz and Timmer²² provide a review of recent approaches to optimal experimental design for these two problems. Our methods here are for optimal experimental design for prediction uncertainty, which generally requires predicting the impact of a proposed experiment on prediction uncertainty. Casey *et al.*²³ measure prediction uncertainty using a linearization of the prediction problem and then show how to predict the impact of a proposed experiment on the approximated prediction uncertainty. Another approach to optimal experiment design is to simulate the outcome of the proposed experiment using the best-fit model, and to measure the corresponding reduction in uncertainty.¹⁶ Useful experiments will themselves have high prediction uncertainty, so there will likely be a large range of possible outcomes, only one of which is the best-fit outcome. As shown in Fig. 8, the impact of the experiment on prediction uncertainty may depend strongly on which of the possible outcomes is realized. The actual reduction of uncertainty from an experiment could be much less than that predicted by the best-fit outcome, potentially wasting a valuable experiment. Solving (7) measures

uncertainty under the worst-case of the possible outcomes of the experiment, ensuring that the experiment will be useful whatever the outcome may be.

VIII. CONCLUSIONS

Two important questions that arise when fitting nonlinear dynamical systems to data are uncertainty quantification and optimal experimental design. We presented in this paper a prediction-centered approach for measuring uncertainty in a dynamical system's fit to data. Prediction deviation is able to directly show, via the pair of prediction deviation models, how much uncertainty remains in the prediction problem, thus answering the uncertainty quantification question. Solving the estimated experiment impact problem provides *a priori* a direct estimate of the impact that a candidate experiment would have on uncertainty. This allows the experimenter to choose the additional experiments that are likely to most reduce uncertainty, thus answering the question of optimal experimental design. We used the estimated experiment impact problem to sequentially choose 4 experiments which produced nearly the same reduction in uncertainty as the full set of 20 candidate experiments. In addition to the sequential experimentation setting that was demonstrated here, estimated experiment impact can also be used to predict the impact of simultaneously running a number of experiments by combining them into a single candidate. Finally, we proved a bound that with high probability provides a direct relationship between prediction deviation and how constrained the underlying true model is, providing a theoretical foundation for using prediction deviation as a metric of uncertainty.

ACKNOWLEDGMENTS

Funding for this project was provided in part by Siemens and the U.S. Army Research Office.

APPENDIX A: IMPLEMENTATION DETAILS

1. Specifying the parameter η

Observations \vec{x}' from candidate experiment \mathcal{P}' would constrain the prediction deviation models according to the two constraints $z_{\text{fit}}(\vec{\theta}^1; \mathcal{P}', \vec{x}') \leq \eta$ and $z_{\text{fit}}(\vec{\theta}^2; \mathcal{P}', \vec{x}') \leq \eta$. The amount that the fit error on the new models would be constrained, η , is a parameter in the estimated experiment impact problem, (7). Assuming normally distributed noise and a reasonable estimate of the experiment noise level σ_{ij}^2 , $z_{\text{fit}}(\vec{\theta}^*; \mathcal{P}', \vec{x}')$ follows a χ^2 distribution whose 95% percentile provides a reasonable choice for η . Alternatively, since observations are normalized by their noise level when computing fit error, if all observations contribute equally to the uncertainty then $\eta = z_u^* |\mathcal{P}'| / |\mathcal{P}|$ provides a reasonable choice, where $|\mathcal{P}'|$ is the number of observations in experiment \mathcal{P}' and z_u^* is the upper end of the 95% confidence interval for the best-fit error. This is the approach we used in our experiments, and the effect of this η through (7d) can be seen in Fig. 5(a).

The following result provides the motivation for constraint (7d) in the estimated experiment impact problem.

Proposition 1. $z_{\text{fit}}(\bar{\theta}^1; \mathcal{P}', \tilde{\mathbf{x}}') \leq \eta$ and $z_{\text{fit}}(\bar{\theta}^2; \mathcal{P}', \tilde{\mathbf{x}}') \leq \eta$ imply $z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{P}') \leq 2\eta$.

Proof.

$$\begin{aligned} z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{P}') &= \sum_{j=1}^J \sum_{i \in I_j} \sum_{t \in T_{i,j}} \left(\frac{x_i(t; \bar{\theta}^1, \mathbf{v}^j) - x_i(t; \bar{\theta}^2, \mathbf{v}^j)}{\sigma_{ijt}} \right)^2 \\ &= \sum_{j=1}^J \sum_{i \in I_j} \sum_{t \in T_{i,j}} \left(\frac{(x_i(t; \bar{\theta}^1, \mathbf{v}^j) - \tilde{x}_i^j(t))}{\sigma_{ijt}} \right. \\ &\quad \left. - \frac{(x_i(t; \bar{\theta}^2, \mathbf{v}^j) - \tilde{x}_i^j(t))}{\sigma_{ijt}} \right)^2 \\ &\leq z_{\text{fit}}(\bar{\theta}^1; \mathcal{P}', \tilde{\mathbf{x}}') + z_{\text{fit}}(\bar{\theta}^2; \mathcal{P}', \tilde{\mathbf{x}}') \leq 2\eta. \end{aligned}$$

The third line uses the triangle inequality, and the last line is by supposition. \square

This result allows for an approximation of the impact of \mathcal{P}' that does not require knowledge of the data $\tilde{\mathbf{x}}'$. The triangle inequality is generally loose, and incorporating this constraint into a maximization problem means that the result is the worst-case impact of $\tilde{\mathbf{x}}'$. These two approximations provide room for the additional approximation made above in choosing η . Ultimately, Fig. 7(a) shows that the approximations involved in estimating experiment impact are good enough to be useful.

2. Simulation and optimization

SloppyCell^{24,25} was used to integrate the model ODE system. In addition to integrating the model, SloppyCell integrates the forward sensitivity system, which provides gradients of the model trajectories with respect to the parameters, $\nabla_{\theta} x_i(t; \theta, \mathbf{v}^j)$. From these gradients, it is a straightforward calculation to obtain the gradients of the objectives and constraints for the three optimization problems in this paper: the data fitting problem, the prediction deviation problem, and the estimated experiment impact problem. All optimization problems were solved using random restarts of gradient-based optimization methods, with each optimization problem solved from 20 random initializations. The data fitting problem is an unconstrained minimization problem and was solved using the Scipy implementation of the Newton conjugate-gradient algorithm.^{26,27} The prediction deviation and estimated experiment impact problems are constrained maximization problems and were solved using the logarithmic barrier method [Ref. 27, Framework 17.2]. This method solves the constrained problem via a sequence of unconstrained problems, each of which was solved using the Newton conjugate-gradient method. The computational difficulty of each of these unconstrained problems is similar to that of the data-fitting problem. Solving (5) and (7) should thus scale in a similar way as the data fitting problem and have similar challenges. Feasible initial values for the prediction deviation and estimated experiment impact optimization problems were obtained using a Gaussian random walk from the best-fit parameters (which are always feasible), rejecting infeasible steps.

3. Experimental data

The data for the experiment on IFN α dynamics were those provided by Browne *et al.*⁹ There, two parameters were measured separately from these data, and we followed and treated these parameters, as well as all initial conditions, as known. One of the known parameters was the IFN α decay rate, and so $I(t)$ was thus known. The estimation done in this paper was then on a space of 7 parameters and 5 variables. The noise variance estimate used for weighted least squares and for prediction deviation, σ_{ijt}^2 , was taken as the average over time of the sample variances across the four replicates at each time point, separately for each set of variables and IFN α level. This is equivalent to the maximum likelihood estimate under a model where the noise is normally distributed with a variance that differs across variables and IFN α levels but is constant across time points.

APPENDIX B: PROOF OF THE THEORETICAL RESULT

The result of Theorem 1 provides a theoretical foundation for using prediction deviation as a metric of uncertainty by showing that it relates directly to bounds on the behavior of the underlying true model. The theorem requires the following assumptions:

Assumption 1. *The observed data are the output of a true model θ^{true} , plus noise: $\tilde{x}_i^j(t) = x_i(t; \theta^{\text{true}}, \mathbf{v}^j) + \epsilon_{ijt}$.*

Assumption 2. *The random variables ϵ_{ijt} are independent.*

Assumption 3. *The probability density function of ϵ_{ijt} is symmetric about 0 and unimodal, meaning the distribution function $F_{\epsilon_{ijt}}(x)$ is convex for $x \leq 0$ and concave for $x \geq 0$.*

Assumption 4. *Let θ^* be the best-fit model under a particular realization of the observations and z_u^* fixed. Then, assume $\mathbb{P}_{\tilde{\mathbf{x}}} (z_{\text{fit}}(\theta^*; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_u^*) \geq 1 - \alpha$.*

Assumption 2 requires independence, but does not require ϵ_{ijt} to be identically distributed, thus the noise level may vary across different observations. Assumption 3 is quite general: it is satisfied by the normal distribution, as well as by other heavy-tailed distributions. In Assumption 4, the model θ^* is held constant and the randomness is over different realizations of ϵ_{ijt} , and thus different realizations of $\tilde{\mathbf{x}}$. This assumption is about how the best-fit to one realization of the data generalizes to other realizations of the data, and requires that z_u^* , used in constraints (5b) and (5c), actually provides a $1 - \alpha$ upper bound for the fit error.

For the proof of Theorem 1, we define notation to describe the squared residuals. Let $R_{ijt}^{\text{true}} = (x_i(t; \theta^{\text{true}}, \mathbf{v}^j) - \tilde{x}_i^j(t))^2$ be the squared residuals under the true model and R_{ijt}^* the squared residuals under the best-fit model, θ^* . Let $b_{ijt} = x_i(t; \theta^*, \mathbf{v}^j) - x_i(t; \theta^{\text{true}}, \mathbf{v}^j)$ be the bias of the best-fit model.

The following result shows that intervals of the noise distribution centered on 0 contain the most probability mass.

Lemma 1. *For $x \geq 0$ and for $a \in \mathbb{R}$, $F_{\epsilon_{ijt}}(x+a) - F_{\epsilon_{ijt}}(-x+a) \leq F_{\epsilon_{ijt}}(x) - F_{\epsilon_{ijt}}(-x)$.*

Proof. This result follows from Assumption 3. When $x = 0$ the result is trivial. For $x > 0$, we first consider the case where $a \geq x$. For all $x \geq 0$, $F_{\epsilon_{ijt}}(x)$ is concave, and thus

$F'_{\epsilon_{ijt}}(x)$ is monotonically non-increasing. This means $\frac{\partial}{\partial a}(F_{\epsilon_{ijt}}(x+a) - F_{\epsilon_{ijt}}(-x+a)) \leq 0 \forall a \geq x$, and this quantity is maximized when $a = x$. Thus,

$$\begin{aligned} F_{\epsilon_{ijt}}(x+a) - F_{\epsilon_{ijt}}(-x+a) &\leq F_{\epsilon_{ijt}}(2x) - F_{\epsilon_{ijt}}(0) \\ &\leq 2(F_{\epsilon_{ijt}}(x) - F_{\epsilon_{ijt}}(0)) \\ &= F_{\epsilon_{ijt}}(x) + 1 - F_{\epsilon_{ijt}}(-x) - 2F_{\epsilon_{ijt}}(0) \\ &= F_{\epsilon_{ijt}}(x) - F_{\epsilon_{ijt}}(-x), \end{aligned}$$

which is the statement of the lemma. The second line follows directly from the concavity of $F_{\epsilon_{ijt}}(x)$ and the third line uses the symmetry $F_{\epsilon_{ijt}}(x) = 1 - F_{\epsilon_{ijt}}(-x)$. When $a \leq -x$, the same argument holds using the convexity of $F_{\epsilon_{ijt}}(x)$ for $x \leq 0$.

For the remaining case, $|a| < x$,

$$F_{\epsilon_{ijt}}(x) \geq \frac{1}{2}(F_{\epsilon_{ijt}}(x+a) + F_{\epsilon_{ijt}}(x-a))$$

by the concavity of $F_{\epsilon_{ijt}}(x)$ on the interval $[x-a, x+a]$. From the symmetry, it then follows that

$$1 + F_{\epsilon_{ijt}}(x) - F_{\epsilon_{ijt}}(-x) \geq F_{\epsilon_{ijt}}(x+a) + 1 - F_{\epsilon_{ijt}}(-x+a).$$

After rearranging, this proves the lemma. □

An important concept for the proof of Theorem 1 is that of a stochastic ordering, which we now define and then use to prove the theorem.

Definition 1. For random variables X and Y , $X \leq Y$ if

$$\mathbb{P}(X > x) \leq \mathbb{P}(Y > x) \forall x.$$

Lemma 2. $R_{ijt}^{\text{true}} \leq R_{ijt}^*$.

Proof.

$$\begin{aligned} \mathbb{P}(R_{ijt}^* \leq x) &= \mathbb{P}((\epsilon_{ijt} - b_{ijt})^2 \leq x) \\ &= F_{\epsilon_{ijt}}(\sqrt{x} + b_{ijt}) - F_{\epsilon_{ijt}}(-\sqrt{x} + b_{ijt}) \\ &\leq F_{\epsilon_{ijt}}(\sqrt{x}) - F_{\epsilon_{ijt}}(-\sqrt{x}) \\ &= \mathbb{P}(R_{ijt}^{\text{true}} \leq x), \end{aligned}$$

using Lemma 1. □

The next result comes from Shaked and Shanthikumar,²⁸ Theorem 1.A.3(b).

Lemma 3. For independent random variables X_1, \dots, X_n and Y_1, \dots, Y_n , let $X = \sum_{i=1}^n w_i X_i$ and $Y = \sum_{i=1}^n w_i Y_i$ with non-negative weights w_1, \dots, w_n . If $X_i \leq Y_i \forall i$, then $X \leq Y$.

Corollary 1. $z_{\text{fit}}(\theta^{\text{true}}; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_{\text{fit}}(\theta^*; \mathcal{P}, \tilde{\mathbf{x}})$.

Proof. The fit error is a weighted sum of the squared residuals, with weights $\frac{1}{\sigma_{ijt}^2}$, so this result follows directly from Lemmas 2 and 3, and Assumption 2. □

Theorem 1. With probability at least $1 - \alpha$

$$\begin{aligned} z_{\text{dev}}(\theta^{\text{true}}, \bar{\theta}^1; \mathcal{Y}) &\leq z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y}) \quad \text{and} \\ z_{\text{dev}}(\theta^{\text{true}}, \bar{\theta}^2; \mathcal{Y}) &\leq z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y}). \end{aligned}$$

Proof. By Corollary 1 and Assumption 4

$$\mathbb{P}(z_{\text{fit}}(\theta^{\text{true}}; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_{\text{fit}}^*) \geq \mathbb{P}(z_{\text{fit}}(\theta^*; \mathcal{P}, \tilde{\mathbf{x}}) \leq z_{\text{fit}}^*) \geq 1 - \alpha.$$

Thus with probability at least $1 - \alpha$, $(\theta^{\text{true}}, \bar{\theta}^1)$ is a feasible solution to problem (5). The proof of the theorem is then by contradiction: If $z_{\text{dev}}(\theta^{\text{true}}, \bar{\theta}^1; \mathcal{Y}) > z_{\text{dev}}(\bar{\theta}^1, \bar{\theta}^2; \mathcal{Y})$, then $(\bar{\theta}^1, \bar{\theta}^2)$ cannot be an optimal solution to problem (5). However, $(\bar{\theta}^1, \bar{\theta}^2)$ are defined to be optimal solutions, and so the theorem holds. The same argument simultaneously holds for $(\theta^{\text{true}}, \bar{\theta}^2)$. □

¹C. A. Brackley, O. Ebenhh, C. Grebogi, J. Kurths, A. de Moura, M. C. Romano, and M. Thiel, "Introduction to focus issue: Dynamics in systems biology," *Chaos* **20**, 045101 (2010).
²A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer, "Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood," *Bioinformatics* **25**, 1923–1929 (2009).
³G. Bellu, M. P. Saccomani, S. Audoly, and L. D'Angiò, "DAISY: A new software tool to test global identifiability of biological and physiological systems," *Comput. Methods Programs Biomed.* **88**, 52–61 (2007).
⁴O. Chiş, J. R. Banga, and E. Balsa-Canto, "GenSSI: A software toolbox for structural identifiability analysis of biological models," *Bioinformatics* **27**, 2610–2611 (2011).
⁵A. Sedoglavic, "A probabilistic algorithm to test local algebraic observability in polynomial time," in Proceedings of the 2001 International Symposium on Symbolic and Algebraic Computation (2001).
⁶R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna, "Universally sloppy parameter sensitivities in systems biology models," *PLoS Comput. Biol.* **3**, e189 (2007).
⁷B. Efron and R. Tibshirani, "Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy," *Stat. Sci.* **1**, 54–77 (1986).
⁸See supplementary material at <http://dx.doi.org/10.1063/1.4953795> for additional figures showing prediction deviation models and estimated experiment impact models for the Lorenz system.
⁹E. P. Browne, B. Letham, and C. Rudin, "A computational model of inhibition of HIV-1 by interferon-alpha," *PLoS One* **11**, e0152316 (2016).
¹⁰C. Kreutz, A. Raue, and J. Timmer, "Likelihood based observability analysis and confidence intervals for predictions of dynamic models," *BMC Syst. Biol.* **6**, 120 (2012).
¹¹P. C. St John and F. J. Doyle, "Estimating confidence intervals in predicted responses for oscillatory biological models," *BMC Syst. Biol.* **7**, 71 (2013).
¹²J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, and N. A. W. van Riel, "A Bayesian approach to targeted experiment design," *Bioinformatics* **28**, 1136–1142 (2012).
¹³J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, and N. A. W. van Riel, "Parameter uncertainty in biochemical models described by ordinary differential equations," *Math. Biosci.* **246**, 305–314 (2013).
¹⁴S. Bandara, J. P. Schlöder, R. Eils, H. G. Bock, and T. Meyer, "Optimal experimental design for parameter estimation of a cell signaling model," *PLoS Comput. Biol.* **5**, e1000558 (2009).
¹⁵A. Raue, V. Becker, U. Klingmüller, and J. Timmer, "Identifiability and observability analysis for experimental design in nonlinear dynamical models," *Chaos* **20**, 045105 (2010).
¹⁶M. K. Transtrum and P. Qiu, "Optimal experiment selection for parameter estimation in biological differential equation models," *BMC Bioinf.* **13**, 181 (2012).
¹⁷R. J. Flassig and K. Sundmacher, "Optimal design of stimulus experiments for robust discrimination of biochemical reaction networks," *Bioinformatics* **28**, 3089–3096 (2012).
¹⁸J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, and N. A. W. van Riel, "Optimal experiment design for model selection in biochemical networks," *BMC Syst. Biol.* **8**, 20 (2014).
¹⁹A. G. Busetto, A. Hauser, G. Krummenacher, M. Sunnåker, S. Dimopoulos, C. S. Ong, J. Stelling, and J. M. Buhmann, "Near-optimal experimental design for model selection in systems biology," *Bioinformatics* **29**, 2625–2632 (2013).
²⁰J. Daunizeau, K. Preuschoff, K. Friston, and K. Stephan, "Optimizing experimental design for comparing models of brain function," *PLoS Comput. Biol.* **7**, e1002280 (2011).
²¹D. Skanda and D. Lebiedz, "An optimal experimental design approach to model discrimination in dynamic biochemical systems," *Bioinformatics* **26**, 939–945 (2010).

- ²²C. Kreutz and J. Timmer, "Systems biology: experimental design," *FEBS J.* **276**, 923–942 (2009).
- ²³F. P. Casey, D. Baird, Q. Feng, R. N. Gutenkunst, J. J. Waterfall, C. R. Myers, K. S. Brown, R. A. Cerione, and J. P. Sethna, "Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model," *IET Syst. B.* **1**, 190–202 (2007).
- ²⁴C. R. Myers, R. N. Gutenkunst, and J. P. Sethna, "Python unleashed on systems biology," *Comput. Sci. Eng.* **9**, 34–37 (2007).
- ²⁵R. N. Gutenkunst, J. C. Atlas, F. P. Casey, B. C. Daniels, R. S. Kuczynski, J. J. Waterfall, C. R. Myers, and J. P. Sethna, "SloppyCell," 2007, see <http://sloppyCell.sourceforge.net/>.
- ²⁶S. G. Nash, "Newton-type minimization via the Lanczos method," *SIAM J. Numer. Anal.* **21**, 770–778 (1984).
- ²⁷J. Nocedal and S. Wright, *Numerical Optimization* (Springer-Verlag, New York, 2006).
- ²⁸M. Shaked and J. G. Shanthikumar, *Stochastic Orders* (Springer-Verlag, New York, 2007).