Title
Bayesian experiment design and estimation for probabilistic modeling of biological systems

Author
Martin Casas, Marc

Publication Date
2019

Bayesian experiment design and estimation for probabilistic modeling of biological systems

by

Marc Martin Casas

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Chemical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ali Mesbah, Chair
Professor Susan J. Muller
Professor Michiko E. Taga

Spring 2019

Bayesian experiment design and estimation for probabilistic modeling of biological systems

Marc Martin Casas

Abstract

Bayesian experiment design and estimation for probabilistic modeling of biological systems

by

Marc Martin Casas

Doctor of Philosophy in Chemical Engineering

University of California, Berkeley

Professor Ali Mesbah, Chair

Computational models have emerged as a key tool to study and characterize the behavior of biological systems. In order to accurately describe biological systems, computational models need to capture inherently complex features of biological behavior, such as nonlinear and probabilistic dynamics. Probabilistic models of biological systems can account for our lack of knowledge in model structure or parameters (uncertainty), as well as the physical sources of probabilistic behavior leading to heterogeneity.

Uncertainty quantification and optimal experiment design frameworks play a key role in the development of probabilistic models for biological systems, as they can enable system learning from informative experimental data by improving the accuracy of model structures and parameter estimates. Despite the extensive implementation of uncertainty quantification and experimental design frameworks in complex engineering systems, their application in biological systems has been lagging. This is because of their high computational cost, which can be further exacerbated when accounting for the large number of parameters and cellular constituents that characterize biological models. There is a need for computationally efficient tools that can capture complex biological properties to enable learning of biological systems in a systematic manner.

In this thesis, we present novel uncertainty quantification and optimal experiment design tools that can capture complex traits of biological behavior, namely uncertainty, heterogeneity, nonlinearities, and large numbers of parameters and

cellular constituents. More specifically, we introduce novel surrogate modeling methods for fast propagation of sources of uncertainty and heterogeneity, which enable the application of uncertainty quantification and optimal experiment design tools in biological systems. Subsequently, we introduce novel Bayesian methods for the estimation of genome-scale model parameters from noisy, sparse, and incomplete biological datasets. Additionally, we introduce computationally efficient methods to design experiments that yield informative data to refine parameter estimates and model structures in an offline and online manner (i.e., in the course of an experiment). The contributions of this thesis will create new opportunities for seamless implementation of uncertainty quantification and optimal experiment design for systematic and hypothesis-based modeling in biological systems.

Per l’avia Teresa i l’avi Joan

Contents

Contents ii

1 Introduction 1
- 1.1 Computational models in biological systems 1
- 1.2 Sources of probabilistic behavior in biological system modeling . . 3
  - 1.2.1 Physical sources: heterogeneity 3
  - 1.2.2 Non-physical sources: system uncertainty 5
- 1.3 Probabilistic modeling of cellular dynamics 5
- 1.4 Uncertainty quantification and experiment design in biological systems 7
- 1.5 Optimal experiment design for parameter estimation and model structure selection 8
  - 1.5.1 OED for model structure selection 9
  - 1.5.2 OED for parameter estimation 10
  - 1.5.3 Challenges of OED implementation 12
- 1.6 Outline of the thesis 12
- 1.7 Bibliography 13

2 Estimation of parameters in genome-scale biological models 17
- 2.1 Introduction 17
- 2.2 Methods 20
  - 2.2.1 Dynamic flux balance analysis (DFBA) 20
  - 2.2.2 Sparse multi-element polynomial chaos 22
  - 2.2.3 Numerical implementation 28
- 2.3 Results 29
  - 2.3.1 Decomposition of parameter space 31
  - 2.3.2 ME-PCE surrogate model validation 36
  - 2.3.3 Bayesian parameter inference using sequential Monte Carlo 39
  - 2.3.4 Forward uncertainty propagation 43
- 2.4 Discussion 45
  - 2.4.1 Potential improvements to surrogate model 47
  - 2.4.2 Alternative methods for uncertainty quantification 47
  - 2.4.3 Extensions to optimal experiment design 48

2.5 Bibliography 49

3 Optimal experiment design for offline model structure selection 56

3.1 Introduction 56
3.2 Modeling population noise in the JAK2/STAT5 signaling pathway 57
3.3 Probabilistic optimal experiment design problem 60
3.4 Model discrimination in the presence of population noise 61
3.5 Discussion 61
3.6 Bibliography 63

4 Optimal experiment design for online model structure selection 66

4.1 Introduction 66
4.2 Background: Model Selection 69
  4.2.1 OED for model structure selection 72
  4.2.2 Challenges of OED for model structure selection 73
4.3 Method 1: OED for online model selection using the Bayes Risk . 74
  4.3.1 Uncertainty Propagation 74
  4.3.2 Tractable Criterion for Multiple Hypothesis Discrimination 79
  4.3.3 Moment-based Chance Constraint Approximation 81
  4.3.4 Tractable OED Problem 82
  4.3.5 OED for online model selection 84
  4.3.6 Simulation Case Study 87
4.4 Method 2: OED for online model selection using the k-nearest neighbors distance 95
  4.4.1 Uncertainty Propagation 95
  4.4.2 Tractable OED Problem 99
  4.4.3 Simulation Case Study 100
4.5 Conclusions 102
4.6 Bibliography 103

5 Bayesian optimal experiment design for parameter estimation 108

5.1 Introduction 108
5.2 Formulation of optimal Bayesian experimental design 111
5.3 Stochastic dynamic optimization with chance constraints 115
  5.3.1 Sample-based estimator for the expected utility 115
  5.3.2 Moment-based approximation of chance constraints 116
  5.3.3 Sample average approximation for chance constrained Bayesian OED 117
5.4 Arbitrary polynomial chaos expansions as surrogates 119
  5.4.1 PCE formulation 120
  5.4.2 Convergence, optimality, and error analysis 122
  5.4.3 Estimation of PCE coefficients 123
  5.4.4 The optimized stochastic collocation method 126

5.4.5 Global versus local PCE with respect to the design space . 128
5.5 Numerical results 130
5.5.1 The dynamic forward model 130
5.5.2 Local and global PCE implementations 131
5.5.3 Stochastic dynamic optimization implementation and results 133
5.5.4 Comparison between local and global PCE surrogates . . . 135
5.5.5 Scaling with respect to number of design variables 138
5.5.6 Chance constraint approximation and tuning 139
5.6 Conclusions 141
5.7 Bibliography 144

6 Conclusions and future work 149
6.1 Conclusions 149
6.2 Future work 151

Acknowledgments

First, I would like to acknowledge my advisor, Prof. Ali Mesbah, for giving me with the opportunity to join his lab. I have learned a lot, and I am extremely grateful for that. Besides my advisor, I would also like to thank my qual and thesis committee members: Prof. Muller, Prof. Schaffer, Prof. Taga, and Prof. Landry. I deeply appreciate their insightful comments and words of encouragement.

I am grateful to the CBE Department; my experiences here have truly made me a better person. A very special thanks to Carlet for providing amazing support and putting out many fires for me. I would also like to thank the learning-a-lot crowd (Eddy, Monica, Parry and Lisa) for helping me maintain my sanity during my first semester at Berkeley. Lastly, a big thank you to Joel and Tor, who have been excellent role models, and provided me with priceless support throughout my PhD.

I would like to acknowledge my college friends Santacani, Silvia, Elisa, Sandra i Albert. From study marathons to lab classes, we endured a lot together in undergrad. I can say that I made it this far because of their encouragement and our 'cotis' in and out of class. My lifelong friends from Palafrugell have always been tremendously supportive - my friendship with them has taught me to love and respect myself as well as others. Silvia, Marina, Tomas, Ricard, Xevi, Irene, Mireia - I owe you a lot. I am also indebted to Maria, who has been an amazing friend and roommate for the past two years. Maria has been incredibly patient with me and, perhaps inadvertently, has also taught me to be more compassionate (and a lot of Catalan words that I did not know). Gracies!

Who I am today has been largely shaped by my family, who have been extremely loving, supporting, and accepting throughout my life. Despite the distance, I know I can count on them for anything I need. Adri, Auri, Mama, avia Teresa i avi Joan, avia Pilar, Dome, tia Rosa i Tere, uncles Lluis i Eduard i Marta - moltes gracies per tot.

I could not have gone through graduate school without the 'rajades', lunches, and shenanigans with Eli - you have truly been the best grad school companion I could have asked for. Being a mate of honor at your wedding was the gift of a lifetime - bring on the post-grad school adventures!

Finally, I truly could not have made it without the unconditional support of Michael, who has been amazingly loving, patient and caring. Thanks for always being there for me. I cannot wait to see what the future has in store for us.

To all of you, thank you from the bottom of my heart.

Chapter 1

Introduction

This chapter discusses the importance of uncertainty quantification and optimal experiment design for building accurate models that can capture and/or reduce uncertainty and heterogeneity in biological systems. The chapter highlights the need to capture probabilistic features of biological systems (i.e., uncertainty and heterogeneity) in our model predictions. Subsequently, the importance of using uncertainty quantification and experiment design tools to build representative models that can quantify uncertainty and/or heterogeneity in biological systems is discussed. The chapter is closed with a discussion on the novel contributions of this dissertation.

1.1 Computational models in biological systems

Computational models have emerged as a key tool to study and characterize the complexity of biological systems. The main goal of computational models is to replicate the true dynamics (i.e., behavior) of the system under study by explaining how components within a cell interact dynamically and predicting quantities of interest over time (e.g., metabolite concentrations, cellular traits). Computational models are also useful to generate and test hypotheses about the underlying system dynamics, as well as designing experiments that can validate such hypotheses. Modeling of biological systems has been used to aid discovery and generate insights, such as elucidating new molecular mechanisms [8, 18], validating hypothesis [21], and mapping biochemical networks [29].

Computational models can take many forms and serve many purposes [15]. Stelling (2009) classified models of biological systems into: i) interaction-based static models with no stoichiometric or parametric information, such as topological network analyses [5], ii) constraint-based static models with stoichiometric

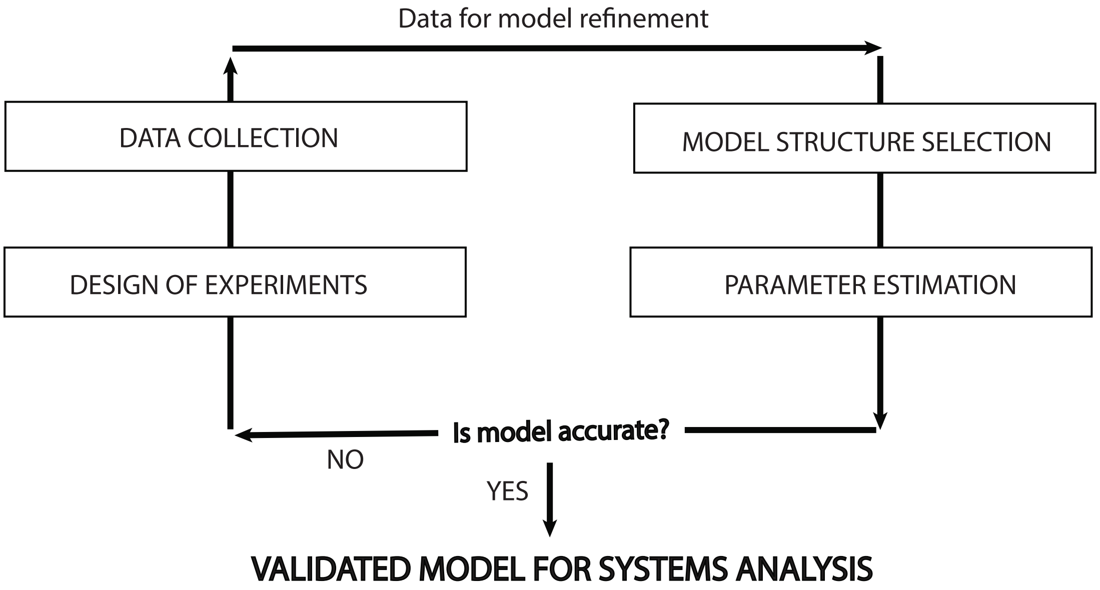

Figure 1.1: The iterative process of building models for biological systems. Available model structures and parameter estimates can be defined from prior mechanistic system information and preliminary experimental data. If the model is not accurate enough, more experimental data is gathered to refine the model structure and the estimates for the parameters. The process is repeated until acceptable accuracy is achieved. Figure adapted from [9]

constraints but no parameters, such as Flux Balance Analysis (FBA) [25] and, iii) mechanism-based dynamic models with stoichiometric constraints and parametric information, such as differential algebraic equation based models [14]. Some efforts have been made to adapt constraint-based models to dynamic frameworks [19] (i.e., dynamic Flux Balance Analysis (dFBA)). We refer the reader to [14] for a comprehensive review on static and dynamic models.

Building dynamic models is a largely iterative process, which requires the interplay of both theoretical and experimental methods (see Figure 1.1). First, a priori mechanistic information and available experimental data can be used to define a model structure and estimate its parameter values. After refining the parameter estimates and validating the model predictions with additional datasets, a working model is generated. The predictions of the working model can then be used to design further experiments. The availability of more data allows for further model refinement, which will likely improve the model predictive capability and hence the quality of the insights it can provide. In addition to quantity of data, its quality will significantly affect model prediction quality. Noisy and sparse system measurement data often leads to non-physical uncertainty in the employed model structure (i.e., biochemical reaction network) and/or model parameter values. An approach to quantify our uncertainty in the range of possible model structures or parameters is by use of probability distributions [32, 34]. As a result, probabilistic models that can account for uncertainty are useful for appraising model prediction quality.

Uncertainty is a non-physical source of probabilistic behavior that originates from incomplete system knowledge and/or noisy measurements. However, physical sources of variability can also lead to probabilistic dynamics and outcomes in individual cells. In the context of cell populations, this probabilistic behavior has been shown to lead to phenotypic heterogeneity, even within clonal cell populations [12]. By accounting for both physical and non-physical sources of probabilistic behavior in our model building process, we can develop more accurate models that capture the heterogeneity of biological systems, while quantifying our degree of belief in model predictions (i.e., uncertainty). We refer the reader to [32, 34] for a thorough review on probabilistic modeling of biological systems.

1.2 Sources of probabilistic behavior in biological system modeling

1.2.1 Physical sources: heterogeneity

Generation of heterogeneity within cell populations is a regular occurrence in biological systems [33]. Varying degrees of cell-to-cell variability can spontaneously arise within cell populations under homogeneous experimental conditions, possibly leading to complex heterogeneous mosaics of phenotypic behavior [30]. The molecular basis for population heterogeneity has been classically attributed to genetic factors; however, clonal cell populations (i.e., cells with identical genetic expression patterns) can show non-trivial levels of cell-to-cell variability due to non-genetic sources of population heterogeneity. Huang et al. (2009) proposed a tiered classification scheme of non-genetic variability sources (Figure 1.2), which considered extrinsic and intrinsic sources of heterogeneity. Extrinsic sources comprised those factors foreign to the cell that have a different impact across a cell population (e.g., temperature gradients), whereas intrinsic sources as those processes that occur within each cell. Huang et al. (2009) further divided intrinsic sources into population noise, when non-genetic traits are broadly distributed across a population of cells, and temporal noise when a cellular trait varies due to local concentration fluctuations of cellular constituents.

The advent of novel experimental tools allowing for quantification of temporal fluctuations in enzymatic reactions and cell-to-cell trait variability has facilitated the characterization of extrinsic and intrinsic sources of noise. Examples of such techniques are single-molecule kinetic assays, which allow for the characterization of enzyme catalytic rate fluctuations (i.e., temporal noise) [35], or single-cell assays that accurately describe the distribution of a cellular trait or constituent

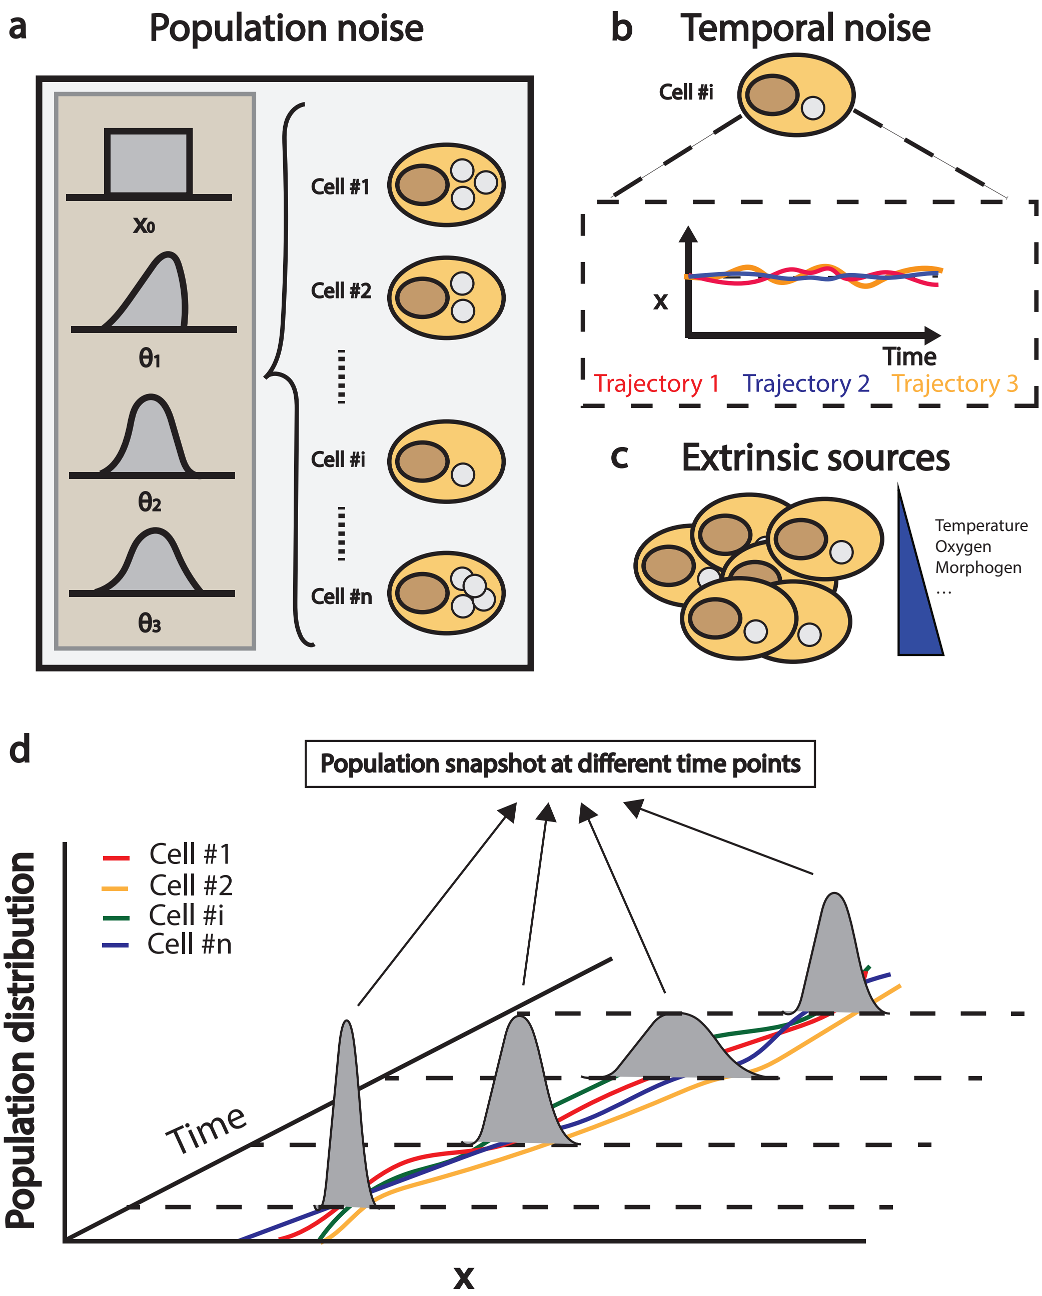

Figure 1.2: Impact and sources of heterogeneity in the dynamics of cell populations. a) Graphic description of the intrinsic sources of heterogeneity in biological systems. Population noise refers to cellular traits and constituents that are probabilistically distributed across a population of cells, for instance, due to asymmetric distribution of molecules upon cell division. b) Temporal noise describes local fluctuations in the concentration of cell constituents over time. c) Extrinsic sources, here depicted as a gradient of temperature or other effectors, can differentially affect cells within a population thus generating heterogeneity. d) The sources of probabilistic behavior will affect the dynamics of heterogeneity within the cell population over time, here described with an arbitrary cell trait or constituent x. The overall aggregation of single cell trajectories in terms of quantities of interest (i.e., x) can be quantified with probability distributions describing population heterogeneity at any given time point.

concentration across a cell population (i.e., population noise) [31]. Models capturing extrinsic and intrinsic sources of noise will output the range of physically possible outcomes associated with cell dynamics, weighted by their probability of occurrence.

1.2.2 Non-physical sources: system uncertainty

In addition to physical, stochastic sources of heterogeneity, probabilistic models can also be employed to capture our lack of knowledge of a system (i.e., system uncertainty). Classically, our uncertainty can be reflected in the model structure when the underlying architecture or network described by the model is partially unknown, model parameters when we are unable to perfectly know the exact value of a model parameter, or initial concentrations of the model, for example, when the amount of cellular constituents in a cell upon division is unknown.

System uncertainty mainly arises from our inability to observe the system fully and at all times. This is due to the fact that, in general, current analytical tools in biology yield incomplete and noisy observations of the system. For example, measurements might be rarely accessible or available (i.e., sparse), or might require the termination of the system under study (e.g., flow cytometry, laser ablation), thus generating irreproducibility. In addition to their sparsity, the noisy nature of experimental data leads to a range of parameter values that can explain the data similarly well. The degree of confidence in a certain parameter value can typically be characterized with a probability distribution.

The interpretation of outputs in probabilistic models of biological systems capturing uncertainty significantly differs from those capturing physical sources of heterogeneity. Upon improvement of our knowledge in the system, uncertainty in model outcomes can be eliminated. Heterogeneity, on the other hand, can be modulated or reduced, but cannot be eliminated upon perfect system knowledge. If a cellular trait is deterministic (i.e., no physical variability), estimation procedures with perfect knowledge of the system would yield a perfect single point estimate. Conversely, if the trait is affected by heterogeneity, perfect knowledge of the system would yield a distribution strictly describing the impact of physical sources.

1.3 Probabilistic modeling of cellular dynamics

The general framework for probabilistic modeling of a biological system relies on the so-called Chemical Master Equation (CME) [10]. The CME is a master equation describing the probability that the system is in a given state at a given time. The CME framework describes the system state as the discrete number of molecules of each cell constituent at any given time. The equation is derived from the probabilities of transition among all feasible system states, which are a function of the rates of biochemical reactions in the system. Given that the CME framework accounts for each individual reactant and reaction in the system, it is prohibitively expensive to implement for large systems (i.e., whole-cell) and large cell populations. We refer the reader to [10], [23] for detailed background and derivation information on the CME, its approximations, and applications.

Gillespie (2000) showed that there is a regime where the CME and its discrete dynamics can be approximated by a continuous-state equation, called the chemical Langevin Equation (CLE). This approximation regime involves assuming large numbers of molecules and fast reaction firing frequency, which allow for the quantification of system constituents as real-valued quantities (i.e., concentration instead of molecular count). The dimensionality of the CME models is proportional to the total number of molecules of all cell constituents in the system. Conversely, the dimensionality of the CLE models is proportional to the number of modeled cell constituents, which can render the CLE a computationally tractable approach to modeling cell dynamics. Furthermore, CLE models are compatible with widely available differential equation solvers and integrators [11].

In this thesis, we describe cellular dynamics using models of the CLE form. In the context of uncertainty or heterogeneity, the CLE models can be expressed as

$$\dot{x}(t) = \mu(x, \theta)dt + \sigma(x, \theta)dw, \tag{1.1}$$

where x is the vector of concentrations of cellular constituents, with initial conditions x_0; θ is the vector of model parameters; t is time; and w is the vector of fluctuations of cellular constituents defined as a Wiener process with zero mean and Σ variance $w \sim \mathsf{N}(0, \Sigma)$. The functions μ and σ denote the so-called drift and diffusion processes, respectively. The drift term μ describes the main driving force of the system (i.e., kinetic reactions), and diffusion denotes stochastic fluctuations affecting the vector of cellular constituents x. When Eq. (1.1) is comprised of only the drift term μ, the model will reduce to a set of deterministic set of ordinary equations (i.e., no stochastic fluctuations). If the diffusion term is included, Eq. (1.1) can be used to model sources of temporal noise as defined in Section 1.2.1.

Eq.(1.1) provides a flexible framework to model the sources of heterogeneity and uncertainty, as both the initial concentrations of cell constituents x_0 and model parameters θ can capture sources of probabilistic behavior. Capturing x_0 as a probability distribution can describe uncertainty in the initial concentration of certain cellular constituents, and/or intrinsic population noise derived from the asymmetric distribution of molecules upon cell division [13]. On the other hand, capturing model parameters as probability distributions can account for uncertainty in the true parameter values, extrinsic (e.g., temperature gradients), as well as intrinsic population sources of heterogeneity. Lastly, intrinsic temporal noise, which arises due to stochastic fluctuations in the concentration of cell constituents, can be captured with the noise process w. In addition to extrinsic and intrinsic non-genetic sources of heterogeneity, the modeling framework defined by Eq. (1.1) has been shown to efficiently capture genetic and epigenetic differences across cells [1, 20].

In the remainder of this thesis, we describe both uncertainty and heterogeneity sources with the vector $\Theta = [x_0; \theta]$, which encompasses initial concentrations and parameters that are probabilistically distributed across the cell population. In addition, uncertainty in the underlying model structure can be accounted for with *competing* drift functions $\mu_1, \mu_2, ...$, which capture different hypothesis regarding the underlying biochemical dynamics of the system.

1.4 Uncertainty quantification and experiment design in biological systems

Numerical simulation of Eq. (1.1) will yield the time evolution of specific cellular traits or properties over time across the population of cells. Three computational challenges stem from Eq. (1.1) with probabilistic Θ. First, the propagation of Θ through the cellular dynamics of interest. Second, the characterization of the statistical properties of Θ from available experimental data, so that the cell population model is representative of the entire cell population. And third, the design of informative experiments that can yield data to generate knowledge about the system, such as improving the model structure μ or the estimates for Θ.

Uncertainty quantification (UQ) tools have been used to quantify, characterize, and reduce uncertainty in a range of experimental settings [24]. Given the analogous mathematical treatment of heterogeneity and uncertainty, UQ tools can be effortlessly extended to heterogeneity quantification, reduction, and modulation. The field of UQ can be mostly categorized into two main problems: the *forward* and the *inverse* problems (see Figure 1.3). The forward problem consists in the propagation of physical or non-physical sources of heterogeneity in terms of measurable outputs. The inverse problem consists in providing estimates for model structures or parameters given system observations, such as metabolite concentration data.

However, more often than not, the quantity and quality of available experimental measurements are insufficient to appropriately characterize model parameters in estimation procedures. This motivates the need to design better experiments that maximize the information content of the data provided by each experiment, with minimal consumption of experimental resources. Model-based Optimal Experiment Design (OED) capitalizes on the use of model predictions to design

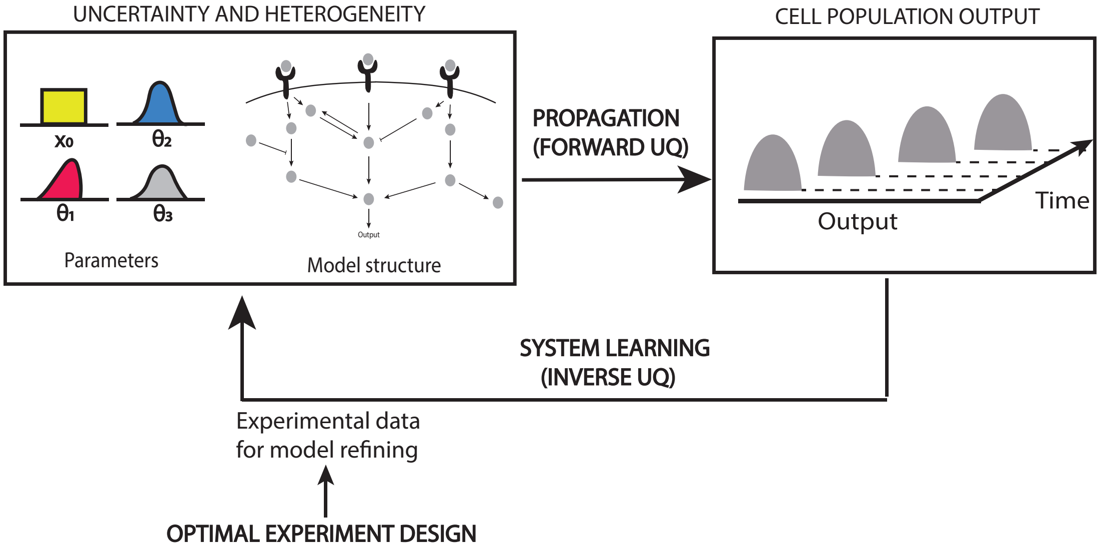

Figure 1.3: Depiction of the relationship between forward and inverse uncertainty quantification (UQ) problems, and optimal experiment design (OED). The forward problem starts from probabilistic Θ describing uncertainty and/or heterogeneity in the parameters, as well as model structure uncertainty, and aims to describe the impact of such probabilistic behavior in the output of interest. Conversely, the inverse problem attempts to infer and characterize, from system observations, the original sources of probabilistic behavior in Θ or the most valid model structure. The more informative the experimental data used for the inverse problem, the better the estimation of Θ can be expected. OED frameworks are important for elucidation of biological systems, as they can significantly improve the efficacy of the inverse problem by designing experiments that yield informative data for learning about the system.

experiments that maximize the information content of the desired metrics, for instance, for parameter estimation or model structure elucidation [3, 4] (see Figure 1.3). The model is employed as a proxy for the real system, and an experimental design that maximizes the predicted data utility is defined as the most optimal experimental design. In an analogous way to uncertainty reduction around parameters or model structures, OED frameworks could also be employed to reduce heterogeneity in the desired model outcomes or quantities of interest.

1.5 Optimal experiment design for parameter estimation and model structure selection

OED frameworks have been defined for a variety of purposes [2]. Generally, OED objectives can be categorized into two main categories: i) selection of a model structure, and ii) improvement of estimates for model structure parameters. The usefulness of measurement data is largely dependent on the objective of the experiment; for example, experiments can be designed to increase sensitivity to model parameters when performing parameter estimation, or to differentiate among competing model output predictions when performing model structure selection. In this section, we introduce formulation for a general OED framework so that it can be adapted for parameter estimation or model structure selection. Reviews on OED and its formulation for biological designs can be found in [3] and [16].

First, let us define the variables of an experimental setting. Following notation from Eq. (1.1), the system observations can be defined as

$$y_i = h_i(x, t, \Theta, d) + \epsilon_i, \quad i = 1, \ldots, n_y \tag{1.2}$$

where y is the vector of measurement outputs, h is the vector of functions with size n_y that describes the relationship between each experimental output and cellular constituents, ϵ is the vector of measurement noise with size n_y, and d are physical variables of the experiment that can be manipulated. d is intentionally defined to be a general design variable, which can comprise measurement times, input to the system, or other controllable experimental settings as will be defined further for each type of experiment.

For the sake of generality, we define the following optimization framework for experimental design

$$\begin{aligned} \max_d \quad & \Phi(\mathrm{x}, \Theta) \\ \text{subject to: } & \text{A set of models of form 1.1} \\ & \text{System constraints} \end{aligned} \tag{1.3}$$

where Φ is a scalar metric of the experiment that measures the desired outcome. An important remark is that, given that there is uncertainty or heterogeneity in the model(s), the objective Φ will be typically represented as an expectation over all potential system outcomes. Upon solving, the problem defined by Eq. (1.3) yields the most optimal design d^* that is able to maximize the scalar metric of information content. In the ensuing sections, we introduce the specific adaptations of Φ in Eq. (1.3) for OED for model structure selection and parameter estimation.

1.5.1 OED for model structure selection

In OED for model structure selection, we are interested in selecting the most representative model out of a set of competing models. Let H denote the set of competing model structures to describe the dynamics of a given system, with known prior probabilities of being the correct model $\mathsf{P}(H_i)$. OED for model discrimination aims to determine, provided certain system measurements y, which model in the set is the most representative of the underlying system dynamics. Making use of conditional probabilities, the problem of model selection can be stated as selecting the model hypothesis $\mathsf{P}(H_i)$ such that the probability of the correct model given experimental data is maximized [26]

$$i \in \text{argmax}_j \ \mathsf{P}(H_j|y, d). \tag{1.4}$$

Bayes rule can mathematically describe the change in our degree of belief (i.e., probability) that a given model is correct, provided experimental data y, as follows [7]

$$\mathsf{P}(H_j|y, d) = \frac{\mathsf{P}(y|H_j, d)\mathsf{P}(H_j)}{\mathsf{P}(y|d)}, \tag{1.5}$$

where $\mathsf{P}(H_j|y, d)$ describes an updated degree of belief in model hypothesis j, provided experimental data y and design d. $\mathsf{P}(y|H_j, d)$ describes the likelihood function, which captures the probability of the observed experimental data belonging to the model hypothesis j in the current experimental setting; $\mathsf{P}(H_j)$ is our degree of belief in hypothesis j prior to gathering data y, and $\mathsf{P}(y|d)$ is the total probability of observing experimental evidence y across all candidate models.

Using Bayes rule, the model selection criterion (i.e., $i \in \text{argmax}_j \quad \mathsf{P}(H_j|y, d)$) yields several regions $\mathcal{R}_i$ within the experimental measurement space where our degree of belief is higher for model hypothesis i than for the remaining hypotheses (see Figure 1.4). If our experimental measurements belong to the region $\mathcal{R}_0$, then H_0 will be selected. If any other model hypothesis other than H_0 is selected, a model selection error will have occurred.

The objective of the OED problem for model selection, Φ, can be related to the probability of model selection error, defined as [6]

$$\mathsf{P}(\text{error}) = \sum_i \sum_{j \neq i} \mathsf{P}(y \in \mathcal{R}_j, H_i|d). \tag{1.6}$$

An optimal experiment design that minimizes the overlap among regions $\mathcal{R}$ will also minimize the probability of model selection error (and maximize Φ), thus leading to more confidence in the selection of the model. The probability of model selection error is a direct function of the experimental design d. As a result, it can be advantageous for selecting the model that is most representative of the underlying system dynamics.

1.5.2 OED for parameter estimation

In OED for parameter estimation, we are interested in selecting the most optimal experiments for estimating model parameters from noisy, indirect, and incomplete measurement datasets. Following the definition of the CLE model in Eq. (1.1),

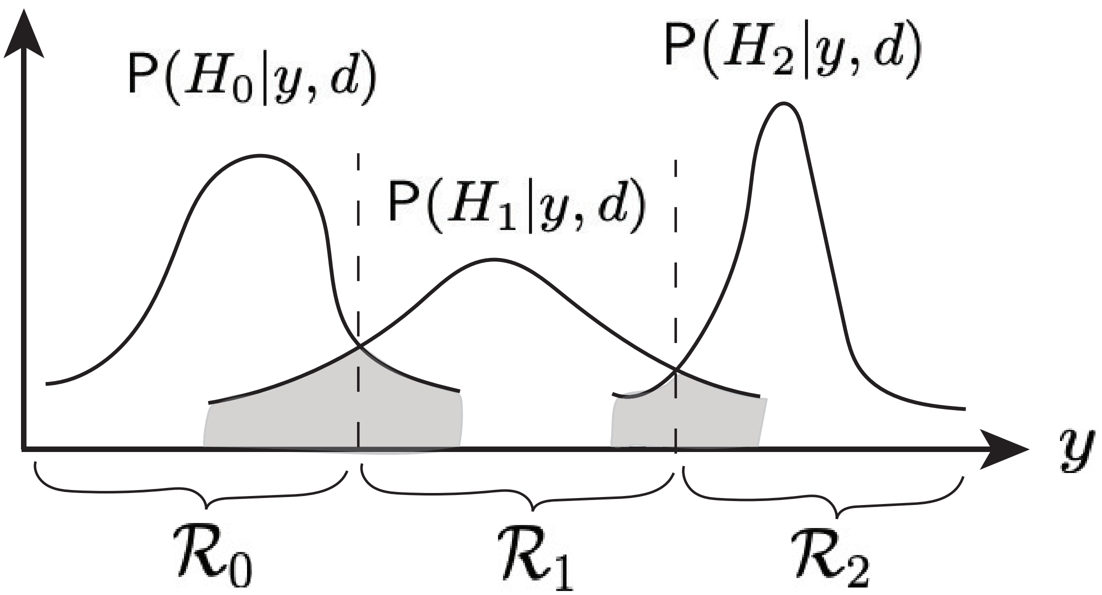

Figure 1.4: Illustration of model selection for a given experimental design d. In the case of multiple competing model structures, model selection results in a number of decision regions in the observation space. If the observation lies in $\mathcal{R}_1$, then the model selection algorithm will select model 1. The overlap between distributions of competing model predictions gives rise to the probability of model selection error (shown by grey area). Graph was adapted from [6].

with system observations y, uncertainty and/or heterogeneity Θ and the Bayes' rule, one can derive the definition of so-called expected experimental utility as follows [17]

$$U(d) = \int_{\mathcal{Y}} \int_{\mathcal{O}} u(d, y, \Theta) \mathsf{P}(\Theta, y|d) d\Theta dy, \tag{1.7}$$

where $\mathcal{Y}$ and $\mathcal{O}$ are the measurement and uncertainty and/or heterogeneity spaces, and $u(d, y, \Theta)$ is the utility function. The utility function is to reflect how useful the experiment is for the estimation of model parameters as tightly as possible. Given that y is noisy and Θ is oftentimes probabilistically distributed, Eq. (1.7) is defined as an expectation over both the measurement and the uncertainty and/or heterogeneity spaces. The objective of the OED framework for parameter estimation, Φ, will be equivalent to the expectation of utility $U(d)$. A suitable example of a choice for the utility function $u(d, y, \Theta)$ include the Kullback-Leibler (KL) distance, which quantifies the difference between the distributions of the parameter estimates before and after the experiment is performed. The experimental design that leads to the largest change in our parameter estimates will be the one with the largest KL distance, and will thus be the most informative.

The performance of OED using Eq. (1.7) as the objective of the optimization, without any further approximations, constitutes a so-called Bayesian OED. Bayesian OED is very computationally challenging to perform because the utility is a function of the full distributions of Θ and the measurement noise, as shown by the double integral in Eq. (1.7). Approximations of measurement noise as zero-mean Gaussian with diagonal covariance, and linearly approximating the system output trajectories y around the optimal value of the distributed parameters Θ, leads to classical OED approaches. Classical OED approaches optimize some scalar metric of the so-called Fisher Information Matrix (FIM), which describes the weighted sensitivity of system outputs with respect to the distributed parameters.

1.5.3 Challenges of OED implementation

There exist several challenges to the implementation of Eq. (1.3) for both parameter estimation and model structure selection. First of all, the propagation of probabilistic behavior in Θ in terms of model predictions is challenging and oftentimes prohibitively expensive. Secondly, no closed-form expression of Eq. (1.6) and Eq.(1.7) are available to easily compute the objective of the OED, Φ. Finally, the possibly probabilistic constraints in the OED problem (see Eq. (1.3)) are intractable. In this thesis, we introduce novel approaches that can circumvent the challenges associated with the implementation of OED for both model structure selection or parameter estimation.

1.6 Outline of the thesis

This dissertation contributes to the area of quantification of uncertainty and heterogeneity in biological systems. Specifically, this thesis introduces novel computational methods and frameworks that allow to circumvent common challenges in uncertainty quantification and experiment design for optimization-based applications. Namely, we introduce new methods for the propagation of sources of uncertainty and/or heterogeneity, as well as the estimation of model parameters and structures from uncertain data. Subsequently, we introduce multiple OED frameworks that allow for the design of experiments or protocols that can perform model selection and classical parameter estimation in an offline or online manner (i.e., in the course of the experiment). The introduced frameworks are based on both Bayesian and Classical OED frameworks. The outline of the thesis is as follows.

In Chapter 2, we introduce a novel framework for the estimation of model parameters in whole-cell models with thousands of biochemical reactions, with the availability of very few data points. In particular, the framework deals with dynamic Flux Balance Analysis models (i.e., DFBA), which are static models coupled with ordinary differential equation systems to capture system dynamics. DFBA models are largely prone to discontinuities, which precludes the use of UQ methods on this type of models. A novel approach that is able to perform efficient parameter estimation in DFBA models is introduced in this chapter. The contributions of Chapter 2 have been submitted as [27].

Chapter 3 presents a method for OED for model structure selection based on the construction of the full distribution of model predictions. The efficacy of the method is demonstrated in a model of JAK2/STAT5 signaling pathway in the presence of intrinsic population noise. Chapter 4, conversely, introduces two OED methods for model structure selection that rely on the statistics of model predictions, instead of their full construction. This allows for their implementation in an online manner (i.e., as the experiment is taking place). The efficacy of these OED methods are demonstrated in so-called fault diagnosis on a continuous bioreactor, where models of faulty and normal operation are selected to identify the state of the system at any given time. The contributions of Chapters 3 and 4 have been published in [21, 22, 26].

Chapter 5 presents a novel Bayesian OED method for parameter estimation in nonlinear dynamic systems, which allows for the use of prior information of parameter estimates. The Bayesian OED framework enables design of experiments that cause the most change from prior to posterior parameter estimates (i.e., the most informative experiment designs). The contributions of Chapter 5 have been published in [28].

Finally, in Chapter 6, we lay out our perspectives on future work that can be performed in the field. We anticipate that the frameworks presented in this thesis will pave the way for real-time implementation of Bayesian OED methods for model structure selection and parameter estimation.

1.7 Bibliography

1. J. Albeck, J. Burke, B. Aldridge, M. Zhang, D. Lauffenburger, and P. Sorger. Quantitative analysis of pathways controlling extrinsic apoptosis in single cells. *Molecular Cell*, 30(1):11–25, 2008.
2. S. Asprey and S. Macchietto. Designing robust optimal dynamic experiments. *Journal of Process Control*, 12:545–556, 2002.
3. E. Balsa-Canto, M. Peifer, J. Banga, J. Timmer, and C. Fleck. Hybrid optimization method with general switching strategy for parameter estimation. *BMC Systems Biology*, 2(26), 2008.
4. S. Bandara, J. Schloder, R. Eils, H. Bock, and T. Meyer. Optimal experimental design for parameter estimation of a cell signaling model. *PLoS Computational Biology*, 2009.

5. A. Barabasi and Z. Oltvai. Network biology: understanding the cell's functional organization. *Nature Reviews Genetics*, 5:101–113, 2004.

6. L. Blackmore and B. Williams. Finite horizon control design for optimal discrimination between several models. In *Proceedings of the* 45^{th} *IEEE Conference on Decision and Control*, pages 1147–1152. San Diego, 2006.

7. T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13:21–27, 1967.

8. G. Curien, O. Bastien, M. Robert-Genthon, A. Cornish-Bowden, M. Cardenas, and R. Dumas. Understanding the regulation of aspartate metabolism using a model based on measured kinetic parameters. *Molecular Systems Biology*, 5(1), 2009.

9. M. Fujiwara, Z. Nagy, J. Chew, and R. Braatz. First-principles and direct design approaches for the control fo pharmaceutical crystallization. *Journal of Process Control*, 15(5):493–504, 2005.

10. D. Gillespie. A rigorous derivation of the chemical master equation. *Statistical mechanics and its applications*, 188(1-3):404–425, 1992.

11. A. Horchler. Sdetools toolbox, matlab. *GitHub*, 2011.

12. S. Huang. Non-genetic heterogeneity of cells in development: More than just noise. *Development*, 136:3853–3862, 2009.

13. D. Huh and J. Paulsson. Random partitioning of molecules at cell division. *Proceedings of the National Academy of Sciences*, 108(36):15004–15009, 2011.

14. H. Kitano. Computational systems biology. *Nature*, 420:206–210, 2002.

15. A. Kremling. Bringing together models from bottom-up and top-down approaches: an application for growth of Escherichia coli on different carbohydrates. *Advanced Experimental Medical Biology*, 736:579–595, 2012.

16. C. Kreutz and J. Timmer. Systems biology: Experimental design. *FEBS Journal*, 276:923–942, 2009.

17. D. Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956.

18. H. MacAdams and A. Arkin. Stochastic mechanisms in gene expression. *Proceedings of the National Academy of Sciences*, 94(3):814–819, 1998.

19. R. Mahadevan, J. Edwards, and F. Doyle. Dynamic flux balance analysis of diauxic growth in Escherichia coli. *Biophysics Journal*, 83:1331–1340, 2002.

20. N. Mantzaris. From single-cell genetic architecture to cell population dynamics: quantitatively decomposing the effects of different population heterogeneity sources for a genetic network with positive feedback architecture. *Biophysics Journal*, 92(12):4271–4288, 2007.

21. M. Martin-Casas and A. Mesbah. Discrimination between competing model structures of biological systems in the presence of population heterogeneity. *IEEE Life Sciences Letters*, 2:23–26, 2016.

22. M. Martin-Casas and A. Mesbah. Active fault diagnosis for stochastic nonlinear systems: online probabilistic model discrimination. *10th IFAC International Symposium on Advanced Control of Chemical Processes*, pages 696–701, 2018.

23. B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The journal of chemical physics*, 124, 2006.

24. H. Najm. Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual review of fluid mechanics*, 41:35–52, 2009.

25. J. Orth, I. Thiele, and B. Palsson. What is flux balance analysis? *Nature Biotechnology*, 28(3):245–248, 2010.

26. J. Paulson, M. Martin-Casas, and A. Mesbah. Input design for online fault diagnosis of nonlinear systems with stochastic uncertainty. *Industrial Engineering and Chemistry Research*, 56:9593–9605, 2017.

27. J. Paulson, M. Martin-Casas, and A. Mesbah. Black-box surrogate modeling of biological systems with non-smooth behavior: Applications in bayesian parameter inference for dynamic flux balance analysis models. *Submitted*, 2019.

28. J. Paulson, M. Martin-Casas, and A. Mesbah. Optimal bayesian experiment design for nonlinear dynamic systems with chance constraints. *Journal of Process Control*, 2019.

29. K. Schallau and B. Junker. Simulating plant metabolic pathways with enzyme-kinetic models. *Plant Physiology*, 152(4):1763–1771, 2010.

30. D. Stockholm, R. Benchaouir, J. Picot, P. Rameau, T. Nelidez, G. landini, C. Laplace-Builhe, and A. Paldi. The origin of phenotpic heterogeneity in a clonal cell population in vitro. *PLoS ONE*, 4:e394, 2007.

31. T. Stuart and R. Satija. Integrative single-cell analysis. *Nature Reviews Genetics*, 2019.

32. T. Szekely and K. Burrage. Stochastic simulation in systems biology. *Computational and Structural Biotechnology Journal*, 12:14–25, 2014.

33. L. Tsimring. Noise in biology. *Reports on Progress in Physic*, 77, 2014.

34. D. Wilkinson. Stochastic modelling for quantitative description of heterogeneous biological systems. *Nature Reviews*, 10:122–133, 2009.

35. Y. Zhang, P. Song, Q. Fu, M. Ruan, and W. Xu. Single-molecule chemical reaction reveals molecular reaction kinetics and dynamics. *Nature communications*, 5, 2014.

Chapter 2

Estimation of parameters in genome-scale biological models

Modeling of large-scale biological systems with thousands of biochemical reactions and metabolites (e.g., whole-cell models) can be prohibitively expensive. This is particularly true when performing computationally intensive procedures, such as parameter estimation from noisy data (i.e., inverse uncertainty quantification), or quantifying the error or heterogeneity in model predictions (i.e., forward uncertainty quantification). In this chapter, we introduce a surrogate modeling approach that enables faster forward and inverse uncertainty quantification methods in genome-scale biological models. The computational efficiency of the surrogate modeling approach enables the performance of Bayesian estimation of parameters in genome-scale metabolic networks for the first time.[*]

2.1 Introduction

The utility of mathematical modeling in biology is on the rise due to computational advancements and the increasing availability of data provided by high-throughput experimental techniques [14]. Mathematical models that capture complex features of cellular behavior, such as whole-cell scale interaction of cellular constituents and probabilistic behavior, can be used to gain mechanistic insight into biological systems by enabling the performance of model parameter estimation from available experimental data. However, parameter estimation procedures can be prohibitively expensive to perform, especially if the employed mathematical model has large dimensionality.

Flux Balance Analysis (FBA) models are an example of a class of models that can capture whole-cell dynamics. FBA is widely used for modeling

---

[*] The content of this chapter was published in [49].

cellular metabolism in a large range of metabolic and biochemical engineering problems [47, 48]. Given a constrained metabolic network, FBA assumes the intracellular fluxes are regulated by the cell to optimize a predefined cellular objective function (e.g., maximizing the biomass growth rate [65]) subject to mass balances of the intracellular metabolites and other feasibility constraints (e.g., bounds on the substrate uptake and product secretion rates). However, FBA only identifies metabolic flux distributions at steady-state and, thus, provides no information on metabolite concentrations or the dynamic behavior of the fluxes. A dynamic extension to FBA, commonly referred to as dynamic FBA (DFBA), was originally developed in [39] and has been subsequently applied in several engineering applications [27, 29, 36, 44]. In DFBA models, the intracellular fluxes are given by the solution of a FBA model, which is coupled to a set of dynamic equations that describe the time-varying nature of the extracellular substrate and product concentrations as a function of the extracellular environment [25]. The key assumption in DFBA is that the intracellular fluxes equilibrate instantaneously such that the sizes of the intracellular metabolite pools remain unchanged. This ”quasi steady-state” assumption is valid as long as the intracellular dynamics are significantly faster than the extracellular dynamics.

Prediction of the behavior of biological systems, such as those described by DFBA models, can be subject to various sources of uncertainty including: (i) unknown model parameters; (ii) unknown model structure; and (iii) experimental uncertainty such as measurement error [34]. Accurate characterization of these uncertainties, as well as their impact on the quality of model predictions, is vital when applying these models in decision-support or optimization tasks such as parameter estimation and optimal experiment design. The task of uncertainty quantification (UQ) can be divided into two major problems: *forward* uncertainty propagation and *inverse* uncertainty analysis. The forward problem focuses on propagating all uncertainties through the model to predict the overall uncertainty in the outputs, whereas the inverse problem aims to calibrate the model with experimental data [23, 61, 70]. However, the most commonly used UQ methods are intractable for expensive-to-evaluate computational models [43, 62], which has severely limited their application to DFBA models. For example, the inverse problem has been tackled in [33] using a variant of maximum likelihood estimation (MLE) to estimate parameters of a yeast DFBA model. The MLE method for DFBA models is formulated as a bilevel optimization that is challenging to solve and, thus, is limited to small metabolic network sizes [59]. To give an idea of scale, this reformulation procedure was recently applied to *in-silico* optimization of the production of recombinant proteins in *Pichia pastoris* using a simplified network with 37 metabolites and 47 reactions [19], which is far from the available genome-scale model (e.g., iPP668) that consists of 1361 reactions and 1177 metabolites [12].

In this work, we present a novel surrogate modeling approach that enables forward and inverse UQ methods to be executed significantly faster on genome-scale DFBA models. We assume that the underlying model structure is accurate, meaning that the uncertain inputs in this work are model parameters or initial conditions (referred to as "parameters" for short). A variety of surrogate modeling techniques have been proposed and used in the context of UQ [4, 46]. We focus on polynomial chaos expansions (PCEs) due to their ability to closely match the exact model response using limited computational resources in a variety of engineering applications [5, 52, 69, 71]. In the context of biological systems, PCE has been applied to both parameter estimation [56] and model discrimination [42]. A prerequisite in PCE is that the model response be a smooth function of the parameters, such that it can be accurately approximated by a collection of polynomials. Whenever the model has non-smooth behavior, the PCE may converge very slowly or even fail to converge altogether [66, 67]. This is a key challenge in DFBA models, as they become singular (i.e., loss of differentiability; a type of non-smoothness) at certain time points due to the underlying quasi steady-state assumption [25, 28, 30]. Thus, we aim to develop an extension of PCE that can adequately capture the non-smooth character of DFBA models. Inspired by [67], we propose a multi-element approach that decomposes the parameter space into a collection of non-overlapping regions so that the model response behaves smoothly within each chosen region. We also present a simple procedure for selecting the elements based on our ability to detect singularities when integrating the DFBA model. Separate PCE models can then be constructed in each of these elements, such that the overall model response is approximated by a piecewise polynomial function.

We adopt a regression-based approach for constructing the PCEs from a finite set of DFBA model evaluations [10, 18, 60]. Consequently, the model can be treated as a black-box, i.e., the original DFBA solver need not be modified in any way and its solution can be directly obtained from readily available toolboxes. Although any regression method can be used in principle, we focus on sparse regression methods, as they are particularly advantageous in problems with a large number of parameters [5, 56]. Sparse regression has the ability to systematically locate the terms in the PCE that have the greatest impact on the model response, meaning they avoid heuristic exclusion of terms in the expansion. This is important in DFBA models since they often involve many parameters due to their multi-scale nature, i.e., the fact that they capture both intracellular and extracellular effects.

To demonstrate the effectiveness of the proposed surrogate modeling framework, we apply it to accelerate Bayesian inference of parameters in the substrate uptake kinetic expressions of diauxic growth of a batch monoculture of *Escherichia coli* on glucose/xylose mixed media. The metabolic network reconstruction used for *E. coli* is iJ904, which is a genome-scale model that contains 1075 reactions and 761 metabolites [55]. Parameter estimation is performed using measurements of the concentrations of extracellular metabolites and biomass that are taken at certain time points throughout the batch. We selected this particular example, not only because experimental data is available in [17], but also due to the fact that reported parameter estimates were determined using a trial-and-error procedure [27], likely due to the complexity of the genome-scale model. Thus, one of our main objectives is to formalize this exercise of estimating parameters in DFBA models using a very general Bayesian framework. We focus on Bayesian methods because they have several advantages over MLE including: (i) avoid the need to solve a large-scale bilevel optimization problem; (ii) provide an explicit representation of the uncertainty in the estimates in terms of a probability distribution, as opposed to a single point estimate; and (iii) avoid numerical challenges due to insensitive parameters or complex measurement noise models. We also demonstrate how the proposed multi-element PCE method can be applied to forward UQ problems such as global sensitivity analysis and estimation of the probability distribution of the system outputs.

2.2 Methods

In this section, we introduce the model class of interest, summarize our proposed multi-element polynomial chaos surrogate modeling approach, and discuss some key aspects of the numerical implementation.

2.2.1 Dynamic flux balance analysis (DFBA)

We are interested in modeling a microbial cultivation process using DFBA, in which the bioreactor is viewed as a combination of the fluid medium (extracellular environment) and the microorganisms (intracellular environment). Cell walls act as physical boundaries between these two phases, through which certain chemical metabolites are exchanged. The DFBA model can then be mathematically formulated as [30]:

$$\dot{\mathbf{q}}(t) = \mathbf{f}(t, \mathbf{q}(t), \mathbf{v}(\mathbf{q}(t))), \quad \mathbf{q}(t_0) = \mathbf{q}_0, \tag{2.1}$$

with $\mathbf{v}(\mathbf{q}(t))$ being an element of the solution set of the flux balance model

$$\begin{aligned} \mathbf{v}(\mathbf{q}) \in \operatorname{argmax}_{\mathbf{v}} \quad & h(\mathbf{v}, \mathbf{q}) \\ \text{s.t.} \quad & \mathbf{S}\mathbf{v} = \mathbf{0}, \\ & \mathbf{v}^{\mathrm{LB}}(\mathbf{q}) \leq \mathbf{v} \leq \mathbf{v}^{\mathrm{UB}}(\mathbf{q}), \end{aligned} \tag{2.2}$$

where $\mathbf{q}$ denotes to the state variables describing the extracellular environment (e.g., concentrations of substrates, biomass, and products) with time derivative $\dot{\mathbf{q}}$ and initial conditions $\mathbf{q}_0$; $\mathbf{v}$ denotes the metabolic fluxes that include both intracellular fluxes and exchange rates; $\mathbf{S}$ is the stoichiometric matrix of the metabolic network; and $\mathbf{v}^{\mathrm{LB}}(\mathbf{q})$ and $\mathbf{v}^{\mathrm{UB}}(\mathbf{q})$ are the lower and upper bounds on the fluxes, respectively, which are functions of the extracellular concentrations. The vector function $\mathbf{f}$, specified by the set mass balances in the extracellular medium, defines the rate of change of each component of $\mathbf{z}_0$ and must be integrated to determine the concentration profiles over time. The scalar function h is the cellular objective that is maximized by the cells. Whenever more than one microbial species is present in the culture, then multiple flux balance models of the form (2.2) must be incorporated into (2.1), as shown in [25].

DFBA models can be classified as ordinary differential equations with embedded optimization (ODEO) wherein the lower-level FBA optimization can either be a linear program (LP) or nonlinear program (NLP). The vast majority of DFBA simulation has focused on LP objectives, though NLP objectives were recently tackled in [72]. DFBA simulation strategies can be broadly categorized as: (i) static optimization approach (SOA) [39], (ii) dynamic optimization approach (DOA) [39], (iii) direct approach (DA) [30], and (iv) interior point reformulation (IPR) [59]. In SOA, a forward Euler scheme is used to integrate (2.1) while the FBA (2.2) is solved at each time step using a suitable solver. SOA is known to be inefficient, especially for stiff problems that require small time steps to ensure convergence, due to the fact that (2.2) must be repeatedly solved. DOA, on the other hand, discretizes the time horizon and then converts the DFBA model into a NLP problem; however, this approach cannot be easily applied to genome-scale metabolic networks due to the large number of variables and constraints that are introduced. DA directly includes the solver for the FBA (2.2) in the right hand side evaulator $\mathbf{f}$. Thus, DA can take advantage of implicit ODE integrators that employ adaptive step sizes and error control, which reduces the number of integration steps when compared to SOA. The IPR method is based on the fact that the optimization problem in ODEO models can be replaced by their Karush-Kuhn-Tucker (KKT) optimality conditions, which can subsequently be relaxed using logarithmic barrier functions so that the DFBA model can be transformed into a set of implicit ODEs. However, IPR introduces a new source of error since the relaxed solution only converges to the true optimal solution as the barrier parameter approaches zero.

Although any one of these methods can be used to integrate DFBA models in this work, the DA approach has the ability to provide highly accurate solutions in a reasonably efficient manner. Another advantage of DA is that the solution set to the FBA (2.2) can be made unique through the use of lexicographic optimization [25, 28], which avoids numerical issues that are known to plague other DFBA simulation methods (e.g., see [24, Chapter 3]). Since DA requires continuous monitoring and identification of any active set changes in (2.2), it constitutes a dynamic simulation involving discrete events (i.e., a hybrid system).

Next, we propose a novel surrogate modeling approach that can systematically handle the hybrid nature of DFBA models.

2.2.2 Sparse multi-element polynomial chaos

We assume that the DFBA model involves a set of $M \geq 1$ input parameters, denoted by $\boldsymbol{x} = (x_1, \ldots, x_M)$, that can appear in the initial conditions $\mathbf{z}_0$, rate of change function $\mathbf{f}$, cellular objective h, and/or the flux limits $\mathbf{v}^{\mathrm{LB}}, \mathbf{v}^{\mathrm{UB}}$. We are interested in an output of the DFBA model (commonly referred to as the *model response*) that can be any desired function of the states or fluxes including, for example, the metabolite or biomass concentrations at specific time points or even post-processed quantities such as the time that a certain metabolite is consumed. The output can be represented by some function $y = \mathcal{M}(\boldsymbol{x})$, which may not be known analytically, but can be evaluated by integrating (2.1)-(2.2) using previously discussed approaches. An exact value for y can be obtained when the parameters $\boldsymbol{x}$ are perfectly known; otherwise, these outputs must be treated as uncertain. Parametric uncertainty can be generally represented by a random vector $\boldsymbol{X}$ with some known probability density function (PDF) $f_{\boldsymbol{X}}(\mathbf{x})$, implying the model response is also a random variable $Y = \mathcal{M}(\boldsymbol{X})$. Although we focus on a scalar output for notational clarity, the developed procedure can easily be applied separately to each component of a random vector $\boldsymbol{Y}$.

Since $\mathcal{M}$ can be computationally expensive to evaulate, we aim to construct a polynomial chaos (PC) *surrogate model* that closely approximates the function $\mathcal{M}$, but is much cheaper to evaluate for any given parameter value. In the PC method, the model response is expanded in terms of a set of multivariate polynomials that are orthonormal with respect to $f_{\boldsymbol{X}}$. PC expansions (PCEs) are known to be efficient and accurate approximations to the model response whenever $\mathcal{M}$ is sufficiently smooth, but can converge very slowly (or even fail to converge) when singularities or discontinuities are present [66]. This issue is inherent to DFBA models because they are defined in terms of non-smooth functions such as (2.2). Inspired by [67], we instead look to develop a *multi-element* (ME) extension of standard PCE. Let $S \subseteq \mathbb{R}^M$ denote the support of the random parameters, such that $\boldsymbol{X} \in S$. The idea behind ME-PCE is then to decompose S into a finite number N_e of non-overlapping elements

$$S = \bigcup_{k=1}^{N_e} S_k, \quad S_{k_1} \cap S_{k_2} = \emptyset, \quad k_1 \neq k_2. \tag{2.3}$$

Based on this decomposition, we define the following indicator random variables

$$I_{S_k}(\boldsymbol{X}) = \begin{cases} 1 & \text{if } \boldsymbol{X} \in S_k \\ 0 & \text{otherwise.} \end{cases} \qquad k = 1, \ldots, N_e, \tag{2.4}$$

as well as a new set of conditional random variables $\boldsymbol{X}_k = \boldsymbol{X}|(I_{S_k}(\boldsymbol{X}) = 1)$ with PDF

$$f_{\boldsymbol{X}_k}(\boldsymbol{x}_k) = \frac{f_{\boldsymbol{X}}(\boldsymbol{x}_k)}{\Pr(I_{S_k}(\boldsymbol{X}) = 1)} = \frac{f_{\boldsymbol{X}}(\boldsymbol{x}_k)}{\int_{S_k} f_{\boldsymbol{X}}(\boldsymbol{x})d\boldsymbol{x}}. \tag{2.5}$$

We look to construct local PCEs $\mathcal{M}_k^p(\boldsymbol{X}_k)$ in element S_k using the locally orthonormal basis $\phi_{k,\boldsymbol{\alpha}}(\boldsymbol{X}_k)$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_M) \in \mathbb{N}^M$ is a multi-index truncated at $|\boldsymbol{\alpha}| = \alpha_1 + \cdots + \alpha_M \leq p$ that represents the total order of the polynomials. The ME-PCE approximation $\mathcal{M}^p(\boldsymbol{X})$ over the entire random field can then be defined as

$$\mathcal{M}^p(\boldsymbol{X}) = \sum_{k=1}^{N_e} \mathcal{M}_k^p(\boldsymbol{X}) I_{S_k}(\boldsymbol{X}) = \sum_{k=1}^{N_e} \sum_{|\boldsymbol{\alpha}| \leq p} y_{k,\boldsymbol{\alpha}} \phi_{k,\boldsymbol{\alpha}}(\boldsymbol{X}) I_{S_k}(\boldsymbol{X}), \tag{2.6}$$

where $y_{k,\boldsymbol{\alpha}} \in \mathbb{R}$ are the coefficients of the expansion in the k^{th} element. Let $L^2_{\boldsymbol{X}}$ be the space of all real-valued, square integrable functions with respect to $f_{\boldsymbol{X}}$. Whenever Y has finite variance (i.e., $\mathcal{M} \in L^2_{\boldsymbol{X}}$) and the support S is compact, the series (2.6) converges to the $\mathcal{M}(\boldsymbol{X})$ in the $L^2_{\boldsymbol{X}}$ sense [67]

$$\begin{aligned} &\lim_{p\to\infty} \mathbb{E}\left\{(\mathcal{M}(\boldsymbol{X}) - \mathcal{M}^p(\boldsymbol{X}))^2\right\} \qquad (2.7)\\ &= \sum_{k=1}^{N_e} \Pr(I_{S_k}(\boldsymbol{X}) = 1) \lim_{p\to\infty} \int_{S_k} \left(\mathcal{M}(\boldsymbol{x}_k) - \mathcal{M}_k^p(\boldsymbol{x}_k)\right)^2 f_{\boldsymbol{X}_k}(\boldsymbol{x}_k) d\boldsymbol{x}_k = 0, \end{aligned}$$

where the equality above follows from the law of total expectation. In the case of a single element $N_e = 1$, the ME-PCE reduces to a standard *global* PCE. After selecting a decomposition of the form (2.3), we must perform the following steps in each of the elements: (i) construct the basis functions and (ii) compute the expansion coefficients. An overview of the calculations involved in these steps is provided next.

The local basis polynomials are assumed to be orthonormal with respect to the conditional PDF (2.5). That is,

$$\int_{S_k} \phi_{k,\boldsymbol{\alpha}}(\boldsymbol{x}_k) \phi_{k,\boldsymbol{\beta}}(\boldsymbol{x}_k) f_{\boldsymbol{X}_k}(\boldsymbol{x}_k) d\boldsymbol{x}_k = \delta_{\boldsymbol{\alpha}\boldsymbol{\beta}}, \tag{2.8}$$

for all $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{N}^M$ where $\delta_{\boldsymbol{\alpha}\boldsymbol{\beta}} = 1$ if $\boldsymbol{\alpha} = \boldsymbol{\beta}$ and 0 otherwise. The complexity of determining the polynomials $\{\phi_{k,\boldsymbol{\alpha}}\}_{|\boldsymbol{\alpha}| \leq p}$ depends fully on the structure of $f_{\boldsymbol{X}_k}$. Whenever the uncertain parameters are statistically independent, i.e.,

$$f_{\boldsymbol{X}_k}(\boldsymbol{x}_k) = \prod_{i=1}^{M} f_{X_{k,i}}(x_{k,i}), \quad S_k = [a_{k,1}, b_{k,1}) \times \cdots \times [a_{k,M}, b_{k,M}), \tag{2.9}$$

where $f_{X_{k,i}}$ denotes the marginal PDF of the i^{th} element of $\boldsymbol{X}_k$, then (2.8) reduces to simply the tensor product of M one-dimensional polynomials that are orthonormal with respect to $f_{X_{k,i}}$. These polynomials are known for certain distributions, which is often referred to as *generalized polynomial chaos* (gPC) [71], and can be generated numerically for random variables with arbitrary marginal distributions using readily available algorithms that are based on three-term recurrence relationships [22]. The most general case of $\boldsymbol{X}_k$ having dependent elements can be tackled in one of the following two ways. The orthonormality conditions (2.8) can be simultaneously imposed in M dimensions using numerical approaches presented in, e.g., [20, 50, 54], which is often referred to as *arbitrary polynomial chaos* (aPC). Alternatively, the dependent random vector $\boldsymbol{X}_k$ can be transformed into a new standard random vector $\boldsymbol{Z}_k$ with independent elements. As long as this is an *isoprobabilistic transformation* that preserves the PDFs of the random vectors, then gPC can be applied to the transformed function $\mathcal{M}(\boldsymbol{X}_k) = \mathcal{M}(\mathcal{T}_k(\boldsymbol{Z}_k))$ where $\mathcal{T}_k$ is most commonly chosen as the Rosenblatt transformation [57]. In this work, we focus on the Rosenblatt transformation approach to ensure that the proposed method can be easily implemented using readily available toolboxes. It should be noted that the convergence rate can drop when $\mathcal{T}_k$ is highly nonlinear, and aPC would be preferred in these situations (see [53] for an example comparing aPC to gPC); however, this was not an issue in the chosen case study.

The second step in constructing the ME-PCE is to determine the coefficients of the expansion in each element. We are now interested in approximating the composition of the model response with the appropriate transformation, i.e., $g_k = \mathcal{M} \circ \mathcal{T}_k$. For notational simplicity, we drop the subscript k as the following steps are the same within each element. The PCE for $g(\boldsymbol{Z})$ is then stated as

$$g^{\text{PCE}}(\boldsymbol{Z}) = \sum_{\boldsymbol{\alpha} \in \mathcal{A}} \hat{a}_{\boldsymbol{\alpha}} \phi_{\boldsymbol{\alpha}}(\boldsymbol{Z}) = \hat{\boldsymbol{a}}^\top \boldsymbol{\phi}(\boldsymbol{Z}), \tag{2.10}$$

where $\mathcal{A}$ is the set of multi-indices retained in the expansion, and $\hat{\boldsymbol{a}}$ and $\boldsymbol{\phi}$ are, respectively, vectors of the coefficients and polynomial basis functions. Let P denote the total number of terms retained in the series (2.10). To mitigate the exponential growth in P when using a total degree truncation, we define $\mathcal{A}$ using the so-called *hyperbolic* truncation scheme, which selects all multi-indices that satisfy [5]

$$\mathcal{A} = \{\boldsymbol{\alpha} \in \mathbb{N}^M : \|\boldsymbol{\alpha}\|_q \leq p\}, \quad \|\boldsymbol{\alpha}\|_q = \left(\sum_{i=1}^{M} \alpha_i^q \right)^{1/q}, \tag{2.11}$$

where $0 < q \leq 1$. Lower values for q limit the number of high-order interaction terms considered, leading to sparse solutions, while $q = 1$ results in the standard truncation scheme $|\boldsymbol{\alpha}| \leq p$ where $P = \frac{(M+p)!}{M!p!}$. The hat on coefficients

$\hat{\boldsymbol{a}}$ signifies that we can only obtain numerical estimates of the true coefficients $\boldsymbol{a} = \mathbb{E}\{g(\boldsymbol{Z})\boldsymbol{\phi}(\boldsymbol{Z})\}$ in general. These estimates can be obtained using either *intrusive* approaches (e.g., Galerkin projection [23]) or *non-intrusive* approaches (e.g., pseudo-spectral projection [13], regression [62]). The term "non-intrusive" indicates that the coefficients are estimated from only a finite set of parameter realizations $\mathcal{Z} = \{\boldsymbol{z}^{(1)}, \dots, \boldsymbol{z}^{(N)}\}$, referred to as the *experimental design* (ED). These samples can be chosen in various ways including Monte Carlo sampling, quasi-random samples derived from Sobol or Halton sequences, or sparse grids to name a few [60]. The computational model is then evaluated at every point in the ED, i.e., $\mathcal{Y} = \{y^{(1)}, \dots, y^{(N)}\}$ with $y^{(i)} = g(\boldsymbol{z}^{(i)})$ for all $i = 1, \dots, N$. As such, non-intrusive approaches are "black-box" in the sense that they can be applied to any function and do not require any modification to the deterministic solver.

To further sparsify the PCE, without discarding potentially important interaction terms, the coefficients are estimated using the regularized least squares problem

$$\hat{\boldsymbol{a}} = \underset{\boldsymbol{a} \in \mathbb{R}^P}{\operatorname{argmin}} \ \|\mathcal{Y} - \boldsymbol{\Phi}\boldsymbol{a}\|_{\boldsymbol{W}}^2 + \lambda \|\boldsymbol{a}\|_1, \tag{2.12}$$

where $\|\cdot\|_{\boldsymbol{W}}$ is the weighted Euclidean norm, $\boldsymbol{\Phi} \in \mathbb{R}^{N \times P}$ contains the values of all polynomial bases evaluated at all ED points, $\boldsymbol{W} \in \mathbb{R}^{N \times N}$ is a diagonal weight matrix with elements $\boldsymbol{W}(i,i) = w(\boldsymbol{z}^{(i)})$ that are functions of the sample points, and $\|\boldsymbol{a}\|_1 = \sum_{\boldsymbol{\alpha} \in \mathcal{A}} |a_{\boldsymbol{\alpha}}|$ is a regularization term that forces the minimization to favor low-rank solutions. It can be shown that $\lim_{N \to \infty} \hat{\boldsymbol{a}} = \boldsymbol{a}$ converges to the true coefficients by increasing the number of samples in the ED for $\lambda = 0$ [53], which corresponds to the weighted least squares solution $\boldsymbol{\Phi}^\top \boldsymbol{W} \boldsymbol{\Phi} \hat{\boldsymbol{a}} = \boldsymbol{\Phi}^\top \boldsymbol{W} \mathcal{Y}$. Since this problem only has a unique solution for $N \geq P$ (typically $N = cP$ with $c \in [2, 3]$), the regularization term is added to ensure a unique solution exists even when $N < P$. This problem is commonly referred to as lasso (least absolute shrinkage and selection operator) in the fields of statistics and machine learning [63].

Methods for solving (2.12) require a proper choice of λ, which specifies the number of non-zero coefficients, and they usually rely upon rigorous estimation of the PCE approximation error. A natural error choice is the relative mean square error, defined as $\varepsilon = \mathbb{E}\{(g(\boldsymbol{Z}) - g^{\text{PCE}}(\boldsymbol{Z}))^2\}/\text{Var}(g(\boldsymbol{Z}))$, that can be estimated using either the normalized empirical error or cross-validation. Since the normalized empirical error suffers from *overfitting* (i.e., always decreases when the number of terms in the series P increases, even when ε actually increases), cross-validation is preferred as it leads to error estimates that are much less sensitive to overfitting. In cross-validation, the surrogate model is evaluated at additional sample points that were not used when estimating the coefficients of the PCE. The true model output is then compared with the polynomial surrogate at these points to estimate the overall error. In K-fold cross validation, the full ED is randomly partitioned into K equally sized subsets. Of these K subsets, only one is used as validation data, while the remaining $K-1$ subsets are used as training data. This procedure is repeated K times, with each of the K subsets used once as validation data, and the resulting K errors are then averaged to produce a single error estimate. The special case of $K = N$ is the so-called leave-one-out (LOO) error. In the context of linearly parametrized regression, the LOO error can be analytically calculated by [45]

$$Err_{LOO} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{g(\boldsymbol{z}^{(i)}) - g^{\mathrm{PCE}}(\boldsymbol{z}^{(i)})}{1-h_i}\right), \tag{2.13}$$

where h_i is the i^{th} diagonal term of the matrix $\boldsymbol{\Phi}(\boldsymbol{\Phi}^\top\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^\top$. The relative LOO error, denoted by ε_{LOO}, is then equal to (2.13) divided by the sample-based estimate of the model response variance. Additional correction factors can be included to further reduce the sensitivity of the error estimates to overfitting. A particular choice of corrected relative LOO error that has been illustrated to be a robust and conservative estimator for ε, even for small sample sizes, is given by [9]

$$\varepsilon^{\star}_{LOO} = \varepsilon_{LOO}\frac{N}{N-P}\left(1 + \frac{\mathrm{tr}(\boldsymbol{C}^{-1}_{emp})}{N}\right), \quad \boldsymbol{C}_{emp} = \frac{1}{N}\boldsymbol{\Phi}^\top\boldsymbol{\Phi}. \tag{2.14}$$

In this work, we adopt the hybrid least angle regression (LAR) method to solve (2.12) for the expansion coefficients [5]. LAR is an efficient procedure for variable selection that is aimed at selecting the predictors (i.e., the basis polynomials $\phi_{k,\boldsymbol{\alpha}}$) that have the greatest impact on the model response among a potentially large set of candidates [16]. Hybrid LAR is a variant of the original LAR that uses a modified cross-validation scheme in terms of $\varepsilon^{\star}_{LOO}$. This modification relies on only a single call to the LAR procedure, which provides significant savings in computational cost when compared to the original method. A potential limitation of hybrid LAR is that the truncation set $\mathcal{A}$ must be fixed *a priori*, which can be mitigated by employing a *basis-adaptive* algorithm. The basis-adaptive strategy involves fitting a sparse PCE for multiple orders and then selecting the one with the lowest error. In this way, the maximum degree is driven by the data directly, as opposed to being heuristically chosen by the user.

The efficiency and accuracy of the ME-PCE surrogate strongly depends on the chosen decomposition (2.3). It is desired to decompose the support such that the model response behaves smoothly in every element, meaning the best decomposition will depend on the type of singularities present in the function. In DFBA models, the singularities represent points in time that the extracellular state variables are not differentiable, and often arise whenever the solution to the FBA (2.2) changes abruptly. Note that these abrupt changes correspond to a switch in the intracellular behavior and may occur, for example, when a metabolite is fully consumed. Since these singularities can be related back to the feasibility of (2.2), they can be straightforwardly detected during integration of the DFBA model equations. However, the time of occurrence of any singularity depends on the value of the parameters, which are uncertain in this work. Therefore, we propose to decompose the support S into two elements S_1 and S_2 that denote, respectively, the set of parameters for which the singularity has not and has occurred. This idea is best illustrated through a simple example.

Consider the following non-smooth ODE system, $\dot{y} = -x$ if $y > 0$ and $\dot{y} = 0$ otherwise, with initial condition $y_0 > 0$, whose solution is given by

$$y(t,x) = \begin{cases} y_0 - tx, & \text{if } y_0 > tx, \\ 0, & \text{otherwise.} \end{cases} \tag{2.15}$$

This function is not differentiable at time $t_s(x) = y_0/x$, which can be thought of as the "singularity manifold" in the parameter space $x \in S$. In other words, this represents the boundary that separates S_1 and S_2. For any given time t, the two elements are given by

$$S_1(t) = \{x \in S : t_s(x) > t\}, \quad S_2(t) = \{x \in S : t_s(x) \le t\}. \tag{2.16}$$

The same decomposition principle applies in multi-dimensional DFBA problems with a couple of caveats. The first issue is that we cannot rely on an analytic solution for $t_s(\boldsymbol{x})$ since DFBA models are most often only solvable numerically. As long as the singularity manifold is sufficiently smooth in the parameters, we can straightforwardly construct a global PCE to represent this function using the tools discussed above. The second issue relates to the fact that S_1 and S_2 are not guaranteed to have simple geometrical representations. Therefore, it can be advantageous to further break down these elements into a set of regions with more standard shapes.

A flowchart that summarizes the main steps required to construct the proposed black-box ME-PCE for DFBA models is shown in Figure 2.1. The ME-PCE approach produces the following piecewise polynomial surrogate model for $\mathcal{M}$

$$\mathcal{M}(\boldsymbol{x}) \approx \mathcal{M}^{\text{ME-PCE}}(\boldsymbol{x}) = \sum_{k=1}^{N_e} g_k^{\text{PCE}}(\mathcal{T}_k^{-1}(\boldsymbol{x})) I_{S_k}(\boldsymbol{x}), \tag{2.17}$$

where $\mathcal{T}_k^{-1}$ denotes the inverse of the isoprobabilistic transform $\mathcal{T}_k$. The steps for evaluating (2.17) for any parameter $\boldsymbol{x} \in S$ can then be summarized as: find the element k in which the parameter lies $\boldsymbol{x} \in S_k$; calculate the standardized parameter values $\boldsymbol{z}_k = \mathcal{T}_k^{-1}(\boldsymbol{x})$; and evaluate the corresponding PCE surrogate model at the standardized point $g_k^{\text{PCE}}(\boldsymbol{z}_k)$. By performing these steps on a collection of Monte

Carlo samples, we can directly approximate statistical properties of the model response $Y = \mathcal{M}(\boldsymbol{X})$ including moments, parametric sensitivities, or even its full distribution.

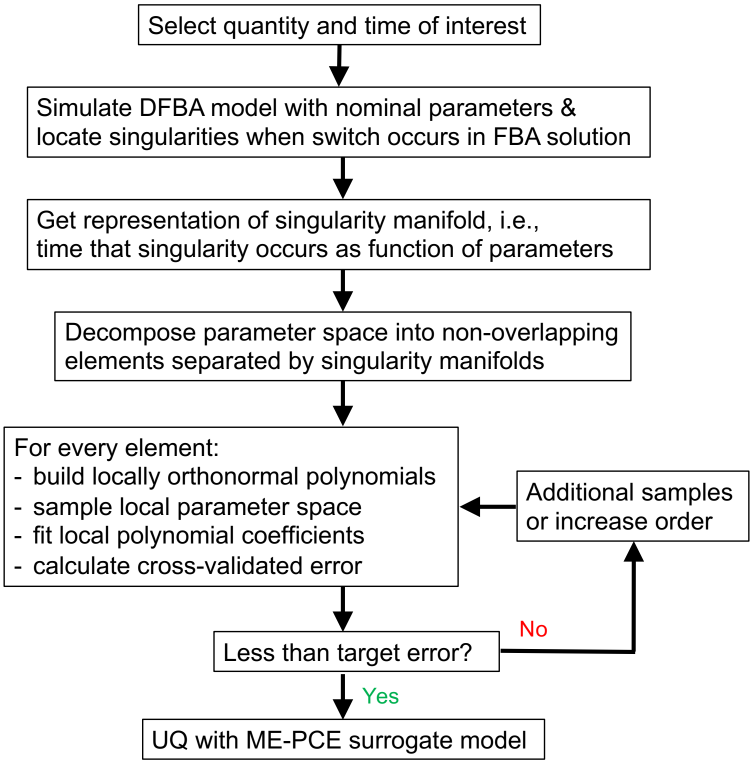

Figure 2.1: Flowchart for the proposed ME-PCE surrogate modeling approach. The quantity and time of interest correspond to a scalar output of the DFBA model and, in general, the decomposition of the parameter space depends on these values. When an exact solution is not available for the singularity manifold, it can be approximated with a global PCE. The coefficients in each element can be fit using a regularized least squares method (e.g., hybrid LAR) and should be made adaptive (w.r.t. samples and/or order) to ensure that a target error level is achieved. The fitted piecewise polynomial surrogate model can then be used for any desired UQ task.

2.2.3 Numerical implementation

The implementation of both DFBA model integration and construction of sparse PCEs are non-trivial and error-prone. To help make the proposed method more easily usable, the overall approach was constructed to be compatible with readily available toolboxes. The simulation of DFBA models can be done with a variety of toolboxes including COBRA [3], ORCA [40] and DFBAlab [25], which can all be executed in MATLAB. We use DFBAlab throughout this work, which is a DA solver that performs lexicographic optimization to render unique exchange fluxes, because it exhibits important numerical advantages when compared to the readily available alternatives (see [24, 28] for more details). The sparse PCE models are estimated using UQLab [41], which includes implementations of the state-of-the-art methods discussed above. UQLab also allows for multiple PCEs with different input distributions to be fit and used simultaneously, such that it is ideal for dealing with many elements at once.

2.3 Results

The case study is based on a DFBA model of a batch fermentation reactor consisting of an *E. coli* monoculture, which has been investigated for the production of valuable chemicals such as ethanol. Here, the simulation represents the initial phase of batch operation of the *E. coli* fermentation reactor under aerobic growth on glucose and xylose mixed media [27]. No ethanol production is observed under aerobic conditions (i.e., this phase is mainly used to increase the biomass), such that the concentration of ethanol can be omitted from the dynamics. This problem is commonly used as a benchmark for comparing DFBA solvers (see, e.g., [24, 28, 59]), as it exhibits stiff dynamics and multiple singularities.

The dynamic mass balance equations for the extracelluar environment of the form (2.1) can be summarized as follows:

$$\begin{aligned}
\dot{b}(t) &= \mu(t)b(t), \\
\dot{g}(t) &= -u_g(t)b(t), \\
\dot{z}(t) &= -u_z(t)b(t),
\end{aligned} \tag{2.18}$$

where $b(t)$, $g(t)$, and $z(t)$ denotes the biomass, glucose, and xlyose concentrations, respectively, at time t. The uptake kinetics for glucose, xylose, and oxygen are given by Michaelis-Menten kinetics

$$u_g(t) = u_{g,max}\frac{g(t)}{K_g + g(t)}, \tag{2.19}$$

$$u_z(t) = u_{z,max}\frac{z(t)}{K_z + z(t)}\frac{1}{1 + \frac{g(t)}{K_{ig}}}, \tag{2.20}$$

$$u_o(t) = u_{o,max}\frac{o(t)}{K_o + o(t)}, \tag{2.21}$$

where $u_{g,max}$, $u_{z,max}$, $u_{o,max}$, K_g, K_z, K_o, and K_{ig} are parameters of the model that correspond to the maximum substrate uptake rates, saturation constants, and inhibition constants. It is assumed that the reactor oxygen concentration, $o(t)$, is controlled and is therefore constant. The growth rate $\mu(t)$, on the other hand, is determined from the metabolic network model of wild-type *E. coli.* The chosen metabolic network reconstruction was iJR904 [55], which contains 1075 reactions and 761 metabolites. The cells are assumed to maximize growth, implying (2.2) is an LP of the form

$$
\begin{aligned}
\mu(t) = \min_{\mathbf{v}} \quad & \mathbf{c}^\top \mathbf{v}, \\
\text{s.t.} \quad & \mathbf{S}\mathbf{v} = \mathbf{0}, \\
& v_{g_{ext}} = u_g(t), \\
& v_{z_{ext}} = u_z(t), \\
& v_{o_{ext}} = u_o(t), \\
& \mathbf{v}^{\mathrm{LB}} \leq \mathbf{v} \leq \mathbf{v}^{\mathrm{UB}},
\end{aligned} \tag{2.22}
$$

where $\mathbf{c}$ is a vector of weights that represent the contribution of each flux to biomass formation while $v_{g_{ext}}$, $v_{z_{ext}}$, and $v_{o_{ext}}$ are, respectively, the exchange fluxes for glucose, xylose, and oxygen (i.e., elements of the flux vector $\mathbf{v}$). Thus, the metabolic network interacts with the extracellular environment through the exchange fluxes in (2.19)-(2.21).

Table 2.1: Parameter estimates obtained from [8]

Parameter	Value	Units
$u_{g,max}$	10.5	mmol/g/hr
K_g	0.0027	g/L
$u_{z,max}$	6	mmol/g/hr
K_z	0.0165	g/L
K_{ig}	0.005	g/L
$u_{o,max}$	15	mmol/g/hr
K_o	0.024	g/L

The initial conditions of the batch are assumed to be fixed at 0.03 g/L of inoculum, 15.5 g/L of glucose, and 8 g/L of xylose, the oxygen concentration is kept constant at 0.24 mmol/L, and $\mathbf{S}$, $\mathbf{c}$, $\mathbf{v}^{\mathrm{LB}}$, and $\mathbf{v}^{\mathrm{UB}}$ are specified by the iJR904 FBA model. However, the parameters in the substrate uptake rates (2.19)-(2.21) should be fit to experimental data since they cannot be easily predicted from first-principles. This problem was partially tackled in [27], where most of the parameters were fixed according to estimates provided in the literature, while $u_{z,max}$ and K_{ig} were adjusted by trial-and-error to match transient measurements of biomass, glucose, and xylose. The reported parameter estimates are given in Table 2.1. Since $o(t)$ is fixed, $u_{o,max}$ and K_o can be lumped into a single parameter u_o. Hence, the six parameters $\boldsymbol{X} = (u_{g,max}, K_g, u_{z,max}, K_z, K_{ig}, u_o)$ are not perfectly known and modeled as a random vector whose elements are independent and uniformly distributed around ±10% of the nominal values. The support S of the uncertain parameters $\boldsymbol{X}$ can then be stated as

$$
\begin{aligned}
S = {} & [9.45{,}11.55]\text{mmol/g/hr} \times [0.0024, 0.0030]\text{g/L} \times [5.4, 6.6]\text{mmol/g/hr} \\
& \times [0.0149, 0.0182]\text{g/L} \times [0.0045, 0.0055]\text{g/L} \times [12.27, 15]\text{1/hr}.
\end{aligned} \tag{2.23}
$$

The problems of inverse and forward UQ are intractable on the full *E. coli* DFBA model due to the genome-scale metabolic network and non-smooth behavior. In the following, we demonstrate how challenging UQ tasks can be enabled with the proposed surrogate modeling approach.

All reported computations are performed in MATLAB R2016a on a MacBook Pro with 8 GB of RAM and a 2.6 GHz Intel i5 processor. The DFBA model is simulated using DFBAlab with default options for integration and LP optimization tolerances. CPLEX was used as the LP solver and MATLAB's ode15s was used as the integrator.

2.3.1 Decomposition of parameter space

Before selecting the element decomposition, we must first simulate the DFBA model to locate any significant singularities. The extracellular glucose, xylose, and biomass concentration profiles are plotted in Figure 2.2 for one hundred randomly sampled parameter values. For a given realization of the parameter, the full simulation requires approximately 1.3 seconds of CPU time.

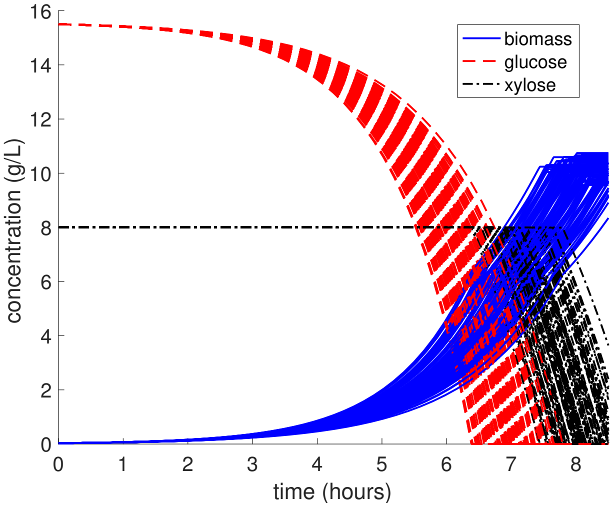

Figure 2.2: Monte Carlo simulation of *E. coli* DFBA model. The genome-scale model is integrated from 0 to 8.5 hours for 100 different parameter realizations that are independently drawn from the uniform prior density. The consumption of xylose only occurs after glucose is fully exhausted, which is a strong function of the parameters.

At the start of the batch, glucose is consumed preferentially over xylose. Once glucose has been depleted, the LP solution switches and xylose becomes the main carbon source. The final batch time is then specified as the time that both glucose and xylose have been fully depleted, at which point the LP becomes infeasible and the solution ceases to exist. The *E. coli* stop growing at this point due to the lack of a carbon source. Although physically the cells would begin to die in this situation, DFBA models cannot directly predict the cell death phase and thus we assume the biomass remains constant for simplicity. The time-to-consumption of glucose t_g and xylose t_z represent the two sources of singularities in this problem, and clearly depend on the value of the model parameters. Since we cannot analytically derive these functions, we look to construct PCE approximations for t_g and t_z instead.

We initially focus on t_g because glucose is consumed first. Since t_g depends smoothly on the parameters, we can fit a global PCE over the entire support (2.23). Two different fitting methods are investigated: classical *full* PCE with coefficients estimated using ordinary least squares, and basis-adaptive *sparse* PCE with coefficients estimated using hybrid LAR. In full PCE, the degree varies from 1 to 6 with $N = 2P$ model evaluations performed where P denotes the basis size. In sparse PCE, the maximum degree of the basis is allowed to vary from 2 to 30 and a hyperbolic truncation scheme (2.11) is used with $q = 0.75$. Both of these computational schemes are carried out using UQLab, and all experimental designs (EDs) are composed of parameter samples drawn from a quasi-random (low-discrepancy) Sobol' sequence. Figure 2.3a plots the RMSE versus the number of model evaluations used to fit the surrogate. We can clearly see that sparse LAR-based PCE consistently outperforms full PCE, achieving more than an order-of-magnitude lower RMSE using about 1000 model evaluations. The resulting sparse polynomial surrogate model is first used to approximate the Sobol' indices of $t_g(\boldsymbol{X})$, which are a commomly used tool in global sensitivity analysis for ranking the random parameters according to their weight in the variance of the model response (Figure 2.3b). The Sobol' indices can be computed *analytically* from the PCE coefficients [62], which requires less than one second of CPU time in this case, and shows that only $u_{g,max}$ and u_o contribute significantly to the variance of $t_g(\boldsymbol{X})$. The surrogate model can also be used to efficiently estimate the density function of $t_g(\boldsymbol{X})$, as plotted in Figure 2.3c. From the set of samples used to estimate this density function, we can determine that t_g ranges from approximately 6.31 to 7.87 hr, suggesting that $g(t; \boldsymbol{X})$ is a non-smooth function of $\boldsymbol{X} \in S$ for any $t \in [6.31, 7.87]$ hr. We can split S into two disjoint regions in a similar fashion to (2.16), which is plotted in Figure 2.3d for the fixed time point $t = 7$ hr. The red regions represent parameters with positive glucose concentrations, while the blue region represents parameters with zero glucose concentrations. Only the projection of the 6d space onto the two sensitive parameters is plotted in Figure 2.3d since we do not need to split the support with respect to the insensitive parameters. As expected, the blue and red regions are approximately equal size because around half of the density of $t_g(\boldsymbol{X})$ is below 7 hours, as shown by the red line in Figure 2.3c.

The same procedure described above for t_g was carried out for t_z, as shown in Figure 2.5. We again see that sparse PCE consistently produces significantly

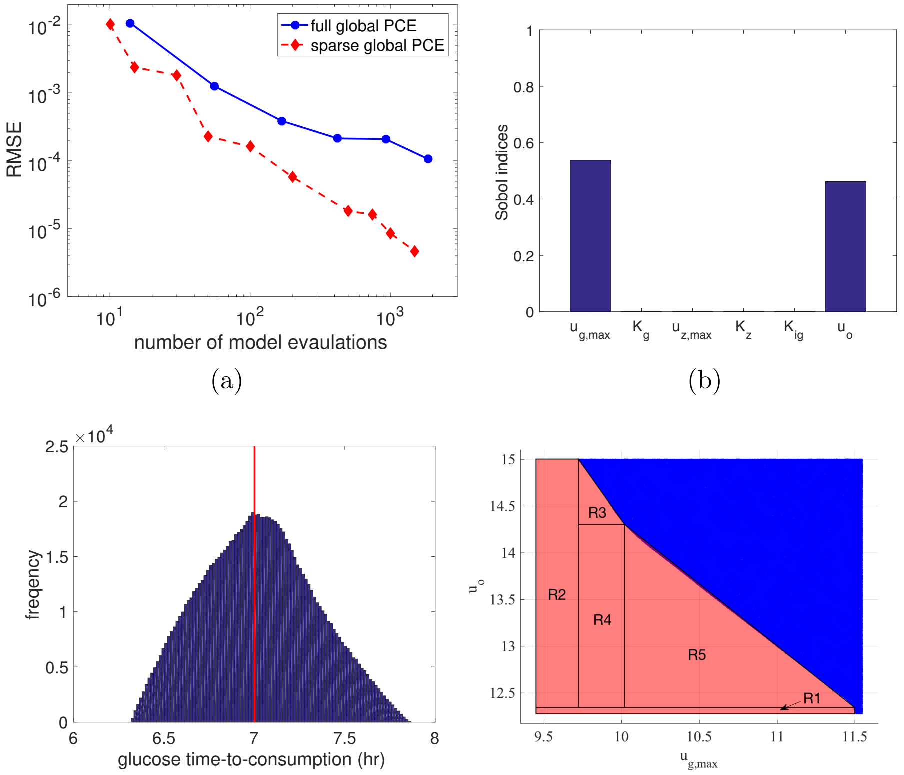

(a) (b)

(c) (d)

Figure 2.3: Parameter space decomposition for glucose. **(a)** Surrogate model validation error for time-to-consumption of glucose t_g versus the number of DFBA evaluations used in model fitting. **(b)** Sensitivity of t_g to the uncertain parameters. Global sensitivity indices are estimated inexpensively using the surrogate model. **(c)** Estimated PDF of t_g based on 1e+6 surrogate model evaluations. **(d)** Decomposition of support of the prior density into six non-overlapping elements at time 7 hr, projected onto the two most sensitive parameters. The blue region is composed of parameters for which glucose is zero, estimated using the surrogate model, while the red region represents parameters with non-zero glucose concentrations.

lower RMSE when compared to full PCE (Figure 2.5a). The global sensitivity analysis on the surrogate model shows that three parameters contribute to the variance of $t_z(\boldsymbol{X})$, mainly $u_{g,max}$, $u_{z,max}$, and u_o (Figure 2.5b). Parameters $u_{g,max}$ and u_o are important for glucose consumption while $u_{z,max}$ is important for xylose consumption. From the estimated density function of $t_z(\boldsymbol{X})$, we observe that its support ranges from approximately 7.33 to 9.12 hr (Figure 2.5c). Before time 7.33 hr, we can use the same element breakdown based on t_g only since the xylose singularity has yet to happen. For times $t \in [7.33, 7.87]$ hr, the supports of $t_g(\boldsymbol{X})$ and $t_z(\boldsymbol{X})$ partially overlap, meaning the xylose singularity must be incorporated into the element breakdown. We thus have three situations that must be represented with separate elements: (i) $\{\boldsymbol{x} \in S : t_g(\boldsymbol{x}) > t\}$ are parameter values for which xylose remains at its initial condition since glucose has yet to be exhausted; (ii) $\{\boldsymbol{x} \in S : t_z(\boldsymbol{x}) > t\}$ are parameter values for which some xylose has been consumed but not fully; and (iii) $\{\boldsymbol{x} \in S : t_z(\boldsymbol{x}) \leq t\}$ are parameter values for which xylose has been fully consumed. An example representation of these regions for 7.75 hr is shown in Figure 2.4.

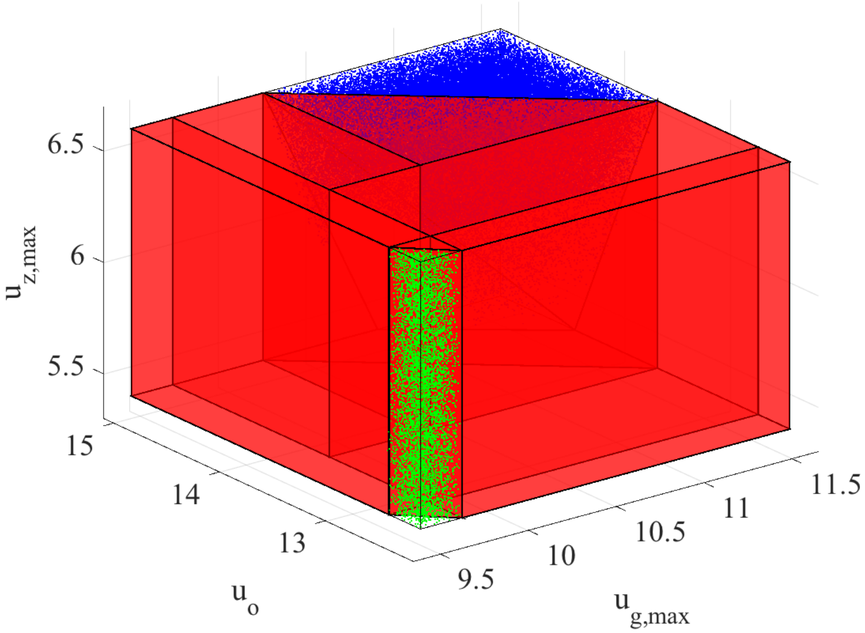

Figure 2.4: Parameter space decomposition for xylose at 7.75 hr. The decomposition is projected onto the three (out of six) most sensitive parameters. The green region represents parameter values for which xylose remains at its initial condition since glucose has yet to be exhausted. The red regions represent parameter values for which some xylose has been consumed but not yet zero. The blue region represents parameter values for which xylose has been fully consumed.

After time 7.87 hr, the first region becomes the empty set and we only have two remaining regions to model. These latter two regions are shown in Fig 2.5d for $t = 8.5$ hr, which depicts the projection of non-zero (red) and zero (blue) xylose values onto the three sensitive parameters. In this case, the blue region is much larger than the red region due to the fact that the majority of the density of $t_z(\boldsymbol{X})$ is below 8.5 hr, as indicated by the red line in Figure 2.5c.

The "true" value of the RMSE that is reported in Figs. 2.3-2.5 was estimated using a large validation set that consists of 20,000 evaluations of the full DFBA model, which requires over 7 hours of CPU time in this case. Ideally, these additional model evaluations could be avoided by directly estimating the RMSE from the ED either empirically or using cross-validation techniques. The empirical estimate of the RMSE is simply based on sample-based approximations to the integral expressions for mean and variance. Cross-validation looks to obtain a more

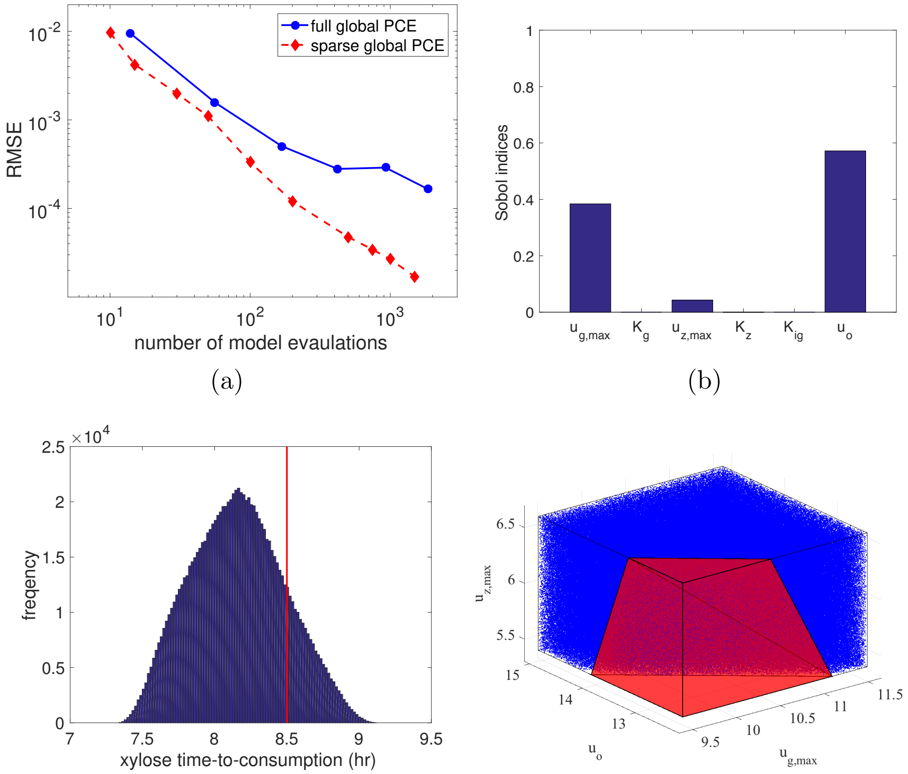

(a) (b)

(c) (d)

Figure 2.5: Parameter space decomposition for xylose. **(a)** Surrogate model validation error for time-to-consumption of xylose t_z versus the number of DFBA evaluations used in model fitting. **(b)** Sensitivity of t_z to the uncertain parameters. Global sensitivity indices are estimated inexpensively using the surrogate model. **(c)** Estimated PDF of t_z based on 1e+6 surrogate model evaluations. **(d)** Decomposition of support of prior density into a two non-overlapping elements at time 8.5 hr, projected onto the three sensitive parameters. The blue region is composed of parameters for which xylose is zero, estimated using the surrogate model, while the red region represents parameters with non-zero xylose concentrations.

robust RMSE estimate by splitting the ED into various training and validation sets, fitting different models with each training set, and averaging the prediction error (over the validation set) of each model. We focus exclusively on the corrected LOO error $\varepsilon^{\star}_{LOO}$ in (2.14). Table 2.2 gives the estimated RMSE values for sparse PCE surrogate models fit using different sized EDs. We observe that the empirical estimator greatly underpredicts the "true" RMSE found from the large validation set. In fact, for the smallest size $N = 10$, the empirical estimate is a factor of 10^4 smaller than the true RMSE. The cross-validated RMSE, on the other hand, predicts the correct order in all considered cases except $N = 10$ where it is only off by a factor of 10 instead of 10^4. Note that $\varepsilon^{\star}_{LOO}$ is used within the hybrid LAR algorithm to select the best surrogate out of all potential candidates.

Table 2.2: Relative mean square error estimates for time-to-consumption of glucose surrogate models under multiple experimental design sizes

N	Validation	Cross-validation	Empirical
10	1.014e-02	1.268e-03	2.601e-06
50	2.230e-04	3.718e-04	2.130e-04
100	1.616e-04	1.347e-04	5.333e-05
150	7.864e-05	6.468e-05	2.820e-05
200	5.787e-05	3.898e-05	2.416e-05
500	1.817e-05	1.273e-05	7.679e-06

The validation error is computed using a large set of samples not used in the fitting procedure, such that it closely matches the true error. Cross-validation and empirical error, however, are computed using only points in the original experimental design. Cross-validation partitions the experimental design into various training and validation sets, such that multiple models can be fit and their prediction errors averaged, in order to compute more robust error estimates than its empirical counterpart. Here, a corrected leave-one-out cross-validation procedure is utilized.

2.3.2 ME-PCE surrogate model validation

Once models for the singularity manifolds, i.e., $t_g(\boldsymbol{X})$ and $t_z(\boldsymbol{X})$ are determined, we can use them to construct ME-PCE surrogate models for the extracellular concentrations based on the strategy summarized in Figure 2.1. In principle this can be done at any time of interest, but we focus on time 7 hr for glucose and time 8.5 hr for xylose and biomass for illustrative purposes. The two key choices are the breakdown of elements and design of the ED. We manually defined the regions to be in terms of simple geometries for convenience sake, implying a variety of other choices would likely further improve the results shown in this section. If the conditional random variable $\boldsymbol{X}_k$ with PDF (2.5) does not have independent elements as per (2.9), then it was transformed to a collection of i.i.d. uniform random variables using the Rosenblatt transformation. The examples of when the transformation is necessary are the red triangular regions R3 and R5 in Figure 2.3d and the red truncated triangular pyramid in Figure 2.5d. Out of a total number of samples N, a fixed number of samples N_k is allocated to every element $k = 1, \dots, N_e$. In this case, $N_k \propto \Pr(\boldsymbol{X} \in S_k)$ is proportional to the probability that the parameter lies in that element, which is equivalent to the ratio of the volume of S_k to that of S for the uniform density considered here. All EDs are populated with quasi-random Sobol' sequences to ensure fair comparisons below.

The convergence properties of the proposed ME-PCE surrogate model and a comparison of predictions to the exact DFBA model on a large validation set for glucose at time 7 hr are shown in Figure 2.6. The RMSE for ME-PCE shows a more than quadratic rate of convergence with respect to number of model evaluations, which is a significant improvement over the first-order and half-order rates observed with sparse and full global PCE, respectively (Figure 2.6a). Thus, the incorporation of both sparse regression and the elemental breakdown of the support provide substantial gains in accuracy. The benefits of much lower RMSE can be observed in parity plot (Figure 2.6b), which shows the ME-PCE and global PCE surrogate model predictions versus the exact glucose value for 2×10^4 validation simulations. The ME-PCE surrogate consistently matches the DFBA model over the entire range of possible values while the global PCE shows significant deviations around 0, 2, 4, and 6.5 g/L even though the same number of samples are used to fit each of these models. Similarly good results were obtained for xylose at time 8.5 hr, which are summarized in Figure 2.7. Although the rate of convergence is not quite as large here, the RMSE for ME-PCE is consistently more than an order-of-magnitude lower than global PCE for every considered ED size (Figure 2.7a). The fact that this differs from glucose suggests there is a delicate balance between the size of the ED and the number/shape of the elements. Additionally, the parity plot for xylose (Figure 2.7b) shows more pronounced differences between the ME-PCE and global PCE surrogates than that for glucose, mainly due to the fact that the global xylose model predicts a large spread around zero that includes many non-physical negative values. The convergence and parity plots for the biomass surrogates at time 8.5 hr, based on the element breakdown shown in Figure 2.8, are summarized in Figure2.9. The only difference in the case of biomass is that the concentrations are non-zero throughout the entire parameter support.

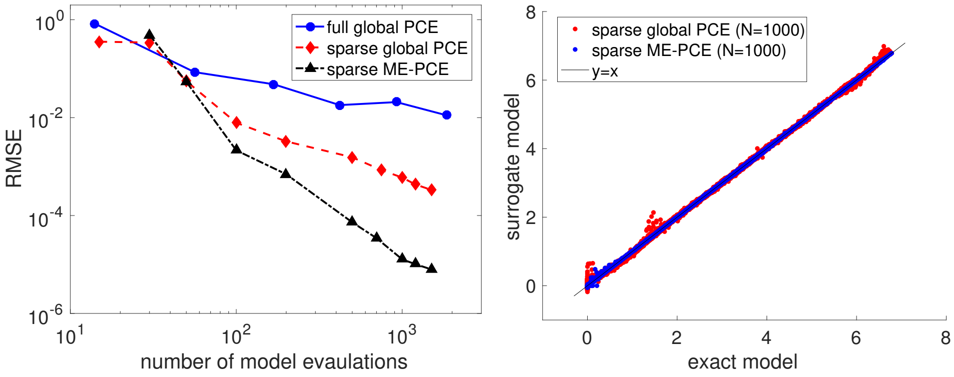

(a) (b)

Figure 2.6: Surrogate model accuracy for glucose. **(a)** Surrogate model validation error for glucose at time 7 hr versus the number of model evaluations used in fitting procedure for sparse ME-PCE, sparse global PCE, and full global PCE. **(b)** Parity plot for the sparse ME-PCE and sparse global PCE surrogate models, fit using 1000 model evaluations, for 2×10^4 validation simulations.

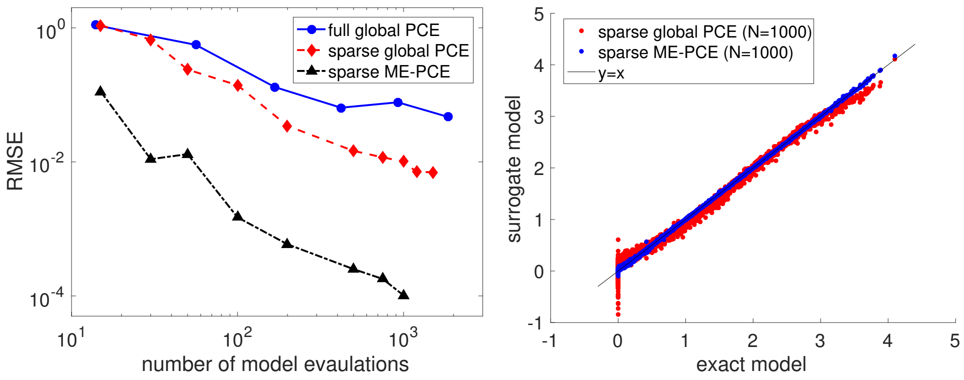

(a) (b)

Figure 2.7: Surrogate model accuracy for xylose. **(a)** Surrogate model validation error for xylose at time 8.5 hr versus the number of model evaluations used in fitting procedure for sparse ME-PCE, sparse global PCE, and full global PCE. **(b)** Parity plot for the sparse ME-PCE and sparse global PCE surrogate models, fit using 1000 model evaluations, for 2×10^4 validation simulations.

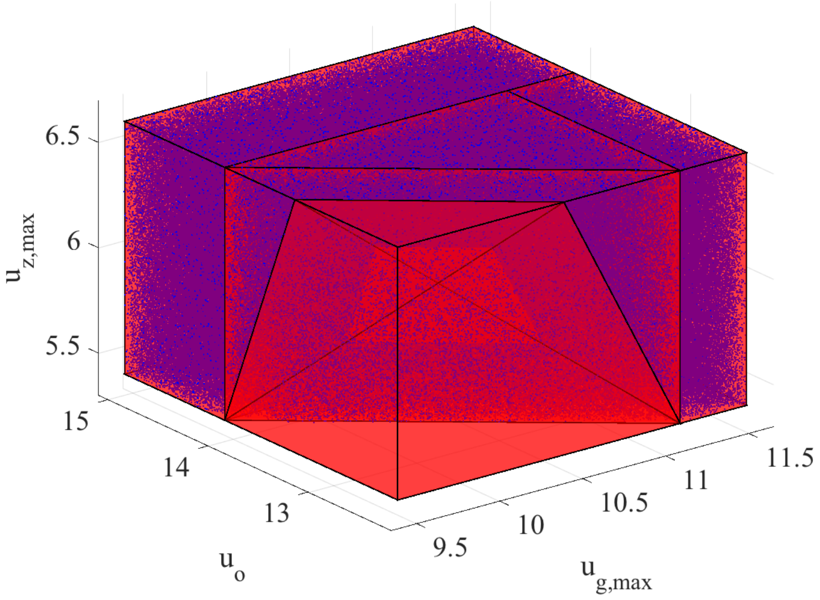

Figure 2.8: Parameter space decomposition for biomass at 8.5 hr. The decomposition is projected onto the three (out of six) most sensitive parameters. The blue dots represent parameters for which xylose is zero, which were estimated using a surrogate model for t_z. The breakdown mimics that for xylose at 8.5 hr, except more elements are required since the biomass concentration is non-zero throughout the entire parameter support.

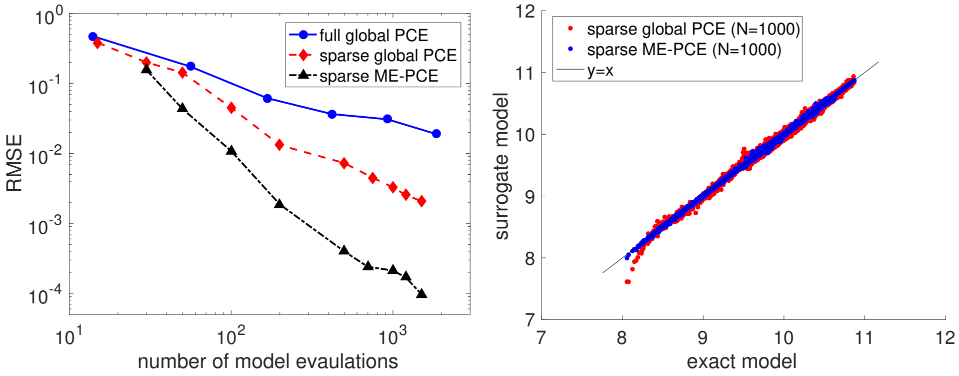

(a) (b)

Figure 2.9: Surrogate model accuracy for biomass. **(a)** Surrogate model validation error for biomass at time 8.5 hr versus the number of model evaluations used in fitting procedure for sparse ME-PCE, sparse global PCE, and full global PCE. **(b)** Parity plot for the sparse ME-PCE and sparse global PCE surrogate models, fit using 1000 model evaluations, for 2×10^4 validation simulations.

2.3.3 Bayesian parameter inference using sequential Monte Carlo

We focus on the inverse UQ problem of estimating parameters from data, which can be greatly accelerated using the ME-PCE strategy. The same data set used in [27] is assumed here, which includes measurements of the extracellular biomass, glucose, and xylose concentrations at $t \in \{0, 5.5, 6, 6.5, 7, 7.25, 7.75, 8, 8.25, 8.5\}$ hr. The measurements are assumed to be corrupted with noise

$$\begin{aligned} D_i^b &= b(t_i; \boldsymbol{X}) + E_i^b, & i &= 1, \ldots, 10, \\ D_i^g &= g(t_i; \boldsymbol{X}) + E_i^g, & i &= 1, \ldots, 10, \\ D_i^z &= z(t_i; \boldsymbol{X}) + E_i^x, & i &= 1, \ldots, 10, \end{aligned} \tag{2.24}$$

where $\boldsymbol{D}_i = (D_i^b, D_i^g, D_i^z)$ and $\boldsymbol{E}_i = (E_i^b, E_i^g, E_i^z)$ are, respectively, vectors of the measured data and noise at the ith time point. The concatenated data (resp. noise) vector is denoted by $\boldsymbol{D} = (\boldsymbol{D}_1, \ldots, \boldsymbol{D}_{10})$ (resp. $\boldsymbol{E} = (\boldsymbol{E}_1, \ldots, \boldsymbol{E}_{10})$). The measurement noise variables are modeled as independent zero-mean Gaussian random variables with state-dependent variance that equals 5% of the measured signal, i.e.,

$$E_i^v \sim \mathcal{N}(0, \sigma_{v,i}^2(\boldsymbol{X})), \quad \sigma_{v,i}(\boldsymbol{X}) = 0.05|v(t_i; \boldsymbol{X})|, \quad v \in \{b, g, z\}. \tag{2.25}$$

Given a set of measurements, the change in the state of information about the parameters is given by Bayes' rule [34]

$$f_{\boldsymbol{X}|\boldsymbol{D}}(\boldsymbol{x}|\boldsymbol{d}) = \frac{f_{\boldsymbol{D}|\boldsymbol{X}}(\boldsymbol{d}|\boldsymbol{x}) f_{\boldsymbol{X}}(\boldsymbol{x})}{f_{\boldsymbol{D}}(\boldsymbol{d})}, \tag{2.26}$$

where $f_{\boldsymbol{X}|\boldsymbol{D}}$ is the posterior density (i.e., parameter density given realized data); $f_{\boldsymbol{D}|\boldsymbol{X}}$ is the likelihood function (i.e., data density given fixed parameters); $f_{\boldsymbol{X}}$ is the prior density (i.e., parameter density before observing data); and $f_{\boldsymbol{D}}$ is the evidence (i.e., marginal data density). As Bayesian inference looks to characterize the full posterior density, it directly provides an explicit representation of the uncertainty in the parameter estimates, which is an important advantage over alternative optimization-based formulations that only provide point estimates.

The prior and likelihood function must be specified before solving (2.26). We assume the same uniform prior used to construct the ME-PCE surrogate models though these can differ in general. The likelihood function describes the discrepancy between the observed data and the model predictions in a probabilistic way. The likelihood function, which is fully specified by the data and noise models (2.24)-(2.25), is given by

$$f_{\boldsymbol{D}|\boldsymbol{X}}(\boldsymbol{d}|\boldsymbol{x}) = \prod_{i=1}^{10} \prod_{v \in \{b,g,z\}} \frac{1}{\sqrt{2\pi\sigma_{v,i}^2(\boldsymbol{x})}} \exp\left(-\frac{(d_i^v - v(t_i;\boldsymbol{x}))^2}{2\sigma_{v,i}^2(\boldsymbol{x})}\right). \tag{2.27}$$

Although we use a Gaussian likelihood here, the same framework can be applied to any choice of likelihood function and thus can be easily modified to incorporate other potentially important factors including sensor bias or asymmetric noise.

Since (2.26) cannot be solved analytically in practical situations, we must resort to sample-based approximations that rely on generating samples from the target posterior distribution. A variety of methods have been developed for sampling from the unknown posterior $f_{\boldsymbol{X}|\boldsymbol{D}}$ including Markov Chain Monte Carlo (MCMC) [2, 7, 68] and sequential Monte Carlo (SMC) [11, 15, 37] algorithms. The proposed surrogate models can be used to accelerate any sampling-based method; however, we focus on SMC since this is a class of completely parallel algorithms. Since MCMC is not inherently parallelizable, it can take on the order of hours to converge even for cheap polynomial models as illustrated in [56]. SMC is based on the concept of *importance sampling*, which can be straightforwardly implemented in an iterative fashion such that the posterior is updated every time a new measurement is available. For a given number of particles N, the SMC approximation to (2.26) can be summarized as follows:

1. Initialization: set $k = 1$ and generate samples and weights $\{\boldsymbol{x}_i, w_i\}_{i=1}^N$ from prior.
2. Reweighting: update the weights $w_i \leftarrow w_i \times w_k(\boldsymbol{x}_i)$ where $w_k(\boldsymbol{x}_i) \propto f_{D_k|\boldsymbol{X}}(d_k|\boldsymbol{x}_i)$.
3. Resampling: resample $\{\boldsymbol{x}_i, w_i\}_{i=1}^N$ for particles with equal weights $\{\boldsymbol{x}_i^r, \frac{1}{N}\}_{i=1}^N$.

4. Loop: set $k \leftarrow k + 1$ and $\{\boldsymbol{x}_i, w_i\}_{i=1}^N \leftarrow \{\boldsymbol{x}_i^r, \frac{1}{N}\}_{i=1}^N$. Return to Step 2 if $k < k_f$.

When the algorithm stops at final time k_f (equal to 10 in this case), the particle system targets the posterior distribution of interest. We use systematic resampling in Step 3 due to its computational simplicity and good empirical performance, though a variety of other methods are available [37]. Step 2 is usually the computational bottleneck because the model must be repeatedly solved in order to evaluate the likelihood weight factors using (2.27). Therefore, we propose to replace the evaluation of $v(t_i; \boldsymbol{x})$ with a ME-PCE surrogate model $v^{\text{ME-PCE}}(t_i; \boldsymbol{x})$ for every $v \in \{b, g, z\}$ and $i = 1, \ldots, 10$. We must then construct a total of 30 surrogates before running the SMC algorithm.

We follow the same basic strategy describe in the previous section for constructing the ME-PCE surrogate models. The main difference is that we now have multiple time points of interest for every concentration. Since the singularity manifolds evolve over time, the element breakdown will be different for any of the time points that fall within the range of $t_g(\boldsymbol{X})$ and/or $t_z(\boldsymbol{X})$. This issue is illustrated graphically for glucose in Figure 2.10, which has four measurement times that overlap with $t_g(\boldsymbol{X})$. The element breakdown for all 30 models can be determined from the global PCE representations of $t_g(\boldsymbol{X})$ and $t_z(\boldsymbol{X})$, which both achieved $\varepsilon^{\star}_{LOO} < 10^{-3}$ using the same ED with 200 parameter realizations. Once the elements have been selected, we must choose how to design the ED and simulate the DFBA model for all 30 quantities of interest. Although treating each measurement time independently is simple, it can be inefficient since the model has to be integrated from the initial condition in every case. If we instead recycle points in the ED, we can avoid repetitive computations by integrating between time points wherein the model is initialized using the solution at the end of the previous interval. Ideally the same ED could be used for all time points, but this is unlikely when the size and shape of the elements change substantially over time. Thus, we recommend to sequentially add points to the ED whenever the target error is not met. Using this methodology, all 30 surrogates were fit using a total of less than 2500 DFBA model evaluations with errors $\varepsilon^{\star}_{LOO}$ consistently lower than 10^{-3}. Note that the cost of estimating the expansion coefficients using the hybrid LAR method was less than 30 seconds in all cases, indicating that the DFBA simulation is the dominant cost during surrogate model construction.

Figure 2.11 shows the posterior density estimated using SMC with $N = 5 \times 10^5$ particles for a synthetic data set wherein the likelihood weights are evaluated using the inexpensive ME-PCE surrogates. The synthetic data was obtained by simulating the genome-scale *E. coli* DFBA model with fixed parameters (red lines in Figure 2.11) and adding random noise realizations from (2.25) to the resulting model outputs, and is shown with 'x' marks in Figure 2.12. The nearly 2500 DFBA

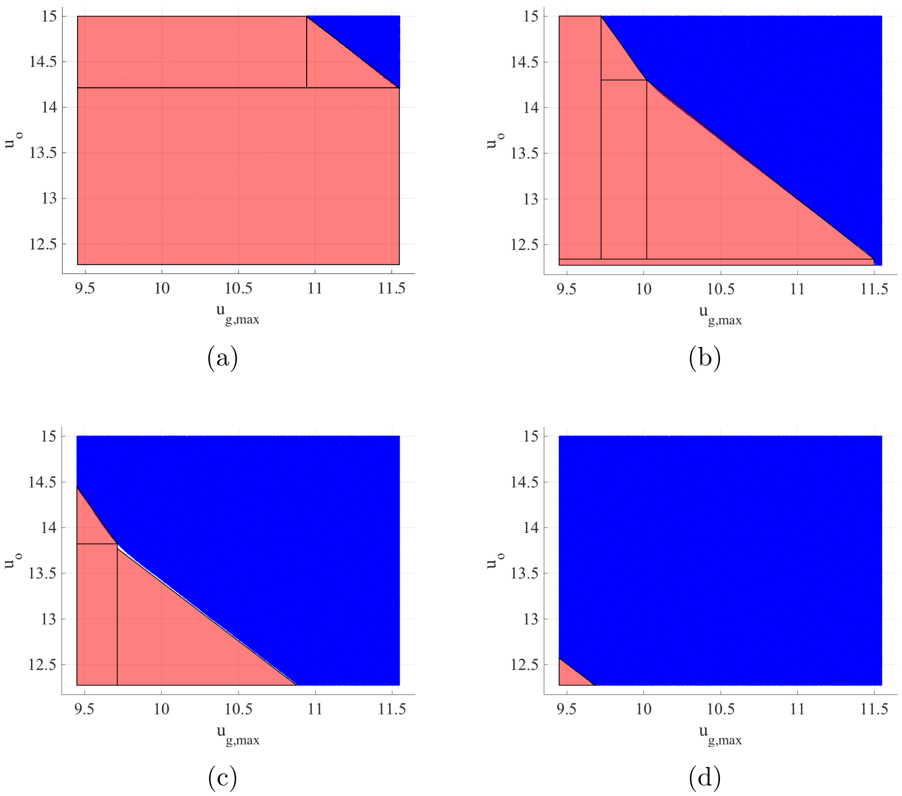

(a) (b)

(c) (d)

Figure 2.10: Element breakdown for glucose concentration over time. **(a)** Four elements at time 6.5 hr. **(b)** Six elements at time 7.0 hr. **(c)** Four elements at time 7.25 hr. **(d)** Two elements at time 7.75 hr.

simulations used to construct the surrogates requires $\sim$ 55 minutes while the surrogate-based SMC algorithm, which takes advantage of vectorization, finishes in $\sim$ 1 minute of CPU time. Therefore, the full Bayesian inference problem can be solved in around 1 hr using the proposed method that represents a roughly 200-fold savings in computational cost when compared to SMC without surrogates that would require $\sim$ 7.5 days of CPU time under the same settings (5×10^5 DFBA model evaluations at a cost of 1.3 seconds per evaluation). To verify that the SMC algorithm provides sensible results, we plotted the DFBA model predictions under the MAP estimates (i.e., parameters that maximize the posterior) in Figure 2.12, which closely matches the synthetic data.

The estimated posterior density in Figure 2.11 also provides interesting physical insights. Three of the parameters (K_g, K_z, K_{ig}) are unobservable with the current data set since their posterior (green) and prior (blue) densities are equivalent. This observation could not easily be made before running the estimation procedure due to the nonlinear and indirect relationship between $\boldsymbol{D}$ and $\boldsymbol{X}$. The data is sensitive to ($u_{g,max}$, $u_{z,max}$, u_o); however, these parameters are highly correlated as seen in off-diagonal plots of their joint densities in Figure 2.11. As such, the currently available data from a single batch is insufficient for providing accurate estimates for all of the parameters of interest. Additionally, the evolution of the marginal posterior densities of the observable parameters over time is shown in Figure 2.13. Since glucose is mostly consumed by 7.25 hr, the densities of $u_{g,max}$ and u_o remain constant for the remaining batch time. The density of $u_{z,max}$, however, is constant before 7.25 hr because xylose remains mostly at its initial condition. Also, notice that at 7.75 hr the density of $u_{z,max}$ is positively skewed away from the true value since the noise realization happened to be a large negative value. The next few measurements quickly correct this behavior, which results in a posterior density that is a fairly accurate representation of the true value for $u_{z,max}$.

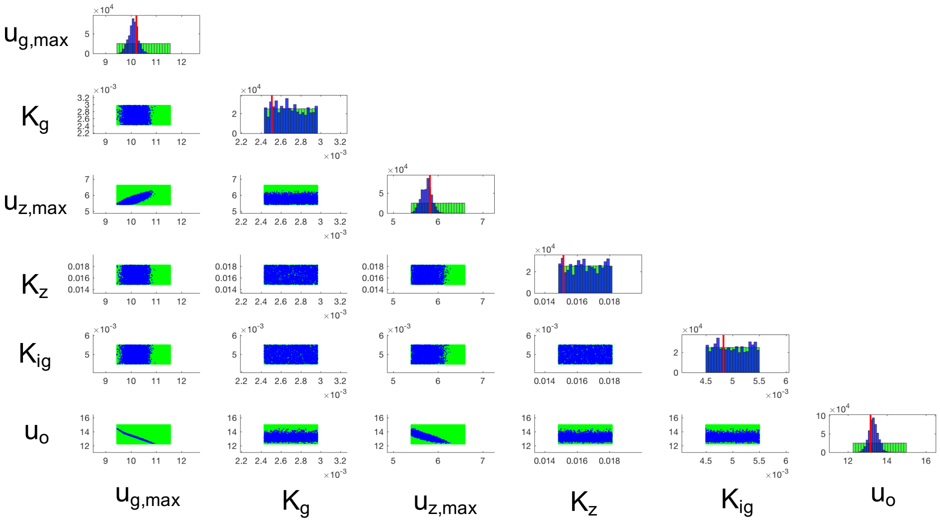

Figure 2.11: Posterior distributions of the estimated model parameters. The diagonal subplots represent marginal densities while the off-diagonal subplots represent two-dimensional projections of samples from the joint density. Blue denotes the posterior density while green denotes the prior density. The red line represents the true parameter values used to generate synthetic data for estimation purposes.

2.3.4 Forward uncertainty propagation

Let $\boldsymbol{Y} = \mathcal{M}(\boldsymbol{X})$ denote a vector of all model responses. The forward UQ problem looks to characterize the uncertainty in the model predictions via propagating uncertainty in the parameters through $\mathcal{M}$. This can involve estimating either the prior predictive distribution $f_{\boldsymbol{Y}}$ (before any data has been collected), or the posterior predictive distribution $f_{\boldsymbol{Y}|\boldsymbol{D}}$ (after data has been obtained). The only difference between these two problems is that $\mathcal{M}$ is evaluated at i.i.d. samples drawn from the prior in the former and the posterior in the latter. The densities

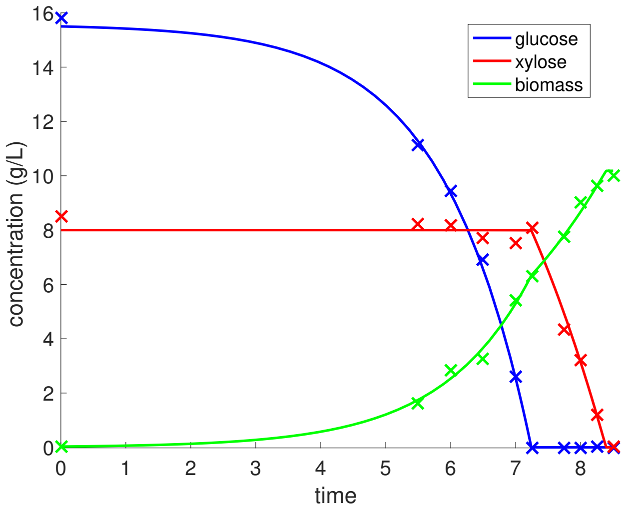

Figure 2.12: Comparison of model predictions and synthetic data. The model predictions, shown with solid lines, were obtained by integrating the DFBA model with the maximum a posteriori (MAP) estimates of the parameters, which correspond to the mode of the posterior density in Figure 2.11. The 'x' marks represent synthetic data generated by corrupting model predictions for the true (unknown) parameters with randomly generated noise.

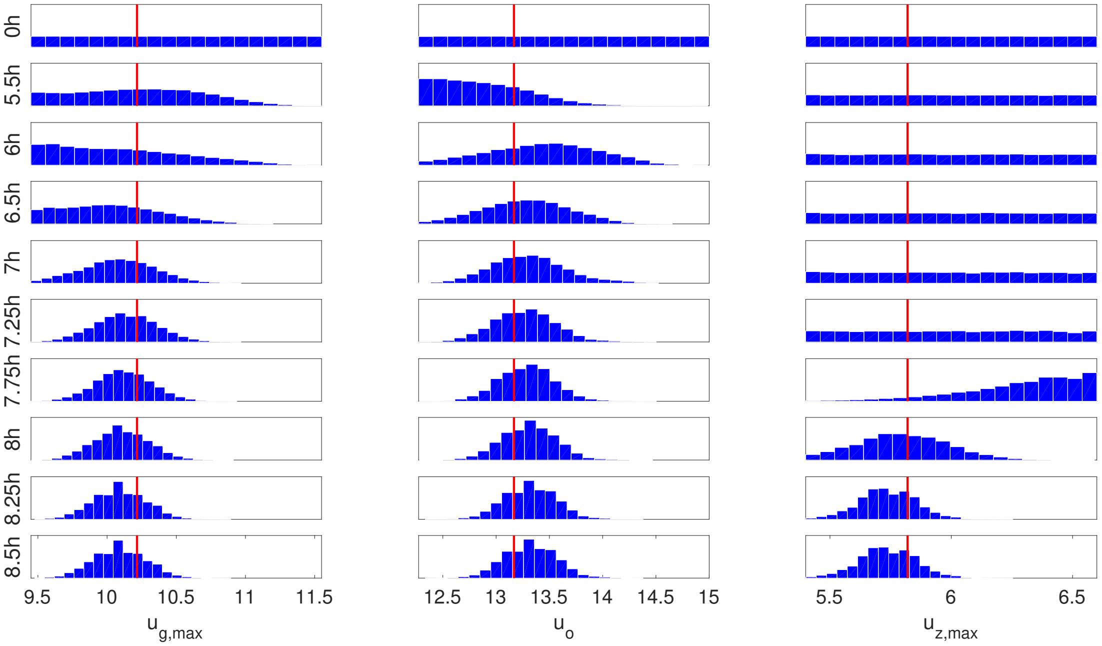

Figure 2.13: Evolution of the posterior marginal densities for the observable model parameters over time. Each subplot shows the histogram of parameter posterior samples estimated using the sequential Monte Carlo method. The x-axis represents the range in values of the parameters and the y-axis represents frequencies. The red line represents the true parameter values.

of the model predictions estimated using 1×10^5 samples are shown in Figure 2.14. By replacing the DFBA model with the ME-PCE surrogate, these histograms were obtained in less than 30 seconds of CPU time. As expected, the prior predictive are much wider than the posterior predictive distributions, indicating there is significant uncertainty before incorporating data. In addition, we see that many of these distributions have sharp changes and long tails due to the non-smooth behavior, which can be accurately captured with the proposed ME-PCE approach. It is also interesting to note that the posterior predictive distributions have low variance, even though the parameters are not perfectly estimated, which highlights the impact that nonlinearity can have on both estimation and uncertainty propagation.

2.4 Discussion

In this work, we develop a novel surrogate modeling approach for handling discontinuous (non-smooth) behavior in large-scale computational models of biological systems. Discontinuities occur in many biological models including bistable switching [21], positive feedback loops [32], and intracellular transport delays [6], to name a few. The developed surrogate model can vastly accelerate uncertainty quantification (UQ) tasks, such as calibrating the model with experimental data (inverse problem) and quantifying confidence in the model predictions (forward problem). The proposed approach is based on a extension of state-of-the-art polynomial chaos expansion (PCE) methods, which is able to accurately reconstruct the discontinuous model response from a limited number of expensive model evaluations. The main idea behind this so-called *multi-element* (ME) extension to PCE is to systematically decompose the set of possible parameters into a collection of non-overlapping regions so that the model response is smooth within each element. We focus on *genome-scale* DFBA models in this chapter, however, the basic concept can be applied to any nonlinear, discontinuous model in the presence of uncertainty.

We demonstrate advantages of the proposed sparse ME-PCE surrogate model on a DFBA model of *E. coli* fermentation reactor under aerobic growth on glucose and xylose mixed media that utilizes a genome-scale metabolic network reconstruction with 1075 reactions and 761 metabolites. We showcase how ME-PCE surrogates can be used to enable both inverse and forward UQ with respect to six uncertain parameters related to the substrate uptake kinetics. These six parameters were estimated with a trial-and-error procedure in past work [27] and, thus, our goal was to repeat this procedure using fully Bayesian methods. The posterior parameter distribution was estimated using sequential Monte Carlo with 5×10^5 samples, which would have required ~ 7 days of CPU time to compute using the full DFBA model, but took around one hour total (including

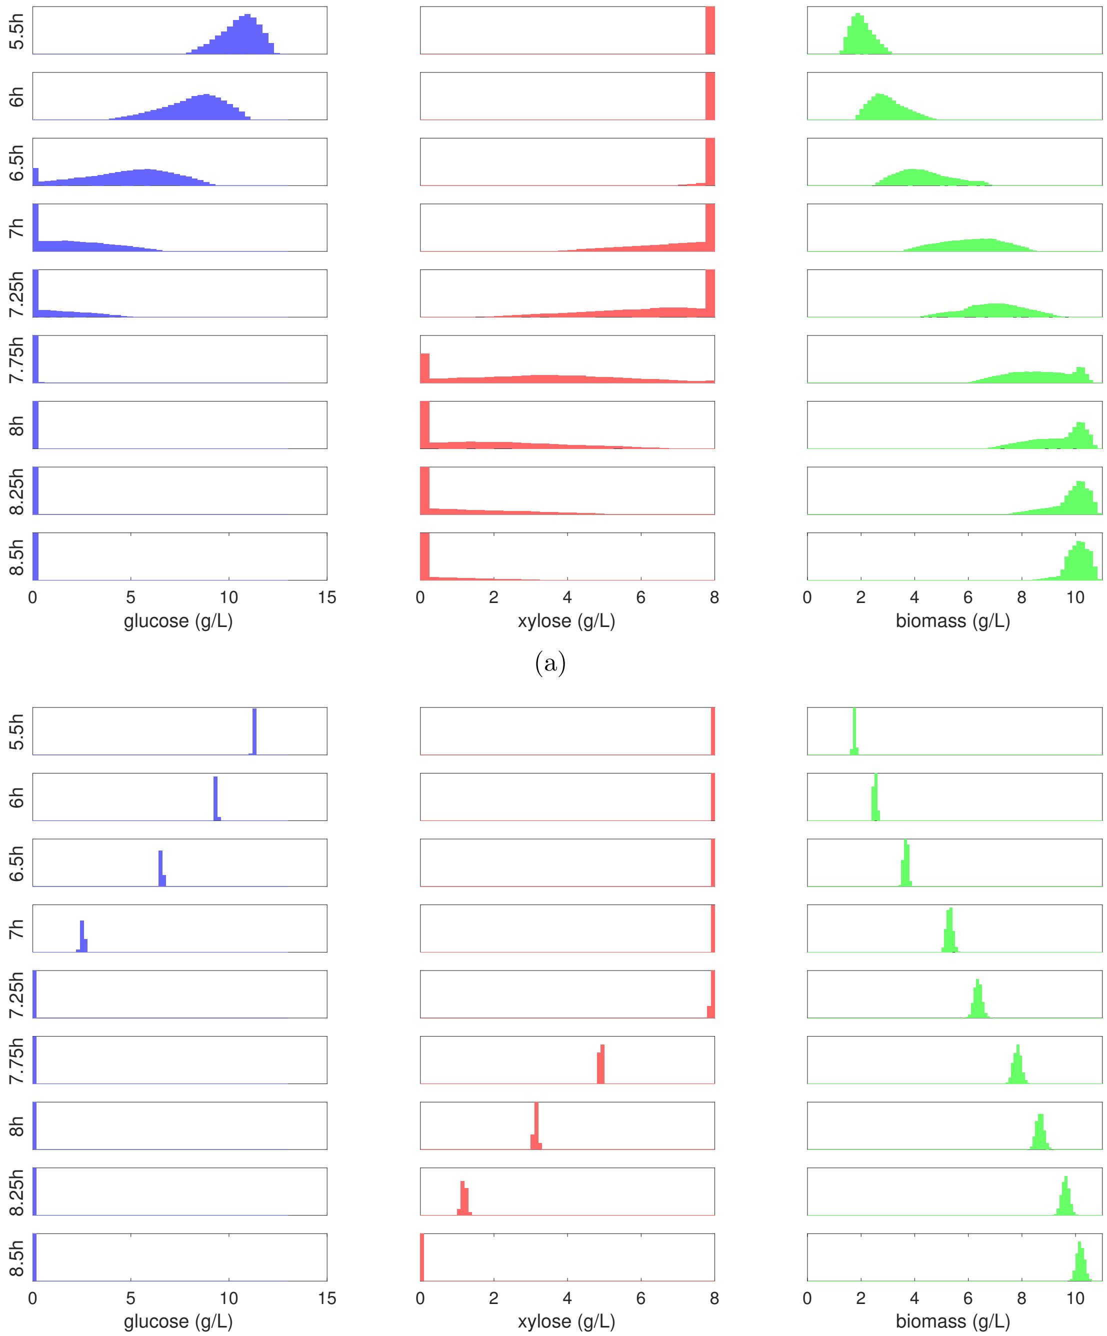

(a)

(b)

Figure 2.14: Predicted probability distributions of extracellular concentrations. **(a)** Model predictions using parameter samples from the prior. **(b)** Model predictions using parameter samples from the posterior. (a-b) Each subplot shows the histogram of samples obtained by substituting i.i.d. samples from the parameter distribution into the corresponding ME-PCE surrogate model. The x-axis represents the range in values of the model outputs and the y-axis represents frequencies.

cost of fitting) when ME-PCE surrogates were used. The resulting posterior distribution yielded significant physical insights including that the available data set is insufficient to estimate the six parameters reliably, with three of the parameters being completely unobservable. We also demonstrate how the prior and posterior predictive distributions can be estimated with negligible computational cost by evaluating the ME-PCE surrogate model for many randomly sampled parameter realizations. Global sensitivity analysis can also easily be performed to determine how much each kinetic parameter influences the variability in the model responses. Note that this observed speed-up is expected to be even greater for more complex DFBA models, such as those with nonlinear cellular objectives, multiple cultures, or even larger metabolic networks due to the increased cost of the simulations.

2.4.1 Potential improvements to surrogate model

The efficacy of the ME-PCE surrogate modeling approach critically relies on the choice of elements as well as the selection of parameter values in the experimental design (ED). We observed large improvements over global PCE using simple element geometries and random sampling for the ED; however, we expect even further gains in the convergence rate by developing procedures for optimizing these choices. For example, adaptive methods have been developed for decomposing the random parameter space using sensitivity information to decide which elements to split [67]. A similar concept could be applied to models of the singularity manifolds in this work (e.g., $t_g(\boldsymbol{X})$) for automatically partitioning the parameter space, even when there are multiple sensitive parameters. These methods would likely benefit from the incorporation of more advanced geometries than simple boxes. After the parameter support has been partitioned, the ED points in each element should be chosen to maximize their information content. Multiple approaches have been recently developed to tackle this challenging problem including coherence-optimal sampling [26] and numerical “moment-matching” optimization [53, 54]. When model evaluations are time consuming, it can be useful to build up the ED sequentially by iteratively enriching the initial ED. Although some work has been done in this area, the choice of elements and the ED are generally coupled in the multi-element version of the problem and remains largely unexplored in the literature.

2.4.2 Alternative methods for uncertainty quantification

The proposed ME-PCE surrogate modeling approach can be used to accelerate a wide-range of UQ methods. This includes alternative parameter estimation methods such as Markov chain Monte Carlo, maximum likelihood estimation, and maximum a posteriori estimation. Although not an issue in the case study, sequential Monte Carlo can suffer from *degeneracy* wherein fewer and fewer particles retain significant weight. This sample degeneracy can be protected against by adding a *rejuvenation* step that “moves” the resampled particles according to a Markov chain transition kernel [11]. This operation does not change the target distribution, but does reduce impoverishment since identical replicates of a single particle are replaced with new values. Even though the move step is computationally demanding, it can still benefit from parallel processing. Another way to reduce sample degeneracy is to sample from a distribution that more closely matches the posterior than the prior, which is especially useful for diffuse priors or concentrated likelihoods. One such example is the so-called Laplace approximation to the posterior, which assumes the posterior is normally distributed around the maximum a posteriori estimate with covariance approximated from a second-order Taylor series of the logarithm of the posterior distribution [38, 64]. The ME-PCE surrogates can also be used to accelerate methods for Bayesian model selection, which involves computing the probability that every model in a set of potential models matches available data. This is often an important problem is systems biology since uncertainty in model structure can significantly impact the conclusions of parameter inference [35].

2.4.3 Extensions to optimal experiment design

The selection of optimal conditions for conducting experiments (e.g., measurement times, initial conditions, and time-varying input profiles) is important for ensuring maximum information is extracted from the observations, especially when the experiments are expensive and time-consuming to perform. For example, it may be useful to change the feed rate or the measurement times in the considered case study so that the data ensures tight parameter estimates are obtained. Optimal experiment design (OED) has been extensively studied in the classical framework wherein the design criteria are defined as some scalar function of the Fisher information matrix (FIM) [1]. More recently, OED has been tackled from a fully Bayesian perspective that replaces the approximated classical design criteria with an *expected utility* function that is rigorously chosen from a decision-theoretic point of view [8, 51, 58].

The ME-PCE surrogate models could be used to efficiently evaluate classical or Bayesian design criteria at any fixed experimental condition. However, the parameter space decomposition depends strongly on the experiment, such that separate surrogates need to be constructed for all experiments of interest. This is not a major challenge when only a small, finite number of experiments (on the order of ten) are considered, but is intractable for continuous design spaces. Developing efficient procedures for both classical and Bayesian OED in genome-scale DFBA models is an important area for future research. One possible direction is to treat the experiment design variables as parameters when constructing the surrogate model, as suggested in [31] for global PCE. However, developing methods for automatically choosing elements, as mentioned above, would be critical to ensure similar ideas would be viable in ME-PCE because the model responses are undoubtedly sensitive to the design variables.

2.5 Bibliography

1. A. C. Atkinson and A. N. Donev. *Optimum experimental designs*. New York: Oxford University Press, 2007.
2. M. A. Beaumont and B. Rannala. The Bayesian revolution in genetics. *Nature Reviews Genetics*, 5:251–261, 2004.
3. S. A. Becker, A. M. Feist, M. L. Mo, G. Hannum, B. Ø. Palsson, and M. J. Herrgard. Quantitative prediction of cellular metabolism with constraint-based models: The COBRA toolbox. *Nature Protocols*, 2:727, 2007.
4. I. Bilionis and N. Zabaras. Multi-output local Gaussian process regression: Applications to uncertainty quantification. *Journal of Computational Physics*, 231:5718–5746, 2012.
5. G. Blatman and B. Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230:2345–2367, 2011.
6. G. A. Bocharov and F. A. Rihan. Numerical modelling in biosciences using delay differential equations. *Journal of Computational and Applied Mathematics*, 125:183–199, 2000.
7. S. Brooks. Markov chain Monte Carlo method and its application. *The Statistician*, 47:69–100, 1998.
8. K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.
9. O. Chapelle, V. Vapnik, and Y. Bengio. Model selection for small sample regression. *Machine Learning*, 48:9–23, 2002.
10. S.-K. Choi, R. V. Grandhi, R. A. Canfield, and C. L. Pettit. Polynomial chaos expansion with latin hypercube sampling for estimating response variability. *AIAA journal*, 42:1191–1198, 2004.
11. N. Chopin. A sequential particle filter method for static models. *Biometrika*, 89:539–552, 2002.

12. B. K. S. Chung, S. Selvarasu, A. Camattari, J. Ryu, H. Lee, J. Ahn, H. Lee, and D. Lee. Genome-scale metabolic reconstruction and in silico analysis of methylotrophic yeast Pichia pastoris for strain improvement. *Microbial Cell Factories*, 9:50, 2010.

13. P. G. Constantine, M. S. Eldred, and E. T. Phipps. Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229:1–12, 2012.

14. M. W. Covert, E. M. Knight, J. L. Reed, M. J. Herrgard, and B. O. Palsson. Integrating high-throughput and computational data elucidates bacterial networks. *Nature*, 429(6987):92–96, 2004.

15. A. Doucet, N. De Freitas, and N. Gordon. An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo methods in practice*, pages 3–14. Springer, 2001.

16. B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32:407–499, 2004.

17. M. A. Eiteman, S. A. Lee, and E. Altman. A co-fermentation strategy to consume sugar mixtures effectively. *Journal of Biological Engineering*, 2:3, 2008.

18. M. Eldred and J. Burkardt. Comparison of non-intrusive polynomial chaos and stochastic collocation methods for uncertainty quantification. In *47th AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition*, page 976, 2009.

19. V. N. Emenike, R. Schenkendorf, and U. Krewer. Model-based optimization of biopharmaceutical manufacturing in Pichia pastoris based on dynamic flux balance analysis. *Computers & Chemical Engineering*, 118:1–13, 2018.

20. J. Feinberg, V. G. Eck, and H. P. Langtangen. Multivariate polynomial chaos expansions with dependent variables. *SIAM Journal on Scientific Computing*, 40(1):A199–A223, 2018.

21. J. E. Ferrell and W. Xiong. Bistability in cell signaling: How to make continuous processes discontinuous and reversible processes irreversible. *Chaos*, 11(1):227–236, 2001.

22. W. Gautschi. On generating orthogonal polynomials. *SIAM Journal on Scientific and Statistical Computing*, 3(3):289–317, 1982.

23. R. Ghanem and P. Spanos. *Stochastic Finite Elements A Spectral Approach*. SpringerVerlag, 1991.

24. J. A. Gomez. *Simulation, sensitivity analysis, and optimization of bioprocesses using dynamic flux balance analysis.* PhD thesis, Massachusetts Institute of Technology, 2018.
25. J. A. Gomez, K. Höffner, and P. I. Barton. Dfbalab: A fast and reliable MATLAB code for dynamic flux balance analysis. *BMC Bioinformatics*, 15:409, 2014.
26. J. Hampton and A. Doostan. Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression. *Computer Methods in Applied Mechanics and Engineering*, 290:73–97, 2015.
27. T. J. Hanly and M. A. Henson. Dynamic flux balance modeling of microbial co-cultures for efficient batch fermentation of glucose and xylose mixtures. *Biotechnology and Bioengineering*, 108:376–385, 2011.
28. S. M. Harwood, K. Höffner, and P. I. Barton. Efficient solution of ordinary differential equations with a parametric lexicographic linear program embedded. *Numerische Mathematik*, 133:623–653, 2016.
29. J. L. Hjersted and M. A. Henson. Optimization of fed-batch Saccharomyces cerevisiae fermentation using dynamic flux balance models. *Biotechnology Progress*, 22:1239–1248, 2006.
30. K. Höffner, S. M. Harwood, and P. I. Barton. A reliable simulator for dynamic flux balance analysis. *Biotechnology and Bioengineering*, 110:792–802, 2013.
31. X. Huan and Y. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232:288–317, 2013.
32. N. T. Ingolia and A. W. Murray. Positive feedback loops as a flexible biological module. *Current Biology*, 17(8):668–677, 2007.
33. U. Kaplan, M. Türkay, L. Biegler, and B. Karasözen. Modeling and simulation of metabolic networks for estimation of biomass accumulation parameters. *Discrete Applied Mathematics*, 157:2483–2493, 2009.
34. M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:425–464, 2001.
35. J. Liepe, P. Kirk, S. Filippi, T. Toni, C. P. Barnes, and M. P. H. Stumpf. A framework for parameter estimation and model selection from experimental data in systems biology using approximate Bayesian computation. *Nature Protocols*, 9:439, 2014.

36. K. P. Lisha and D. Sarkar. Dynamic flux balance analysis of batch fermentation: Effect of genetic manipulations on ethanol production. *Bioprocess and Biosystems Engineering*, 37:617–627, 2014.

37. J. S. Liu and R. Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93:1032–1044, 1998.

38. Q. Long, M. Scavino, R. Tempone, and S. Wang. Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations. *Computer Methods in Applied Mechanics and Engineering*, 259:24–39, 2013.

39. R. Mahadevan, J. S. Edwards, and F. J. Doyle III. Dynamic flux balance analysis of diauxic growth in Escherichia coli. *Biophysical Journal*, 83:1331–1340, 2002.

40. L. Mao and W. S. Verwoerd. Orca: a COBRA toolbox extension for model-driven discovery and analysis. *Bioinformatics*, 30:584–585, 2013.

41. S. Marelli and B. Sudret. Uqlab: A framework for uncertainty quantification in Matlab. In *Vulnerability, Uncertainty, and Risk: Quantification, Mitigation, and Management*, pages 2554–2563. 2014.

42. M. Martin-Casas and A. Mesbah. Discrimination between competing model structures of biological systems in the presence of population heterogeneity. *IEEE Life Science Letters*, 2:23–26, 2016.

43. Y. M. Marzouk and H. N. Najm. Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics*, 228:1862–1902, 2009.

44. A. L. Meadows, R. Karnik, H. Lam, S. Forestell, and B. Snedecor. Application of dynamic flux balance analysis to an industrial Escherichia coli fermentation. *Metabolic engineering*, 12:150–160, 2010.

45. A. M. Molinaro, R. Simon, and R. M. Pfeiffer. Prediction error estimation: a comparison of resampling methods. *Bioinformatics*, 21:3301–3307, 2005.

46. R. Noori, H.-D. Yeh, M. Abbasi, F. T. Kachoosangi, and S. Moazami. Uncertainty analysis of support vector machine for online prediction of five-day biochemical oxygen demand. *Journal of Hydrology*, 527:833–843, 2015.

47. J. D. Orth, I. Thiele, and B. Ø. Palsson. What is flux balance analysis? *Nature Biotechnology*, 28:245–248, 2010.

48. E. J. O'Brien, J. M. Monk, and B. O. Palsson. Using genome-scale models to predict biological capabilities. *Cell*, 161(5):971–987, 2015.

49. J. Paulson, M. Martin-Casas, and A. Mesbah. Black-box surrogate modeling of biological systems with non-smooth behavior: Applications in bayesian parameter inference for dynamic flux balance analysis models. *Submitted*, 2019.

50. J. A. Paulson, E. A. Buehler, and A. Mesbah. Arbitrary polynomial chaos for uncertainty propagation of correlated random variables in dynamic systems. *IFAC-PapersOnLine*, 50:3548–3553, 2017.

51. J. A. Paulson, M. Martin-Casas, and A. Mesbah. Optimal Bayesian experiment design for nonlinear dynamic systems with chance constraints. *Journal of Process Control*, https://doi.org/10.1016/j.jprocont.2019.01.010, 2019.

52. J. A. Paulson and A. Mesbah. An efficient method for stochastic optimal control with joint chance constraints for nonlinear systems. *International Journal of Robust and Nonlinear Control*, pages 1–21, 2017.

53. J. A. Paulson and A. Mesbah. Arbitrary polynomial chaos for quantification of general probabilistic uncertainties: Shaping closed-loop behavior of nonlinear systems. In *Proceedings of the 57th IEEE Conference on Decision and Control*, page Accepted, Miami, 2018.

54. J. A. Paulson and A. Mesbah. Nonlinear model predictive control with explicit backoffs for stochastic systems under arbitrary uncertainty. In *Proceedings of the 6th IFAC Conference on Nonlinear Model Predictive Control*, pages 622–633. Madison, WI, 2018.

55. J. L. Reed, T. D. Vo, C. H. Schilling, and B. O. Palsson. An expanded genome-scale model of Escherichia coli K-12 (i JR904 GSM/GPR). *Genome Biology*, 4:R54, 2003.

56. M. Renardy, T.-M. Yi, D. Xiu, and C.-S. Chou. Parameter uncertainty quantification using surrogate models applied to a spatial model of yeast mating polarization. *PLoS Computational Biology*, 14:e1006181, 2018.

57. M. Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23:470–472, 1952.

58. E. G. Ryan, C. C. Drovandi, J. M. McGree, and A. N. Pettitt. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84:128–154, 2016.

59. F. Scott, P. Wilson, R. Conejeros, and V. S. Vassiliadis. Simulation and optimization of dynamic flux balance analysis models using an interior point method reformulation. *Computers & Chemical Engineering*, 119:152–170, 2018.

60. M. Sinsbeck and W. Nowak. An optimal sampling rule for nonintrusive polynomial chaos expansions of expensive models. *International Journal for Uncertainty Quantification*, 5:275–295, 2015.

61. R. C. Smith. *Uncertainty quantification: Theory, implementation, and applications*, volume 12. SIAM, 2013.

62. B. Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93:964–979, 2008.

63. R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58:267–288, 1996.

64. L. Tierney and J. B. Kadane. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81:82–86, 1986.

65. W. M. Van Gulik and J. J. Heijnen. A metabolic network stoichiometry analysis of microbial growth and product formation. *Biotechnology and Bioengineering*, 48:681–698, 1995.

66. X. Wan and G. E. Karniadakis. An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *Journal of Computational Physics*, 209:617–642, 2005.

67. X. Wan and G. E. Karniadakis. Multi-element generalized polynomial chaos for arbitrary probability measures. *SIAM Journal on Scientific Computing*, 28:901–928, 2006.

68. D. J. Wilkinson. Bayesian methods in bioinformatics and computational systems biology. *Briefings in Bioinformatics*, 8:109–116, 2007.

69. D. Xiu. Efficient collocational approach for parametric uncertainty analysis. *Communications in Computational Physics*, 2:293–309, 2007.

70. D. Xiu. Fast numerical methods for stochastic computations: A review. *Communications in Computational Physics*, 5:242–272, 2009.

71. D. Xiu and G. E. Karniadakis. The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM Journal of Scientific Computing*, 24:619–644, 2002.

72. X. Zhao, S. Noack, W. Wiechert, and E. von Lieres. Dynamic flux balance analysis with nonlinear objective function. *Journal of Mathematical Biology*, 75:1487–1515, 2017.

Chapter 3

Optimal experiment design for offline model structure selection

The ability to discriminate between computational model structures depicting competing hypotheses and subsequently selecting the most representative one can be useful in a variety of contexts. Oftentimes, several competing hypothesis are proposed to describe the underlying molecular mechanisms of a biological system. Selecting the most representative model is imperative for obtaining meaningful insights into the true system dynamics. In this chapter, we introduce an OED method to elucidate biochemical network structures in an offline setting.[*]

3.1 Introduction

Computational models can play a central role in unraveling the fundamental mechanisms of biological systems, as they facilitate systematic design of experiments and dynamical analysis of complex networks of biochemical reactions [2]. The predictive quality of computational models critically hinges on selection of adequate descriptions of molecular mechanisms (i.e., model structures) that characterize the system dynamics. Multiple competing model structures, which correspond to different hypotheses for molecular mechanisms, typically exist for describing the dynamics of the same cellular processes. Optimal experiment design (OED) tools for model discrimination can assist in the invalidation of model structures that cannot adequately describe experimental data (e.g., [13, 17]). However, a key challenge in OED for biological systems arises from cell-to-cell variability in a cell population, i.e., cell traits may differ within a cell population [8]. The heterogeneity of biological systems can significantly increase the complexity of model discrimination.

---

[*] The content of this chapter was published in [12]

Numerous studies have elucidated the critical role of population heterogeneity in major biological processes such as embryogenesis, tumorigenesis or survival [1, 6, 19]. Biological sources of cell population variability can be broadly classified into genetic and non-genetic. Genetic sources arise from spontaneous or induced mutation of the genetic material within a cell population. On the other hand, generation of variability in isogenic cell populations has revealed the importance of non-genetic factors that give rise to phenotypic heterogeneity [8]. Non-genetic heterogeneity mainly results from extrinsic or intrinsic sources, which can in their turn be classified into temporal and population noise (see Chapter 1). In this chapter, modeling of the population variability exclusively focuses on time-invariant cell-to-cell differences or population noise.

The goal of this chapter is to demonstrate the effectiveness of using a *probabilistic* OED framework [14, 18] for model structure selection (i.e., discrimination) in the presence of population heterogeneity. Probabilistic OED allows for designing system input(s) that can discriminate between distributions of cellular outputs (within a cell population) predicted by the competing models. Thus, upon application of the optimal input(s) to the biological system, invalidation of inadequate model structures can be achieved. In this work, probabilistic OED is performed for the erythropoietin-induced JAK2/STAT5 signaling pathway, for which several competing model structures can describe experimental data equally well under nominal stimulation levels [20]. The population heterogeneity that arises from population noise is modeled by time-invariant, probabilistic distribution of kinetic parameters of the underlying biochemical reactions of the signaling pathway.

3.2 Modeling population noise in the JAK2/STAT5 signaling pathway

The family of JAK/STAT signaling pathways is ubiquitous in mammalian cells, and is of relevance to the fields of stem cell and cancer research - its activation leads to modulation of expression of genes involved in growth, differentiation, migration, apoptosis, and other vital cellular processes [4, 5, 7, 9, 15]. JAK/STAT signaling can be stimulated with a large range of hormones, cytokines, and growth factors [16]. In mammals, four different JAKs and seven different STATs exist [16]. This work focuses on the eryhropoietin-stimulated pathway involving JAK2 and STAT5 (i.e., JAK2/STAT5 pathway), which is critical for growth and differentiation of hematopoietic progenitor cells [3, 21]. The key components of the signaling system include: (i) Erythropoietin Receptor (EpoR), which is a transmembrane receptor that interacts with the extracellular pathway ligand, (ii) Janus Kinase 2 (JAK2), which interacts with the cytosolic domain of EpoR, and (iii) Signal Transducer and Activator of Transcription 5 (STAT5), which can be in monomeric or dimeric form.

A schematic diagram of the JAK2/STAT5 signaling system is shown in Figure 3.1. EpoR subunits undergo multimerization upon stimulation by erythropoietin [16]. The cytoplasmic domain of each EpoR monomer interacts with a JAK2 tyrosine kinase. When two JAK2 proteins transphosphorylate each other, they become functionally active and induce phosphorylation of EpoR and other downstream signaling biomolecules, including STAT5 monomers. Phosphorylated STAT5 monomers dimerize in the cytoplasm and translocate into the nucleus, where they act as a transcription factor for upregulation of target genes. The dynamics of the JAK2/STAT5 signaling system are typically described by the four forms of the STAT5 molecules in the pathway [20]: unphosphorylated cytoplasmic STAT5 monomer x_1, phosphorylated cytoplasmic STAT5 monomer x_2, cytoplasmic STAT5 dimer x_3, and the transcriptionally active nuclear STAT5 dimer x_4. The input to the pathway, u, consists of the level of phosphorylation of EpoR, which can be stimulated by extracellular erythropoietin (Epo).

Swameye et al. [20] have reported several model structures for the JAK2/STAT5 signaling pathway in murine B cells. This work considers two competing models of the signaling pathway (denoted by superscripts [1] and [2]) that could describe the population-average experimental results reported in [20] with comparable statistical significance. Model 1, which assumes nucleocytoplasmic shuttling of STAT5 in the cell, takes the form

$$\begin{aligned}
\dot{x}_1^{[1]} &= -k_1 x_1^{[1]} u(t) + 2k_4 x_3^{[1]}(t-\tau), & x_1^{[1]}(0) &= x_{1,0}\\
\dot{x}_2^{[1]} &= -k_2 (x_2^{[1]})^2 + k_1 x_1^{[1]} u(t), & x_2^{[1]}(0) &= x_{2,0}\\
\dot{x}_3^{[1]} &= -k_3 x_3^{[1]} + \frac{1}{2} k_2 (x_2^{[1]})^2, & x_3^{[1]}(0) &= x_{3,0}\\
\dot{x}_4^{[1]} &= k_3 x_3^{[1]} - k_4 x_3^{[1]}(t-\tau), & x_4^{[1]}(0) &= x_{4,0}.
\end{aligned} \tag{3.1}$$

Model 2, which also assumes nucleocytoplasmic shuttling while accounting for a bidirectional kinetic flow in the dimerization of STAT, is described by

$$\begin{aligned}
\dot{x}_1^{[2]} &= -k_1 x_1^{[2]} u(t) + 2k_4 x_3^{[2]}(t-\tau), & x_1^{[2]}(0) &= x_{1,0}\\
\dot{x}_2^{[2]} &= -k_2 (x_2^{[2]})^2 + k_1 x_1^{[2]} u(t) + 2k_{-2} x_3^{[2]}, & x_2^{[2]}(0) &= x_{2,0}\\
\dot{x}_3^{[2]} &= -k_3 x_3^{[2]} + \frac{1}{2} k_2 (x_2^{[2]})^2 - k_{-2} x_3^{[2]}, & x_3^{[2]}(0) &= x_{3,0}\\
\dot{x}_4^{[2]} &= k_3 x_3^{[2]} - k_4 x_3^{[2]}(t-\tau), & x_4^{[2]}(0) &= x_{4,0}.
\end{aligned} \tag{3.2}$$

In (3.1) and (3.2), k_1 and k_2 are the rates of phosphorylation and dimerization of STAT, k_{-2} is the rate of spontaneous disintegration of STAT dimers into monomers, k_3 and k_4 are the rates of transport in and out of the nucleus, respectively, and τ is the residence time of STAT dimers in the nucleus. The (measurable) outputs of the pathway are phosphorylated STAT5 and total STAT5 in the cytoplasm, which

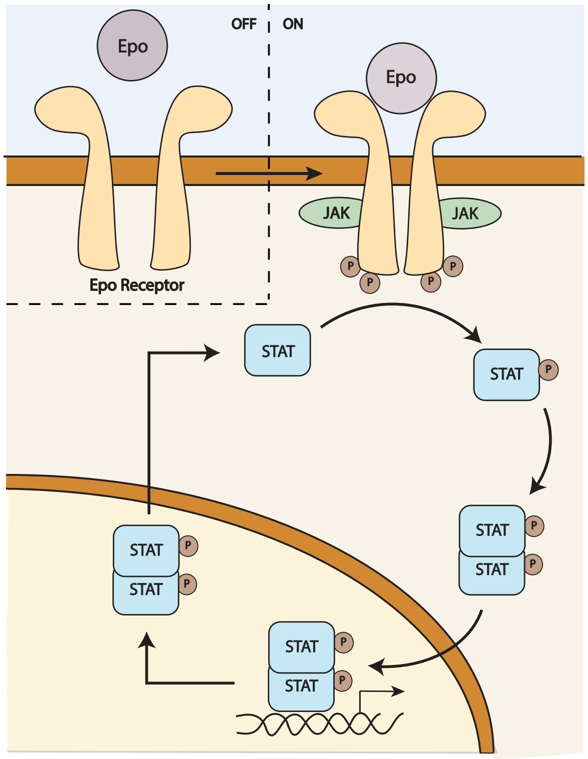

Figure 3.1: Depiction of the JAK2/STAT5 signaling pathway [16].

are denoted by $y_1 = k_5(x_2 + 2x_3)$ and $y_2 = k_6(x_1 + x_2 + 2x_3)$, respectively; k_5 and k_6 being scaling parameters. In a population of cells, cell-to-cell differences in kinetic rates or initial conditions can give rise to population noise. To capture population noise in the competing models (3.1) and (3.2), the kinetic rates and time-delay parameter are assumed to take beta distributions across the population of cells. The bounds of the beta distributions were adopted from the confidence intervals provided in [20]. The time-invariant differences in the kinetic and time-delay characteristics of individual cells in a population will give rise to noise in constituents of the signaling pathway (and thus, in the outputs y_1 and y_2). The impact of each kinetic and time-delay mechanism on heterogeneity of pathway outputs is likely to be different, reflecting their unequal contributions to the pathway dynamics. An analysis of the impact of variability in the kinetic rates and time-delay parameter on population noise can provide insight into molecular mechanisms of generation of population noise.

For illustrative purposes, the effect of probabilistic variability in k_3 on population noise is presented in Figure 3.2, which shows snapshots of population noise in the output y_1 at time instant 60 min (predicted by Model 1). The distributions of y_1 are constructed based on propagating the probability distributions of the kinetic rates and time-delay parameter through the pathway dynamics using Monte Carlo

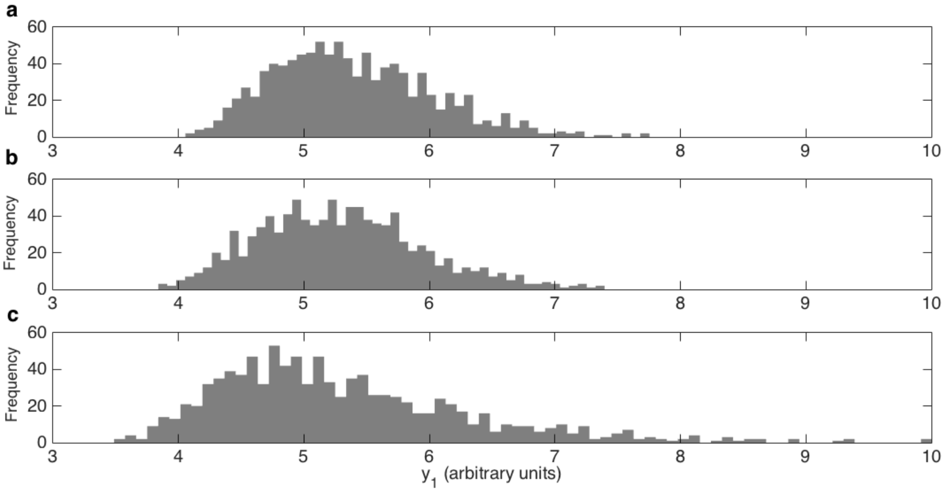

Figure 3.2: Distribution of the cytoplasmic phoshporylated STAT5 (y_1) within a population of cells at 60 min when k_3 takes a point estimate (a) and k_3 is distributed with the same (b) and triple (c) the standard deviation reported in [20]. The variability of other kinetic parameters among individual cells is accounted for.

simulations. In Figure 3.2a, k_3 is assigned a population-average value, whereas in Figs. 3.2b and 3.2c, k_3 is distributed respectively with the same and triple the standard deviation reported in [20]. It is evident that the variability of k_3 across individual cells of a population affects the population noise of the output y_1.

3.3 Probabilistic optimal experiment design problem

The probabilistic OED problem is formulated as a nonlinear optimization problem that aims to maximize the dissimilarity between probability distributions of outputs of competing models [14, 18]. The optimization problem is subject to input and state constraints. The Kolmogorov distance [10] is used as a metric to quantify (dis)similarity between outputs of competing models. The generalized polynomial chaos (gPC) framework is used for propagation of population heterogeneity over pathway dynamics [23]. In gPC, each (probabilistic) system state is approximated by an expansion of orthogonal polynomial basis functions, which are defined based on the known descriptions of probabilistic uncertainties. The statistical moments of stochastic states can be efficiently computed from the coefficients of PC expansions or, alternatively, PC expansions can be used as a surrogate for the nonlinear system model to efficiently perform Monte Carlo simulations.

3.4 Model discrimination in the presence of population noise

This work seeks to discriminate between predictions of the competing models (3.1) and (3.2) in the presence of population noise. Figure 3.3 shows the dynamics of cytoplasmic phosphorylated STAT5 (y_1) predicted by the competing models upon EpoR stimulation with a nominal level of extracellular erythropoietin. The predictions of population-average y_1 by the competing models (depicted by continuous profiles) are very similar, which is consistent with the results reported in [20]. More importantly, the predicted distributions of y_1 (due to population noise), which can be seen from distribution snapshots in Figure 3.3 (at times 30, 45, and 60 min), overlap significantly. This makes the discrimination between predictions of models (3.1) and (3.2) impractical under the nominal stimulation level.

A probabilistic optimal experiment design approach [14, 18] is employed for discriminating between the competing models (3.1) and (3.2) in the presence of population noise. The distinct feature of the probabilistic OED approach lies in the ability to systematically account for population noise generation in the signaling pathway when the optimal input for model discrimination is designed. The probabilistic OED approach aims to separate the predicted distributions of pathway outputs within a population of cells. If the predicted distributions of one of the outputs are separated at least at one time instant, the competing models can be discriminated. Figure 3.4a shows the evolution of distribution of y_1 when the optimal input designed by probabilistic OED is applied to the competing models (3.1) and (3.2) under 1,000 realizations of the sources of population noise generation. The optimal input allows for discriminating between the predicted distributions of y_1 within a cell population at 60 min.

The performance of the optimal input designed by probabilistic OED is compared to that of standard OED (e.g., [11]), which disregards the time-invariant cell-to-cell differences within a population of cells. The optimal input designed by standard OED is able to separate the population-average predictions of the competing models (compare Figure 3.4b with Figure 3.3). However, it fails to effectively discriminate between the predicted distributions of y_1 across the cell population.

3.5 Discussion

The use of computational models in combination with state-of-the-art measurement techniques can significantly improve our understanding of cellular processes that are prone to cell-to-cell variability. The above discussed OED approaches are intended to complement experimental efforts, in particular when collecting rich experimental data is limited by various technical constraints and/or availability of resources. OED tools are invaluable for reducing the experimental effort required for elucidation of competing hypotheses of molecular mechanisms in biological systems. Accounting for the presence of cell-to-cell variability in a population of cells can be critical to achieving accurate quantitative representations of the underlying dynamics of a cell population. Traditional quantitative methods (such as western blotting and qPCR) typically yield end-point, population-average information and, consequently, can mask heterogeneity within a population [22]. On the other hand, real-time, single-cell resolution techniques can surface cell-to-cell differences and, if high throughput, can provide a complete picture of population heterogeneity [22]. The above discussed probabilistic OED approach enables discrimination between competing models regardless of the employed measurement technique (population-average or single-cell measurement). Comparison of the results of standard and probabilistic OED reveals that accounting for population heterogeneity allows for more effective experiment designs to achieve model discrimination with high confidence. In the presented case study, probabilistic OED not only achieved a better discrimination of the population-average behavior of the outputs as compared to standard OED, but also led to separation of the distributions of outputs of the competing models.

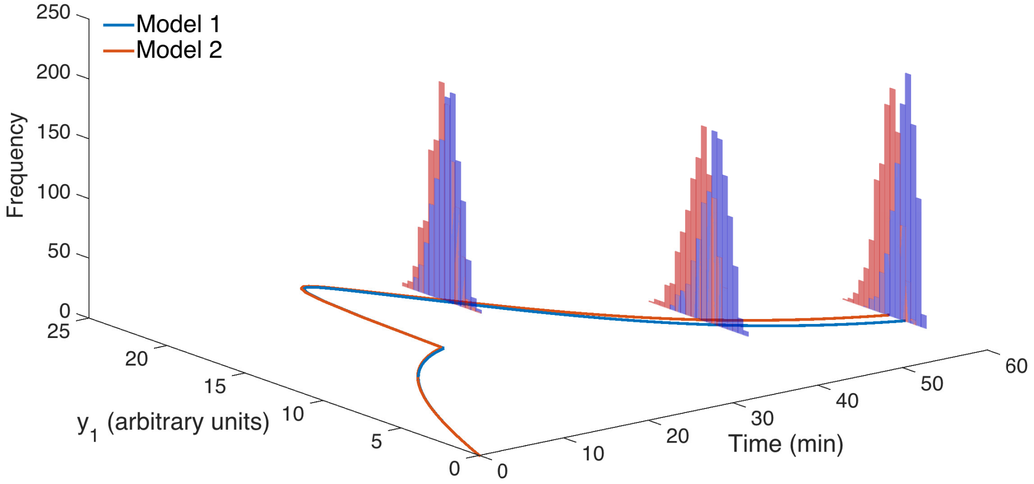

Figure 3.3: Evolution of distributions of cytoplasmic phosphorylated STAT5 (y_1) within a cell population under nominal stimulation as predicted by the competing models 3.1 and 3.2. The continuous profiles depict the population-average behavior and the histograms depict the distribution of y_1 within a population of cells. Distributions correspond to time instants 30, 45, and 60 min.

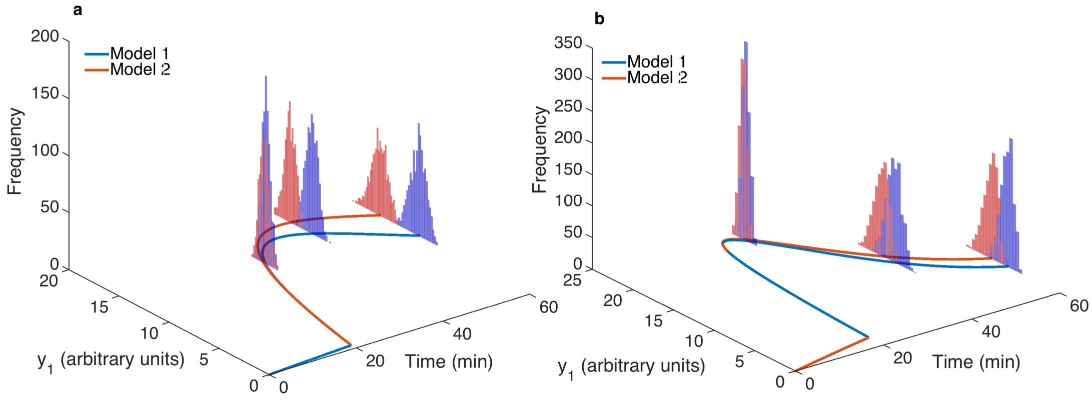

Figure 3.4: Evolution of distributions of cytoplasmic phoshporylated STAT5 (y_1) within a cell population (predicted by the competing models (3.1) and (3.2)) when the signaling pathway is stimulated with the optimal input designed by (a) probabilistic OED and (b) standard OED. The histograms are generated by applying the optimal inputs to (3.1) and (3.2) under 1,000 realizations of the sources of population noise generation. The continuous profiles depict the average y_1 within the cell population. Distributions correspond to time instants 30, 45, and 60 min after initiation of stimulation.

3.6 Bibliography

1. A. Brock, H. Chang, and S. Huang. Non-genetic heterogeneity - A mutation-independent driving force for the somatic evolution of tumours. *Nature Perspectives*, 10:336–342, 2009.

2. H. Chuang, M. Hofree, and T. Ideker. A decade of systems biology. *Annual Review of Cell and Developmental Biology*, 26:721–744, 2010.

3. S. Constantinescu, S. Ghaffari, and H. Lodish. The Erythropoietin receptor: Structure, activation and intracellular signal transduction. *Trends in Endocrinology and Metabolism*, 10:18–23, 1999.

4. K. Ghoreschi, M. Jesson, X. Li, J. Lee, S. Ghosh, J. Alsup, J. Warner, M. Tanaka, S. Steward-Tharp, M. Gadina, et al. Modulation of innate and adaptive immune responses by tofacitinib. *Journal of Immunology*, 186:4234–4243, 2011.

5. K. Ghoreschi, A. Laurence, and J. O'Shea. Janus kinases in immune cell signaling. *Immunology Review*, 228:273–287, 2009.

6. T. Graf and M. Stadtfeld. Heterogeneity of embryonic and adult stem cells. *Cell Stem Cell*, 3:480–483, 2008.

7. D. Harrison. The JAK/STAT pathway. *Cold Spring Harbor Perspectives in Biology*, 2012.

8. S. Huang. Non-genetic heterogeneity of cells in development: More than just noise. *Development*, 136:3853–3862, 2009.

9. S. Jatiani, S. Baker, L. Silverman, and E. Reddy. JAK/STAT pathways in cytokine signaling and myeloproliferative disorders: Approaches for targeted therapies. *Genes and Cancer*, pages 979–993, 2010.

10. A. Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell'Istituto Italiano degli Attuari*, 4:83–91, 1933.

11. C. Kreutz and J. Timmer. Systems biology: Experimental design. *FEBS Journal*, 276:923–942, 2008.

12. M. Martin-Casas and A. Mesbah. Discrimination between competing model structures of biological systems in the presence of population heterogeneity. *IEEE Life Sciences Letters*, 2:23–26, 2016.

13. B. Melykuti, E. August, A. Papachristodoulou, and H. El-Samad. Discriminating between rival biochemical network models: Three approaches to optimal experiment design. *BMC Systems Biology*, 4, 2010.

14. A. Mesbah and S. Streif. A probabilistic approach to robust optimal experiment design with chance constraints. In *Proceedings of the 9th International Symposium on Advanced Control of Chemical Processes (ADCHEM)*, pages 100–105. Whistler, Canada, 2015.

15. J. O'Shea and R. Plenge. JAKs and STATs in immunoregulation and immune-mediated disease. *Immunity*, 36:542–550, 2012.

16. J. Rawlings, K. Rosler, and D. Harrison. The JAK/STAT signaling pathway. *Journal of Cell Science*, 117:1281–1283, 2004.

17. D. Skanda and D. Lebiedz. An optimal experimental design approach to model discrimination in dynamic biochemical systems. *Bioinformatics*, 26:939–945, 2010.

18. S. Streif, F. Petzke, A. Mesbah, R. Findeisen, and R. Braatz. Optimal experimental design for probabilistic model discrimination using polynomial chaos. *In Proceedings of the IFAC World Congress*, pages 4103–4109, 2014, Cape Town.

19. G. Suel, R. Kulkarni, J. Dworkin, J. Garcia-Ojalvo, and M. Elowitz. Tunability and noise dependence in differentiation dynamics. *Science*, 315:1716–1719, 2007.

20. I. Swameye, T. M¸ller, J. Timmer, O. Sandra, and U. Klingmuller. Identification of nucleocytoplasmic cycling as a remote sensor in cellular signaling by databased modeling. *Proceedings of the National Academy of Sciences*, 100:1028–1033, 2003.

21. W. Vainchenker and S. Constantinescu. JAK/STAT signaling in hematological malignancies. *Oncogene*, 32:2601–2613, 2013.

22. J. Wu and E. Tzanakakis. Deconstructing stem cell population heterogeneity: Single-cell analysis and modeling approaches. *Biotechnology Advances*, 31:1047–1062, 2013.

23. D. Xiu and G. E. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal of Scientific Computation*, 24:619–644, 2002.

Chapter 4

Optimal experiment design for online model structure selection

Online model structure selection is imperative for gaining insight into experiment execution or operation as it takes place. In this chapter, we present two computationally efficient optimal experiment design methods that can discriminate between competing model structures in online settings. A useful application of the presented online optimal experiment design methods is fault diagnosis, which consists in the discrimination of models describing normal and faulty operation in experimental or industrial settings. Fault diagnosis is achieved by designing an experiment that minimizes tractable optimization objectives, which allows to maximize the probability of selecting the most representative model structure (i.e., diagnose the fault) in real time. Simulation results on a continuous bioreactor demonstrate the effectiveness of both methods for online model structure selection for fault diagnosis in practical settings.[*]

4.1 Introduction

In Chapter 3, we introduced a method to design experiments that enables elucidation of biochemical networks in an offline manner. The offline nature of the optimal experiment design (OED) method allowed for full reconstruction of all model outcomes as a probability distribution, given that computational cost was not a limitation. Conversely, computational efficiency is a key feature of online OED frameworks, as they must allow for real-time learning of the system under study. An emerging application of online OED is fault detection and diagnosis, where several plausible models describing normal and faulty modes of system operation are discriminated in real time.

---

[*] The content of this chapter was published in [35] and [26]

Online OED for model selection in the context of fault detection and diagnosis allows for maintaining stable, reliable, and profitable operation of technical systems in the presence of component malfunctions, drifting system parameters, and other abnormal events [32]. However, nonlinear system dynamics as well as system uncertainty, disturbances, and measurement noise often render model selection a challenging task in practical applications. OED for online model selection in the context of fault diagnosis is further compounded by the growing complexity of technical systems and the increasingly stringent requirements on their high-performance operation, which necessitate retaining the system within admissible operational constraints during fault diagnosis.

To address the aforementioned challenges, a wide range of methods for fault detection and diagnosis have been developed including residual- and observer-based methods [13, 16, 34], set-based approaches [32], and data-based methods [13]. The vast majority of these methods are *passive*, where the status of the system is deduced only from measurements acquired during nominal operation, which are then compared with model predictions or historical data. However, system uncertainties or the corrective actions of a controller, when the system is under feedback control, can often obscure faulty behavior to an extent that impairs reliable fault diagnosis. In these cases, auxiliary input signals can be applied to the system to enhance the detectability and isolability of faults in the system measurements [11]. This notion has led to development of OED for online model selection for *active* fault diagnosis in which an input sequence to the system is designed to diagnose the system faults within a given time window in the presence of uncertainty. [8]

For systems with deterministic, bounded uncertainty, OED methods have been reported for designing input sequences that are “robust” to *worst-case* realizations of system uncertainty. These methods, which are typically developed for linear systems, look to minimize the energy of the input sequence while ensuring separation of the reachable sets of all model hypotheses [4, 11, 40]. In [41], a tractable approach was presented for solving the online problem for linear systems subject to bounded disturbances and measurement noise by using zonotopes for efficient computation. OED for nonlinear systems with sufficiently small nonlinearity was addressed in [3]. In [38], a set-based approach for providing ”robustness certificates” was used for model selection of nonlinear systems.

When probabilistic descriptions of system uncertainty are available, a more natural approach to OED for online model selection involves accounting for the stochastic system uncertainty in designing the input sequence. In a stochastic setting for OED, the uncertainty in model parameters and initial conditions, as well as system disturbances and measurement noise, can be characterized in terms of probability distributions instead of sets [47]. By accounting for distributional information in the OED problem, the designed input sequences can handle system uncertainty in a (possibly) less conservative manner than OED methods that adopt a set-based description of uncertainty. In the context of fault diagnosis, OED for model selection for linear systems subject to additive stochastic disturbances and measurement noise has been explored in Refs. [6, 7], and further extended to jump Markov linear systems in [5]. These methods use an upper bound on the probability of model discrimination error, defined in terms of the Bhattacharyya coefficient [22], as the optimization criteria [22]. In [28], an OED method was presented for nonlinear systems with probabilistic uncertainty in model parameters and initial conditions, where generalized polynomial chaos [45] was used for obtaining a tractable surrogate problem. A tractable OED approach for nonlinear systems was presented in [29], where results from randomized optimization were used to provide theoretical guarantees for fault detection.

This chapter addresses the problem of OED for online model selection for nonlinear systems subject to white-noise stochastic disturbances and measurement noise, with probabilistic model uncertainty in parameters and initial conditions. The probabilistic model uncertainty represents our incomplete knowledge of the system. Two computationally efficient methods are introduced.

In the first method, and similarly to [7], the probability of model hypothesis discrimination error is used as the tractable optimization criterion. Input and state constraints are enforced to guarantee safe system operation that is minimally intrusive during the fault diagnosis experiment. State constraints are enforced as a *chance constraint* to provide flexibility and minimal conservatism in dealing with stochastic system uncertainty. A key challenge in solving the proposed OED problem arises from efficient propagation of the stochastic system uncertainty through the nonlinear system dynamics. There exists no closed-form expression for evaluating the probability distribution of the predicted states variables as a function of the input sequence. This method adopts two uncertainty propagation methods, based on linearization and unscented transformation (UT), for tractable propagation of the moments of the state variables.

In the second method, non-intrusive generalized polynomial chaos (gPC) is combined with UT using conditional probability rules to jointly propagate both model uncertainty and stochastic disturbances. The proposed sample-based uncertainty propagation method, which draws from the benefits of gPC and UT for handling model uncertainty and disturbances, respectively, is particularly suitable for optimization, since it can approximate the statistics of system variables with a fairly small number of samples that are chosen systematically. Inspired by the k-nearest neighbors scheme, classically used for data clustering [14, 17], the optimization criterion is defined in terms of a sample-based measure of distance between probabilistic predictions of model hypotheses for the normal and faulty system operation.

The optimization criteria for both methods do not require the generation of histograms and thus circumvent the need for full probability distribution reconstruction. The performance of the proposed OED methods for online model selection in the context of fault diagnosis is demonstrated on a continuous bioreactor case study for online diagnosis of multiple operation scenarios.

The chapter is organized as follows. First, a general background on model selection is provided. The OED problem formulation is then presented, including the details of the model structure, the OED problem in a stochastic setting, and the main challenges in solving the latter problem. This is followed by the description of each presented OED for online model selection method, which includes the summary of the adopted uncertainty propagation techniques for approximating moments of the state/output probability distributions, the description and derivation of the optimization criteria, and discussions on the offline and online implementation of the tractable OED problems. Lastly, the performance of the proposed OED methods is demonstrated on a continuous bioreactor with multiple fault scenarios including operational and structural changes to the process.

Notation. Hereafter, $\mathbb{R}^n$ denotes the set of real vectors with n elements and $\mathbb{R}^{n\times m}$ denotes the set of real matrices with n rows and m columns. $\mathbb{N} = \{1, 2, \ldots\}$ is the set of natural numbers and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. $x_{a:b} = [x_a^\top, x_{a+1}^\top, \ldots, x_b^\top]^\top$ denotes a concatenated vector of the sequence of values of x from discrete time a to b. $\mathrm{diag}(\cdot)$ represents a block diagonal matrix. $P(A)$ denotes the probability of event A. For a random vector X, $p(X)$ denotes the probability density function (pdf) of X, $\mathbb{E}[X]$ or $\bar{X}$ denotes the expected value of X, and $\mathrm{cov}(X)$ or $\mathit{\Sigma}_X$ denotes the covariance matrix of X. For random vectors X and Y, the cross covariance between X and Y is denoted by $\mathrm{cov}(X, Y) = \mathit{\Sigma}_{XY} = \mathbb{E}[(X - \bar{X})(Y - \bar{Y})]$. $X \mid Y$ denotes that the random vector X is conditioned on realization Y, with conditional pdf $p(X \mid Y)$. The hat accent $\hat{\cdot}$ is used to denote estimated or approximate variables.

4.2 Background: Model Selection

Consider a finite number of possible model hypotheses $\{H_i\}$ with known prior probabilities $P(H_i)$. We look to investigate the decision-theory problem of classifying a system observation $\mathbf{y}$, under a given input $\mathbf{u}$, as coming from one of the model hypotheses $\{H_i\}$. It is well-known that the decision rule that minimizes the probability of error, denoted by $P(error)$, in hypothesis selection is the Bayes decision rule, which aims at selecting the hypothesis with the largest *a posteriori* probability [14]. The problem of Bayesian hypothesis selection (aka model selection) can be expressed as follows

$$\text{select } H_i \text{ such that } i = \text{argmax}_j \; P(H_j \mid \mathbf{y}, \mathbf{u}). \tag{4.1}$$

According to Bayes rule, $P(H_j \mid \mathbf{y}, \mathbf{u}) = p(H_j, \mathbf{y} \mid \mathbf{u})/p(\mathbf{y} \mid \mathbf{u}) = p(\mathbf{y} \mid H_j, \mathbf{u})P(H_j)/p(\mathbf{y} \mid \mathbf{u})$, where the denominator can be neglected in the optimization Problem (4.1) since it is independent of H_j. As shown in Figure 4.1, the Bayes decision rule yields a number of decision regions $\mathfrak{R}_i$ that are defined with respect to the observation space, i.e.,

$$\mathfrak{R}_i = \{\mathbf{y} : p(H_i, \mathbf{y} \mid \mathbf{u}) > p(H_l, \mathbf{y} \mid \mathbf{u}), \; \forall l \neq i\}. \tag{4.2}$$

If $\mathbf{y} \in \mathfrak{R}_i$, then hypothesis H_i is selected. This implies that $P(\mathbf{y} \in \mathfrak{R}_j, H_i \mid \mathbf{u})$ is the probability of incorrectly selecting the model hypothesis j when the true hypothesis is i. By adding all these probabilities across all possible model pairs, the following expression is derived for the probability of selecting the incorrect hypothesis, known as the *Bayes risk*

$$\begin{aligned} P(error) &= \sum_i \sum_{j \neq i} P(\mathbf{y} \in \mathfrak{R}_j, H_i \mid \mathbf{u}) \\ &= \sum_i \sum_{j \neq i} P(\mathbf{y} \in \mathfrak{R}_j \mid H_i, \mathbf{u}) P(H_i) \\ &= \sum_i \sum_{j \neq i} \int_{\mathfrak{R}_j} p(\mathbf{y} \mid H_i, \mathbf{u}) P(H_i) d\mathbf{y}. \end{aligned} \tag{4.3}$$

In fact, the Bayes risk can be used for quantifying the probability of error of the optimal hypothesis selection. The Bayes risk is dependent on the system inputs $\mathbf{u}$. As such, it is clearly advantageous to select an input trajectory that minimizes the Bayes risk. The latter will be the focus of this chapter for nonlinear systems with stochastic uncertainty.

Consider an uncertain, nonlinear system described by a discrete-time model of general form

$$H : \begin{cases} x_{k+1} = f(x_k, u_k, w_k, \theta) \\ y_k = h(x_k, v_k, \theta), \end{cases} \tag{4.4}$$

where $k \in \mathbb{N}_0$ is the discrete time index; $x \in \mathbb{R}^{n_x}$, $u \in \mathbb{R}^{n_u}$, and $y \in \mathbb{R}^{n_y}$ denote the system states, manipulated inputs, and measured outputs, respectively; $w \in \mathbb{R}^{n_w}$ denotes the system disturbances; $v \in \mathbb{R}^{n_v}$ denotes measurement noise; $\theta \in \mathbb{R}^{n_\theta}$ denotes model parameters; and $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \times \mathbb{R}^{n_\theta} \to \mathbb{R}^{n_x}$ and $h : \mathbb{R}^{n_x} \times \mathbb{R}^{n_v} \times \mathbb{R}^{n_\theta} \to \mathbb{R}^{n_y}$ describe the (possibly) nonlinear system dynamics and measurement functions, respectively.

Due to imperfect knowledge of the system, the initial conditions x_0 and parameters θ in (4.4) are modeled as probabilistic uncertainties with (possibly) joint pdf $p(x_0, \theta)$. The system disturbances $\{w_k\}_{k\in\mathbb{N}_0}$ and measurement noise $\{v_k\}_{k\in\mathbb{N}_0}$ satisfy the following conditions:

1. w_k and v_k are independent and identically distributed (i.i.d.) sequences with known pdfs $p(w_k)$ and $p(v_k)$, respectively;
2. w_l, and v_t are mutually independent (i.e., white-noise random sequences) for all $l, t \in \mathbb{N}_0$, and are independent of x_0 and θ.

The system is subject to hard input constraints $u_k \in \mathcal{U}$, where the set $\mathcal{U} \subset \mathbb{R}^{n_u}$ is assumed to be closed and bounded. The system is also subject to (possibly nonlinear) state constraints $x_k \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$. Since the system evolves as a stochastic process, the state constraints are enforced as a *chance constraint* of the form

$$P(x_k \in \mathcal{X} \mid H) \geq 1 - \epsilon, \qquad k \in \mathbb{N}, \tag{4.5}$$

where $\mathcal{X}$ is a collection of n_c inequality state constraints; $P(x \in X \mid y) = \int_X p(x \mid y)dx$ denotes the probability that the event $x \in X$ occurs given a realization of the random variable y; and ϵ is the maximum allowed probability of state constraint violation.

Now, consider that a finite number n_f of models, with known dynamics of the form (4.4), exist to describe potential fault scenarios of the system. Denote the set of models by $\mathcal{H} \triangleq \{H_0, H_1, \dots, H_{n_f}\}$, where H_i represents the system information

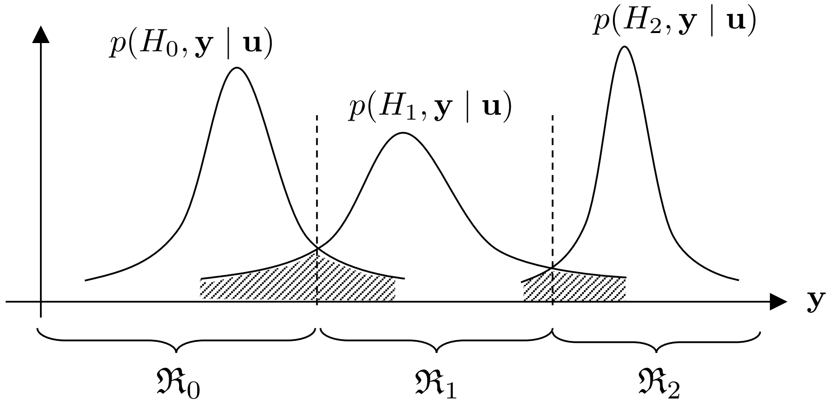

Figure 4.1: Illustration of Bayesian hypothesis selection for a given input $\mathbf{u}$. In the case of multiple models, hypothesis selection generally results in a number of decision regions in the observation space. If the observation $\mathbf{y}$ lies in the region $\mathfrak{R}_i$ (i.e., $\mathbf{y} \in \mathfrak{R}_i$), then the hypothesis selection yields hypothesis H_i. Note that even when the Bayes hypothesis selection is optimal, there is a finite probability of error related to overlap of the likelihood distributions (shown by the shaded region). Adapted from [7].

required for specifying the i^{th} model

$$H_i = \{f^{(i)}, h^{(i)}, p^{(i)}(x_0, \theta), p^{(i)}(w_k), p^{(i)}(v_k), \mathcal{X}^{(i)}, \epsilon^{(i)}\}.$$

The superscript (i) is used to distinguish variables in (4.4) for each model $H_i \in \mathcal{H}$. The *nominal* system model (i.e., no faulty behavior) is denoted by the superscript $i = 0$, while $i = 1, \ldots, n_f$ refers to the different *fault* models. Note that all models must have the same manipulated inputs u and measured outputs y, suggesting that n_u and n_y are the same for all $H_i \in \mathcal{H}$. However, the other model attributes such as the state and output functions, state constraints, and model uncertainty descriptions can be different for the various $H_i \in \mathcal{H}$. For notational convenience, the superscripts (i) are dropped in the remainder of the chapter unless otherwise stated.

Remark 1 *The prior $P(H_i)$ for each model can be determined in a number of ways. For example, there may exist explicit knowledge about the likelihood of the different model hypotheses, or the priors can represent the system belief from an estimator. If no a-priori knowledge is available, an uninformative prior can be utilized such that all models are considered to have an equal probability of occurrence, i.e., $P(H_i) = 1/(n_f + 1)$ for all $H_i \in \mathcal{H}$.*

The model selection framework presented above readily applies by defining the observations $\mathbf{y}$ to be a collection of system measurements and the input trajectory $\mathbf{u}$ as the manipulated inputs.

4.2.1 OED for model structure selection

The goal of this work is to present an OED framework for designing an input sequence that minimizes the probability of model selection error in $\mathcal{H}$, while satisfying the input and state constraints of the system. Due to the stochastic nature of system uncertainty in (4.4), the optimal input sequence should be designed such that the system evolves through a trajectory of state *probability distributions* that minimizes the Bayes risk [6]. For a finite (prespecified) planning horizon $N \in \mathbb{N}$, the OED problem can be stated as follows.

Problem 1 (OED for model structure selection) *Determine the input sequence $u_{0:N-1}$ that minimizes the probability of selecting the wrong model $P(error)$ while satisfying input constraints $u_k \in \mathcal{U}$, $k = 0, \ldots, N-1$ and chance constraints* (4.5) *for all possible models $H_i \in \mathcal{H}$.*

The designed input sequence should then be applied to the system to observe its behavior and, accordingly, select the model hypothesis that best (in a probabilistic sense) describes the measurements. The constraints $\mathcal{U}$ and $\mathcal{X}$ must be defined such that the designed input sequence ensures safe and least intrusive system operation during the experiment. Note that the chance constraints (4.5) can be converted to hard state constraints by setting $\epsilon = 0$ when the stochastic system uncertainty is bounded. In addition, the input constraints can be straightforwardly modified to limit the rate of input changes or the total input energy [18, 39].

OED can be performed in an *open-loop* fashion where the optimization in Problem 1 is solved only once, *offline*, yielding an input sequence that is subsequently applied to the system. In offline OED, system measurements collected during the implementation of the optimal input sequence are not used for re-planning of the experiment. This can reduce the effectiveness of the designed input sequence due to unmodeled system uncertainties and disturbances in the models $H_i \in \mathcal{H}$. Alternatively, the OED problem can be solved *online* in a *closed-loop* fashion. In this case, Problem 1 is solved repeatedly at every measurement sampling time, when new system information becomes available. At each sampling time, only the first element of the optimal input sequence is applied to the system, and the new measurements are used to re-plan the experiment via recomputing the optimal input sequence. The so-called *receding-horizon* implementation of the optimal input sequence allows for incorporating measurement feedback into the OED problem. Thus, online OED can provide some degree of robustness to unmodeled system uncertainties and disturbances. Note that online OED requires solving Problem 1 quicker than the measurement sampling intervals so that the designed input sequences can be implemented in real time.

4.2.2 Challenges of OED for model structure selection

To solve Problem 1, the following challenges must be addressed:

1. **Stochastic uncertainty propagation**: Stochastic uncertainties x_0, θ, and $w_{0:N-1}$ must be propagated through the nonlinear dynamics of each model $H_i \in \mathcal{H}$ to evaluate the distribution of the concatenated observation vector $p(y_{0:N} \mid H_i, u_{0:N-1})$. This is required for computing the Bayes risk $P(error)$ in (4.3). Traditional sample-based approaches to uncertainty propagation, such as Markov-Chain Monte Carlo [15], are prohibitively expensive and ill-suited for gradient-based optimization methods.
2. **Computation of the Bayes risk** $P(error)$: There exist no closed-form expressions for the distribution $p(y_{0:N} \mid H_i, u_{0:N-1})$ or the decision regions $\mathfrak{R}_i$, making evaluation of $P(error)$ particularly challenging. This is further compounded by the multivariate integral in the Bayes risk.

3. **Chance constraint evaluation**: Chance constraints are generally nonconvex and intractable [10, 31]. This is because chance constraint evaluation involves solving a multivariate integral of state pdfs, for which no closed-form expression exists.

Although construction of the multivariate distribution $p(y_{0:N} \mid H_i, u_{0:N-1})$ is inherently difficult, $p(y_{0:N} \mid H_i, u_{0:N-1})$ can be efficiently approximated in terms of its moments obtained from (i) linearization [24], (ii) the unscented transform [44], or (iii) generalized polynomial chaos (gPC) [45]. The first method presented in this chapter adopts linearization and UT for moment-based uncertainty propagation, whereas the second method adopts UT and gPC. To address the second challenge stated above, an upper bound of the Bayes risk $P(error)$ and a metric of the k-nearest neighbors distance are used as tractable objectives amenable to online optimization. Lastly, an approximate, moment-based surrogate is derived for the chance constraint (4.5) for both methods to avoid sample-based integration or construction of the state pdfs. In the remainder of the chapter, the methods used for obtaining a tractable surrogate for Problem 1 are outlined. The tractable OED problems are then presented both in their offline and online implementations.

4.3 Method 1: OED for online model selection using the Bayes Risk

This section presents a computationally tractable surrogate for Problem 1 and the used uncertainty propagation methods to approximate state and output statistics. The proposed OED problem relies on a bound on the Bayes risk, $P(error)$, which is derived for discrimination between multiple model hypotheses. The bound on $P(error)$ is used as the fault diagnosis criterion for posing the optimization in Problem 1, where the chance constraint is approximated using a moment-based surrogate. Subsequently, an algorithm is presented for receding-horizon implementation of the optimization problem, which enables online re-planning of the fault diagnosis experiment. Lastly, the OED method is demonstrated on a bioreactor case study.

4.3.1 Uncertainty Propagation

A closed-form solution for exact propagation of stochastic system uncertainty (i.e., uncertain parameters and/or disturbances) through general nonlinear models does not exist [12]. Sampled-based approaches to uncertainty propagation can be expensive for optimization applications since they require many model simulation runs that must be repeated for any candidate decision variables. In this work, we look to propagate moments of $p(y_{0:N} \mid H_i, u_{0:N-1})$ to avoid sample-based construction of the distribution. To reduce the computational complexity of Problem 1, $p(y_{0:N} \mid H_i, u_{0:N-1})$ is approximated as a multivariate Gaussian distribution so that it can be fully represented in terms of its first two moments. Two approaches are presented for obtaining closed-form expressions for the moments of $p(y_{0:N} \mid H_i, u_{0:N-1})$ as a function of the input sequence. The first approach relies on linearization of the nonlinear model equations, which is a common approach used in, for example, the development of the extended Kalman filter (EKF) [24]. On the other hand, in the second approach linearization is circumvented by using the *unscented transformation* (UT) for calculating statistics of a random variable that undergoes a nonlinear transformation [44]. The UT method is particularly useful when the system dynamics are highly nonlinear, so that the linearization method provides a poor approximation of the nonlinear dynamics. In the UT method, a set of samples, called *sigma points*, are selected around the mean of the random variables. The sigma points are then propagated through the nonlinear model equations to construct new mean and covariance estimates for the random variables. A key advantage of the UT method is that Jacobian calculations, which can be cumbersome for complex functions or even impractical for non-smooth/non-differentiable functions, are no longer required. In the following, the linearization and UT methods for propagating the moments of $p(y_{0:N} \mid H_i, u_{0:N-1})$ are presented for the augmented state and parameter system dynamics.

Augmenting States with Parameters

To simplify presentation, a new "augmented" state vector is defined as

$$z_k = \left[x_k^\top,\ \theta_k^\top\right]^\top,$$

where the system dynamics (4.4) are augmented with artificial "parameter dynamics" $\theta_{k+1} = \theta_k$ that enforce the time-invariant nature of the unknown parameters. The augmented system dynamics are described by

$$\begin{aligned} z_{k+1} = \begin{bmatrix} x_{k+1} \\ \theta_{k+1} \end{bmatrix} = \begin{bmatrix} f(x_k, u_k, w_k, \theta_k) \\ \theta_k \end{bmatrix} &= f_z(z_k, u_k, w_k) && \text{(4.6a)} \\ y_k = h(x_k, v_k, \theta_k) &= h_z(z_k, v_k), && \text{(4.6b)} \end{aligned}$$

$f_z : \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_z}$ and $h_z : \mathbb{R}^{n_z} \times \mathbb{R}^{n_v} \to \mathbb{R}^{n_y}$ are defined appropriately, with $n_z = n_x + n_\theta$. $p(z_0) = p(x_0, \theta)$ denotes the pdf of the initial augmented state.

Linearization

The nonlinear dynamics of the augmented system (4.6a) can be linearized using a first-order Taylor series[‡]

$$z_{k+1} \approx f_z(\hat{z}_k, u_k, \bar{w}) + \left.\frac{\partial f_z}{\partial z_k}\right|_{\hat{z}_k, u_k, \bar{w}} (z_k - \hat{z}_k) + \left.\frac{\partial f_z}{\partial w_k}\right|_{\hat{z}_k, u_k, \bar{w}} (w_k - \bar{w}), \tag{4.7}$$

where $\hat{z}_k \approx \bar{z}_k = \mathbb{E}[z_k]$ denotes the approximated mean of the augmented states and $\bar{w} = \mathbb{E}[w_k]$. Taking the expectation of (4.7) yields the following recursion for the mean

$$\hat{z}_{k+1} = f_z(\hat{z}_k, u_k, \bar{w}), \qquad \hat{z}_0 = \mathbb{E}[z_0]. \tag{4.8}$$

Define $\delta z_k = z_k - \hat{z}_k$, so that (4.7) can be rewritten as a linear time-varying (LTV) system

$$\delta z_{k+1} = A_k \delta z_k + E_k \delta w_k, \tag{4.9}$$

where δz_0 has a known pdf; $\delta w_k = w_k - \bar{w}$; and

$$A_k = \left.\frac{\partial f_z}{\partial z_k}\right|_{\hat{z}_k, u_k, \bar{w}}, \qquad E_k = \left.\frac{\partial f_z}{\partial w_k}\right|_{\hat{z}_k, u_k, \bar{w}}.$$

Let $\hat{\Sigma}_{z_k} = \mathbb{E}[\delta z_k \delta z_k^\top] \approx \Sigma_{z_k}$ denote the approximate covariance of the augmented states. A recursive expression for this covariance can now be derived from (4.9)

$$\hat{\Sigma}_{z_{k+1}} = A_k \hat{\Sigma}_{z_k} A_k^\top + E_k \Sigma_w E_k^\top, \qquad \hat{\Sigma}_{z_0} = \Sigma_{z_0}. \tag{4.10}$$

Accordingly, the covariance of the augmented states over the planning horizon N (including cross-correlation in time) can be approximated by stacking (4.9) into

$$\delta z_{0:N} = \mathbf{A} \delta z_0 + \mathbf{E} \delta w_{0:N-1}, \tag{4.11}$$

where

$$\mathbf{A} = \begin{bmatrix} I \\ A_1 \\ A_1 A_2 \\ \vdots \\ \prod_{i=1}^{N-1} A_i \\ \prod_{i=1}^{N} A_i \end{bmatrix}, \qquad \mathbf{E} = \begin{bmatrix} 0 & 0 & \cdots & \cdots & 0 \\ I & 0 & & & 0 \\ A_1 & I & \ddots & & \vdots \\ A_2 A_1 & A_2 & I & \ddots & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ \prod_{i=1}^{N-1} A_i & \prod_{i=2}^{N-1} A_i & \cdots & A_{N-1} & I \end{bmatrix} \begin{bmatrix} E_0 \\ E_1 \\ \vdots \\ E_{N-2} \\ E_{N-1} \end{bmatrix}$$

[‡] The dynamics are linearized only with respect to the states and the disturbances (not the inputs) for improved accuracy. In this way, the dynamics are locally linearized around any candidate input sequence that captures variability in the covariance of the predicted outputs through changes in the Jacobian matrices.

such that

$$\hat{\Sigma}_{z_{0:N}} = \mathbf{A}\Sigma_{z_0}\mathbf{A}^\top + \mathbf{E}\Sigma_{w_{0:N-1}}\mathbf{E}^\top, \tag{4.12}$$

with $\Sigma_{w_{0:N-1}} = \mathrm{diag}(\Sigma_w, \ldots, \Sigma_w)$. Note that $\mathbf{A}$ and $\mathbf{E}$ are functions of the input sequence $u_{0:N-1}$ (since the linearization is performed locally with respect to an input trajectory) as well as the moments $\bar{z}_0$ and Σ_{z_0} of the initial augmented states.

Likewise, the measurement function (4.6b) can be linearized around $\hat{z}_k$ and $\bar{v} = \mathbb{E}[v_k]$

$$y_k \approx h_z(\hat{z}_k, \bar{v}) + \underbrace{\left.\frac{\partial h_z}{\partial z_k}\right|_{\hat{z}_k,\bar{v}}}_{C_k}(z_k - \hat{z}_k) + \underbrace{\left.\frac{\partial h_z}{\partial v_k}\right|_{\hat{z}_k,\bar{v}}}_{M_k}(v_k - \bar{v}), \tag{4.13}$$

which leads to the following approximations for the mean and covariance of the system outputs

$$\hat{y}_k = h_z(\hat{z}_k, \bar{v}) \tag{4.14a}$$

$$\hat{\Sigma}_{y_k} = C_k\hat{\Sigma}_{z_k}C_k^\top + M_k\Sigma_v M_k^\top. \tag{4.14b}$$

Stacking (4.13) over the planning horizon N leads to

$$\delta y_{0:N} = \mathbf{C}\delta z_{0:N} + \mathbf{M}\delta v_{0:N}, \tag{4.15}$$

with

$$\mathbf{C} = \mathrm{diag}(C_0, \ldots, C_N), \qquad \mathbf{M} = \mathrm{diag}(M_0, \ldots, M_N).$$

The predicted covariance of the outputs over the horizon N is then given by

$$\begin{aligned}\hat{\Sigma}_{y_{0:N}} &= \mathbf{C}\hat{\Sigma}_{z_{0:N}}\mathbf{C}^\top + \mathbf{M}\Sigma_{v_{0:N}}\mathbf{M}^\top \\ &= \mathbf{C}\mathbf{A}\Sigma_{z_0}\mathbf{A}^\top\mathbf{C}^\top + \mathbf{C}\mathbf{E}\Sigma_{w_{0:N-1}}\mathbf{E}^\top\mathbf{C}^\top + \mathbf{M}\Sigma_{v_{0:N}}\mathbf{M}^\top,\end{aligned} \tag{4.16}$$

where $\Sigma_{v_{0:N}} = \mathrm{diag}(\Sigma_v, \ldots, \Sigma_v)$. Note that $\hat{\Sigma}_{y_{0:N}}$ is a symmetric matrix that accounts for cross-correlation in the system outputs at different times, i.e.,

$$\hat{\Sigma}_{y_{0:N}} \approx \Sigma_{y_{0:N}} = \begin{bmatrix} \Sigma_{y_0} & \Sigma_{y_0y_1} & \cdots & \Sigma_{y_0y_N} \\ \Sigma_{y_1y_0} & \Sigma_{y_1} & \cdots & \Sigma_{y_1y_N} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{y_Ny_0} & \Sigma_{y_Ny_1} & \cdots & \Sigma_{y_N} \end{bmatrix}.$$

The accuracy of the above discussed uncertainty propagation method based on model linearization can be improved by retaining the higher-order terms in (4.7). In addition, the method can be straightforwardly extended for computing higher-order moments of the augmented states (such as skewness and kurtosis) through taking expectations of polynomial expressions of (4.11) and (4.15).

Unscented Transform

UT method. The general notion of the UT method is first described for propagation of an n-dimensional random vector $x \in \mathbb{R}^n$ through a nonlinear function $y = g(x)$. Let x have a known mean $\bar{x}$ and covariance matrix $\mathit{\Sigma}_x$. Define a set of $2n+1$ sigma points $\{\mathcal{X}_i\}_{i=0}^{2n}$ such that

$$\begin{aligned}
\mathcal{X}_0 &= \bar{x} \\
\mathcal{X}_i &= \bar{x} + \left(\sqrt{(n+\lambda)\mathit{\Sigma}_x}\right)_i && i = 1, \dots, n \\
\mathcal{X}_i &= \bar{x} - \left(\sqrt{(n+\lambda)\mathit{\Sigma}_x}\right)_i && i = n+1, \dots, 2n,
\end{aligned}$$

with the mean weights $\{W_i^{(m)}\}_{i=0}^{2n}$ and covariance weights $\{W_i^{(c)}\}_{i=0}^{2n}$ given by

$$\begin{aligned}
W_0^{(m)} &= \frac{\lambda}{(n+\lambda)} \\
W_0^{(c)} &= \frac{\lambda}{(n+\lambda)} + (1 - \alpha^2 + \beta) \\
W_i^{(m)} = W_i^{(c)} &= \frac{1}{2(n+\lambda)}, && i = 1, \dots, 2n.
\end{aligned}$$

In the above expressions, $(A)_i$ is the i^{th} column of the matrix A; $\lambda = \alpha^2(n+\kappa) - n$ is a scaling parameter; α specifies the spread of the sigma points around the mean; κ is another scaling parameter; and β is used to account for prior knowledge of the pdf of x. Some common values for these parameters are $\alpha = 1$, $\kappa = 2$, and $\beta = 0$ [44]. Note that $\beta = 2$ is the optimal choice for β when x has a Gaussian distribution. The UT method relies on propagating these sigma points through the nonlinear function $y = g(x)$, i.e.,

$$\mathcal{Y}_i = g(\mathcal{X}_i), \quad i = 0, \dots, 2n,$$

and approximating the mean and covariance of y based on weighted sample mean and covariance of $\mathcal{Y}_i$

$$\begin{aligned}
\bar{y} &\approx \sum_{i=0}^{2n} W_i^{(m)} \mathcal{Y}_i \\
\mathit{\Sigma}_y &\approx \sum_{i=0}^{2n} W_i^{(c)} (\mathcal{Y}_i - \bar{y})(\mathcal{Y}_i - \bar{y})^\top.
\end{aligned}$$

Note that the UT method requires $2n+1$ function evaluations, which scale linearly with respect to the dimension of the uncertainty. Therefore, this method can be substantially cheaper than Monte Carlo-based uncertainty propagation methods, which typically require several orders of magnitude more sample points for accurate computation of $\bar{y}$ and $\mathit{\Sigma}_y$ [44].

Prediction scheme using UT. The UT method can be adopted to approximate the mean and covariance of the predictions of (4.4) over the planning horizon N. Define the concatenated vector of augmented states, disturbances, and measurement noise as

$$x_k^a = \left[z_k^\top,\ w_k^\top,\ v_k^\top\right]^\top .$$

For $k = 0, 1, \ldots, N-1$, perform the following steps: (i) generate the sigma points $\mathcal{X}_k^a = [(\mathcal{X}_k^z)^\top,\ (\mathcal{X}_k^w)^\top,\ (\mathcal{X}_k^v)^\top]^\top$ using $\hat{x}_k^a = \left[\hat{z}_k^\top,\ \bar{w}^\top,\ \bar{v}^\top\right]^\top$ and $\hat{\Sigma}_{x_k^a} = \operatorname{diag}(\hat{\Sigma}_{z_k}, \Sigma_w, \Sigma_v)$ and compute the weights $W_i^{(m)}$, $W_i^{(c)}$ as described above, (ii) propagate the latter sigma points through the dynamics $\mathcal{X}_{k+1}^z = f_z(\mathcal{X}_k^z, \mathcal{X}_k^w)$, and (iii) compute the mean and covariance estimates of the augmented states from

$$\hat{z}_{k+1} = \sum_{i=0}^{2n} W_i^{(m)} \mathcal{X}_{i,k+1}^z \tag{4.17a}$$

$$\hat{\Sigma}_{z_{k+1}} = \sum_{i=0}^{2n} W_i^{(c)} (\mathcal{X}_{i,k+1}^z - \hat{z}_{k+1})(\mathcal{X}_{i,k+1}^z - \hat{z}_{k+1})^\top, \tag{4.17b}$$

where $n = n_z + n_w + n_v$. Similarly, the sigma points for the system outputs can be defined using the measurement function $\mathcal{Y}_k = h_z(\mathcal{X}_k^z, \mathcal{X}_k^v)$ at each $k = 0, \ldots, N$. The sigma points are used to approximate the mean and covariance of the outputs by

$$\hat{y}_k = \sum_{i=0}^{2n} W_i^{(m)} \mathcal{Y}_{i,k} \tag{4.18a}$$

$$\hat{\Sigma}_{y_k y_l} = \sum_{i=0}^{2n} W_i^{(c)} (\mathcal{Y}_{i,k} - \hat{y}_k)(\mathcal{Y}_{i,l} - \hat{y}_l)^\top, \quad \forall k, l = 0, 1, \ldots, N. \tag{4.18b}$$

The covariance matrix of the system outputs over the planning horizon, $\hat{\Sigma}_{y_{0:N}}$, can then be constructed from these blocks, of which only $N(N+1)/2$ elements should be calculated due to symmetry.

4.3.2 Tractable Criterion for Multiple Hypothesis Discrimination

Bound on the Bayes risk. There exists no closed-form expression for the Bayes risk due to the multivariate nature of the integral in (4.3). This poses a key challenge to direct application of the Bayes risk in Problem 1. In [7], a tractable upper bound on the Bayes risk, known as the *Bhattacharrya bound* [14], has been derived for the case of multiple hypotheses. For a general system observation $\mathbf{y}$ and manipulated input $\mathbf{u}$, the result is summarized as follows.

Theorem 1 *For hypothesis discrimination between a finite number of hypotheses, the Bayes risk is upper bounded by*

$$P(error) \leq \sum_{i} \sum_{j>i} P(H_i)^{\frac{1}{2}} P(H_j)^{\frac{1}{2}} \mathscr{B}_{ij}(\mathbf{u}), \tag{4.19}$$

where

$$\mathscr{B}_{ij}(\mathbf{u}) = \int \sqrt{p(\mathbf{y} \mid H_i, \mathbf{u}) p(\mathbf{y} \mid H_j, \mathbf{u})} d\mathbf{y} \tag{4.20}$$

is the Bhattacharrya coefficient between likelihoods of the i^{th} and j^{th} hypotheses.

Proof 1 *The key notion is to express* (4.3) *in terms of the Bayes risk for each of the pairs of hypotheses, and then bound the Bayes risk between each pair using the established Bhattacharrya bound. See [7] for details.*

The Bhattacharrya coefficient as optimization criterion. The bound (4.19) alleviates the need to compute the decision regions $\mathfrak{R}_i$ in (4.3). However, calculating the Bhattacharrya coefficient $\mathscr{B}_{ij}$ still hinges on (possibly expensive) high-dimensional numerical integration. In the case of Gaussian system observations, $\mathscr{B}_{ij}$ can be evaluated analytically, circumventing the numerical integration in (4.20). This motivates the following approximation

$$p(y_{0:N} \mid H_i, u_{0:N-1}) \approx \mathcal{N}(\bar{y}^{(i)}_{0:N}, \Sigma^{(i)}_{y_{0:N}}), \qquad \forall H_i \in \mathcal{H},$$

where the mean $\bar{y}^{(i)}_{0:N}$ and covariance $\Sigma^{(i)}_{y_{0:N}}$ can be obtained using the uncertainty propagation methods described in the previous section. Under this approximation, the Bhattacharrya coefficient (4.20) simplifies to

$$\mathscr{B}_{ij}(u_{0:N-1}) = e^{-d(i,j)}, \tag{4.21}$$

where

$$\begin{aligned} d(i,j) =& \frac{1}{4}(\bar{y}^{(i)}_{0:N} - \bar{y}^{(j)}_{0:N})^\top [\Sigma^{(i)}_{y_{0:N}} + \Sigma^{(j)}_{y_{0:N}}]^{-1} (\bar{y}^{(i)}_{0:N} - \bar{y}^{(j)}_{0:N}) \\ &+ \frac{1}{2}\ln \left(\frac{\det\left(\frac{\Sigma^{(i)}_{y_{0:N}} + \Sigma^{(j)}_{y_{0:N}}}{2}\right)}{\sqrt{\det\left(\Sigma^{(i)}_{y_{0:N}}\right) \det\left(\Sigma^{(j)}_{y_{0:N}}\right)}} \right). \end{aligned} \tag{4.22}$$

Using the error bound (4.19) on the Bayes risk $P(error)$, a tractable objective function for Problem 1 can now be defined as

$$J(u_{0:N-1}) = \sum_{i=0}^{n_f} \sum_{j=i+1}^{n_f} P(H_i)^{\frac{1}{2}} P(H_j)^{\frac{1}{2}} e^{-d(i,j)}. \tag{4.23}$$

4.3.3 Moment-based Chance Constraint Approximation

A key feature of Problem 1 is to enforce the system inputs and states to satisfy their constraints during the fault diagnosis experiment. Input constraints typically arise from practical considerations such as hard bounds on the manipulated variables (i.e., $u_{min} \leq u_k \leq u_{max}$) or the maximum allowable system perturbation for fault diagnosis, which is commonly defined in terms of the "energy" of the input sequence [18, 39]. On the other hand, state constraints are generally intended to ensure that the system remains in a "safe" operating region (or in a region for which the model hypotheses have been validated) as well as to ensure that certain system performance requirements are met during the fault diagnosis experiment. In the system (4.4), however, the state constraints can only be enforced in probability as the chance constraint (4.5) because of the stochastic system uncertainty.

Without loss of generality, assume that the feasible region $\mathcal{X} \triangleq \{x : A^\top x \leq b\}$ is a polytope defined by matrix $A \in \mathbb{R}^{n_x \times n_c}$ and vector $b \in \mathbb{R}^{n_c}$. Since evaluating the probability distribution $p(x_k \mid H, u_{0:k-1})$ can be computationally prohibitive, we look to derive a moment-based approximation for the chance constraint (4.5) that is not overly conservative. To this end, $p(x_k \mid H, u_{0:k-1})$ is approximated to have a Gaussian distribution, so that it can be represented in terms of its first two moments. The latter moments can be computed directly from $\bar{z}_{0:N}$ and $\Sigma_{z_{0:N}}$ by defining $x_k = T_k z_{0:N}$, where T_k is a matrix of all zeros except for I_{n_x} in the position of the k^{th} step states

$$\begin{aligned}\mathbb{E}[x_k \mid H] &= \bar{x}_k = T_k \bar{z}_{0:N} \\ \operatorname{cov}(x_k \mid H) &= \Sigma_{x_k} = T_k \Sigma_{z_{0:N}} T_k^\top .\end{aligned}$$

For any normally distributed random variable $x \sim \mathcal{N}(\bar{x}, \Sigma_x)$ of dimension n, enforcing $x \in \mathcal{X}$ with at least probability $1 - \epsilon$ can be expressed exactly in terms of the integral condition

$$P(x \in \mathcal{X}) = \frac{1}{\sqrt{(2\pi)^n \det(\Sigma_x)}} \int_{\mathcal{X}} e^{-\frac{1}{2}(\zeta-\bar{x})^\top \Sigma_x^{-1}(\zeta-\bar{x})} d\zeta \; \geq \; 1 - \epsilon. \tag{4.24}$$

Since it is challenging to enforce the above constraint directly, the approach taken in Refs. [36, 42] is adopted to derive a simple and efficient relaxation for (4.24). First, define the ellipsoid

$$\mathcal{E}_r \triangleq \left\{x : x^\top \Sigma_x^{-1} x \leq r^2\right\},$$

for any radius $r > 0$. If it can be ensured that $\bar{x} \oplus \mathcal{E}_r \subset \mathcal{X}$, then $P(x \in \mathcal{X}) > P(x \in \bar{x} \oplus \mathcal{E}_r)$ must hold by the axioms of probability. The expression $P(x \in \bar{x} \oplus \mathcal{E}_r)$ can be stated in terms of the chi-squared distribution

$$P(x \in \bar{x} \oplus \mathcal{E}_r) = P((x - \bar{x})^\top \Sigma_x^{-1}(x - \bar{x}) \leq r^2) = F_{\chi^2}(r^2;\ n),$$

where $F_{\chi^2}(\cdot;\ n)$ is the cdf of the chi-squared distribution with n degrees of freedom. Therefore, the radius r can be selected such that $F_{\chi^2}(r^2;\ n) = 1 - \epsilon$ in order to guarantee that the chance constraint holds. The constraint $P(x \in \mathcal{X})$ can now be replaced with the stronger constraint $\bar{x} \oplus \mathcal{E}_r \subset \mathcal{X}$ for the properly chosen radius. The constraint $\bar{x} \oplus \mathcal{E}_r \subset \mathcal{X}$ is in fact equivalent to requiring that the ellipsoid lies within a collection of half-spaces defined as $\mathcal{H}_j = \{x : a_j^\top x \leq b_j\}$, where a_j is the j^{th} column of A and b_j is the j^{th} element of b. The latter inclusion can be expressed in terms of simpler constraints using the result of the following lemma.

Lemma 1 *Let* $\mathcal{E} = \{y \in \mathbb{R}^n : y^\top Y^{-1} y \leq 1\}$, $b \in \mathbb{R}^n$, *and* $Y = Y^\top > 0$. *Consider the problem of maximizing a linear function over an ellipsoidal set*

$$\max_{y \in \mathcal{E}} \ b^\top y.$$

Then, the solution and optimal objective value to this problem are given by $y^\star = Yb/\sqrt{b^\top Y b}$ *and* $b^\top y^\star = \sqrt{b^\top Y b}$, *respectively.*

Proof 2 *The proof of this result follows from solving the KKT conditions derived using the method of Langrange multipliers. See [30] for a full solution.*

By defining $y = \delta x_j = x - \bar{x}$, $Y = r^2 \Sigma_x$, and $b = a_j$, the result of Lemma 1 yields the "worst-case" vector such that $\bar{x} + \delta x_j^\star$ of length $a_j^\top \delta x_j^\star = r(a_j^\top \Sigma_x a_j)^{\frac{1}{2}}$ is contained in the hyperplane $\mathcal{H}_j$. This directly leads to the assertion that $\bar{x} \oplus \mathcal{E}_r \subset \mathcal{X}$ can be equivalently stated as [42]

$$a_j^\top \bar{x} + r\sqrt{a_j^\top \Sigma_x a_j} \leq b_j, \qquad j = 1, \ldots, n_c. \tag{4.25}$$

This concept is visually illustrated in Figure 4.2. In this work, the expression (4.25) is readily used to obtain a moment-based surrogate for the chance constraint (4.5) given the approximation $p(x_k \mid H_i, u_{0:k-1}) \approx \mathcal{N}(\bar{x}_k^{(i)}, \Sigma_{x_k}^{(i)})$.

4.3.4 Tractable OED Problem

The above discussed methods are used to obtain a tractable surrogate for Problem 1. For any initial values of the augmented states, the set of feasible input sequences over the planning horizon N that ensure the fulfillment of the input constraints

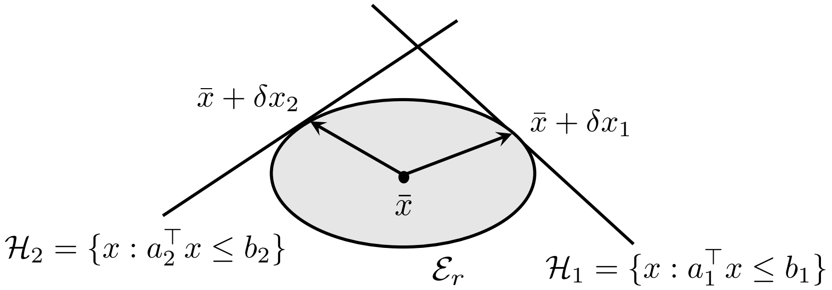

Figure 4.2: Illustration of method for guaranteeing ellipse lies in feasible polytopic region defined by multiple half-spaces. Adapted from [42].

and the moment-based surrogate for the chance constraint (4.5) is defined by

$$\Pi_{0:N}(z,Z) = \left\{ u_{0:N-1}\in\mathbb{R}^{Nn_u} \;\middle|\; \begin{array}{l} \bar{z}_0^{(i)} = z^{(i)},\ \Sigma_{z_0}^{(i)} = Z^{(i)}, \\ \hat{z}_{0:N}^{(i)} \text{ determined from (4.8) or (4.17a)},\\ \hat{\Sigma}_{z_{0:N}}^{(i)} \text{ determined from (4.10) or (4.17b)},\\ \hat{y}_{0:N}^{(i)} \text{ determined from (4.14a) or (4.18a)},\\ \hat{\Sigma}_{y_{0:N}}^{(i)} \text{ determined from (4.14b) or (4.18b)},\\ \hat{x}_{k+1}^{(i)} = T_{k+1}^{(i)}\hat{z}_{0:N}^{(i)},\ \hat{\Sigma}_{x_k}^{(i)} = T_{k+1}^{(i)}\hat{\Sigma}_{z_{0:N}}^{(i)}\left(T_{k+1}^{(i)}\right)^\top,\\ u_k\in\mathcal{U},\ a_j^\top \hat{x}_{k+1}^{(i)} + r\left(a_j^\top\hat{\Sigma}_{x_{k+1}}^{(i)}a_j\right)^{\frac{1}{2}}\le b_j,\\ \forall i=0,\dots,n_f,\\ \forall j=1,\dots,n_c,\\ \forall k=0,\dots,N-1, \end{array}\right\}. \tag{4.26}$$

The proposed tractable OED problem can now be stated as

$$\min_{u_{0:N-1}\in\Pi_{0:N}(\bar{z}_0,\Sigma_{z_0})} J(u_{0:N-1}) = \sum_{i=0}^{n_f}\sum_{j=i+1}^{n_f} P(H_i)^{\frac{1}{2}}P(H_j)^{\frac{1}{2}}e^{-d(i,j)}, \tag{4.27}$$

where $d(i,j)$ is defined in (4.22), which is a nonlinear function of $u_{0:N-1}$. Note that either the linearization method or the UT method can be used in (4.26) for approximating the mean and covariance of the augmented states. A key feature of the proposed formulation is its flexibility in terms of the choice of the uncertainty propagation method and the type of input and state constraints that can be incorporated in $\Pi_{0:N}$. Problem (4.27) is a nonlinear programming (NLP) problem that can be solved efficiently to local optimality using various nonlinear optimization methods such as interior point methods [27] and sequential quadratic programming [33]. The optimal input sequence, denoted by $u_{0:N-1}^{\star}$, will enable discrimination between the model hypotheses in a probabilistic sense since the system uncertainty is of stochastic nature.

In OED for offline model selection, once the measurements $y_0^\star, , \dots, y_N^\star$ that correspond to the system evolution under the optimal input sequence $u_{0:N-1}^\star$ have been collected, the probability of each model hypothesis can be evaluated using Bayes rule [7]

$$P(H_i \mid y_{0:N}^\star, u_{0:N-1}^\star) = \frac{p(y_{0:N}^\star \mid H_i, u_{0:N-1}^\star)P(H_i)}{\sum_{i=0}^{n_f} p(y_{0:N}^\star \mid H_i, u_{0:N-1}^\star)P(H_i)}. \tag{4.28}$$

The main challenge in (4.28) arises from computing the likelihood function $p(y_{0:N}^\star \mid H_i, u_{0:N-1}^\star)$. The likelihood function can be approximated using sample-based methods such as particle filtering [19]. Alternatively, approximating the likelihood function as a Gaussian distribution, which is also done for deriving a bound on the Bayes risk in (4.19), greatly simplifies the computation of (4.28). In the next section, an algorithm is presented for online OED, which relies on solving the problem (4.27) in tandem with Bayes rule (4.28) in real time as the fault diagnosis experiment is conducted.

4.3.5 OED for online model selection

The unmodeled system uncertainties and disturbances in the problem (4.27) can hamper effective fault diagnosis when the optimal input sequence $u_{0:N-1}^\star$ is designed completely offline. On the other hand, as the input sequence is applied to the true (but unknown) system, system measurements can be used to inform fault diagnosis in real time via updating the estimates of the augmented states and the probabilities of each model $H_i \in \mathcal{H}$. The measurement feedback to the problem (4.27) is likely to give some degree of robustness to unmodeled system uncertainties and disturbances when designing the input sequence. The receding-horizon implementation of the OED problem hinges on recursively updating the model probabilities online based on Bayes rule (4.28) and, subsequently, redesigning the input sequence in (4.27) using the most recent system information.

Recursive Calculation of Model Probabilities Online

Given system measurements $y_{0:k}$ and inputs $u_{0:k-1}$ up until sampling time $k \in \mathbb{N}_0$, stacked into a single *information vector* $I^k = \{y_{0:k}, u_{0:k-1}\}$, the probability of each model $H_i \in \mathcal{H}$ at every $k \in \mathbb{N}_0$ can be evaluated recursively using Bayes rule

$$P(H_i \mid I^k) = \frac{p(y_k \mid y_{0:k-1}, u_{0:k-1}, H_i)P(H_i \mid I^{k-1})}{\sum_{i=0}^{n_f} p(y_k \mid y_{0:k-1}, u_{0:k-1}, H_i)P(H_i \mid I^{k-1})}, \qquad P(H_i \mid I^{-1}) = P(H_i). \tag{4.29}$$

The likelihood function $p(y_k \mid y_{0:k-1}, u_{0:k-1}, H_i)$ in fact represents the so-called *evidence* (aka innovation), which can be obtained from state estimation for the given model H_i. As is well-known from the Bayesian estimation literature [12], no closed-form solution exists for state estimation of nonlinear systems. A variety of nonlinear Bayesian estimation techniques such as the extended Kalman filter (EKF) [24], the unscented Kalman filter (UKF) [44], the ensemble Kalman filter [21], and the particle filter [19] have been developed that alleviate the construction of the evidence distribution through recursively computing the mean and covariance (and possibly higher moments) of the posterior state distribution $p(x_k \mid y_{0:k}, u_{0:k-1} H)$ and the evidence $p(y_k \mid y_{0:k-1}, u_{0:k-1}, H)$. Note that the parameter uncertainty in (4.4) poses an additional challenge to evaluating the evidence in (4.29). This challenge can be addressed by converting the combined state and parameter estimation problem into a standard state estimation problem based on the augmented state system in (4.6).

In this work, the evidence is approximated as a multivariate Gaussian distribution $p(y_k \mid y_{0:k-1}, u_{0:k-1}, H) \approx \mathcal{N}(\bar{y}_{k|k-1}, \Sigma_{y_{k|k-1}})$, where $y_{k|k-1} = y_k \mid y_{0:k-1}, u_{0:k-1}, H$ is a conditional random variable representing the evidence with mean $\bar{y}_{k|k-1}$ and covariance $\Sigma_{y_{k|k-1}}$. In principle, any of the aforementioned nonlinear estimation methods can be used to determine the mean and covariance of the evidence. The EKF and UKF methods are straightforward extensions of the linearization and UT uncertainty propagation methods described previously. These filters both rely on a model prediction and a measurement update step. The prediction step calculates $\hat{z}_{k|k-1}$, $\hat{\Sigma}_{z_{k|k-1}}$, $\hat{y}_{k|k-1}$, and $\hat{\Sigma}_{y_{k|k-1}}$ using either (4.8), (4.10), (4.14a), and (4.14b) in the EKF or (4.17a), (4.17b), (4.18a), and (4.18b) in the UKF. In the update step, the latter model predictions are updated using the system observations. To this end, the filter gain is defined as

$$K_k = \hat{\Sigma}_{z_{k|k-1} y_{k|k-1}} \hat{\Sigma}^{-1}_{y_{k|k-1}}, \tag{4.30}$$

where in the EKF the covariance matrices are given by

$$\begin{aligned} \hat{\Sigma}_{z_{k|k-1} y_{k|k-1}} &= \hat{\Sigma}_{z_{k|k-1}} C^\top_{k|k-1} \\ \hat{\Sigma}_{y_{k|k-1}} &= C_{k|k-1} \hat{\Sigma}_{z_{k|k-1}} C^\top_{k|k-1} + M_{k|k-1} \Sigma_v M_{k|k-1}, \end{aligned}$$

and in the UKF the covariance matrices are given by

$$\begin{aligned} \hat{\Sigma}_{z_{k|k-1} y_{k|k-1}} &= \sum_{i=0}^{2n} W_i^{(c)} (\mathcal{X}^z_{i,k|k-1} - \hat{z}_{k|k-1})(\mathcal{Y}_{i,k|k-1} - \hat{y}_{k|k-1})^\top \\ \hat{\Sigma}_{y_{k|k-1}} &= \sum_{i=0}^{2n} W_i^{(c)} (\mathcal{Y}_{i,k|k-1} - \hat{y}_{k|k-1})(\mathcal{Y}_{i,k|k-1} - \hat{y}_{k|k-1})^\top. \end{aligned}$$

The estimates of the mean and covariance of the posterior distribution of the augmented states can now be updated as

$$\hat{z}_{k|k} = \hat{z}_{k|k-1} + K_k (y_k - \hat{y}_{k|k-1}) \tag{4.31a}$$

$$\hat{\Sigma}_{z_{k|k}} = \hat{\Sigma}_{z_{k|k-1}} - K_k \hat{\Sigma}_{y_{k|k-1}} K_k^\top. \tag{4.31b}$$

Algorithm 1 Recursive estimation of states, parameters, and model probabilities.

```
Require: Specify model $H$ and priors $P(H)$ for all $H \in \mathcal{H}$.
 1: Offline:
 2: For all $H \in \mathcal{H}$, calculate the initial estimates $\hat{z}_{0|-1}$ and $\hat{\Sigma}_{z_{0|-1}}$ from $p(x_0, \theta)$.
 3: Online:
 4: for $k = 0, 1, 2, \ldots$ do
 5:     Obtain measurements $y_k$ from the true system.
 6:     For each $H \in \mathcal{H}$, calculate evidence $\hat{y}_{k|k-1}$, $\hat{\Sigma}_{y_{k|k-1}}$ from the state estimates $\hat{z}_{k|k-1}$, $\hat{\Sigma}_{z_{k|k-1}}$ using either the linearization or UT method for uncertainty propagation.
 7:     For each $H \in \mathcal{H}$, update the model probabilities $P(H \mid I^k)$ using the Bayes rule (4.29) and measurements $y_k$ under the approximation $p(y_k \mid y_{0:k-1}, u_{0:k-1}, H) \approx \mathcal{N}(\hat{y}_{k|k-1}, \hat{\Sigma}_{y_{k|k-1}})$.
 8:     For each $H \in \mathcal{H}$, calculate posterior state estimates $\hat{z}_{k|k}$, $\hat{\Sigma}_{z_{k|k}}$ using (4.31).
 9:     Apply input $u_k$ to the true system.
10:     For each $H \in \mathcal{H}$, calculate predictions $\hat{z}_{k+1|k}$, $\hat{\Sigma}_{z_{k+1|k}}$ from $\hat{z}_{k|k}$, $\hat{\Sigma}_{z_{k|k}}$ and $u_k$ using either the linearization or UT method.
11: end for
```

Note that the prediction and update steps in the estimator must be repeated at every time instant k. This procedure yields the estimates of $\bar{y}_{k|k-1}$ and $\Sigma_{y_{k|k-1}}$, which are required for approximating the evidence in (4.29). The recursive estimation procedure is summarized in Algorithm 1.

Online Design of Input Sequence

In online OED, the surrogate problem (4.27) should be solved recursively based on available system information I^k at every measurement sampling time $k = 0, \ldots, N-1$. This amounts to minimizing $P(error \mid I^k)$ subject to system constraints. Hence, for online OED, the priors $P(H_i)$ and the initial conditions $\bar{z}_0$, Σ_{z_0} in (4.27) must be replaced with $P(H_i \mid I^k)$ and $\hat{z}_{k|k}$, $\hat{\Sigma}_{z_{k|k}}$, respectively, which are readily available from the recursive evaluation of Bayes rule (4.29). At each sampling time during the planning horizon $k = 0, \ldots, N-1$, online OED involves solving

$$u^{\star}_{k:N-1}(\hat{z}_{k|k}, \hat{\Sigma}_{z_{k|k}}) \triangleq \underset{u_{k:N-1} \in \Pi_{k:N}(\hat{z}_{k|k}, \hat{\Sigma}_{z_{k|k}})}{\operatorname{argmin}} J(u_{k:N-1} \mid I^k), \tag{4.32}$$

where $J(u_{k:N-1} \mid I^k)$ is defined as in (4.23) except that all values are conditioned on the current system information I^k. The receding-horizon implementation of the online OED problem (4.32) then follows the same recursive estimation procedure detailed in Algorithm 1 with the input u_k in Step 9 being the first element of the optimal input sequence $u^{\star}_{k:N-1}(\hat{z}_{k|k}, \hat{\Sigma}_{z_{k|k}})$. Note that the input at each sampling time is a function of the mean and covariance of the posterior distribution of the augmented states for each model.

4.3.6 Simulation Case Study

System Description

The performance of the proposed OED framework for fault diagnosis is demonstrated on a biochemical reaction occurring in a continuous well-stirred bioreactor. Assuming constant volume, the continuous-time process dynamics are described by [1, 20]

$$dX = (-DX + \mu X)\, dt + \sigma_X dw_X(t) \tag{4.33a}$$

$$dS = \left(D(S_f - S) - \frac{1}{Y_{X/S}} \mu X \right) dt + \sigma_S dw_S(t) \tag{4.33b}$$

$$dP = (-DP + (\alpha\mu + \beta)X)\, dt + \sigma_P dw_P(t), \tag{4.33c}$$

where X, S, and P are, respectively, the concentration of biomass, substrate, and product in the bioreactor; D is the dilution rate (equal to the inlet feed divided by the reactor volume); S_f is the concentration of substrate in the inlet feed; $Y_{X/S}$ is the yield of biomass per substrate consumed; α and β are the yield parameters for the production of P; and $w_X(t)$, $w_S(t)$, and $w_P(t)$ are independent, zero-mean unit-variance Weiner processes scaled by standard deviations σ_X, σ_S, and σ_P, respectively. The quantity μ represents the (substrate-dependent) growth rate of biomass, given by

$$\mu = \frac{\mu_{\max} S}{K_M + S}, \tag{4.34}$$

where $\mu_{\max}$ denotes the maximum growth rate; and K_M is an affinity constant. The dilution rate D is the only manipulated input of the process. The nominal parameter values and operating conditions are listed in Table 4.1.

Normal Operation and Fault Scenarios

Normal operation is described by (4.33) with parameters and operating conditions (aka initial conditions) listed in Table 4.1. The nominal model is the same as that used for normal operation, however, the parameter $\mu_{\max}$ is treated as an uncertainty since it cannot be perfectly estimated from data. The prior distribution $\mu_{\max} \sim \mathcal{N}(\bar{\mu}_{\max}, \Sigma_{\mu_{\max}})$ is assumed to be normal with mean $\bar{\mu}_{\max} = 0.6$ hr^{-1} and variance $\Sigma_{\mu_{\max}} = 0.05$ hr^{-1}. Note that the mean of $\mu_{\max}$ is different from that of its true value.

Under continuous operation of the bioreactor, faults must be swiftly identified in order to prevent product loss, which can lead to negative economic consequences.

Table 4.1: Nominal parameters and operating conditions of the continuous bioreactor [1, 20].

Variable	Nominal Value	Units
$Y_{X/S}$	0.4	g/g
α	2.2	g/g
β	0.2	hr^{-1}
$\mu_{\max}$	0.48	hr^{-1}
K_M	1.2	g/L
S_f	20	g/L
σ_X	0.25	g/L
σ_S	0.25	g/L
σ_P	0.25	g/L
X	7.038	g/L
S	2.404	g/L
P	24.87	g/L

To demonstrate the efficacy of the proposed OED framework, two different types of fault scenarios relative to normal operation are considered. First, a change in the model structure is considered in which inhibition of the biochemical reaction takes place due to the presence of excess substrate (commonly referred to as substrate inhibition). In this fault scenario, the specific growth rate μ is described by

$$\mu = \frac{\mu_{\max} S}{K_M + S\left(1 + \frac{S}{K_I}\right)}, \tag{4.35}$$

instead of (4.34), where K_I is the rate constant that specifies the strength of inhibition. The parameters for the substrate inhibition fault model are the same as those used for the nominal model with $K_I = 1$ g/L. The second fault scenario is an operation fault in which the substrate concentration in the inlet feed drops by 20%, i.e., $S_f \leftarrow 0.8 S_f$, resulting in $S_f = 16$ g/L.

Three model hypotheses based on the model structure (4.33) are used to represent normal operation and the two fault scenarios. The continuous-time models are discretized to obtain discrete-time models of the form (4.4) with states $x = [X,\ S,\ P]^\top$ and input $u = D$. Measurements of S and P are taken every 0.2 hr. Both measurements are corrupted with zero-mean Gaussian white noise with standard deviations of 0.03 g/L and 0.2 g/L, respectively. The initial states are also not known exactly and represented by the prior $x_0 \sim \mathcal{N}(\bar{x}_0, \Sigma_{x_0})$ with mean $\bar{x}_0 = [6.5,\ 2.75,\ 26]^\top$ g/L and covariance $\Sigma_{x_0} = \mathrm{diag}(0.5, 0.5, 0.5)$ g^2/L^2.

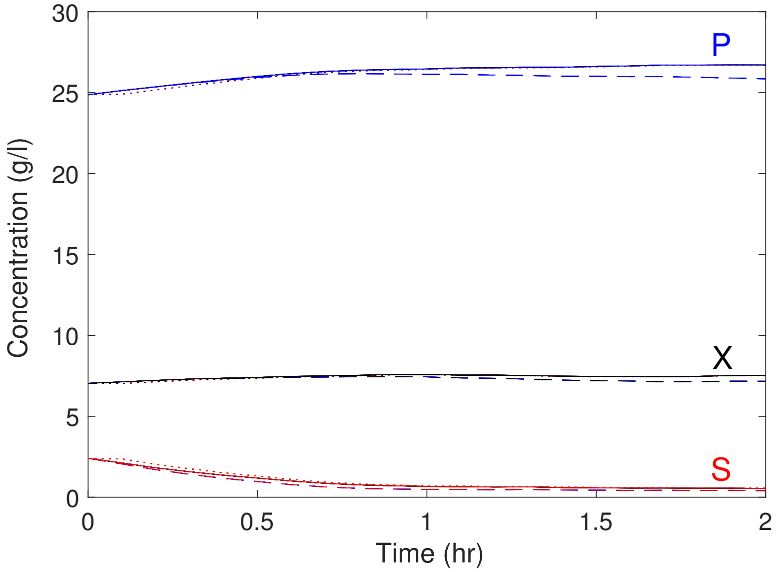

Figure 4.3: Comparison of the expected value of the process dynamics for the three models. The solid, dotted, and dashed lines correspond to the models of normal operation, substrate inhibition, and substrate offset in the feed, respectively.

Performance Evaluation of Offline OED

A comparison of the expected value of the states for normal operation and fault scenarios is shown in Figure 4.3 under a nominal input sequence (adapted from [20]). It is evident that the dynamics of the three operation scenarios are very similar. As such, fault diagnosis is expected to be more challenging due to the fact that there is uncertainty in the maximum growth rate as well as process and measurement noise.

Figure 4.4 shows ten Monte Carlo runs of the evolution of the probability of each model $P(H_i \mid I^k)$ under the nominal input sequence when normal operation model, substrate inhibition model, and substrate offset model are active. The active model is misidentified in all three cases shown in Figure 4.4a-c since the substrate inhibition model appears to be consistently active in the first hour of process operation. In fact, the nominal model is not properly identified throughout the entire process run in Figure 4.4a.

To better discriminate between the three model hypotheses that correspond to the different operation scenarios, we look to determine the optimal dilution rate trajectory, denoted by $D^\star$, that minimizes a bound on the Bayes risk $P(error)$ by solving the OED problem (4.27). The optimization problem is solved using the Matlab solver `fmincon`. Figure 4.5 depicts ten Monte Carlo runs of the evolution of the model probabilities $P(H_i \mid I^k)$, with $D^\star$ applied to the process over a span of two hours, when the normal operation model, substrate inhibition model, and

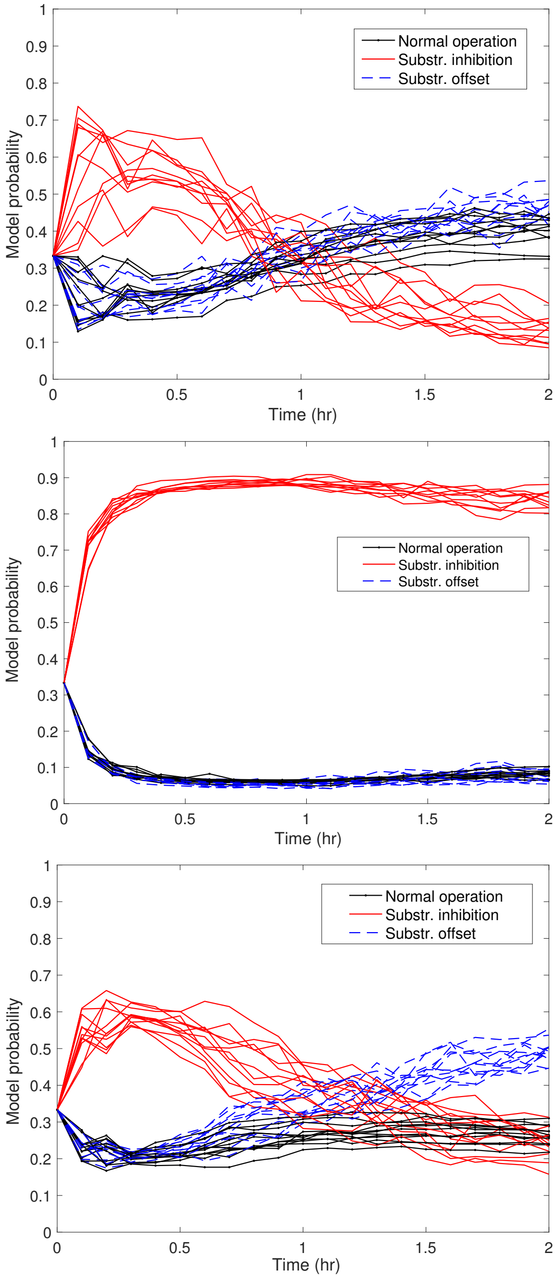

Figure 4.4: Ten Monte Carlo runs of the evolution of probability of each model being active using nominal input sequence under (top) normal operation, (middle) substrate inhibition, and (bottom) substrate concentration offset in the feed.

substrate offset model are active. In all three cases, the active model is correctly identified reasonably fast. Note that subsequent measurements further improve upon the confidence in the true model since the probability of the active model gradually increases. This directly contrasts the poor discriminating power of the nominal input sequence, which illustrates the usefulness of input design for fault diagnosis.

Comparison between Online OED and Offline OED

The above results revealed the efficacy of offline OED, where the optimal input sequence was designed only once, offline. As such, when a deviation from normal operation is detected, for example using process monitoring methods [25, 46], the optimal dilution rate $D^\star$ sequence should be designed and subsequently applied to the process in order to diagnose the current mode of operation using the process observations collected during the fault diagnosis experiment. When the tractable OED problem (4.27) can be solved within the measurement sampling intervals, it will often be advantageous to redesign the input sequence based on the most recent system observations. Due to regular adaptation of the model probabilities and state/parameter estimates (see (4.32)), online fault diagnosis partly accounts for unmodeled system uncertainties and disturbances, which can be detrimental to the efficacy of input sequences designed offline.

To demonstrate the effectiveness of the proposed online OED method (Algorithm 1 combined with the optimization (4.32)), its performance is compared to that of the offline OED problem (4.27). A shrinking-horizon implementation is utilized for online OED in which the final time of the planning horizon is fixed at all iterations of the Algorithm 1. Figure 4.6 compares the optimal dilution rate sequence computed online through repeated solution of (4.32) to that obtained from solving the offline OED problem under the same uncertainty realizations when the second fault scenario is active. The corresponding model probability evolutions are shown in Figure 4.7. Figure 4.6 suggests that the online optimized dilution rate sequence exhibits significant fluctuations between the upper and lower bounds of the manipulated input, whereas the offline computed input sequence exhibits a smoother behavior. The online and offline input sequences are significantly different due to the fact that the model probabilities are changing during the fault diagnosis experiment, which are accounted for when the input sequence is redesigned at each sampling time in online OED. For example, the substrate inhibition model can be ruled out after the first few sampling times as its probability goes to zero (see Figure 4.7). This allows the online OED method to direct its efforts at discriminating between the normal operation and substrate feed offset models, in contrast to considering all three models as in offline OED. Figure 4.7 suggests that the correct model is identified faster in online OED, while leading to higher probability of the true model than the case of offline OED. In

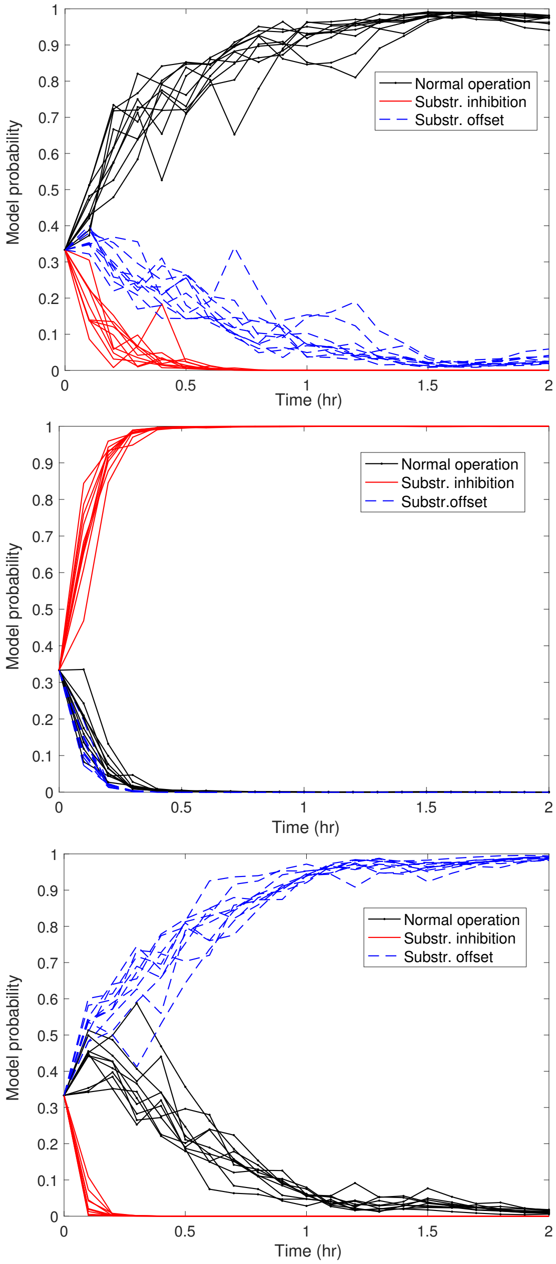

Figure 4.5: Ten Monte Carlo runs of the evolution of probability of each model being active using optimal input sequence $D^{\star}$ under (top) normal operation, (middle) substrate inhibition, and (bottom) substrate concentration offset in the feed.

addition, the pulse-like dilution rate sequence computed online is likely to be less intrusive to the process since less substrate is fed to the bioreactor (see Figure 4.6).

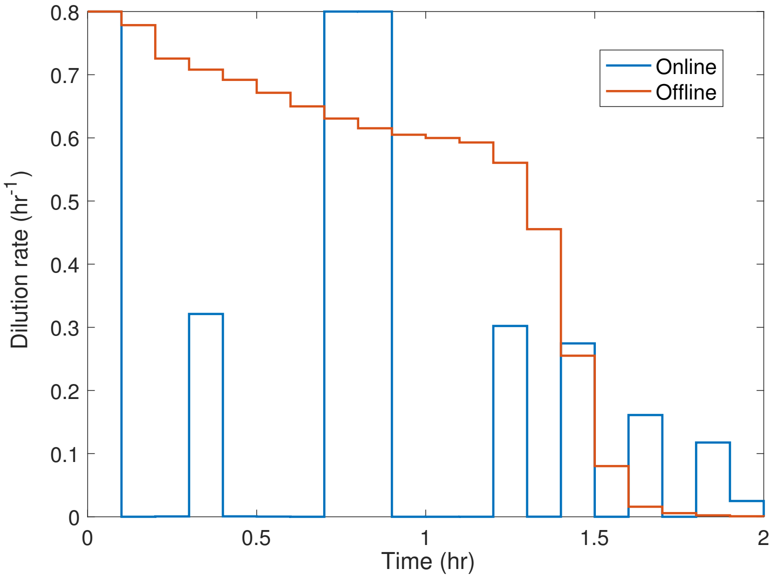

Figure 4.6: Comparison of the optimal input sequences for online OED (blue) and offline OED (orange) computed under the same uncertainty realizations when the second fault scenario is active.

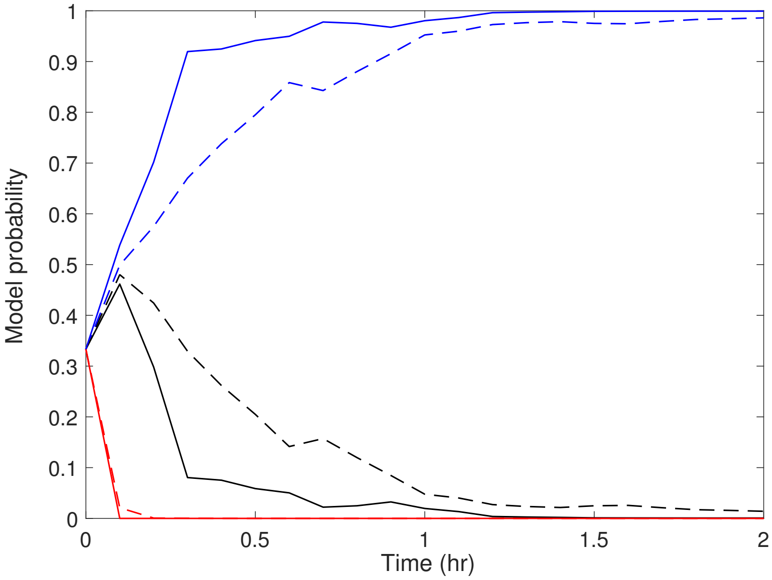

Figure 4.7: Evolution of the probabilities for the normal operation (black), substrate inhibition (red), and substrate concentration offset in the feed (blue) models under the optimal input sequences of online OED (solid line) and offline OED (dashed line) shown in Figure 4.6.

Online OED in Practice

We now demonstrate how the proposed online OED method can be applied in practice. Consider three distinct operational stages of a system: (i) normal system operation during which a performance loss can be detected using standard performance monitoring and fault detection methods, (ii) online OED to diagnose the current system operation mode, and (iii) performance recovery of system to desired conditions. For the bioreactor case study at hand, the there operational stages are illustrated in Figure 4.8, which shows the product concentration, normal and fault model probabilities evaluated online, and the dilution rate sequence (i.e., manipulated input). The first three hours of operation correspond to normal operation wherein the dilution rate is fixed and the product concentration remains relatively constant. Towards the end of the third hour of operation, the product concentration starts to decrease (Figure 4.8a), which indicates that a process change/fault has occurred. Note that the probability of the normal operation model decreases during this period, but changes relatively slowly (Figure 4.8c). The proposed online OED method is then executed during the next two hours of operation (from 3-5 hours). During this period, the dilution rate is manipulated by OED to generate informative measurements that can aid identifying the model that corresponds to the correct mode of operation with high confidence (Figure 4.8d). The latter model can then be used to recover the operation. In this case study, the dilution rate is set to a value that is known to yield a steady-state product concentration at the setpoint of 27 g/L. In general, advanced control strategies, such as model predictive control, can be used in this stage to enable high-performance system operation.

Twenty Monte Carlo runs were performed to evaluate the performance of online OED under various uncertainty realizations (not shown here). In all of the Monte Carlo runs, the correct model was identified in only few sampling times, suggesting the robustness of the online OED method to system uncertainties. In addition, the product concentration eventually returned to the desired setpoint in all runs. To evaluate the effectiveness of the proposed moment-based surrogate for chance constraint approximation, a chance constraint on the product concentration was included in the OED problem. Even though up to 15% violation of the state constraint $P > 23$ g/L was allowed in the chance constraint, a maximum of 10% constraint violation was observed in the Monte Carlo runs, suggesting a non-conservative chance constraint approximation. It should be noted that in this work the online OED problem was solved in Matlab using `fmincon`. The CPU time for solving the optimization problem was approximately 5 min on a machine with 8 GB of RAM (the measurement sampling time in this case study was 12 min). For systems with faster sampling times, the optimization in the OED problem could be substantially sped up by using state-of-the-art solvers such as IPOPT [43] combined with automatic differentiation tools such as CasADi [2] that provide exact gradients to the optimizer.

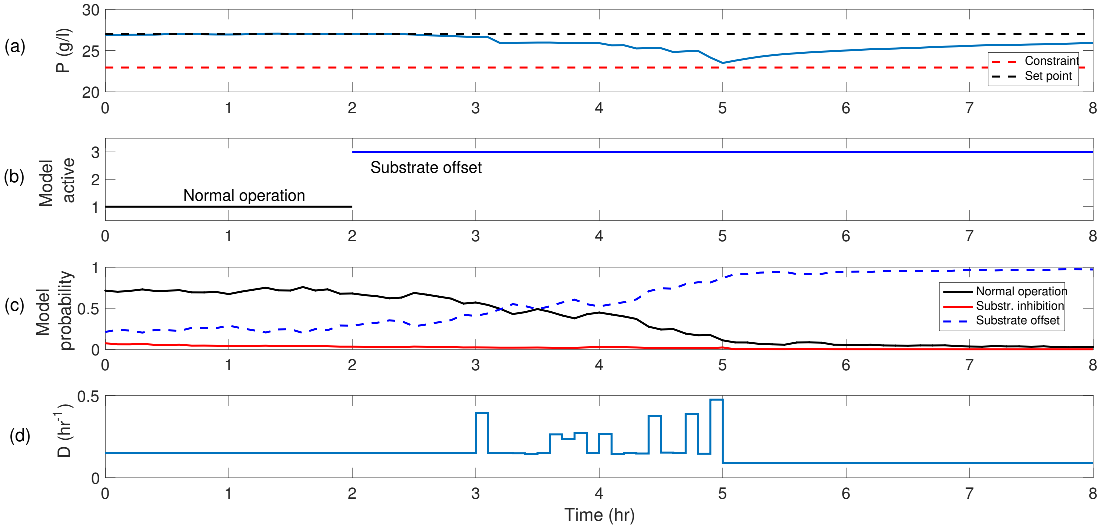

Figure 4.8: One run of online OED during system operation. Different segments of time correspond to different operational stages: 0-3 hours is normal operation (fault occurs at start of second hour), 3-5 hours is the fault diagnosis period during which online OED is performed, and 5-8 hours is the performance recovery stage. The plots show (a) the product concentration (dotted black line is the setpoint and solid red line is the constraint), (b) the active model in different stages of operation, (c) the online model probabilities (black line is nominal operation, red line is substrate inhibition fault, and blue line is substrate feed offset fault), and (d) the dilution rate sequence.

4.4 Method 2: OED for online model selection using the k-nearest neighbors distance

This section presents a computationally tractable, sample-based OED formulation, based on non-intrusive generalized polynomial chaos (gPC) and UT. The proposed sample-based uncertainty propagation method, which draws from the benefits of gPC and UT for handling model uncertainty and disturbances, respectively, is particularly suitable for optimization, since it can approximate the statistics of system variables with a fairly small number of samples that are chosen systematically. The optimization criterion is defined in terms of a sample-based measure of k-nearest neighbors distance between probabilistic predictions of model hypotheses for the normal and faulty system operation. Lastly, the OED method is demonstrated on a bioreactor case study.

4.4.1 Uncertainty Propagation

We present an efficient uncertainty propagation method that considers both model uncertainty and stochastic disturbances in (4.4). For ease of notation, we will rename the (possibly) joint probability distribution of x_0 and θ, $p(x_0, \theta)$, as ξ.The key notion of the proposed method is to decouple the propagation of probabilistic uncertainty in the model parameters and initial conditions from the propagation of stochastic disturbances using conditional probability rules, so that different methods that are best-suited for handling each uncertainty source can be adopted. To this end, UT [44] is used for propagation of stochastic disturbances conditioned on a realization of model uncertainty ξ, which is then integrated over different realizations of model uncertainty using non-intrusive gPC [45]. This results in a sample-based method for joint propagation of stochastic disturbances and model uncertainty, as described below.

Unscented transform for propagation of stochastic disturbances conditioned on model uncertainty

The UT method can effectively deal with nonlinearities, as it does not require linearization of the nonlinear system equations and thus computation of Jacobian. As stated for in Section 4.3, the UT method consists in the propagation of a relatively small number of deterministically-chosen samples (known as *sigma points*), centered around the mean of states, through the system dynamics. In this method, the UT method is used for the propagation of stochastic disturbances $w(t)$ in (4.4) given a realization of the model uncertainty ξ.

Denote the sigma points conditioned on a given realization of model uncertainty by $\mathcal{S}(t;\xi)$. Let $m(t;\xi)$ and $V(t;\xi)$ be the mean and covariance of the states conditioned on a realization of the standard random variables ξ. A set of $2n+1$ sigma points $\{\mathcal{S}_i(t;\xi)\}_{i=0}^{2n}$ is defined as

$$\begin{aligned}
&\mathcal{S}_0(t;\xi) = m(t;\xi), \\
&\mathcal{S}_i(t;\xi) = m(t;\xi) + \left(\sqrt{(n+\lambda)V(t;\xi)}\right)_i, \quad i = 1,\dots,n, \\
&\mathcal{S}_i(t;\xi) = m(t;\xi) - \left(\sqrt{(n+\lambda)V(t;\xi)}\right)_i, \quad i = n+1,\dots,2n,
\end{aligned}$$

where $n = n_x + n_w$; $(V)_i$ is the ith column of matrix V; and λ is a scaling parameter. The sigma points $\{\mathcal{S}_i(t;\xi)\}_{i=0}^{2n}$ are propagated through the nonlinear system dynamics (4.4) to obtain the propagated points $\{\mathcal{P}_i(t;\xi)\}_{i=0}^{2n}$ that characterize the distribution of the states, conditioned on ξ, in terms of the mean and covariance

$$\begin{aligned}
m_{\mathcal{P}}(t;\xi) &\approx \sum_{i=0}^{2n} W_i^m \mathcal{P}_i(t;\xi), \\
V_{\mathcal{P}}(t;\xi) &\approx \sum_{i=0}^{2n} W_i^c \big(\mathcal{P}_i(t;\xi) - m_{\mathcal{P}}(t;\xi)\big)\big(\mathcal{P}_i(t;\xi) - m_{\mathcal{P}}(t;\xi)\big).
\end{aligned}$$

The mean weights $\{W_i^m\}_{i=0}^{2n}$ and covariance weights $\{W_i^c\}_{i=0}^{2n}$ are defined as

$$\begin{aligned} W_0^m &= \frac{\lambda}{n+\lambda}, \\ W_0^c &= \frac{\lambda}{n+\lambda} + (1-\alpha^2+\beta), \\ W_i^m &= W_i^c = \frac{1}{2(n+\lambda)}, \quad i = 1, \ldots, 2n, \end{aligned}$$

where α determines the spread of sigma points around the mean and β accounts for prior knowledge of the state distribution [44].

The computational complexity of the UT method scales linearly with the dimension of the stochastic disturbances n_w, as UT relies on $2n+1$ function evaluations; that is, one function evaluation per sigma point. This can result in considerable computational speed-up in comparison with Monte Carlo-based sampling methods [9].

Non-intrusive polynomial chaos for propagation of model uncertainty of sigma points

Disturbance propagation based on UT yields the (propagated) sigma points $\{\mathcal{P}_i(t;\xi)\}_{i=0}^{2n}$ that are conditioned on model uncertainty. We now employ gPC to integrate the sigma points over the probabilistic model uncertainty ξ. gPC consists in approximating a stochastic variable, the sigma points $\mathcal{P}_i$, with the truncated expansion

$$\mathcal{P}_i(t;\xi) \approx \sum_{k=0}^{L} \tilde{\mathcal{P}}_{i,k}(t)\Phi_k(\xi), \tag{4.36}$$

where $\tilde{\mathcal{P}}_{i,k}$ denotes the expansion coefficients of sigma point $\mathcal{P}_i$, which evolve as a function of system dynamics; and Φ_k denotes multivariate polynomial basis functions constructed from the univariate polynomial basis functions of the individual random variables ξ_j

$$\Phi_k(\xi) = \prod_{j=1}^{m} \phi_{\alpha_j^{(i)}}^{(j)}(\xi_j), \qquad \alpha_j^{(i)} \in \{0,1,\ldots\}, \quad \forall j = \{1,\ldots,m\},$$

with $\alpha_j^{(i)}$ being the jth element of a multi-index whose value corresponds to the order of the basis of the jth random variable in the ith multivariate polynomial basis. The univariate polynomials ϕ belong to the Askey-Wiener scheme of polynomials, so that each univariate polynomial has an optimal convergence rate with respect to ξ_j [45]. The truncation order in the gPC expansion (4.36) is defined by $L+1 = \frac{(n_\xi+m)}{m!n_\xi!}$, where m denotes the prespecified degree of the multivariate polynomial basis and n_ξ denotes the dimension of ξ. An important property of the polynomial basis functions in the Askey-Wiener scheme, key to the computational efficiency of gPC, is their orthogonality with respect to the multivariate distribution of ξ. Note that the polynomials $\Phi_k(\xi)$ are constructed only once, merely based on the known distribution of ξ.

Evaluation of the gPC expansions for each of the $2n+1$ sigma points requires computation of the coefficients $\tilde{\mathcal{P}}_{i,k}$ in (4.36), which can be done using intrusive or non-intrusive methods [23]. In this work, we adopt the non-intrusive method, which hinges on evaluating the system model at given samples of model uncertainty $\{\xi^{(i)}\}_{j=1}^{n_s}$. This enables estimating the gPC coefficients as a weighted sum of the samples $\{\mathcal{P}_i^{(j)}(t,\xi^{(j)})\}_{j=1}^{n_s}$

$$\begin{bmatrix} \tilde{\mathcal{P}}_{i,0}(t) \\ \vdots \\ \tilde{\mathcal{P}}_{i,L}(t) \end{bmatrix} = \begin{bmatrix} \Pi_{0,1},\ldots,\Pi_{0,n_s} \\ \vdots \quad \ddots \quad \vdots \\ \Pi_{L,1},\ldots,\Pi_{L,n_s} \end{bmatrix} \begin{bmatrix} \mathcal{P}_i^{(1)}(t;\xi^{(1)}) \\ \vdots \\ \mathcal{P}_i^{(n_s)}(t;\xi^{(n_s)}) \end{bmatrix}, \tag{4.37}$$

where the weights $\Pi_{k,j}$ describe the effect of the sigma point samples $\{\mathcal{P}_i^{(j)}\}_{j=1}^{n_s}$ on the expansion coefficients $\tilde{\mathcal{P}}_{i,k}$ for $k = 0,...,L$. Note that the sigma point samples are simply determined by evaluating the sigma points $\{\mathcal{P}_i(t;\xi)\}_{i=0}^{2n}$ for the uncertainty realizations $\{\xi^{(j)}\}_{j=1}^{n_s}$. The non-intrusive method can be used irrespective of the form and complexity of the nonlinear model equations.

In this method, the weight matrix in (4.37) is defined in terms of least-squares estimation of the expansion coefficients

$$\mathbf{\Pi} = (\Lambda^\top \Lambda)^{-1}\Lambda^\top,$$

with Λ given by (see [37])

$$\Lambda = \begin{bmatrix} \Phi_0(\xi^{(1)}) & \Phi_1(\xi^{(1)}) & \cdots & \Phi_L(\xi^{(1)}) \\ \Phi_0(\xi^{(2)}) & \Phi_1(\xi^{(2)}) & \cdots & \Phi_L(\xi^{(2)}) \\ \vdots & \vdots & \cdots & \vdots \\ \Phi_0(\xi^{(n_s)}) & \Phi_1(\xi^{(n_s)}) & \cdots & \Phi_L(\xi^{(n_s)}) \end{bmatrix}.$$

The model uncertainty sample set $\{\xi^{(j)}\}_{j=1}^{n_s}$ can be selected as the roots of the polynomial basis of one degree higher $(m+1)$, so that the number of uncertainty samples is $n_s = \frac{(n_\xi+m+1)!}{n_\xi!(m+1)!}$. Due to the optimal choice of the polynomials in (4.36) with respect to the multivariate distribution of ξ, non-intrusive gPC requires a relatively small number of samples to propagate the model uncertainty ξ.

Joint propagation of model uncertainty and disturbances

Joint propagation of the probabilistic model uncertainty and stochastic disturbances entails expressing each one of the $2n+1$ sigma points $\{\mathcal{P}_i(t;\xi)\}_{i=0}^{2n}$ as the gPC expansion (4.36). Evaluating the coefficients $\tilde{\mathcal{P}}_{i,k}(t)$ using the non-intrusive method in (4.37) requires computing the conditional sigma points $\{\mathcal{P}_i(t;\xi)\}_{i=0}^{2n}$ at n_s samples of the model uncertainty, i.e., $\{\xi^{(j)}\}_{j=1}^{n_s}$. Thus, the proposed sample-based uncertainty propagation method uses a total of $n_{\text{tot}} = (2n+1) \times n_s$ samples for the joint propagation of probabilistic model uncertainty and stochastic disturbances, yielding the *total sigma points* $\{\mathsf{P}_s(t)\}_{s=1}^{n_{\text{tot}}} = \{\mathcal{P}_i(t;\{\xi^{(j)}\}_{j=1}^{n_s})\}_{i=0}^{2n}$.

4.4.2 Tractable OED Problem

We now present a tractable criterion for OED that readily uses the total sigma points $\{\mathsf{P}_s(t)\}_{s=1}^{n_{\text{tot}}}$, alleviating the need to either construct the distribution of the model outputs, or evaluate their statistical moments. To enable probabilistic model discrimination, the optimization criterion is defined as a measure of the distance between the sigma points belonging to the competing model hypotheses in the model hypothesis set $\mathcal{H}$. The distance measure used here is inspired by the k-nearest neighbors (KNN) algorithm [14], which computes the distance between the sigma points of a model H_i to the k closest sigma points of other models in $\mathcal{H}$. Maximization of the kNN distance between the sigma points of the model hypotheses in $\mathcal{H}$ will lead to separation of the output distributions, thus enhancing fault diagnosability.

Let $\{\mathsf{P}_i^{(l)}(t)\}_{i=1}^{n_{\text{tot}}}$ and $\{\mathsf{P}_j^{(o)}(t)\}_{j=1}^{n_{\text{tot}}}$ denote the total sigma points corresponding to the model hypotheses H_l and H_o, respectively. The L_2 distance between each sigma point of H_l to every sigma point of H_o is computed, and the indices of the k closest sigma points of H_o to each sigma point of model H_l are stored. After iterating over all points of model H_l, a total of $t_s = kn_{\text{tot}}$ pairs of sigma points are obtained. Let $\{(p_1,q_1),...,(p_{t_s},q_{t_s})\}$ denote the stacked indices of the k closest sigma points of model M_o to each sigma point of model H_l. The kNN measure of probabilistic discrimination between the two models is defined as

$$d^{(l,o)}(t) = \sum_{(p=p_1,q=q_1)}^{(p_{t_s},q_{t_s})} \Big(\|\mathsf{P}_p^{(l)}(t) - \mathsf{P}_q^{(o)}(t)\|_2 \Big),$$

which quantifies the distance between t_s pairs of sigma points predicted by models H_l and H_o. For the multiple model hypotheses in the model set $\mathcal{H}$, the optimization criterion can now be defined as

$$J = \sum_{l=0}^{n_f} \sum_{o=0,l\neq o}^{n_f} w_{l,o} d^{(l,o)}(t), \tag{4.38}$$

where the weights $w_{l,o}$ are user-specified. For example, $w_{m,n} = (\frac{1}{n_f})$ when uniform weights are used for discrimination between all models. The optimization criterion (4.38) allows for handling multiple model hypotheses and tuning the importance of probabilistic discrimination between different model-hypothesis pairs. The tractable surrogate problem for Problem 1 can now be formulated as follows.

Problem 3 (Tractable OED). For the fault diagnosis horizon $t \in [0, T]$ and the (possibly uncertain) initial states $\boldsymbol{x_o}$, the active fault diagnosis for the stochastic nonlinear system described by (4.4) with model hypotheses $\mathcal{H}$ involves solving the nonlinear optimization problem

$$\max_{u(t)} \; J$$

subject to the propagation of the total sigma points $\{\mathsf{P}_i^{(l)}(t)\}_{i=1}^{n_{\text{tot}}}$ for all model hypotheses $l = 0, 1, \dots, n_f$ in $\mathcal{H}$, the input constraints $u(t) \in \mathcal{U}$, and an approximation of the chance constraint $\mathbb{P}(\boldsymbol{x}(t) \leq \mathcal{X}) \geq \alpha$.

Remark 2. The tractable Problem 3 can be solved online using system measurements obtained at every sampling point to account for the effects of unmodeled system uncertainties. This requires recursive estimation of the initial states as well as the model probabilities for all model hypotheses in $\mathcal{H}$ at every sampling point via Bayesian estimation [35].

4.4.3 Simulation Case Study

The proposed OED method is demonstrated on the same continuous bioreactor employed in 4.3.6.

The performance of the OED method is benchmarked against a deterministic OED method that merely minimizes the L_2-norm between the mean of the model outputs. Figure 4.9 shows the distributions of the product concentration at time 2 h predicted by the three process models under the input profiles designed by the proposed OED and the deterministic OED methods. The distributions are constructed based on 500 Monte Carlo runs using the designed input profiles. As can be seen, the proposed OED method enables discriminating between the three model hypotheses more effectively. This is because the OED method can reduce the variance of the outputs, which can in turn reduce the overlap between the outputs of the competing models in the presence of process uncertainties. On the other hand, even though deterministic OED can increase the absolute distance between the means of the output distributions, it cannot effectively reduce the overlap between the distributions to enhance fault diagnosability. The optimal input profiles designed by both OED methods are shown in Figure 4.10. The input designed by the proposed method maintains near maximal dilution rate until approximately 0.5 h, after which it gradually decreases until reaching zero dilution rate at the end of the diagnosis horizon. The input designed via deterministic OED, on the other hand, maintains maximum dilution rate until it is abruptly

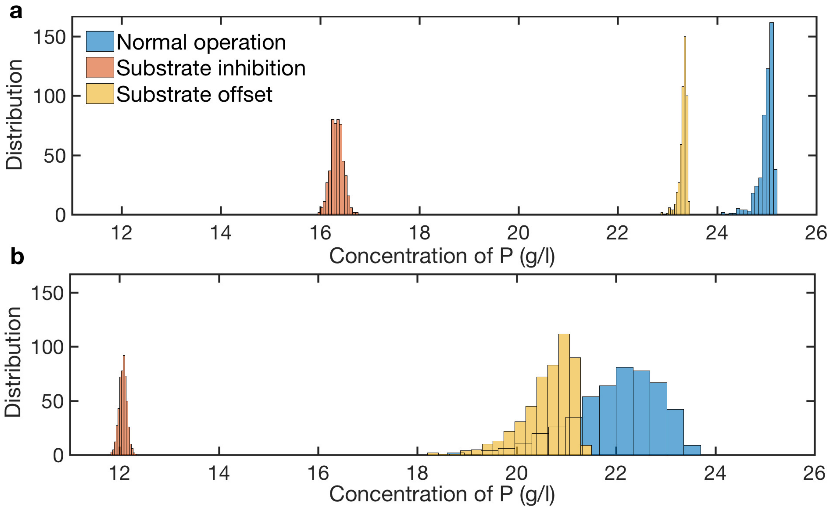

Figure 4.9: Predicted distributions of product concentration at time 2 h obtained using the input designed by (a) the proposed OED method and (b) the deterministic OED method (based on 500 Monte Carlo runs).

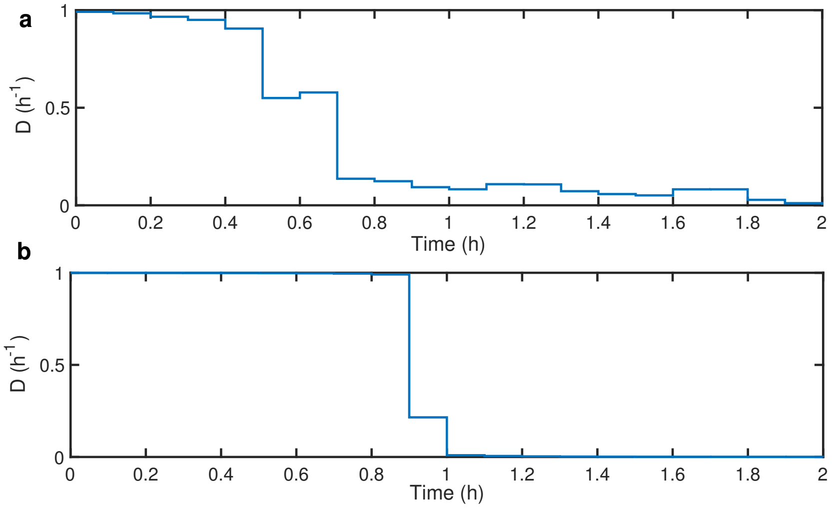

Figure 4.10: Input profiles designed by (a) the proposed OED method and (b) the deterministic OED method.

decreased to near zero values at around time 1 h. Figure 4.10 suggests that deterministic OED results in more process stimulation, while the fault diagnosis is less effective due to the overlap of the output distributions.

Inspired by industrial practice, the proposed OED method is implemented in an online setting in which Problem 3 is solved repeatedly every 6 min when the product concentration is measured (see [35] for the online OED algorithm).

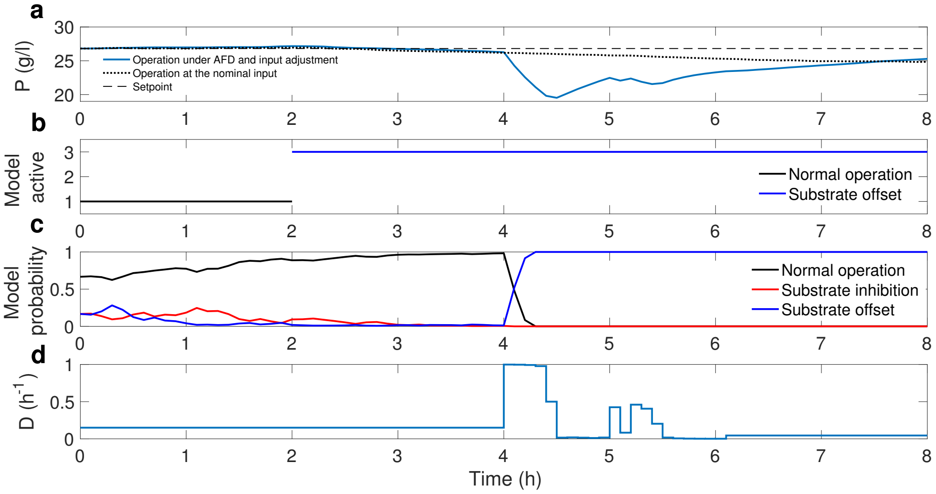

Figure 4.11: Online active fault diagnosis. Different segments of time correspond to different operation stages: 0 to 2 h is normal operation, 2 to 4 h is the fault detection period during which the product concentration starts to deviate from its setpoint; 4 to 6 h is the period during which online OED is performed, and 6 to 8 h is the process recovery stage to return to the desired steady-state operation. The plots show (a) the product concentration in the OED run with subsequent adjustment of the dilution rate (solid line) and the product concentration under the nominal dilution rate (dashed line), (b) the active process model over time, (c) the probabilities of each model estimated via Bayesian recursion (see [35]), and (d) the dilution rate.

The results of online OED are shown in Figure 4.11. The process is initially at steady state. At time 2 h, a fault in the concentration of substrate in the inlet feed becomes active. The fault results in deviation of the product concentration from its desired setpoint ($P = 27$ g/l). At time 4 h, online OED is initiated for a period of approximately 2 h. As can be seen, the probability of detecting the *active* model increases dramatically during this period until reaching a probability of approximately 1. When the correct scenario of process operation is diagnosed at time 6 h, the process input is adjusted based on the diagnosed operational scenario, which is different from the nominal process input (Figure 4.11d). The adjusted process input enables recovering the product concentration to its desired setpoint (Figure 4.11a). If the fault remained undiagnosed and the nominal process input continued to be applied, the concentration of product (dashed line in Figure 4.11a) would continue to deviate from the desired setpoint. The results of this case study clearly demonstrate the significance of online OED for effective process operation in the event of faults.

4.5 Conclusions

Two tractable OED methods for online model selection in the context of active fault diagnosis that can handle probabilistic model uncertainty and stochastic disturbances are presented. The proposed methods aim to design input trajectories that can optimally discriminate between uncertain models of normal and faulty operation in a probabilistic sense, while ensuring safe and least intrusive system operation during the fault diagnosis experiment via enforcing input constraints and state chance constraints. In the first method, tractable approximations for evaluation of chance constraints and the Bayes risk for multiple hypotheses discrimination are introduced to enable online solution of the active fault diagnosis problem. For the second method, a sample-based uncertainty propagation method is proposed for joint propagation of time-invariant model uncertainty and time-varying disturbances using non-intrusive generalized polynomial chaos and unscented transformation. The efficacy of the proposed methods for offline and online active fault diagnosis is demonstrated on a continuous bioreactor. The simulation results indicate that the designed input trajectories achieve correct fault diagnosis in all the considered cases despite the presence of stochastic system uncertainty. The online implementation of the active fault diagnosis method, which relies on repeated design of the input trajectories based on system observations, shows superior performance compared to that of the offline implementation. The simulation results also demonstrate that the proposed active fault diagnosis method can be seamlessly integrated into practical settings.

4.6 Bibliography

1. P. Agrawal, G. Koshy, and M. Ramseier. An algorithm for operating a fed-batch fermentor at optimum specific-growth rate. *Biotechnology and Bioengineering*, 33:115–125, 1989.

2. J. Andersson, J. Åkesson, and M. Diehl. CasADi: A symbolic package for automatic differentiation and optimal control. In *Recent Advances in Algorithmic Differentiation*, pages 297–307. Springer, 2012.

3. I. Andjelkovic, K. Sweetingham, and S. L. Campbell. Active fault detection in nonlinear systems using auxiliary signals. In *Proceedings of the American Control Conference*, pages 2142–2147. Seattle, 2008.

4. A. E. Ashari, R. Nikoukhah, and S. L. Campbell. Auxiliary signal design for robust active fault detection of linear discrete-time systems. *Automatica*, 47:1887–1895, 2011.

5. L. Blackmore, S. Rajamanoharan, and B. C. Williams. Active estimation for jump Markov linear systems. *IEEE Transactions on Automatic Control*, 53:2223–2236, 2008.

6. L. Blackmore and B. Williams. Finite horizon control design for optimal model discrimination. In *Proceedings of the* 45*th* *IEEE Conference on Decision and Control and European Control Conference*, pages 3795–3802. Seville, 2005.
7. L. Blackmore and B. Williams. Finite horizon control design for optimal discrimination between several models. In *Proceedings of the* 45*th* *IEEE Conference on Decision and Control*, pages 1147–1152. San Diego, 2006.
8. M. Blanke, M. Kinnaert, J. Lunze, M. Staroswiecki, and J. Schröder. *Diagnosis and Fault-tolerant Control*. Springer, Berlin, 2006.
9. R. E. Caflisch. Monte Carlo and Quasi-Monte Carlo methods. *Acta Numerica*, 7:1–49, 1998.
10. G. C. Calafiore and L. El Ghaoui. On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130:1–22, 2006.
11. S. L. Campbell and R. Nikoukhah. *Auxiliary Signal Design for Failure Detection*. Princeton University Press, 2015.
12. Z. Chen. Bayesian filtering: From Kalman filters to particle filters, and beyond. *Statistics*, 182:1–69, 2003.
13. L. H. Chiang, E. L. Russell, and R. D. Braatz. *Fault Detection and Diagnosis in Industrial Systems*. Springer, 2000.
14. T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13:21–27, 1967.
15. M. K. Cowles and B. P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91:883–904, 1996.
16. S. Ding. *Model-based Fault Diagnosis Techniques: Design Schemes, Algorithms, and Tools*. Springer, 2008.
17. E. Fix and J. Hodges. Discriminatory analysis-nonparametric discrimination: consistency properties. Technical report, 1951.
18. G. C. Goodwin and R. L. Payne. Dynamic system identification: Experiment design and data analysis. *Mathematics in Science and Engineering*, 136, 1977.
19. N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings on Radar and Signal Processing*, volume 140, pages 107–113, 1993.

20. M. A. Henson and D. E. Seborg. Nonlinear control strategies for continuous fermenters. *Chemical Engineering Science*, 47:821–835, 1992.
21. P. L. Houtekamer and H. L. Mitchell. Ensemble Kalman filtering. *Quarterly Journal of the Royal Meteorological Society*, 131:3269–3289, 2005.
22. T. Kailath. The divergence and Bhattacharyya distance measures in signal selection. *IEEE Transactions on Communication Technology*, 15:52–60, 1967.
23. K. K. Kim, D. E. Shen, Z. K. Nagy, and R. D. Braatz. Wiener's polynomial chaos for the analysis and control of nonlinear dynamical systems with probabilistic uncertainties. *IEEE Control Systems Magazine*, 33:58–67, 2013.
24. L. Ljung. Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems. *IEEE Transactions on Automatic Control*, 24:36–50, 1979.
25. J. MacGregor and A. Cinar. Monitoring, fault diagnosis, fault-tolerant control and optimization: Data driven methods. *Computers & Chemical Engineering*, 47:111–120, 2012.
26. M. Martin-Casas and A. Mesbah. Active fault diagnosis for stochastic nonlinear systems: online probabilistic model discrimination. *10th IFAC International Symposium on Advanced Control of Chemical Processes*, pages 696–701, 2018.
27. S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2:575–601, 1992.
28. A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz. Active fault diagnosis for nonlinear systems with probabilistic uncertainties. In *Proceedings of the IFAC World Congress Conference*, pages 7079–7084. Cape Town, 2014.
29. P. Mohajerin Esfahani and J. Lygeros. A tractable fault detection and isolation approach for nonlinear systems with probabilistic performance. *IEEE Transactions on Automatic Control*, 61:633–647, 2016.
30. A. Nemirovski. Convex optimization in engineering: Modeling, analysis, algorithms. *Lecture Notes Israel Institute of Technology*, 1998.
31. A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17:969–996, 2006.
32. R. Nikoukhah. Guaranteed active failure detection and isolation for linear dynamical systems. *Automatica*, 34:1345–1358, 1998.

33. J. Nocedal and S. J. Wright. *Sequential Quadratic Programming.* Springer, 2006.

34. R. J. Patton and J. Chen. Observer-based fault detection and isolation: Robustness and applications. *Control Engineering Practice*, 5:671–682, 1997.

35. J. Paulson, M. Martin-Casas, and A. Mesbah. Input design for online fault diagnosis of nonlinear systems with stochastic uncertainty. *Industrial Engineering and Chemistry Research*, 56:9593–9605, 2017.

36. J. A. Paulson, V. A. Bavdekar, and A. Mesbah. Output-feedback stochastic model predictive control with joint chance constraints for nonlinear systems. *International Journal of Robust and Nonlinear Control*, submitted 2016.

37. J. A. Paulson and A. Mesbah. An efficient method for stochastic optimal control with joint chance constraints for nonlinear systems. *International Journal of Robust and Nonlinear Control*, pages 1–21, 2017.

38. J. A. Paulson, D. M. Raimondo, R. Findeisen, R. D. Braatz, and S. Streif. Guaranteed active fault diagnosis for uncertain nonlinear systems. In *Proceedings of the European Control Conference*, pages 926–931. Strasbourg, 2014.

39. L. Pronzato. Optimal experimental design and some related control problems. *Automatica*, 44:303–325, 2008.

40. M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43:908–929, 1998.

41. J. K. Scott, R. Findeisen, R. D. Braatz, and D. M. Raimondo. Input design for guaranteed fault diagnosis using zonotopes. *Automatica*, 50:1580–1589, 2014.

42. D. H. van Hessem, C. W. Scherer, and O. H. Bosgra. LMI-based closed-loop economic optimization of stochastic process operation under state and input constraints. In *Proceedings of the 40^{th} IEEE Conference on Decision and Control*, pages 4228–4233. Orlando, 2001.

43. A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.

44. E. A. Wan and R. van der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, pages 153–158. Alberta, 2000.

45. D. Xiu and G. E. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal of Scientific Computation*, 24:619–644, 2002.

46. M. Zagrobelny, L. Ji, and J. B. Rawlings. Quis custodiet ipsos custodes? *Annual Reviews in Control*, 37:260–270, 2013.

47. X. J. Zhang. *Auxiliary Signal Design in Fault Detection and Diagnosis.* Springer, 1989.

Chapter 5

Bayesian optimal experiment design for parameter estimation

Optimal design of experiments is crucial for maximizing information content of experimental data, which can be used for the accurate estimation of model parameters. Classical optimal experiment design approaches are computationally tractable, but they can be inefficient for systems that present nonlinearities and non-gaussian uncertainty and/or heterogeneity. Bayesian optimal experiment design maximizes an expected utility objective, which fully accounts for prior and posterior uncertainty in the model parameters from an information-theoretic standpoint. However, due to the complicated form of the expected utility, it must be estimated using sample-based methods and is expensive to solve using the full dynamic model. In this chapter, we introduce a novel Bayesian optimal experiment design approach for nonlinear models that uses so-called arbitrary polynomial chaos, which readily applies to any type of prior distribution. Numerical simulations indicate that the proposed polynomial chaos-based surrogate can significantly lower the computational cost of the Bayesian optimal experiment design. As such, the methodology presented in this chapter appears to have the potential to pave the way for real-time or sequential dynamic experiment design in a fully Bayesian setting.[*]

5.1 Introduction

The selection of optimal conditions (or designs) for conducting experiments is crucial for improving the information content of observations, especially when experiments are expensive and/or time-consuming to perform. Optimal experiment design (OED) uses a mathematical relationship between the design variables, parameters, and observables of a system to systematically select experimental

---

[*] The content of this chapter was published in [41]

conditions that maximize some design metric that is relevant to, for example, parameter inference or model discrimination [8].

Experiment design has been extensively studied in the classical (or frequentist) framework, in both theory and practice [3]. Classical OED design criteria for parameter inference are generally defined in terms of some scalar metric of the Fisher information matrix (FIM) such as the *alphabetic optimality* criteria A-, D-, and E-optimality [12]. Alternatively, OED can adopt a Bayesian perspective, where the design criteria are expressed in terms of an *expected utility* quantity that accounts for both prior and posterior uncertainty in the model parameters from a decision-theoretic point of view [6]. In the case of linear models subject to Gaussian uncertainties, the Bayesian alphabetic optimality criteria reduce to mathematical forms that are equivalent to their classical FIM counterparts [12]. For example, Bayesian D-optimality corresponds to a utility function equal to the Shannon information of the parameter estimates.

For nonlinear models, however, analytic expressions do not exist for Bayesian design criteria. Thus, extensions of Bayesian OED to nonlinear models are commonly based on linearization of the model and Gaussian approximations of the posterior distribution in order to derive tractable design criteria in terms of the FIM [9, 32]. Many OED approximations involve prior expectations of the FIM, which can be interpreted as robust (or stochastic) versions of classical OED and have also been referred to as *pseudo-Bayesian* OED in the literature [13, 46]. Note that this is in contrast to classical (or "locally optimal") OED approaches that maximize some function of the FIM evaluated around a current "best guess" for the unknown model parameters, which directly leads to deterministic dynamic optimization problems [4]. In both cases, the resulting design criteria can be interpreted as *approximations of an underlying expected utility* in the Bayesian setting. None of these approximations, however, are suitable when the prior distribution is broad (i.e., has large variance) or deviates from normality (i.e., is non-Gaussian), and are known to lead to increasingly suboptimal designs.

This chapter investigates a Bayesian approach to OED for constrained nonlinear systems with continuous (or high-dimensional) design spaces, with the goal of designing experiments that are optimal for parameter inference. Bayesian methods provide the most general framework for experiment design and inference in nonlinear systems with noisy, incomplete, and indirect data [29]. As discussed in [25], the expected utility framework can accommodate a wide variety of information-theoretic criteria. Most often, the expected utility function is explicitly defined in terms of the posterior parameter distribution. One of the most common choices for the expected utility is the mutual information between parameters and observations (equivalent to the expected information gain from the prior to the posterior), which can be expressed in terms of the Kullback-

Leibler (KL) divergence from the posterior to the prior [34]. As such, the main downside of Bayesian OED is its high computational cost relative to classical approaches, which is a direct consequence of numerical evaluation of the expected utility. In general, the expected utility must be approximated using Monte Carlo (MC) integration over the joint observation and parameter space, which can be a high-dimensional space. Thus, early work on Bayesian OED focused on evaluating the expected utility over each element of a small, finite number of designs (on the order of ten) to avoid the challenge of *optimizing* the expected utility over continuous design spaces [48].

Due to the sample-based evaluation of the expected utility, Bayesian OED is naturally formulated as a stochastic optimization problem. In [14], a Markov chain Monte Carlo (MCMC) sampler of the joint design, parameter, and data space is developed such that the marginal distribution of all sampled designs is proportional to the expected utility. Here, the design that leads to the joint mode of the marginal distribution is optimal. Since finding the joint mode is increasingly difficult as the number of design variables increases, a simulated-annealing optimization method was used to achieve faster convergence [36]. However, this approach does not appear to be easily applicable for design dimensions larger than four [49]. Alternative optimization methods, including those based on the Nelder-Mead [37] and simultaneous perturbation stochastic approximation [54], have also been used for Bayesian OED in [25]. The main drawback of these methods is they require many iterations to converge, even for small problems, suggesting they could become excessively expensive for larger design spaces commonly encountered in dynamic OED problems. This is primarily due to the fact that stochastic optimization methods ignore gradient information. Alternatively, so-called "gradient-based" optimization techniques use gradient evaluations to improve the rate of convergence to a local optimum, thus requiring fewer iterations and potentially less computational cost. When applied to problems with stochastic objectives, gradient-based optimization methods can be broadly categorized as stochastic approximation (SA) [33] or sample average approximation (SAA) [51]. Hybrids of these two approaches are also possible. The main practical difference between SA and SAA is that the i.i.d. samples are updated at each iteration in the former while they are treated as fixed in the latter. In either case, when the model used for OED is computationally intensive, evaluating the expected utility and/or its gradients can be computationally prohibitive. To address this challenge, [26] proposed the use of a *surrogate model* for fast estimation of the expected utility, where polynomial approximations (in particular, polynomial chaos expansions (PCEs) [61]) of the model outputs are constructed to capture their dependence on the uncertain parameters and design variables. The main downside of this approach is that the size of the surrogate grows exponentially with respect to the number of design variables, making it impractical for dynamic systems.

This chapter extends the Bayesian OED approach of [26] in several directions. First, we introduce a PCE-based surrogate model that is particularly advantageous for dynamic systems. The proposed approach is based on developing local PCEs for the outputs around each design visited during optimization such that the exponential growth in the size of the PCEs with respect to the number of design variables is avoided. Second, we leverage the theory of orthogonal polynomials to construct the PCEs with respect to *arbitrary* probability measures of uncertain parameters (e.g., priors can be correlated or multi-modal). Thus, the proposed approach is not restricted to particular prior types and also ensures the PCEs are most accurate in high probability regions of the parameter space. Third, the proposed Bayesian OED approach can handle nonlinear probabilistic path and terminal constraints, which can be enforced to ensure safety and/or quality of the experiment. We show how probabilistic constraints can be readily incorporated into the OED problem using the PCE-based surrogate model. A key feature of the proposed Bayesian OED approach is that it can be implemented using state-of-the-art dynamic optimization methods (e.g., multiple shooting or collocation on finite elements [7]), so that the underlying structure of the optimization problem can be exploited for computational efficiency as in classical OED. The Bayesian OED approach is demonstrated on a benchmark dynamic predator-prey problem. To the best of our knowledge, this is the first study on Bayesian OED for nonlinear dynamic systems in the presence of a fairly general class of constraints.

5.2 Formulation of optimal Bayesian experimental design

We are interested in choosing the best experiments from a continuous design space, for the purpose of estimating model parameters from noisy and indirect measurements. In other words, we are interested in experiments that are "optimal" for parameter inference performed in the Bayesian setting, without the need for limiting assumptions such as linear models or strong observability.

Let $(\mathit{\Omega}, \mathfrak{F}, P)$ be a probability space, where $\mathit{\Omega}$ is the sample space (or abstract set of outcomes), $\mathfrak{F}$ is a σ-algebra of the subsets of $\mathit{\Omega}$, and $P : \mathfrak{F} \rightarrow [0, 1]$ is a probability measure. Let the vector of real-valued random variables $\theta : \mathit{\Omega} \rightarrow \mathit{\Theta} \subseteq \mathbb{R}^{n_\theta}$ denote the uncertain model parameters of interest, i.e., these are parameters that we desire to estimate from experimental data. A probability measure μ_θ is induced by the random variables θ, such that $\mu_\theta(A) = P(\theta^{-1}(A))$ for all $A \in \mathbb{R}^{n_\theta}$ (often referred to as the induced or pushforward measure). We can then define $p_\theta(\theta) = d\mu_\theta/d\theta$ as the probability density of θ with respect to the Lebesgue measure. This density is guaranteed to exist as long as the random variables are continuous, which we will assume throughout this work. For simplicity of notation, we shall use $p(\cdot)$ to represent all density functions, and which specific distribution it corresponds to is reflected by its arguments, e.g., $p(\theta)$ denotes $p_\theta(\theta)$. When needed for clarity, we will explicitly include a subscript that denotes the associated random variable.

In a similar fashion, we treat the observations from the experiment $y \in \mathcal{Y}$ (also referred to as "observations", "noisy measurements", or "data") as a real-valued random vector with an appropriate probability density, and $d \in \mathcal{D}$ as the design (also referred to as "control" or "input") variables. Since we are particularly interested in the dynamic evolution of the experiment, we focus on systems modeled by a collection of nonlinear ordinary differential equations (ODE) for ease of presentation

$$\dot{x}(t) = f(t, x(t), d(t), \theta), \qquad \forall t \in [0, t_f] \tag{5.1}$$

where $x : [0, t_f] \to \mathbb{R}^{n_x}$ are the state variables with time derivatives $\dot{x}$ and initial conditions $x(0) = x_0$ and $d : [0, t_f] \to \mathbb{R}^{n_{\text{input}}}$ are the design variables. As such, the state evolution $x(t; d, \theta)$ is implicitly a function of the input trajectory and the model parameters. We assume that the dynamic evolution of (5.1) is constrained so that it must satisfy hard input constraints $d(t) \in \mathbb{D} \subset \mathbb{R}^{n_d}$ and probabilistic (or chance) state constraints of the form

$$P(x(t; d, \theta) \in \mathbb{X}) \geq 1 - \beta, \qquad \forall t \in [0, t_f], \tag{5.2}$$

where $\beta \in [0, 1]$ is the allowed probability of constraint violation. The constraints (5.2) can be interpreted as a generalization of nominal ($\beta = 0.5$) or worst-case ($\beta = 0$) enforcement of state constraints, and can be used to ensure safety or quality throughout the experiment. Terminal state constraints can be handled in a similar fashion to the path constraints (5.2) so we neglect them here to limit the notational complexity. We also assume that measurements can be taken throughout the experiment at discrete times $t_1, \ldots, t_T$ and can be modeled as

$$y_i = g(t_i, x(t_i; d, \theta)) + \epsilon_i, \qquad i = 1, \ldots, T, \tag{5.3}$$

where $y_i \in \mathbb{R}^{n_{y_i}}$ and ϵ_i denote the measurement and the noise in the measurement at time t_i, respectively. The set of observations is then given by $y = (y_1, \ldots, y_T)$ while the corresponding noise vector is given by $\epsilon = (\epsilon_1, \ldots, \epsilon_T)$. Note that we have not made any assumptions on the noise model so that it can have any distribution.

It is important to note that the observation space $\mathcal{Y} \subseteq \mathbb{R}^{n_y}$, $n_y = n_{y_1} + \cdots + n_{y_T}$ is represented by a finite number of dimensions. The design space $\mathcal{D}$, on the other hand, is represented by the set of all continuous-time *trajectories* that satisfy $d(t) \in \mathbb{D}$, which is an infinite-dimensional space. Because of this fact, the design space must be discretized in order to approximate $\mathcal{D} \subseteq \mathbb{R}^{n_d}$ numerically on a computer with finite n_d. Some common approximations include piecewise constant and piecewise linear though, in theory, any finite-dimensional parametrization can be utilized. Letting N_T denote the number of parameters used to approximate the continuous-time trajectories, the total number of design variables becomes $n_d = n_{\text{input}} N_T$, which can be quite large in practice.

If an experiment is performed under a given design d and a realization of the data y is then measured, the change in the state of knowledge/information about the parameters is given by Bayes' rule:

$$p(\theta|y,d) = \frac{p(y|\theta,d)p(\theta|d)}{p(y|d)}, \tag{5.4}$$

where $p(\theta|d)$ is the prior density, $p(y|\theta,d)$ is the likelihood function, $p(\theta|y,d)$ is the posterior density of interest, and $p(y|d) = \int_{\Theta} p(y|\theta,d)p(\theta|d)d\theta$ is the evidence, which represents a normalizing constant that is a function of the design and data. In most practical applications, the prior knowledge on θ does not vary with the choice of design, leading to the simplification $p(\theta|d) = p(\theta)$, i.e., knowing the design of the current experiment without knowing its observations does not affect our belief about the parameters. The likelihood function is assumed to be given, and describes the discrepancy between the observations and a forward model prediction in a probabilistic way. Note that the likelihood function has a one-to-one relationship with the noise model. The forward model $G : \Theta \times \mathcal{D} \to \mathcal{Y}$, generally denoted as $G(\theta, d)$, in this case is implicitly defined by (5.1) and (5.3). Using this notation, we have $y = G(\theta,d) + \epsilon$ with a corresponding likelihood function $p(y|\theta,d) = p_\epsilon(y - G(\theta,d))$. While the majority of classical OED approaches are developed around the assumption that $\epsilon \sim \mathcal{N}(0, \Sigma_\epsilon)$, i.e., a Gaussian likelihood, fully Bayesian approaches are not restricted by this choice and can handle any choice of p_ϵ including noise distributions that depend on the design and parameter values.

We take a decision-theoretic approach to define the *expected utility* (or expected reward) to quantify the value of experiments. As suggested originally in [34], the expected utility can take on the following general form:

$$U(d) = \int_{\mathcal{Y}} \int_{\Theta} u(d,y,\theta)p(\theta,y|d)d\theta dy = \int_{\mathcal{Y}} \int_{\Theta} u(d,y,\theta)p(\theta|y,d)p(y|d)d\theta dy, \tag{5.5}$$

where $u(d,y,\theta)$ denotes the utility function. The utility function should be chosen to reflect the usefulness of an experiment at conditions d, given a particular value of the parameters θ and outcome y. Since the precise values of θ and y are unknown when the experiment is performed, the objective is defined as the expectation of $u(d,y,\theta)$ over the joint distribution of θ and y. It is important that the utility function incorporates the experimental aims and is specific to the application of interest [25]. For example, designs that result in efficient estimation of the model parameters may not be useful for the prediction of future outcomes. A key advantage of Bayesian OED, however, is that a wide variety of experimental goals can be accommodated through the proper choice of utility function including parameter estimation, model discrimination, and the prediction of future observations [48].

Although utility functions are quite flexible and can be tailored to specific goals, in order to derive useful and illustrative results, we focus on utility functions that lead to valid measures of *information gain* on the estimated parameters from the experimental data. In particular, we use the relative entropy, also known as the Kullback-Leibler (KL) divergence, from the posterior to the prior [34]:

$$u(d, y, \theta) = D_{\mathrm{KL}}(p_{\theta|y,d}(\cdot)||p_{\theta}(\cdot)) = \int_{\Theta} \ln\left[\frac{p(\theta|y,d)}{p(\theta)}\right] p(\theta|y,d)d\theta = u(d,y). \quad (5.6)$$

The intuition behind this expression is that a large KL divergence from the posterior to the prior implies that the data y decreases the entropy in θ by a large amount and, hence, those data are more informative for parameter estimation. Note that this choice of utility function integrates over the parameter space and is therefore not a function of the parameters θ. As a result, substituting (5.6) into (5.5) produces the following expression [25]:

$$U(d) = \int_{\mathcal{Y}} \int_{\Theta} \ln\left[\frac{p(\theta|y,d)}{p(\theta)}\right] p(\theta|y,d)d\theta p(y|d)dy. \quad (5.7)$$

Thus, $U(d) = \mathbb{E}_{y|d}\left\{D_{\mathrm{KL}}(p_{\theta|y,d}(\cdot)||p_{\theta}(\cdot))\right\}$ represents the *expected information gain* on the parameters θ where the expectation is taken over the prior predictive distribution $p(y|d)$. Note that $U(d)$ is also equivalent to the *mutual information* between the parameters θ and data y conditioned on d. As discussed in [26], the KL divergence has several desirable properties that warrants its use as a general-purpose utility function for parameter inference, which we briefly summarize here: (i) it satisfies minimal requirements to be a valid measure of information on a set of experiments, such as “always at least informative” ordering; (ii) it gives an indication of information gain in the sense of Shannon information; (iii) it applies to general nonlinear models $G(\theta, d)$ and is consistent with linear optimal design theory based on the FIM; and (iv) it has been shown to perform well for a wide range of tasks, as it provides general guidance for learning in an uncertain environment.

Finally, the optimal design is defined as the design that maximizes the expected utility $U(d)$ subject to constraints:

$$d^{\star} = \underset{d\in\mathbb{D}}{\operatorname{argmax}}\ U(d), \quad \text{s.t.} \quad P(x(t;d,\theta) \in \mathbb{X}) \geq 1-\beta. \quad (5.8)$$

There are a number of challenges that must be overcome when solving (5.8). The biggest difficulty is related to the probabilistic operators that define the expected utility (5.7) and the state chance constraints (5.2), which cannot be evaluated in closed-form even when the model is a simple polynomial function of θ. These challenges are addressed in the next section using different types of approximations.

5.3 Stochastic dynamic optimization with chance constraints

In this section, we formulate the proposed approximated form of the Bayesian OED problem (5.8). Although many different approaches are available for approximating the integrals in (5.2) and (5.7), we focus on two particular choices that lead to a smooth optimization problem, which can be readily solved with state-of-the-art methods for dynamic optimization. The procedure for approximately solving (5.8) is then summarized at the end of this section.

5.3.1 Sample-based estimator for the expected utility

The expected utility in (5.7) does not have a closed-form solution, except when the forward model is a linear function of θ. Instead, this expression must be numerically approximated. By applying Bayes' rule to the quantities inside and outside of the logarithm, and then approximating the integrals using MC, we obtain the following nested MC estimator for the expected utility [49]

$$U(d) \approx \hat{U}_{N,M}(d, \theta_s, y_s) = \frac{1}{N} \sum_{i=1}^{N} \ln \left[\frac{p_{y|\theta,d}(y^{(i)}|\theta^{(i)}, d)}{\frac{1}{M} \sum_{j=1}^{M} p_{y|\theta,d}(y^{(i)}|\tilde{\theta}^{(i,j)}, d)} \right], \tag{5.9}$$

where $\theta_s = \{\theta^{(i)}\}_{i=1}^{N} \cup \{\tilde{\theta}^{(i,j)}\}_{i,j=1}^{N,M}$ are i.i.d. samples from the prior $p(\theta)$ and $y_s = \{y^{(i)}\}_{i=1}^{N}$ are independent samples from the likelihood $p(y|\theta^{(i)}, d)$. The inner sum is needed to approximate the evidence evaluated at $y^{(i)}$, i.e., $p(y^{(i)}|d)$, which typically does not have an analytic form. The variance of this estimator is proportional to $A(d)/N + B(d)/NM$ and its bias (to leading order) is $C(d)/M$, where A, B, and C are constant terms that depend only on the distributions at hand [49]. Hence, the size of N controls variance while the size of M controls the bias. Note that the estimator $\hat{U}_{N,M}$ is asymptotically unbiased, but is biased for finite M. Although alternative numerical integration schemes can replace MC in (5.9), MC is likely the method of choice since its convergence properties are independent of dimension [10] and $n_y + n_\theta$ will often by large in practice. Additionally, MC can be directly applied to any form of prior or likelihood function.

5.3.2 Moment-based approximation of chance constraints

In general, the state constraints $x(t;d,\theta) \in \mathbb{X}$ in (5.2) can be described by a set of nonlinear inequality constraints of the form

$$\mathbb{X} = \{x \in \mathbb{R}^{n_x} : h(x) \leq 0\}, \tag{5.10}$$

where $h : \mathbb{R}^{n_x} \to \mathbb{R}^{n_c}$. Letting $c(x) = \max_{1\leq j \leq n_c} h_j(x)$ where h_j is the j^{th} element of h, a MC estimator can also be developed for the state chance constraints [40]

$$P(h(x(t;d,\theta)) \leq 0) = \mathbb{E}\{\mathbf{1}_{[0,\infty)}(c(x(t;d,\theta)))\} \approx \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{[0,\infty)}(c(x(t;d,\theta^{(i)}))), \tag{5.11}$$

where $\mathbf{1}_{[0,\infty)}$ denotes the indicator function. It is important to note that since the indicator function and the max operator are non-smooth functions, they must be implemented with binary variables. Therefore, under this approximation, (5.8) must be recast as a mixed-integer nonlinear program (MINLP) with a potentially large number of binary variables, i.e., the optimization can be difficult and expensive to solve. In addition, the feasible region of the problem depends on the particular set of samples used in the approximation and can change in a non-smooth way whenever new samples are drawn.

A simpler alternative is to develop a *moment-based* approximation of the chance constraint. The most common example is the mean-variance representation [38]

$$\mathbb{E}\{H_j(t,\theta,d)\} + r(t)\sqrt{\text{Var}\{H_j(t,\theta,d)\}} \leq 0, \quad j = 1,\dots,n_c, \quad \forall t \in [0,t_f], \tag{5.12}$$

where $r(t) \in [0,\infty)$ is a backoff parameter that can vary over time and $H_j(t,\theta,d) = h_j(x(t;\theta,d))$ is the random variable associated with the j^{th} constraint function. The parametrized constraints (5.12) will be smooth functions of time and the design variables whenever the original system equations are smooth in direct contrast to (5.11). Although these constraints can be straightforwardly derived whenever $H(\theta,d)$ is Gaussian [43], this is not a necessary condition for (5.12) to guarantee satisfaction of the original chance constraints (5.2), as discussed next.

Whenever $H(t,\theta,d)$ is normally distributed, $r(t) = r$ is constant and its smallest possible value can be determined exactly from a quantile of the chi-squared distribution. Finding a suitable $r(t)$ becomes more challenging in the case of general distributions, and this is even further complicated by the fact that the shape of the distribution of $H(t,\theta,d)$ can change over time and with the design variable. One promising method for overcoming these challenges in the nonlinear setting is to use an iterative simulation-based procedure that requires $P(x(t; d^\star(r), \theta) \in \mathbb{X})$ to be estimated under the optimal design $d^\star(r)$ found by solving the Bayesian OED problem (5.8) subject to (5.12) in place of (5.2). If the probability of violation is greater (or less) than β, then r should be increased (or decreased). A lower bound of $r = 0$ (corresponding to nominal constraints) can be used, while an upper bound for r can be determined from the "distributionally robust" Cantelli-Chebyshev inequality [5] that ensures the original chance constraints hold for every possible distribution that shares the same mean and covariance as $H(t, \theta, d)$. Note that the violation probability must still be estimated using, for example, MC, but this only needs to be done at the optimal design, as opposed to all of the designs visited during the optimization. Readers are referred to [45] for a similar methodology that has been applied in the context of nonlinear model predictive control. Note that alternative methods have been developed to approximate chance constraints using smoothing functions (e.g., [20]), which can be used in place of (5.12) in this work.

5.3.3 Sample average approximation for chance constrained Bayesian OED

The Bayesian OED problem with moment-based chance constraint approximation can be solved using either gradient-based or non-gradient-based methods. Although gradient-based methods require additional information, they are generally more efficient than their non-gradient-based counterparts. This is especially true in dynamic systems of the form (5.1) that are known to have certain structures that can be efficiently exploited. An important consideration, however, is that U can only be approximated by MC-based estimators such as $\hat{U}_{N,M}$, meaning that optimization methods for stochastic objective functions are needed. The Robbins-Monro (RM) stochastic approximation is one such gradient-based stochastic optimization method. RM is based on an iterative update that resembles steepest descent, except for the fact that it uses an unbiased estimator of the gradient, i.e., the samples used to estimate $U(d)$ may be different than those used to estimate $\nabla_d U(d)$. The choice of the step size sequence is often viewed as a key weakness of RM, as the performance of the algorithm can be very sensitive to the step size [33]. Sample average approximation (SAA), on the other hand, reduces the stochastic optimization problem to a deterministic one by fixing the noise throughout the entire optimization [51]. The main advantage of SAA over RM is that deterministic optimization methods can be directly applied, including state-of-the-art solvers that efficiently handle nonlinear objectives and constraints. For this reason, we focus on SAA exclusively in this work. It is worth noting, however, that the proposed approach is not restricted to SAA and could be solved using RM in a similar manner to that shown in [26], with the main difference being that the nonlinear constraints must be handled using either projection or barrier methods [11].

SAA requires all the "noise" samples θ_s and y_s to be fixed. However, the samples y_s are design-dependent, as they are distributed according to the likelihood function that in turn depends on the given design d. This issue can be addressed in practice by transforming y to be in terms of only design-independent random variables. One example of this transformation for a Gaussian likelihood function is as follows

$$y = G(\theta, d) + \epsilon = G(\theta, d) + C(\theta, d)z, \tag{5.13}$$

where C is a diagonal matrix with the non-zero entries representing the standard deviation of the noise that can generally depend on the parameters and design, and z is a vector of i.i.d. standard normal random variables. For example, the choice of elements $C_{i,j} = 0.1|G_i(\theta, d)|\delta_{ij}$ corresponds to "10% Gaussian noise on the i^{th} component of the model" where δ_{ij} is the Kronecker delta. Other forms of the likelihood can be easily accommodated by replacing the right-hand size of (5.13) by a generic function of θ, d, and some design-independent random vector z.

Let $\mathcal{D}_r$ denotes the set of design variables that satisfy $d(t) \in \mathbb{D}$ and the mean-variance constraints (5.12) for a given backoff radius parameter r. We can then state the proposed SAA approximation to the Bayesian OED problem (5.8) as

$$\hat{d}_s = \underset{d \in \mathcal{D}_r}{\operatorname{argmax}} \; \hat{U}_{N,M}(d, \theta_s, z_s), \tag{5.14}$$

where $\hat{d}_s$ and $\hat{U}_{N,M}(\hat{d}_s, \theta_s, z_s)$ are, respectively, the optimal design and objective values under a particular set of realizations of the random variables θ and z. A deterministic optimization algorithm can then be used to find $\hat{d}_s$ as an approximation to $d^\star$. Estimates of $U(\hat{d}_s)$ can be improved by applying the estimator $\hat{U}_{N',M'}(\hat{d}_s, \theta_{s'}, z_{s'})$ under a larger number of realizations $N' > N$ and $M' > M$ in order to attain a lower variance. Furthermore, multiple optimization runs $B > 0$ can be performed (often referred to as a *bootstrap*) to obtain a sampling distribution for the optimal design values and the optimal objective values, i.e., $\hat{d}_s^b$ and $\hat{U}_{N,M}(\hat{d}_s^b, \theta_s^b, z_s^b)$ for $b = 1, \dots, B$. The sets θ_s^b and z_s^b are independently chosen for each optimization run, but remained fixed within each run. It has been shown that the optimal design and objective estimates converge *in distribution* to their respective true values under certain assumptions [31, 51]. Lastly, stochastic bounds on the true optimal value can be constructed by estimating the optimality gap from the set of B replicate runs. Using the optimality gap estimator and/or its variance estimated from the MC standard error formula, one can decide whether more runs are required or which of the B optimal designs are most trustworthy.

The approximated Bayesian OED problem is still a challenging problem to solve due to the dynamic forward model. In fact, $NM + N$ separate ODEs of the form (5.1) must be integrated in order to evaluate (5.9) at a single design point.

There are a number of ways to handle the infinite-dimensional nature of these ODE constraints (due to the continuous time variable t), including variational, sequential, and simultaneous approaches [7]. Simultaneous methods discretize both the state and design/control profiles in time using, for example, collocation of finite elements, and have the advantage of only solving the ODEs once at the optimal point, i.e., can avoid intermediate solutions that may not exist or require excessive computational effort. Even with this efficient implementation, the huge number of constraints needed to account for the $NM + N$ forward model evaluations can make (5.14) impractical to solve. In addition, we do not have a closed-form expression for the mean and covariance of $H(t, \theta, d)$, which is needed to evaluate the constraint function. Therefore, in the next section, we develop surrogate models for G and H that can greatly reduce the computational cost at each iteration of the optimization while still ensuring that the solutions found are accurate approximations to (5.14).

5.4 Arbitrary polynomial chaos expansions as surrogates

The main challenge in applying the aforementioned stochastic optimization algorithms to the constrained Bayesian OED problem is the complexity of the forward model and its gradients. In fact, only a single evaluation of $\hat{U}_{N,M}(d, \theta_s, z_s)$ requires $O(NM)$ separate solutions of the forward model while an even larger number of equations must be solved to calculate $\nabla_d \hat{U}_{N,M}(d, \theta_s, z_s)$ as the gradient is defined in terms of the sensitivities $\nabla_d G(\theta, d)$. Here, we address this challenge by replacing G with a simple surrogate model based on polynomial expansions (PCEs). This cheaper "surrogate" must be accurate over the entire support of the prior $\mathit{\Theta}$ and the entire design space $\mathcal{D}$. Not only does the surrogate model allow the nested MC estimator in (5.9) to be evaluated in a computationally tractable manner, but its polynomial form greatly simplifies the structure and complexity of the gradient of the expected utility.

These gains come at the cost of introducing a new source of error due to the polynomial approximation of the forward model; however, this error can often be kept low in practice. In fact, the error can always be decreased by increasing the order of the expansion for reasonably smooth functions, as discussed in this section. It is worth noting that we focus on PCE-based surrogates as they have been demonstrated to be effective in the context of Bayesian OED, yielding high accuracy and multiple order-of-magnitude speedups over direct evaluation of the forward model [25, 26]. The proposed surrogate in this work has some important differences to that used in [25, 26], including that our method readily applies to arbitrary prior distributions and the size of the surrogate does not scale with the number of design variables. We first present our proposed surrogate and then describe these differences in more detail at the end of this section.

5.4.1 PCE formulation

With slight abuse of notation, we describe the proposed PCE approximation in the context of a generic scalar function G. Whenever this function is multivariate, the procedure is simply applied to each component of G. As such, the developed procedure can be separately applied to each component of the forward model $G(\theta, d)$ and constraint function $H(\theta, d)$ in the Bayesian OED problem.

The truncated PCE approximation (in terms of the uncertain parameters only) is then defined in the following manner

$$G_L(\theta, d) = \sum_{i=1}^{L} a_i(d) \Psi_i(\theta), \tag{5.15}$$

where L is the total number of terms retained in the expansion; $a_i(d)$ are the expansion coefficients that depend on the design variables; and $\Psi_1, \ldots, \Psi_L$ are polynomial basis functions. We again highlight the fact that θ has an associated probability density $p(\theta)$ on which we have made no restrictions. We can define an inner product $\langle \cdot, \cdot \rangle_\theta$ operator with respect to $p(\theta)$ as

$$\langle f, g \rangle_\theta = \mathbb{E}\{f(\theta) g(\theta)\} = \int_\Theta f(\theta) g(\theta) p(\theta) d\theta, \tag{5.16}$$

for any functions f and g. We can also define a corresponding norm $\|f\|_\theta = \langle f, f \rangle_\theta^{1/2}$ using the definition of the inner product. Now, let $L_\theta^2 = \{f : \|f\|_\theta < \infty\}$ denote the Hilbert space of square integrable functions with respect to density $p(\theta)$. Thus, $G \in L_\theta^2$ is a necessary condition for (5.15) to converge as $L \to \infty$ and is equivalent to the random variable $G(\theta, d)$ having finite variance.

The basis functions can be any complete basis of L_θ^2; however, the computation of the expansion coefficients can be simplified by choosing the basis to be orthogonal with respect to $p(\theta)$. Thus, let $\Psi_1, \Psi_2, \ldots$ be a polynomial orthonormal basis (ONB) of L_θ^2, i.e., each element Ψ_i is a polynomial and for all $i, j \geq 1$ we have

$$\langle \Psi_i, \Psi_j \rangle_\theta = \delta_{ij}, \tag{5.17}$$

where δ_{ij} is the Kronecker delta. In practice, the ONB is constructed to have the following properties: (i) the first polynomial is a constant $\Psi_1(\theta) = 1$ meaning $\mathbb{E}\{\Psi_i(\theta)\} = \delta_{1i}$ is a convenient expression for the expectation of polynomials and (ii) each polynomial Ψ_i contains exactly one monomial θ^α that is not contained in the previous set of polynomials $\Psi_1, \ldots, \Psi_{i-1}$. Most often the polynomials are ordered by degree. Therefore, when approximating G as in (5.15), we first select the number of terms L and define the ansatz space $\mathcal{P}$ as the span of the first L polynomials

$$\mathcal{P} = \{\Psi_1, \dots, \Psi_L\}. \tag{5.18}$$

Note that the size of the expansion L can be chosen to include polynomials of any order, but L is most often chosen according to a "total order" truncation in which all polynomials with degree less than or equal to n_o are retained. This results in the total number of terms in (5.15) being equal to

$$L = \binom{n_\theta + n_o}{n_o} = \frac{(n_\theta + n_o)!}{n_\theta! n_o!}, \tag{5.19}$$

which grows exponentially with respect to the number of uncertainties and the maximum order of polynomials in the expansion.

A variety of methods exist for numerically constructing the polynomial ONB. As mentioned earlier, the approach in [61] (that is applied in the context of Bayesian OED in [25, 26]) assumes separable $p(\theta) = p(\theta_1) \cdots p(\theta_{n_\theta})$, i.e., statistically independent elements of $\theta = (\theta_1, \dots, \theta_{n_\theta})$, so that the construction of the ONB can be done for each dimension separately. These polynomials have been analytically derived for certain scalar probability densities coming from the Askey scheme, and can be derived numerically for generic distributions using algorithms based on three-term recurrence relations for univariate orthogonal polynomials [19]. Whenever θ is composed of statistically dependent elements, a more sophisticated numerical procedure is required to construct the ONB. One example is the Gram-Schmidt process, which is capable of orthonormalizing any starting basis of $\mathcal{P}$, such as the set of monic polynomials (see, e.g., [42] for details). An alternative method is based on a modified Cholesky decomposition of the Gram moment matrix [44], which has shown to be reasonably stable on a variety of examples in [18]. In any case, it is sufficient to know the statistical moments of θ up to a certain order to construct the polynomial ONB. This is an advantage of expanding in θ directly (as opposed to transforming θ into a set of independent random variables), as we are only required to know moments of θ as opposed to an exact expression for $p(\theta)$ [39].

Note that there do exist density functions for which L^2_θ does not admit an ONB of polynomials, i.e., the space of polynomials is not dense in L^2_θ. Interested readers are referred to [17] for more details on this aspect and a list of sufficient conditions to verify the denseness of polynomials. However, since knowing θ is continuous with finite support is sufficient for the space of polynomials to be dense in L^2_θ [17, Theorem 3.4], this will rarely be an issue in practice.

5.4.2 Convergence, optimality, and error analysis

Since $G \in L^2_\theta$ by assumption, we are able to expand it with respect to the ONB of polynomials $\{\varPsi_1, \varPsi_2, \ldots\}$

$$G(\theta, d) = \sum_{i=1}^{\infty} a_i(d) \varPsi_i(\theta), \tag{5.20}$$

where the equality sign in (5.20) should be interpreted in the mean-square sense [17], such that

$$\lim_{L\to\infty} \mathbb{E}\{(G(\theta, d) - G_L(\theta, d))^2\} = \lim_{L\to\infty} \|G - G_L\|^2_{L^2_\theta} = 0. \tag{5.21}$$

In other words, the PCE (5.15) exhibits mean-square convergence. Standard probability theory states that mean-square convergence implies convergence in probability and also convergence in distribution, i.e., $F_{G(\theta,d)}(x) = \lim_{L\to\infty} F_{G_L(\theta,d)}(x)$ for all $x \in \mathbb{R}$ where F_X denotes the cumulative distribution function (CDF) of any random variable X. The rate of convergence depends on the regularity of G with respect to θ and, when G is a smooth function of θ, the convergence rate can be quite large. This means that high accuracy can be achieved in practice with a relatively low order expansion.

According to the Hilbert projection theorem, the best $\|\cdot\|_\theta$ approximation of G in the polynomial space $\mathcal{P}$ is the *orthogonal projection* of G onto $\mathcal{P}$ [53]. This statement can be given mathematically in terms of the optimality condition

$$G_L = \underset{P\in\mathcal{P}}{\operatorname{argmin}} \ \|G - P\|^2_\theta, \tag{5.22}$$

such that no other choice of coefficients $a_1, \ldots, a_L$ will result in a smaller weighted L^2_θ norm. Since the weight function in (5.22) is the density $p(\theta)$, the optimal expansion G_L must more closely match G in regions of $\varTheta$ where the parameter has high probability in order to ensure this norm is small.

Whenever G_L is numerically calculated, we only find approximations to the expansion coefficients and thus obtain the following approximated polynomial

$$\tilde{G}_L(\theta, d) = \sum_{i=1}^{L} \tilde{a}_i(d) \varPsi_i(\theta). \tag{5.23}$$

As such, the difference between G and $\tilde{G}_L$ can be split into two terms: a truncation error and an aliasing error [59]

$$G - \tilde{G}_L = \underbrace{G - G_L}_{\text{truncation error}} + \underbrace{G_L - \tilde{G}_L}_{\text{aliasing error}} = \sum_{i=L+1}^{\infty} a_i \varPsi_i + \sum_{i=1}^{L} (a_i - \tilde{a}_i) \varPsi_i. \tag{5.24}$$

According to the orthogonality property of the ONB (5.17), these two sources of error are orthogonal such that their squared L^2_θ norm is additive

$$\mathbb{E}\{(G(\theta,d)-G_L(\theta,d))^2\} = \|G-G_L\|^2_{L^2_\theta} = \sum_{i=L+1}^{\infty} a_i^2 + \sum_{i=1}^{L}(a_i-\tilde{a}_i)^2. \tag{5.25}$$

Since the ansatz space $\mathcal{P}$ is fixed, the truncation error is constant for a fixed forward model. This is directly controlled by the choice of L, i.e., larger L leads to lower truncation error and solely depends on the nonlinearity of G. Thus, different methods for approximating these expansion coefficients can easily be compared by the aliasing error that they introduce.

Another important property of PCE is that the moments of the random variable $G(\theta,d)$ can be easily computed from only the expansion coefficients

$$\mathbb{E}\{G(\theta,d)\} = a_1(d) \approx \tilde{a}_1(d), \tag{5.26}$$

$$\mathrm{Var}\{G(\theta,d)\} = \sum_{i=2}^{\infty} a_i(d)^2 \approx \sum_{i=2}^{L} \tilde{a}_i(d)^2. \tag{5.27}$$

These equations can be straightforwardly substituted into the mean-variance chance constraint approximation (5.12) so that it can be expressed simply in terms of the PCE coefficients for $H(t,\theta,d)$.

5.4.3 Estimation of PCE coefficients

There are two main approaches for approximating the expansion coefficients: intrusive and non-intrusive [30]. The intrusive approach derives a new system of equations for the coefficients that is larger than the original deterministic system [22]. The difficulty of the intrusive approach strongly depends on the character of the original equations and is often prohibitive (or even impossible to derive) for nonlinear systems [16]. Non-intrusive methods, on the other hand, compute the expansion coefficients from only a finite number of parameter realizations [60]. The main advantage of these approaches is that a deterministic solver for $G(\theta,d)$ can be reused and treated as a black box. Non-intrusive methods also offer flexibility in choosing any function of the state trajectory as the model output, which may depend more smoothly on θ even when the state itself has less regular dependence. In other words, we can avoid representing $x(t;d,\theta)$ with a PCE and instead directly apply the method to $G(\theta,d)$.

Here, we use a non-intrusive approach for estimating the PCE coefficients. The derivation of the method starts from the fact that, by taking the inner product of the expansion (5.20) with one of the basis functions Ψ_i and applying the orthogonality property of the ONB, we obtain an analytic expression for the coefficients

$$a_i(d) = \langle G, \Psi_i \rangle_{L^2_\theta} = \int_\Theta G(\theta, d)\Psi_i(\theta)p(\theta)d\theta, \quad i = 1, \dots, L. \tag{5.28}$$

Then, a set of n sample points $\theta^{(1)}, \dots, \theta^{(n)} \in \Theta$ is chosen and the integrals in (5.28) are approximated using a finite number of forward model evaluations according to some chosen quadrature (or integration) rule [59]

$$\tilde{a}_i(d) = \sum_{j=1}^{n} w_j G(\theta^{(j)}, d)\Psi_i(\theta^{(j)}), \tag{5.29}$$

where $w_1, \dots, w_n$ are corresponding weight values in the quadrature rule. The resulting approach is termed pseudo-spectral projection as it defines a mapping between the forward model G and polynomial $\tilde{G}_L$ that is a *discretized projection operator*.† If a convergent integration rule is employed such that $\lim_{Q\to\infty} \tilde{a}_i = a_i$, then $\lim_{Q\to\infty} \tilde{G}_L(\theta, d) = G_L(\theta, d)$ for all $\theta \in \Theta$ and convergence of $\tilde{G}_L$ to the true forward model G follows naturally.

The key step in any non-intrusive PCE method is the selection of integration points and weights to be used to approximate the coefficients. The number of points n should be as small as possible to achieve a desired level of accuracy in the PCE approximation (5.23). A wide variety of integration (or sampling) rules for multidimensional spaces have been proposed and applied in the context of PCE. Broadly speaking these methods can be categorized as follows: (i) grid-based, (ii) randomized, (iii) monomial cubature rules, or (iv) optimization-based.

Grid-based methods such as tensor and sparse grids [21] are the most commonly used integration rules since they can be easily derived from univariate Gaussian quadrature rules, which are optimal in one dimension [55]. However, tensor-grid quadrature suffers from the *curse of dimensionality* due to the exponential growth of the number of points with dimension of the parameter space. Sparse grids are directly constructed from tensor grids and are built to accurately capture functional features in each separate parameter dimension while investing fewer points in the cross terms between parameters. Although sparse grids have fewer points than the full tensor grid, they have increasingly large error with increasing dimension and are known to produce negative weight values. Another key limitation of tensor and sparse grids is that they require the uncertain parameters to be statistically independent. If the parameters are dependent, then a transformation must be applied, which may place integration points in low probability regions of

---

[†] Regression methods are an alternative class of non-intrusive PCE in which the discrete quadrature rule is directly applied to the optimality condition in (5.22), and these can straightforwardly be used in place of pseudo-spectral methods in this work [52].

Θ that contribute only a very small amount to the PCE projection. Randomized integration rules, on the other hand, select points by randomly sampling from the parameter distribution $p(\theta)$ via MC methods [23]. MC is often the method of choice for approximation of high-dimensional integrals, but are known to require a large number of points to achieve low error due to their relatively slow rates of convergence.

Monomial cubature rules are nongrid-based methods that can be more effective than sparse grids when integrating functions that are well represented by total-degree polynomials [58]. These can be thought of as efficient multivariate extensions of Gaussian quadrature. Their main downside, however, is that effective cubature rules have only been constructed for a very specific set of probability distributions, integration domains, and polynomial degree of exactness. Optimization-based integration rules are based on the same idea as monomial cubature rules, with the main difference being that the quadrature rule is not selected manually. Instead, the integration points and weights are determined numerically through the use of some optimization procedure. In this way, efficient quadrature rules can be constructed for any distribution $p(\theta)$ and any desired polynomial degree. Also, constraints on the position of the points and value of the weights can readily be incorporated.

The main cost of non-intrusive PCE arises from the forward model simulations at fixed nodes n, and these simulations must be repeated for every d visited when numerically solving the Bayesian OED problem (5.8). Thus, we adopt the optimization-based methodology here so that n can be minimized without compromising accuracy of the integration rule. There are two main types of optimization-based methods available: moment matching and optimized stochastic collocation (OSC). The moment-matching rule corresponds to a non-negative measure on Θ that minimizes a sensitivity function subject to the constraints that the measure matches moments of $p(\theta)$ up to a certain finite order [50]. This corresponds to an infinite-dimensional linear program (LP) that must be heuristically solved in practice. One approach, presented in [44], is based on three steps: (i) solve a finite-dimensional LP wherein moments are matched based on a fine grid of Θ, (ii) locate the “clusters” of sample points obtained from the LP solution, and (iii) refine this solution by locally solving a nonlinear least-squares problem with initial guess corresponding to the clustered integration rule. However, the derived moment matching rule can be sensitive to the choice of the initial grid and the clustering step. The OSC method, on the other hand, derives the optimal points and weights through the formal minimization of an integration operator error norm [52]. Here, we adopt the OSC method since it limits the number of heuristic choices by the user and has been shown to effectively handle uncertainty dimensions up to around ten.

5.4.4 The optimized stochastic collocation method

The OSC method is summarized in this subsection. OSC is formulated as a polynomial optimization problem with an objective function that is adapted to be able to efficiently and accurately approximate the PCE coefficients. First, define the exact integral operator for a generic function $f \in L^2_\theta$ as

$$I : L^2_\theta \to \mathbb{R} : f \mapsto \int_\Theta f(\theta)p(\theta)d\theta, \tag{5.30}$$

while, for a given list of points $\boldsymbol{\theta} = (\theta^{(1)}, \dots, \theta^{(n)})$ and weights $w = (w_1, \dots, w_n)$, the discrete quadrature operator is defined as

$$Q_{(\boldsymbol{\theta},w)} : L^2_\theta \to \mathbb{R} : f \mapsto \sum_{j=1}^{n} w_j f(\theta^{(j)}). \tag{5.31}$$

The operators I and $Q_{(\boldsymbol{\theta},w)}$ must be bounded in order to define an operator norm to measure the distance between them. Thus, we restrict I and $Q_{(\boldsymbol{\theta},w)}$ to a finite-dimensional test space $\mathcal{T} \subseteq L^2_\theta$. Let $\mathfrak{L}(\mathcal{T}, \mathbb{R})$ denote the space of all bounded linear operators from $\mathcal{T}$ to $\mathbb{R}$. For any operator $A \in \mathfrak{L}(\mathcal{T}, \mathbb{R})$, the induced operator norm on the space $\mathfrak{L}(\mathcal{T}, \mathbb{R})$ is defined as:

$$\|A\|_{\mathfrak{L}(\mathcal{T},\mathbb{R})} = \sup_{f \in \mathcal{T}} \frac{\|Af\|_\mathbb{R}}{\|f\|_\theta}. \tag{5.32}$$

The OSC method can then be summarized using this induced norm as follows:

1. Choose a finite-dimensional test space $\mathcal{T}$ and number of integration points n.
2. Find the optimal integration points and weights by solving

$$(\boldsymbol{\theta}_{\text{osc}}, w_{\text{osc}}) = \underset{\substack{\boldsymbol{\theta} \in \Theta^n \\ w \in [0,\infty)^n}}{\operatorname{argmin}} \|I - Q_{(\boldsymbol{\theta},w)}\|^2_{\mathfrak{L}(\mathcal{T},\mathbb{R})}. \tag{5.33}$$

Note that there are some basic relationships between $\mathcal{T}$ and n. Mainly, the number of integration points is bounded by $\dim(\mathcal{P}) = L \le n \le t = \dim(\mathcal{T})$ since $n < L$ points cannot even distinguish the L different ansatz functions and t points are always capable of reducing the operator error norm to zero [52].

A good choice for the test space $\mathcal{T}$ can be derived from the integrals that we want to approximate in (5.29). Whenever $G \in \mathcal{P}$, then we would like $\tilde{G}_L = G$, i.e., there is no truncation or aliasing error for polynomial models within the ansatz space $\mathcal{P}$ (5.18). This means that the integral of all products of two elements in $\mathcal{P}$ have to be exact, which corresponds to the test space

$$\mathcal{T} = \operatorname{span}\{\Psi_i \Psi_j,\ 1 \le i, j \le L\}. \tag{5.34}$$

Based on this choice of $\mathcal{T}$, it is then desired to choose n large enough so that the operator norm is reduced to zero. A simple procedure can be derived from the degrees of freedom (DOF) in the optimization problem (5.33). The number of DOF in the optimization is $n(n_\theta + 1)$. In order to reduce the objective function in (5.33) to zero, t equations must be satisfied. Therefore, if we choose $\mathcal{T}$ and n such that

$$t = n(n_\theta + 1), \tag{5.35}$$

then we may have enough integration points to be able to satisfy all t conditions. Since t is fixed according to (5.34), we should initially select $n = t/(n_\theta + 1)$, which is much lower than the upper bound of t. However, it is important to note that this is a heuristic choice and cases exist that n has to be larger or can be chosen smaller. Thus, a practical approach is to first choose n according to the DOF condition (5.35) and then numerically perform the optimization (5.33). If the minimum objective value is not small enough, then n can be increased by one and the optimization repeated until the objective has been reduced to a sufficiently low value. Note that alternative choices of $\mathcal{T}$ and n are discussed in [52].

Remark 2 *Whenever L is chosen using the "total order" truncation method with maximum order n_o, then the choice of test space in* (5.34) *effectively doubles the PCE order such that the dimensionality of the test space is $t = \frac{(n_\theta + 2n_o)!}{n_\theta! 2n_o!}$.*

We can now derive an expression for the operator norm in (5.33) explicitly in terms of the integration points $\boldsymbol{\theta}$ and weights w. Since the elements of $\mathcal{T}$ in (5.34) are polynomials, they can be represented as coordinate vectors with respect to the ONB $\Psi_1, \dots, \Psi_t$, i.e., any function $f \in \mathcal{T}$ can be written as $f = \sum_{i=1}^{t} c_i \Psi_i$. Based on this representation, the numerator of the induced norm can be written as

$$\|Af\|_{\mathbb{R}} = \left\|A\left(\textstyle\sum_{i=1}^{t} c_i \Psi_i\right)\right\|_2 = \left\|\textstyle\sum_{i=1}^{t} c_i A\Psi_i\right\|_2 \leq \|c\|_2 \|A\Psi\|_2, \tag{5.36}$$

where $c = (c_1, \dots, c_t)$ and $A\Psi = (A\Psi_1, \dots, A\Psi_t)$ denotes the vector representation of any operator A with respect to the ONB Ψ. The inequality above directly follows from the well-known Cauchy-Schwartz inequality. Similarly, we can apply this representation to the squared denominator of the induced norm to derive

$$\|f\|_\theta^2 = \int_\Theta \left(\sum_{i=1}^{t} c_i \Psi_i(\theta)\right)^2 p(\theta) d\theta = \sum_{i=1}^{t}\sum_{j=1}^{t} c_i c_j \langle \Psi_i, \Psi_j \rangle_\theta = \sum_{i=1}^{t} c_i^2 = \|c\|_2^2. \tag{5.37}$$

Since the supremum is achieved when the inequality exactly holds, we can combine these two expressions to derive a finite-dimensional representation of the operator norm as the 2-norm of its vector representation, i.e., $\|A\|_{\mathfrak{L}(\mathcal{T},\mathbb{R})} = \|A\Psi\|_2$.

We recall (5.17) to find that $I\Psi_i = \delta_{1i}$, meaning that the vector representation of I can be written as

$$I\Psi = \mathrm{e}_1 = (1, 0, \ldots, 0). \tag{5.38}$$

For $Q_{(\boldsymbol{\theta},w)}$, from (5.31), we find that

$$Q_{(\boldsymbol{\theta},w)}\Psi_i = \sum_{j=1}^{n} w_j \Psi_i(\theta^{(j)}), \tag{5.39}$$

which can be easily converted into its vector representation that is an explicit function of the integration points and weights:

$$Q_{(\boldsymbol{\theta},w)}\Psi = \boldsymbol{\Psi}(\boldsymbol{\theta})w, \tag{5.40}$$

where $\boldsymbol{\Psi}(\boldsymbol{\theta})$ is a $t \times n$ matrix:

$$\boldsymbol{\Psi}(\boldsymbol{\theta}) = \begin{bmatrix} \Psi_1(\theta^{(1)}) & \cdots & \Psi_1(\theta^{(n)}) \\ \vdots & \ddots & \vdots \\ \Psi_t(\theta^{(1)}) & \cdots & \Psi_t(\theta^{(n)}) \end{bmatrix}. \tag{5.41}$$

We can then calculate the squared operator norm of $I - Q_{(\boldsymbol{\theta},w)}$ as

$$\|I - Q_{(\boldsymbol{\theta},w)}\|^2_{\mathfrak{L}(\mathcal{T},\mathbb{R})} = \|\mathrm{e}_1 - \boldsymbol{\Psi}(\boldsymbol{\theta})w\|_2^2, \tag{5.42}$$

which is a sum of squares of polynomial functions. This is a smooth function with structure that can be easily exploited by gradient-based optimization algorithms. An important practical issue in the OSC method (5.33) is finding the global minimum. Since the lowest attainable value of the objective is known to be zero, we are guaranteed to have found a global optimum as long as this bound is reached.

5.4.5 Global versus local PCE with respect to the design space

PCE is simply an orthogonal polynomial approximation to random functions and thus can be applied to G in various ways. For example, in [25] a *single* (or global) PCE is constructed for $G(\theta, d)$ over the entire product of the parameter and design spaces. In this way, a random vector $\xi \in \mathbb{R}^{n_s}$, $n_s = n_\theta + n_d$ is defined to have one dimension associated to each component of θ and one to each component of d. The density $p(\xi)$ is required to be separable and is assumed to map to the joint space $(\theta, d) = T(\xi)$ based on some (possibly) nonlinear diffeomorphism $T : \mathbb{R}^{n_s} \to \Theta \times \mathcal{D}$ that preserves the probability density functions of ξ and (θ, d).

The global PCE can then be defined similarly to (5.15), except now in terms of this new random vector ξ

$$G(\theta(\xi), d(\xi)) \approx \sum_{i=1}^{L} b_i \varPhi_i(\xi), \tag{5.43}$$

where $b_1, \ldots, b_L$ are the global expansion coefficients and $\varPhi_1, \ldots, \varPhi_L$ are polynomials that are orthogonal with respect to $p(\xi)$.

The main advantage of this approach is that the coefficients are constant and therefore only need to be computed once before solving the Bayesian OED problem; however, there are two important limitations. First, the number of terms in (5.43) increases exponentially with $n_s = n_\theta + n_d$ and the truncated order n_o. The effect of this growth is twofold: time-varying trajectories $d(t)$ must be heavily discretized in order to keep n_d small, and accuracy must be sacrificed when G is highly nonlinear in either θ or d to keep n_o small. Second, we must select some probability distribution for the design variables. This distribution represents the weight function that governs what regions of $\mathcal{D}$ that the PCE should be most accurate. Therefore, the probability distribution should be proportional to how often values of d are visited during the optimization algorithm. Since this quantity is too complex to extract in practice, a heuristic strategy must be applied instead. For example, in [25], a uniform weight function over the bounded design space is chosen. As a result, the surrogate can be inaccurate near the unknown optimal design.

The proposed PCE-based surrogate (5.15) avoids both of these issues by developing a *local* surrogate around each design encountered during the optimization. Therefore, the size of the surrogate is completely independent of n_d, and we do not need to artificially define a distribution over the design space. Although the coefficients must be updated at every iteration in the proposed approach, the OSC rule used to define the quadrature operator in (5.29) ensures that this process only requires a minimal number of forward model simulations. This means that we can significantly reduce the number of full model evaluations in (5.14) from $O(NM)$ to merely n. These features suggest that the proposed approach is especially advantageous in dynamic systems, which can very easily result in OED problems with n_d on the order of tens to hundreds of independent variables.

The final important difference between (5.15) and (5.43) is related to the essence of the so-called germ ξ. The *generalized polynomial chaos* (gPC) method requires the stochastic parameters to be statistically independent in order to simplify the basis construction. As long as T is a density-preserving transformation, then ξ can be chosen as any set of independent random variables. The Rosenblatt transformation is the most common example as it applies to any collection of continuous random variables [47]. However, a known problem with the Rosenblatt transformation is that, even for simple problems, T can quickly become discontinuous and highly nonlinear. In fact, it has been shown that transformations between some standard scalar random variables exhibit Gibbs phenomena and thus deteriorate the convergence rate of the expansion [59]. Even when this transformation is reasonably well-behaved, it can be complicated to determine and expensive to evaluate. Therefore, it is preferred to expand in terms of θ when possible. This implementation of PCE has been referred to as *arbitrary polynomial chaos* (aPC) since there are no restrictions on $p(\theta)$, and can be interpreted as a generalization of gPC. We explicitly represent (5.15) using aPC because this helps keep n_o small, which directly results in lower values for n due to smaller-sized test spaces (5.34).

5.5 Numerical results

5.5.1 The dynamic forward model

The well-known Lotka-Volterra (LV) system has been used to model the nonlinear and oscillatory dynamics of interacting predator and prey populations, and is a commonly used benchmark problem in the dynamic OED literature, e.g., [56]. The time-evolution of LV system is governed by the ODEs

$$\dot{x}_1(t) = x_1(t) - (1 + 0.25\theta_1)x_2(t)x_1(t) - 0.4x_1(t)d(t) \tag{5.44a}$$

$$\dot{x}_2(t) = -x_1(t) + (1 + 0.25\theta_2)x_1(t)x_1(t) - 0.2x_2(t)d(t), \tag{5.44b}$$

where $t \in [0, t_f]$ is the time variable with $t_f = 12$, $x_1(t)$ is the normalized prey population, $x_2(t)$ is the normalized predator population, and θ_1 and θ_2 are the unknown parameters for which we have limited information. The design profile $d(t)$ can be manipulated throughout the experiment, and is constrained to the domain $d(t) \in [0, 1]$ for all $t \in [0, t_f]$. We also assume a noisy measurement of the predator population can be made at the final time, i.e., $y = x_2(t_f) + \epsilon$. The noise is modeled as a zero-mean Gaussian random variable with standard deviation $\sigma = 0.1|x_2(t_f)|$, i.e., the noise variance is *state-dependent* and equals 10% of the signal. For this study, we select a statistically-dependent prior in terms of two coupled beta distributions

$$\theta_1 \sim \mathcal{B}(2, 2), \quad \theta_2|\theta_1 \sim \mathcal{B}(\theta_1 + 3, -\theta_1 + 2). \tag{5.45}$$

A contour plot of this joint distribution on the support $(\theta_1, \theta_2) \in [-1, 1]^2$ is shown in Figure 5.1. This prior was chosen as an example of one that is able to capture potential relationships between parameters.

The OED formulation (5.8) seeks the design $d^\star(t)$ such that, when the experiment is performed, on average the predator signal yields the greatest information gain from prior to posterior, i.e., the information gain is averaged over all possible prior parameters and over all possible resulting measured predator populations. State chance constraints are added to the problem after the initial comparisons.

Note that all NLP optimization problems discussed in this case study are numerically solved using CasADi [2] that automatically passes the required derivatives (based on a symbolic implementation of the equations) to the interior point solver IPOPT [57]. In addition, all computations were performed on a MacBook Pro with 8 GB of RAM and a 2.6 GHz Intel i5 processor.

5.5.2 Local and global PCE implementations

Evaluating the forward model requires solving the ODE in (5.44) at fixed realizations of θ and extracting the predator population at the final time. These equations are integrated with CVODE [15] set to a tolerance of 10^{-8}, which is an error-controlled solver for stiff and non-stiff initial value problems. As discussed in Section 5.4, the full forward model can be replaced with a PCE-based surrogate to improve computational efficiency of the MC estimator for the expected utility. In this example, we compare two surrogate modeling approaches: (i) the proposed method that we will refer to as "local" or "design-dependent" PCE for short and (ii) the "global" PCE method proposed in [25, 26].

The local PCE method expands in terms of θ directly, meaning the polynomials must be constructed to be orthogonal to (5.45). This was done by applying the modified Cholesky decomposition to the Gram moment matrix, composed of moments of a finite number of moments of θ. The coefficients of the expansion are estimated using a quadrature rule chosen as the solution to the OSC problem (5.33) with a test space $\mathcal{T}$ composed of all polynomials up to degree 8, which corresponds to $t = \dim(\mathcal{T}) = \frac{(8+2)!}{8!2!} = 45$. Based on the DOF condition in (5.35), n was initially set equal to 15; however, this did not produce satisfactory objective values near zero and so was increased by one. For $n = 16$ points, we were able to consistently find solutions to (5.33) that result in a norm of zero (global optimum) starting from an initial condition with equal weights and nodes sampled randomly from $p(\theta)$. This indicates the OSC rule can integrate 45 polynomials exactly using only 16 points. One such example of a converged OSC rule is shown in Figure 5.1. This rule matches intuition as points with larger weights are concentrated in regions of the density function with higher values. Note that the OSC problem took about 2 seconds on average to solve and more than 90% of the runs converged to the global optimum.

Global PCE, on the other hand, expands with respect to the joint parameter and design space, meaning the design space must be discretized before it can be applied. To this end, the design profile is discretized into N_T piecewise constant

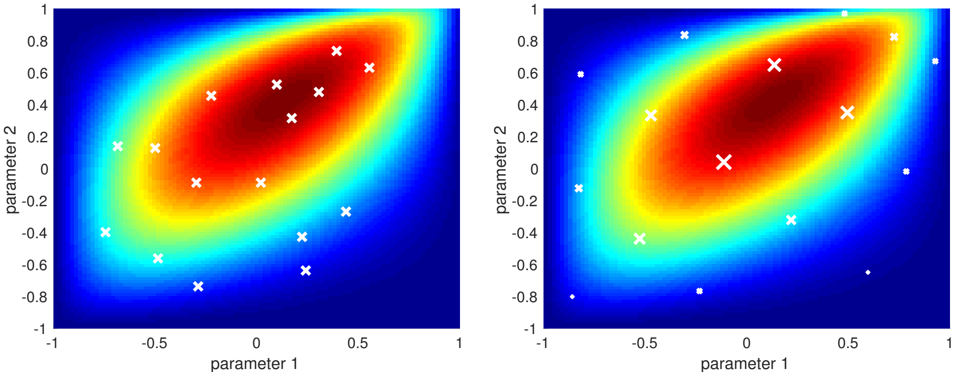

Figure 5.1: Randomly sampled initial condition provided to OSC optimization (5.33) (left) and derived optimal OSC quadrature rule with $n = 16$ nodes that can exactly integrate $t = 45$ polynomials (right). The nodes are shown with a white 'x', and the size of each node is proportional to its weight. The colors represent contours of the joint parameter distribution in (5.45).

intervals $[0, t_f] = [t_0, t_1) \cup \cdots \cup [t_{N_T-1}, t_{N_T})$ such that $d(t) = d_i$ for all $t \in [t_i, t_{i+1})$ and $i = 1, \ldots, N_T$. For comparison purposes, we fix the number of intervals at $N_T = 2$. We thus have two additional variables (d_1, d_2) to include in the PCE, which we assume are independent and uniformly distributed in order to build the surrogate (5.43). We again note that this is heuristic choice, as we do not know the distribution of designs that will be visited during the optimization procedure. The parameter and design variables must then be mapped to a 4-dimensional germ $\xi = T(\theta_1, \theta_2, d_1, d_2)$ that has statistically independent elements. In this case, we choose $\xi_1 \sim \beta(2,2)$, $\xi_2 \sim \beta(2,2)$, $\xi_3 \sim \mathcal{U}(-1,1)$, $\xi_4 \sim \mathcal{U}(-1,1)$, and T according to the Rosenblatt transformation. The coefficients of the global PCE were determined with a tensor product of Gaussian quadrature rules of order 10 that resulted in a total of 10^4 forward model evaluations.

Both the local and global PCE method are implemented using total-order polynomial truncation. In order to select this truncation order, we calculated the L^2_θ mean-squared error (MSE) for various truncation orders, which is plotted in Figure 5.2. We clearly observe that the error decreases as order increases for both methods; however, the local expansion exhibits a faster rate of convergence and has errors more than an order-of-magnitude lower than the global approach. This is not surprising as local PCE does not expand in the design space so that it can directly capture nonlinear effects with respect to d. For global PCE, we selected order 9 as this resulted in reasonably small MSE while retaining less than one thousand terms in the expansion, i.e., $L = 715$. Based on this choice, we selected order 4 for local PCE, which corresponds to $L = 15$, as this is the smallest order with lower MSE than global PCE of order 9.

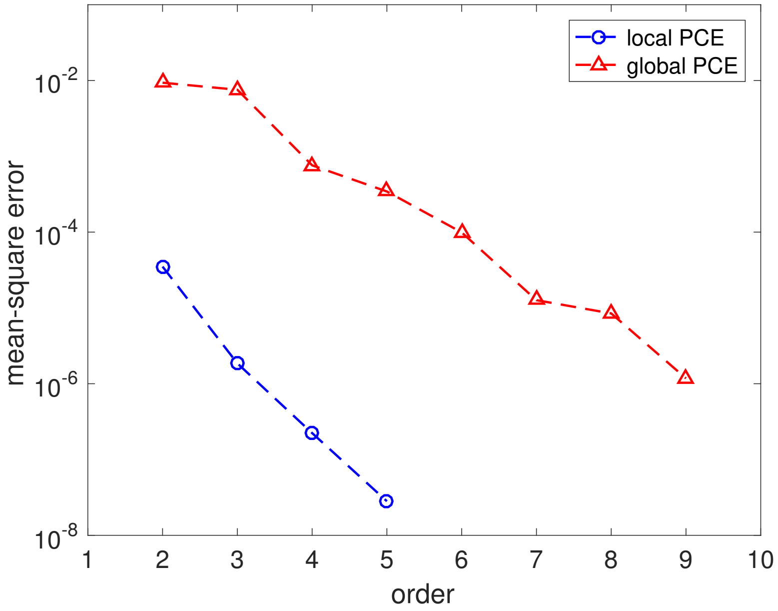

Figure 5.2: The mean-square error (MSE), i.e., L^2_θ norm versus truncation order of the local and global PCE surrogates for the forward model in the LV system.

5.5.3 Stochastic dynamic optimization implementation and results

We now discuss the implementation of the proposed SAA optimization (5.14) and analyze the results for varying number of samples. We first focus on the local PCE method and then compare performance to the global PCE method.

Regardless of the choice of sample sizes N and M, the local PCE method must impose n separate ODE constraints corresponding to the nodes of the OSC rule. A direct transcription approach was used to discretize the state profile in time using collocation of finite elements. We chose 20 elements and used a third-order collocation scheme within each element. Note that the design profile is only discretized into 2 elements in order for the local method to be more fairly compared with the global approach. As such, the dynamic optimization problem has been converted to a large-scale NLP that again can readily be solved using CasADi and IPOPT. We chose to use the limited-memory BFGS estimate of the Hessian, as opposed to exact evaluation of the Hessian, since this provided computational savings in this problem (i.e., cost per iteration decreased more than number of iterations increased).

Under SAA, each choice of sample sets θ_s and z_s yields a different deterministic objective. Example realizations of this objective surface are shown in Figures 5.3–5.5. For each realization, a local optimum is found efficiently in only a relatively few (usually less than 15) iterations. Note that for low N, the objective realizations

Table 5.1: Histograms of the optimal design variables $\hat{d}_s$ from 1000 independent bootstrap runs of SAA over a matrix of N and M sample sizes. For each histogram, the bottom-right axis represents $d(t) = d_1$ for $t \in [0, t_f/2)$, the bottom-left axis represents $d(t) = d_2$ for $t \in [t_f/2, t_f)$, and the vertical axis represents frequency.

M N	2	11	101	1001
1	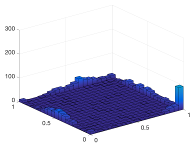	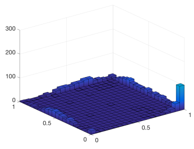	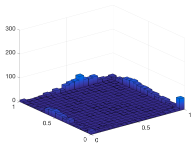	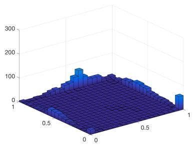
11	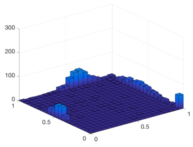	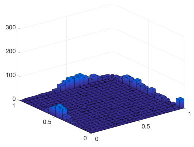	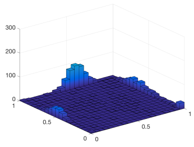	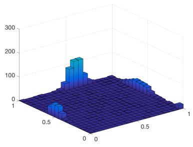
101	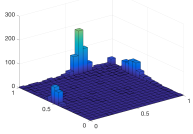	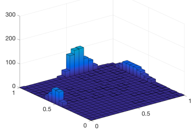	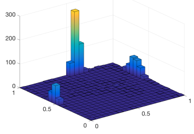	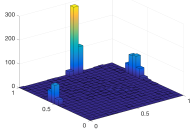

Table 5.2: High-quality estimates of the expected utility (or information gain in this case) at the optimal designs resulting from 1000 independent runs of SAA. For each histogram, the horizontal axis represents values of $\hat{U}_{1001,1001}$ and the vertical axis represents frequency.

M N	2	11	101	1001
1	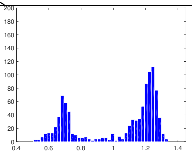	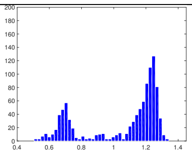	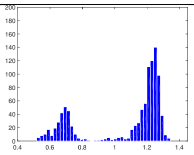	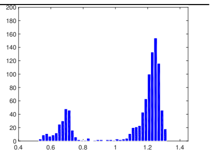
11	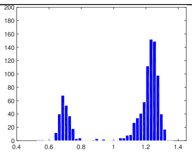	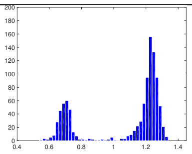	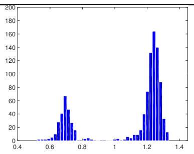	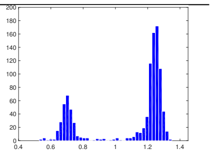
101	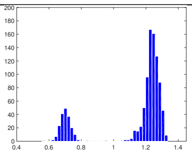	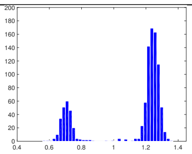	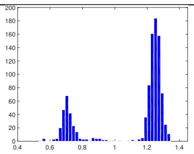	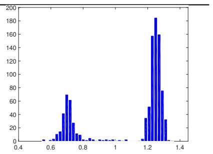

can be extremely different including the location of the optimum points as well as the estimated maximum value of the expected utility. In general, however, the objectives have less variability as N is increased. Looking at $N = 101$ in Figure 5.5, we consistently see two designs that (locally) maximize the objective and one design that minimizes the objective. To better understand the performance of the proposed method, we conducted 1000 independent bootstrap runs, over a matrix of N and M values. Each optimization is initialized with a uniformly distributed random design to assess the performance of the method on average. Histograms of the optimal design variables resulting from each set of 1000 optimization runs are shown in Table 5.1. We can immediately recognize that more designs cluster around the three local optima as N and M are increased. The distribution of final designs is not enough to understand the robustness of the optimization results. For example, if U is flat near the optimum, then the suboptimal designs need not be close to the true optimal design to be considered good designs in practice. A "high-quality" estimate of the objective $\hat{U}_{1001,1001}$ is computed for each of the 1000 designs in Table 5.1 to evaluate robustness, and the resulting histograms are shown in Table 5.2. We can again see that performance improves as N and M increase. It is interesting to note that all histograms in Table 5.2 are bimodal. The higher mode reflects a mixture of the two maxima while the lower mode corresponds to the minimum design. Although the variance in these modes decrease with increasing N and M, both modes are always present. Around 70% of runs converge to the high expected utility mode while 30% of the runs converge to the low mode. Note that similar features are observed when using global PCE as the surrogate model.

5.5.4 Comparison between local and global PCE surrogates

To compare the local and global PCE surrogates, we develop a single integrated measure of the quality of the solutions from the SAA optimization. As suggested in [26], we use the following MSE expression as this metric

$$\mathrm{MSE} = \frac{1}{B}\sum_{b=1}^{B}\left(\hat{U}_{1001,1001}(\hat{d}^b, \theta_{s'}^b, z_{s'}^b) - U^{\mathrm{ref}}\right)^2 \tag{5.46}$$

where $\hat{d}^b$, $b = 1, \ldots, B$ are the final designs from the optimization algorithm for $B = 1000$ bootstraps (using either local or global PCE) and U^{ref} is the true maximum expected information gain. Since the true maximum is not available in this case study, we take U^{ref} to be the maximum value of the objective over all runs. An important issue in the evaluation of the MSE is that it will be significantly biased by the designs that lead to the local minimum. To avoid this bias, we initialized all optimizations at $[1, 1]$ in the design space that consistently produced designs that globally maximize the objective. Figure 5.6 describes the

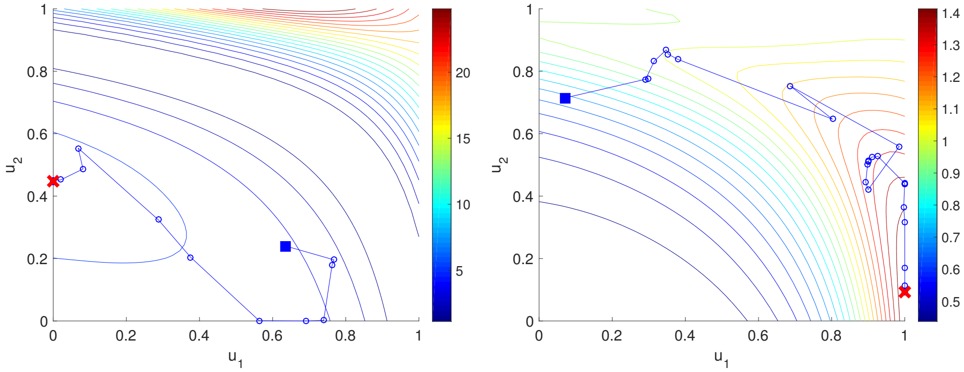

(a) $N = 1$, $M = 2$ (b) $N = 1$, $M = 11$

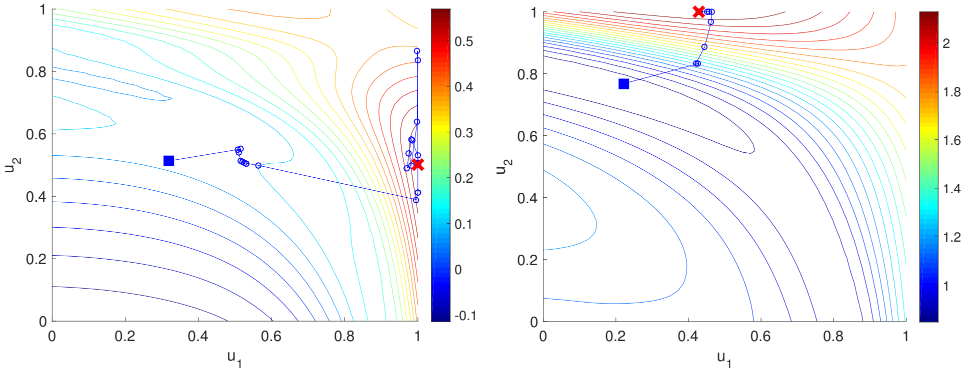

(c) $N = 1$, $M = 101$ (d) $N = 1$, $M = 1001$

Figure 5.3: Realizations of the objective surface using SAA and the corresponding iterations of IPOPT, with $N = 1$ and four separate M values. The blue □ is the starting point and the red × is the final converged point.

solution quality in relation to computational effort by plotting the MSE against average computational time per run for both the local and global PCE methods. Each symbol represents a particular value of N, i.e., ×, ◯ and □ represent $N = 1$, $N = 11$, and $N = 101$, respectively, and the four different M values are shown through average run times.

There are a number of interesting observations within Figure 5.6. Both methods generally result in lower MSE as N and M increase. We also observe that N has a much larger effect than M for this particular example. However, in the $N = 101$ case for the global method, we see that MSE is nearly constant as M increases. This is likely due to the fact that the global PCE has a larger L^2_θ error than the local method and thus hits its lowest achievable error around this value for N. As expected, the global PCE method has solution times that are directly proportional to the total number of samples used to evaluate the expected utility since the cost of the model (5.43) is the same per sample. The

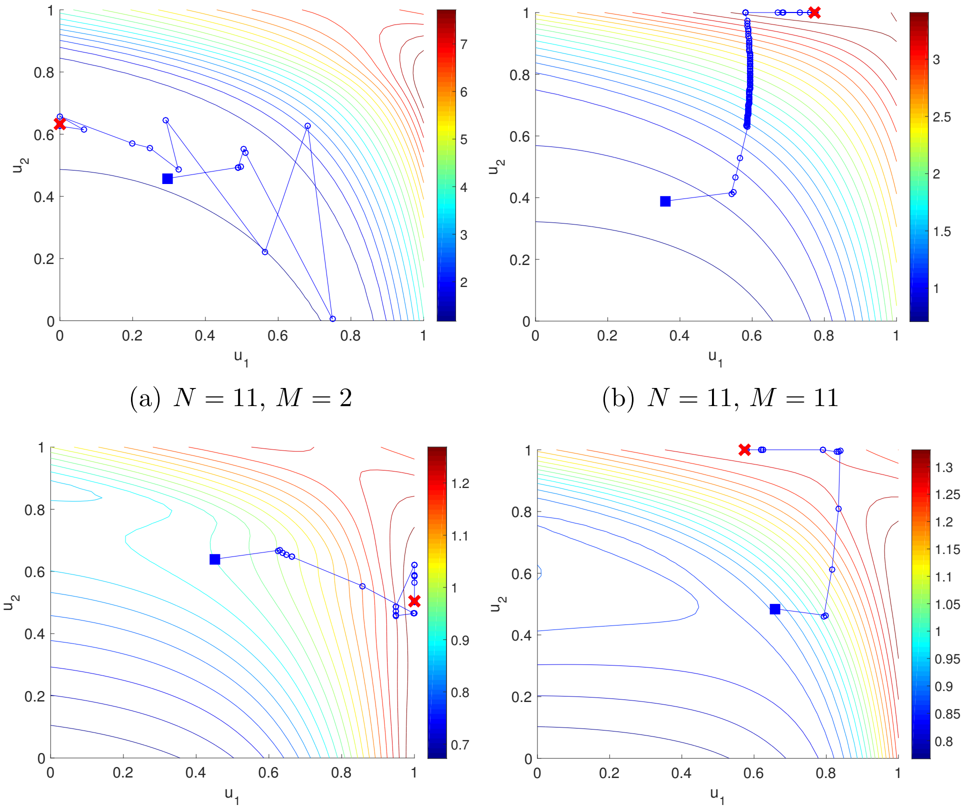

(a) $N = 11,\ M = 2$ (b) $N = 11,\ M = 11$

(c) $N = 11,\ M = 101$ (d) $N = 11,\ M = 1001$

Figure 5.4: Realizations of the objective surface using SAA and the corresponding iterations of IPOPT, with $N = 11$ and four separate M values. The blue □ is the starting point and the red × is the final converged point.

local PCE method, on the other hand, requires a minimum of approximately 1 second to find a solution, regardless of N and M. This is due to the fact that the local method has a fixed cost corresponding to the $n = 16$ discretized ODE constraints. These nonlinear constraints are the dominant cost in the optimization when $N = 1, M = 2$ all the way to $N = 101,\ M = 101$. In other words, we see no increase in the approximately 1 second solution time when there are $NM + N = 3$ versus $NM + N = 10,302$ total evaluations of the polynomial (5.15) needed at each iteration. This highlights the importance of the minimal OSC rule as 16 forward model evaluations are more expensive than over 10,000 surrogate evaluations, meaning we can expect the solution time to massively increase if the full model is evaluated at each sample instead of the surrogate (if even possible to store all of the constraints in memory). We do see that the cost of the polynomial evaluations overtake the ODE cost for the largest considered case of $N = 101,\ M = 1001$. The local PCE method, however, is increasingly cheaper to evaluate than the global method as the number of samples increases. This is a direct consequence

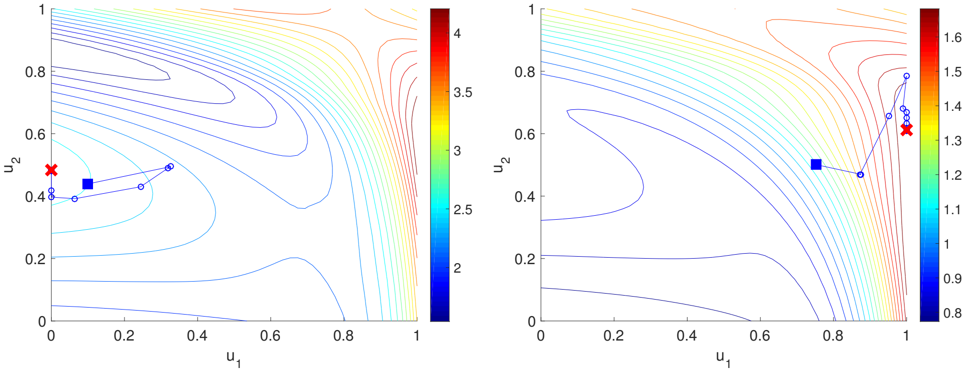

(a) $N = 101$, $M = 2$ (b) $N = 101$, $M = 11$

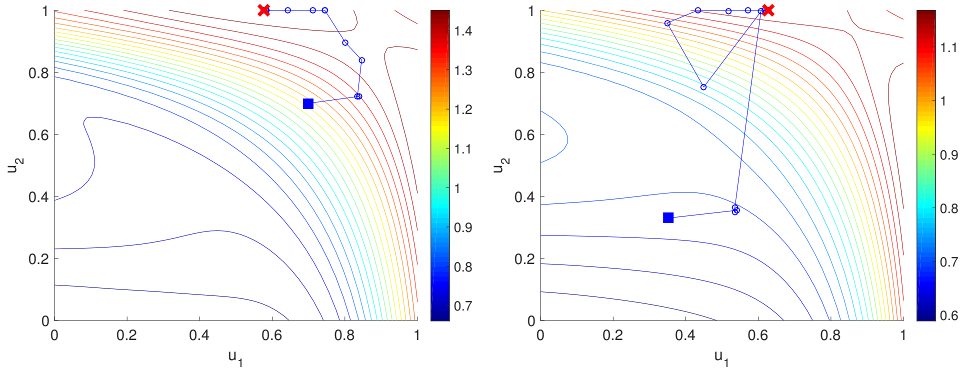

(c) $N = 101$, $M = 101$ (d) $N = 101$, $M = 1001$

Figure 5.5: Realizations of the objective surface using SAA and the corresponding iterations of IPOPT, with $N = 101$ and four separate M values. The blue □ is the starting point and the red × is the final converged point.

of the local expansion having a factor of 50 less terms in the expansion than the global method. In fact, the local method is 7.5 times cheaper than the global method in the largest case considered, while also producing solutions with two order-of-magnitudes lower MSE.

5.5.5 Scaling with respect to number of design variables

The previous analysis focused on the case of $N_T = 2$ discretized design variables. To understand the effect that N_T has on the optimization when using the proposed local PCE method, the average computational time to solve (5.14) (over ten independent runs with $N = 101$ and $M = 101$) is plotted against the number of design variables N_T in Figure 5.7. We can clearly see that the cost scales sublinearly with respect N_T, which is mainly due to the fact that the state discretization level is fixed at 20 elements in all cases as well as sparsity being exploited in the gradient computation. It is important to note that, as N_T increases, the size of the

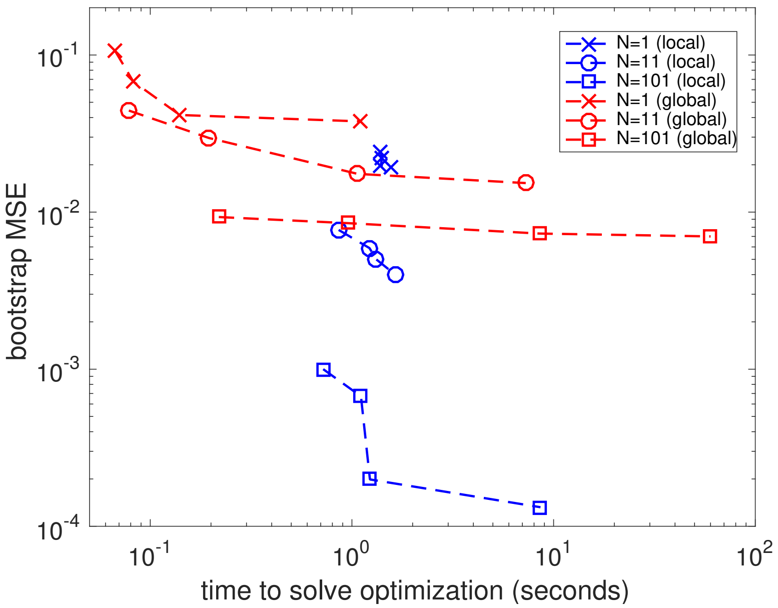

Figure 5.6: Mean-square error (MSE) for the set of "bootstrap" optimization runs, defined by (5.46), versus the average run time for SAA under both the local and global PCE surrogate models and various choices of inner and outer sample sizes N and M.

surrogate model remains fixed and cost only increases due to the larger number of decision variables. This is in sharp contrast to the global PCE model, which grows exponentially in size as N_T increases. For example, the global expansion has more than two million terms for 20 design variables, 2 parameters, and a truncation order of 9. As such, we were unable to apply the global method for $N_T = 20$, while the local method only took approximately 65 seconds to find a solution.

High-quality expected utility estimates for each discretation level are also shown in Figure 5.7. As expected, the optimal objective value increases as the number of design variables increases due to the fact that the design profile has more freedom (i.e., is less constrained). We see a large increase (more than 50%) in the optimal objective for $N_T = 5$ whereas fairly minor increases for larger N_T. This suggests that five intervals provide enough freedom in this problem to find a solution that nearly maximizes the original dynamic problem (5.8), and these solutions can be found in around 25 seconds whereas the original problem is unsolvable.

5.5.6 Chance constraint approximation and tuning

An important consideration in this work is chance constraints of the form (5.2). Here, we consider box state constraints $(x_1(t), x_2(t)) \in [0, 3] \times [0, 1.5]$ and $\beta = 0.1$, meaning 10% of the state profiles are allowed to violate the box constraints.

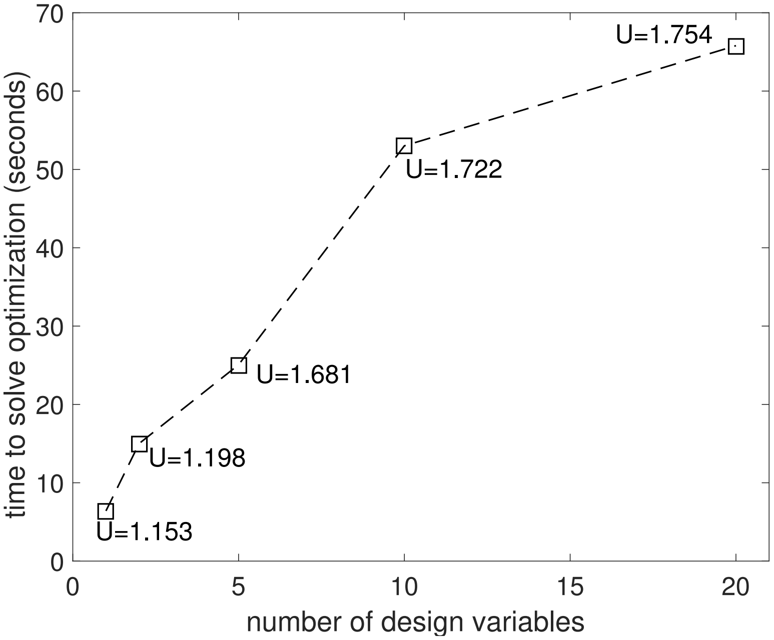

Figure 5.7: Average run time of SAA with the local PCE surrogate versus the number of discretized design variables. The maximum expected utility achieved for each discretization level is also shown.

However, it appears that the only active state constraint for the Bayesian OED problem at hand is $x_2(t) \leq 1.5$, such that $P(x(t; d, \theta) \in \mathbb{X}) = P(x_2(t; d, \theta) \leq 1.5)$. For tractability purposes, we enforce the smooth moment-based approximation (5.12) instead of the original chance constraint. The required mean and variance terms can be expressed as a function of the design variables using the coefficients of the local PCE surrogate as shown in (5.26) and (5.27).

As discussed in Section 5.3, we cannot directly select a good value for the backoff radius parameter r as the distribution of x_2 is unknown at the optimal design. Therefore, we explore a simple simulation-based procedure to determine what value of r results in the largest possible objective. All that is required is to solve the OED problem with $\mathbb{E}\{x_2(t; d, \theta)\} + r\sqrt{\text{Var}\{x_2(t; d, \theta)\}} \leq 1.5$ added for a given r. Once the optimal design for this problem has been found, denoted by $d^\star(r)$, MC simulations are performed on the forward model to empirically estimate $P(x_2(t; d^\star(r), \theta))$. Note that the surrogate model can be used to further speed up the estimation of this violation probability, as established in [45]. If the violation is greater than β, then r should be increased and, if on the other hand, the observed violation is less than β, then r should be decreased. This procedure is repeated until satisfactory convergence is achieved. The basic methodology is presented in Figure 5.8, which plots the estimated violation probability under 1000 MC runs versus the backoff parameter r, with other parameters set to $N = M = 101$ and $N_T = 5$. We see that r has a nonlinear effect on $P(x_2(t; d^\star(r), \theta) \leq 1.5)$. We also see that the constraint violation equals the desired level of 10% when the radius is $r = 1.5$. Due to the simple structure of the constraints combined with the cheap surrogate model, these average solution times are virtually identical with or without (5.12) included.

We emphasize the fact that the violation probability only had to be estimated at the optimal design, as opposed to being calculated at every iteration of the optimization and expressed with binary variables. High-quality estimates of the expected utility are also shown for each backoff radius in Figure 5.8 and, as expected, it decreases with increasing radius due to the reduced feasible region. The expected utility shows almost a linear decrease with increasing β for this problem, but in general this effect can be highly nonlinear. This type of Pareto analysis is useful to perform to probe the tradeoff between performance and "robustness" to parameter uncertainty. It is worth noting that this type of Pareto curve can be easily traced over $\beta \in [0, 1]$ by finding the r that yields $P(x_2(t; d^\star(r), \theta) \leq 1.5) = 1 - \beta$.

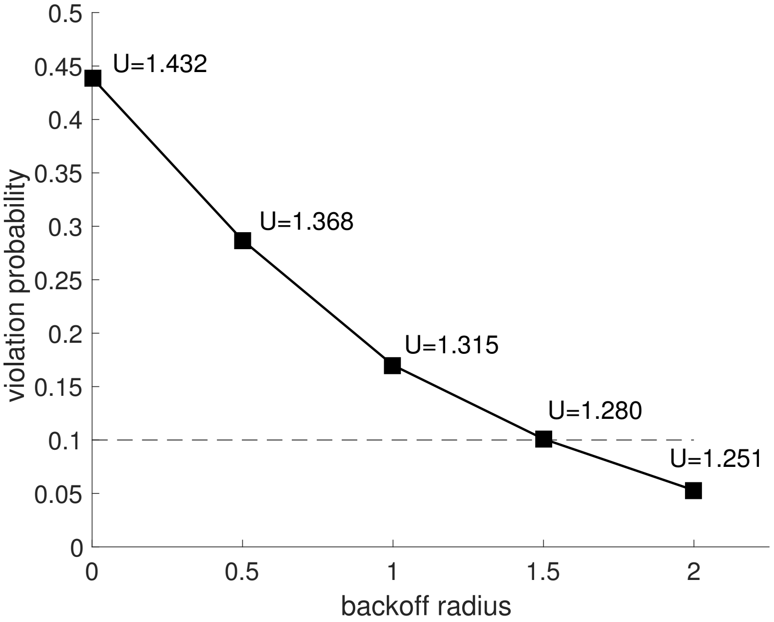

Figure 5.8: The estimated state constraint violation probability under the optimal design profile versus the backoff radius. The maximum expected utility for each backoff radius is also shown.

5.6 Conclusions

This chapter studies the stochastic optimization problem arising from a general nonlinear formulation for Bayesian optimal experiment design (OED), with a particular focus on dynamic systems subject to state chance constraints. The main objective function of interest is the *expected information gain* in the model parameters due to an experiment, which can be written in terms of the KL divergence from the posterior to the distribution. Since this expected information gain cannot be evaluated exactly, we must resort to a finite-sample Monte Carlo (MC) approximation to the objective function using the well-known sample average approximation (SAA) method. Although a similar SAA approximation can be applied to the chance constraints, these constraints require integer variables (non-smooth) and produces a feasible region that changes for every set of realizations used. Therefore, we propose a smooth *moment-based* approximation to the chance constraints that has a tunable backoff parameter, which can be determined through a limited number of simulations to ensure the original constraints are satisfied. Methods for computing the objective must also deal with the fact that the estimator of the expected information gain is not a simple MC sum, but involves a nested sum of MC estimates. It is therefore expensive to evaluate the objective and/or its gradients, as each sample in the estimator requires the dynamic forward model to be integrated over time. As suggested in previous work [25, 26], we look to circumvent these challenges by approximating the forward model with polynomial chaos expansions (PCEs) and subsequently computing the expected information gain with PCEs instead.

The main contribution of this chapter is to develop a new PCE-based surrogate model that is *design-dependent*, i.e., coefficients are updated locally at each design. In this way, the exponential growth with respect to the number of design variables that occurs in the global PCE method of [25, 26] can be avoided. Another key feature of the proposed local PCE method is that the expansions are defined in terms of polynomials that are orthogonal with respect to arbitrary priors, meaning they most accurately approximate the forward model in high probability regions of the parameter space. Since the expansion coefficients are defined as weighted multidimensional integrals, the cost of estimating these coefficients is directly proportional to the number of nodes in a chosen discrete quadrature rule. Therefore, we apply an optimization-based procedure to numerically derive a rule that ensures high accuracy integration with a minimal number of nodes. We compare the performance of the proposed local PCE method to the global method of [25, 26] on the problem of estimating parameters from noisy data in a dynamically evolving predator-prey system. Numerical experiments are performed over a matrix of inner- and outer-loop sample sizes to examine their impact on bias and variance of the objective function. Unsurprisingly, we see that the solution quality improves as the sample sizes increase, but observe that the outer sample size has a larger effect than inner sample size. We also note that multiple local solutions exist, though the global solution is found approximately 70% of the time. When comparing the local and global PCE methods for only two design variables, we observe that the proposed local method can provide a significant speedup and lower error, especially as the sample size increases.

Another important observation is that the average solution time with the local PCE surrogate scales *sublinearly* with respect to the design profile discretization level. This is a result of the coefficients being updated at each design, which ensures the size of the expansion is independent of the design profile. This is in sharp contrast to the global PCE model, which grows exponentially in size as the discretization level increases. We also show that the moment-based constraints can be tuned to guarantee satisfaction of the original chance constraints, and how the surrogate model ensures these added computations are negligible compared to the main cost of estimating the expected information gain.

The developed approach is based on a nested MC estimator for the expected information gain. Some known issues with this estimator are that it requires a large number of samples due to the double-loop sample-average structure and the inner loop can suffer from arithmetic underflow for small sample sizes, diffuse priors, or concentrated posteriors. Future work should focus on developing methods that can avoid these issues. One such example is to replace samples from the prior with those drawn from an importance sampling distribution. A good candidate importance sampler can be derived from the Laplace approximation (LA) [35], which expands the posterior distribution in terms of a second-order Taylor series around the maximum a posteriori (MAP) estimate. Although this reduces the inner-loop sample size, it come at the cost of constructing the LA, so it is not obvious this will reduce the overall computational burden. Note that LA has also been used to directly approximate certain expected utilities in large-scale OED problems [1]. It would also be interesting to explore methods, such as parallelized solvers [28] or distributed optimization algorithms [27], that are capable of exploiting the structure of the NLPs derived from stochastic dynamic optimization problems in order to achieve further reductions in the solution time.

Lastly, this chapter focuses on batch (or open-loop) OED, where the experiment is fully designed before any data are actually collected. An important area of future research is sequential (or closed-loop) OED, where data from previous experiments can be used to guide the design of future experiments. The closed-loop OED problem can be rigorously formulated using dynamic programming [24, Chapter 3], but significant computational challenges must be overcome for this to be practically solvable, especially since the state must be represented in terms of the posterior distribution of the parameters. The proposed PCE-based surrogate model could potentially be used to address some of these challenges and, thus, help pave the way for OED to be solved in real-time in a fully Bayesian setting.

5.7 Bibliography

1. A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas. A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems. *SIAM Journal on Scientific Computing*, 38:A243–A272, 2016.
2. J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, pages 1–36, 2018.
3. A. C. Atkinson and A. N. Donev. *Optimum experimental designs.* New York: Oxford University Press, 2007.
4. I. Bauer, H. G. Bock, S. Körkel, and J. P. Schlöder. Numerical methods for optimum experimental design in dae systems. *Journal of Computational and Applied mathematics*, 120:1–25, 2000.
5. A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski. *Robust Optimization.* Princeton University Press, 2009.
6. J. O. Berger. *Statistical decision theory and Bayesian analysis.* Springer New York, New York, NY, 1985.
7. L. T. Biegler. An overview of simultaneous strategies for dynamic optimization. *Chemical Engineering and Processing: Process Intensification*, 46:1043–1053, 2007.
8. X. Bombois, G. Scorletti, M. Gevers, M. J. Van den Hof, P, and R. Hildebrand. Least costly identification experiment for control. *Automatica*, 42:1651–1662, 2006.
9. G. E. P. Box and H. L. Lucas. Design of experiments in non-linear situations. *Biometrika*, 46:77–90, 1959.
10. R. E. Caflisch. Monte Carlo and Quasi-Monte Carlo methods. *Acta Numerica*, 7:1–49, 1998.
11. P. Carbonetto, M. Schmidt, and N. D. Freitas. An interior-point stochastic approximation method and an L1-regularized delta rule. In *Advances in neural information processing systems*, pages 233–240, 2009.
12. K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.

13. Y. Chu and J. Hahn. Integrating parameter selection with experimental design under uncertainty for nonlinear dynamic systems. *AIChE Journal*, 54(9):2310–2320, 2008.

14. M. A. Clyde, P. Müller, and G. Parmigiani. Exploring expected utility surfaces by markov chains. *Technical report, Duke University*, 1996.

15. S. D. Cohen, A. C. Hindmarsh, and P. F. Dubois. Cvode, a stiff/nonstiff ODE solver in C. *Computers in Physics*, 10:138–143, 1996.

16. B. J. Debusschere, H. N. Najm, P. P. Pébay, M. Knio, O, R. G. Ghanem, and O. P. Le Maitre. Numerical challenges in the use of polynomial chaos representations for stochastic processes. *SIAM Journal on Scientific Computing*, 26:698–719, 2004.

17. O. Ernst, A. Mugler, H. Starkloff, and E. Ullmann. On the convergence of generalized polynomial chaos expansions. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46:317–339, 2012.

18. J. Feinberg, V. G. Eck, and H. P. Langtangen. Multivariate polynomial chaos expansions with dependent variables. *SIAM Journal on Scientific Computing*, 40(1):A199–A223, 2018.

19. W. Gautschi. On generating orthogonal polynomials. *SIAM Journal on Scientific and Statistical Computing*, 3(3):289–317, 1982.

20. A. Geletu, A. Hoffmann, M. Kloppel, and P. Li. An inner-outer approximation approach to chance constrained optimization. *SIAM Journal on Optimization*, 27:1834–1857, 2017.

21. T. Gerstner and M. Griebel. Dimension–adaptive tensor–product quadrature. *Computing*, 71:65–87, 2003.

22. R. G. Ghanem and P. D. Spanos. Stochastic finite element method: Response statistics. In *Stochastic Finite Elements: A Spectral Approach*, pages 101–119. Springer, 1991.

23. S. Hosder, R. Walters, and M. Balch. Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables. In *48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, page 1939, 2007.

24. X. Huan. *Numerical approaches for sequential Bayesian optimal experimental design*. PhD thesis, Massachusetts Institute of Technology, 2015.

25. X. Huan and Y. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232:288–317, 2013.

26. X. Huan and Y. Marzouk. Gradient-based stochastic optimization methods in Bayesian experimental design. *International Journal for Uncertainty Quantification*, 4, 2014.

27. Y. Jiang, P. Nimmegeers, D. Telen, J. Van Impe, and B. Houska. A distributed optimization algorithm for stochastic optimal control. *IFAC-PapersOnLine*, 50(1):11263–11268, 2017.

28. J. Kang, N. Chiang, C. D. Laird, and V. M. Zavala. Nonlinear programming strategies on high-performance computers. In *Proceedings of the 54^{th} IEEE Conference on Decision and Control*, pages 4612–4620, Osaka, 2015. IEEE.

29. M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:425–464, 2001.

30. K.-K. K. Kim, D. E. Shen, Z. K. Nagy, and R. D. Braatz. Wiener's polynomial chaos for the analysis and control of nonlinear dynamical systems with probabilistic uncertainties [historical perspectives]. *IEEE Control Systems*, 33:58–67, 2013.

31. A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12:479–502, 2002.

32. S. Körkel, E. Kostina, H. G. Bock, and J. P. Schlöder. Numerical methods for optimal control problems in design of robust optimal experiments for nonlinear dynamic processes. *Optimization Methods and Software*, 19:327–338, 2004.

33. H. Kushner and G. G. Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Applications of Mathematics, Springer, 2003.

34. D. V. Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956.

35. Q. Long, M. Scavino, R. Tempone, and S. Wang. Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations. *Computer Methods in Applied Mechanics and Engineering*, 259:24–39, 2013.

36. P. Müller, B. Sansó, and M. De Iorio. Optimal bayesian design by inhomogeneous markov chain simulation. *Journal of the American Statistical Association*, 99:788–798, 2004.

37. J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7:308–313, 1965.

38. A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17:969–996, 2006.

39. S. Oladyshkin and W. Nowak. Data-driver uncertainty quantification using the arbitrary polynomial chaos expansion. *Reliability Engineering and System Safety*, 106:179–190, 2012.

40. K. Pagnoncelli, B, S. Ahmed, and A. Shapiro. Sample average approximation method for chance constrained programming: Theory and applications. *Journal of Optimization Theory and Applications*, 142:399–416, 2009.

41. J. Paulson, M. Martin-Casas, and A. Mesbah. Optimal bayesian experiment design for nonlinear dynamic systems with chance constraints. *Journal of Process Control*, 2019.

42. J. A. Paulson, E. A. Buehler, and A. Mesbah. Arbitrary polynomial chaos for uncertainty propagation of correlated random variables in dynamic systems. *IFAC-PapersOnLine*, 50:3548–3553, 2017.

43. J. A. Paulson and A. Mesbah. An efficient method for stochastic optimal control with joint chance constraints for nonlinear systems. *International Journal of Robust and Nonlinear Control*, pages 1–21, 2017.

44. J. A. Paulson and A. Mesbah. Arbitrary polynomial chaos for quantification of general probabilistic uncertainties: Shaping closed-loop behavior of nonlinear systems. In *Proceedings of the 57th IEEE Conference on Decision and Control*, page Accepted, Miami, 2018.

45. J. A. Paulson and A. Mesbah. Nonlinear model predictive control with explicit backoffs for stochastic systems under arbitrary uncertainty. In *Proceedings of the 6th IFAC Conference on Nonlinear Model Predictive Control*, pages 622–633. Madison, WI, 2018.

46. L. Pronzato and E. Walter. Robust experiment design via stochastic approximation. *Mathematical Biosciences*, 75:103–120, 1985.

47. M. Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23:470–472, 1952.

48. E. G. Ryan, C. C. Drovandi, J. M. McGree, and A. N. Pettitt. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84:128–154, 2016.

49. K. J. Ryan. Estimating expected information gains for experimental designs with application to the random fatigue-limit model. *Journal of Computational and Graphical Statistics*, 12:585–603, 2003.

50. E. K. Ryu and S. P. Boyd. Extensions of Gauss quadrature via linear programming. *Foundations of Computational Mathematics*, 15:953–971, 2015.

51. A. Shapiro. Asymptotic analysis of stochastic programs. *Annals of Operations Research*, 30:169–186, 1991.

52. M. Sinsbeck and W. Nowak. An optimal sampling rule for nonintrusive polynomial chaos expansions of expensive models. *International Journal for Uncertainty Quantification*, 5:275–295, 2015.

53. C. Soize and R. Ghanem. Physical systems with random uncertainties: Chaos representations with arbitrary probability measure. *SIAM Journal on Scientific Computing*, 26:395–410, 2004.

54. J. C. Spall. An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL technical digest*, 19(4):482–492, 1998.

55. B. Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93:964–979, 2008.

56. D. Telen, F. Logist, E. Van Derlinden, I. Tack, and J. Van Impe. Optimal experiment design for dynamic bioprocesses: A multi-objective approach. *Chemical Engineering Science*, 78:82–97, 2012.

57. A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming Series A*, 18:25–57, 2005.

58. D. L. Wei, Z. S. Cui, and J. Chen. Uncertainty quantification using polynomial chaos expansion with points of monomial cubature rules. *Computers & Structures*, 86:2102–2108, 2008.

59. D. Xiu. Efficient collocational approach for parametric uncertainty analysis. *Communications in Computational Physics*, 2:293–309, 2007.

60. D. Xiu. Fast numerical methods for stochastic computations: A review. *Communications in Computational Physics*, 5:242–272, 2009.

61. D. Xiu and G. E. Karniadakis. The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM Journal of Scientific Computing*, 24:619–644, 2002.

Chapter 6

Conclusions and future work

As computing power and biological data become more available, future modeling and optimal experiment design approaches can be expected to capture more complexity of biological systems. We expect the field to move toward Bayesian estimation and experiment design approaches with no restrictive assumptions. We also expect the field to adapt and integrate model predictions in the real-time design and control of experiments, so that outcomes of an experiment can be modulated as it takes place. In this chapter, we summarize the main contributions and findings of this thesis. We also elaborate on our thoughts on future work and trends in the field of uncertainty quantification and optimal experiment design for biological systems.

6.1 Conclusions

In this thesis, we demonstrated the importance of probabilistic modeling and experiment design methods that can quantify and reduce uncertainty and/or heterogeneity in biological systems. We showed how the iterative process of building complex models can benefit from the development of uncertainty quantification (UQ) and optimal experiment design (OED) tools that can maximize the information content of experimental data in a computationally efficient manner. Despite the increasing use of UQ and OED methods in complex engineering systems, their use in biological systems is lagging. In this thesis, we have introduced approaches that can circumvent the challenges of implementation of UQ and OED for large nonlinear biological systems with probabilistic behavior arising from uncertainty and/or heterogeneity. We expect the contributions of this thesis to create new opportunities for seamless implementation of UQ and OED for systematic and hypothesis-based modeling in biological systems.

A key aspect in the process of building models is learning about the system from experimental data, that is, the inverse UQ problem for parameter estimation and hypothesis testing. Estimating model parameters from experimental data can be very computationally intensive, particularly for large biological models. Chapter 2 introduces a novel surrogate modeling method that significantly cuts down the cost of simulation of computationally expensive biological models, such as dynamic Flux Balance Analysis (DFBA), whilst remaining faithful to the true system behavior. We demonstrated that surrogate models can radically decrease the computational cost of parameter estimation approaches. In particular, the presented surrogate modeling method can be used to perform Bayesian estimation of genome-scale model parameters using, for example, sequential Monte Carlo and Markov Chain Monte Carlo methods.

The ability to discriminate between predictions of competing model structures is key for determining the model structure that most correctly describes the underlying physical properties of a biological system. In Chapter 3, we introduce an OED method, based on polynomial chaos theory, that is able to discern between two competing model structures by maximizing the separation between the full probably distributions of competing model outcomes. Conversely, Chapter 4 introduces two additional methods for OED for model structure selection. These methods are amenable to real-time implementation (i.e., during operation or experimental execution). This is achieved by enforcing separation between the statistics or properties of the outcome distribution (i.e., mean, variance), as opposed to the full description of the full probability distribution. By not capturing the full distribution, the computational cost is significantly reduced so that the approach can be applied in real-time. Finally, Chapter 4 introduces a Bayesian method that allows for online estimation of the probability of each competing model being true based on system observations. By tracking how well each model describes newly obtained system observations, the method can be used to perform online model separation.

The model selection methods presented in Chapters 3 and 4 offer a trade-off between modeling accuracy and computational speed, which provides flexibility of choice to accommodate a variety of scenarios and needs. There is value in building the entire range of outcomes of a given model in order to be as accurate as possible, if the experiment can be designed in a fully offline manner. Oftentimes, however, it is convenient to identify the system status quickly, for instance, in industrial operation. In an analogous way, the employed methods will be subject to the types and frequency of system observations that are available. Overall, the choice of a specific framework will depend on the characteristics of the system at hand and the designed experiment. The availability of methods tailored to each experimental scenario is imperative for the wide adoption of OED methods in practical settings.

In Chapter 5, we introduce a novel Bayesian OED method for parameter estimation. This method intends to maximize the information content of experimental data for parameter estimation purposes. Bayesian approaches make no assumptions about the system and model parameter priors, but can be very computationally intensive, often intractable for nonlinear complex systems. Upon application of a few restrictive assumptions and approximations, Bayesian OED approaches can be reduced to classical OED approaches, which maximize a scalar metric of the Fisher Information Matrix. Implementation of classical OED for biological systems can be challenging due to the fact that their dynamics can be highly nonlinear and their probabilistic behavior can deviate from normality. To address this, Chapter 5 introduces a novel Bayesian OED approach for nonlinear systems subject to experimental constraints. The Bayesian approach defines a utility function that maximizes the change of the posterior distribution versus the prior knowledge of the parameters. In our method, we proposed an arbitrary polynomial chaos approach to model the utility function, which was defined locally with respect to the design space. We expect this approach to be a new paradigm in the field of OED for parameter estimation.

6.2 Future work

The complexity inherent to the dynamics of biological systems has generally hindered the implementation of UQ and OED tools. Such complexity is compounded by probabilistic non-linear behavior, non-observability, and large system sizes with thousands of constituents. In this thesis, we have offered multiple methods that circumvent some of these issues to enable novel implementations of UQ and OED to biological systems. However, many challenges remain to be overcome until computationally efficient solutions that capture sufficient complexity in biological systems are widely available.

In this thesis, we presented a method for Bayesian learning of genome-scale models of biological systems. We believe that an important area of future research will be the development of efficient procedures to perform OED in genome-scale DFBA models. Designing biological experiments using DFBA models that account for uncertainty and/or heterogeneity in a classical or Bayesian sense can lead to significant learning from the system. Key computational challenges include i) the presence of discontinuities and singularities, and ii) scalability of UQ/OED methods with respect to number of sources of probabilistic behavior in the system.

Throughout this thesis, we have outlined key distinctions between classical and OED approaches. It is important to acknowledge that classical OED approaches can be successful when the system is observable and the system's probabilistic behavior is nearly gaussian. This is preferable from an implementation standpoint because classical OED approaches are straightforward and computationally efficient as compared to Bayesian approaches. However, biological systems exhibit characteristics that suggest classical approaches can be suboptimal, such as departure from normality in probabilistic behavior, weakly or poorly observable problems, large parameter sets that are possibly correlated, and large system dimensionality due to numerous system constituents. As a result, we believe that the field of biological system modeling will continue to benefit from progress in Bayesian OED, which will allow designing experiments whilst accounting for all prior system information.

Important future work on Bayesian OED will revolve around the generalization of the method for model discrimination, as well as handling of more complex systems (e.g., mixed integer systems). The presented Bayesian approach, which is the first-of-its-kind for nonlinear systems, heavily relies on the choice of surrogate model for the computation of expected utility. More work remains to be done in the surrogate modeling front, as it poses significant challenges when dealing with large numbers of parameters and constituents. In addition, we believe that it would be beneficial to explore methods like distributed numerical solution and optimization algorithms, which will allow to parallelize sampling of the system to reduce the overall computational cost of the method.

A potentially transformative application of the proposed Bayesian OED method is its implementation in a closed-loop setting. The proposed Bayesian approach was conceived to be applied as an open-loop (or batch) OED, where the experiment can be fully designed before the collection of any data. A promising area of research is the development of closed-loop (or sequential) OED, where previous experimental data can be used to guide future experiments. However, significant computational challenges are to be overcome for this to be solvable in a practical and systematic manner.

Throughout this thesis, we have presented extensive validation of our methods with respect to historical data and state-of-the-art approaches. We believe that the next step is the integration of these computational contributions within established experimental workflows, so that experimental decisions or designs can be directly influenced by the presented modeling, UQ, and OED methods. To attain this, the employed methods should be made sufficiently flexible and adaptable so that they can be seamlessly implemented into experimentalists protocols. A step towards that goal would be the development of ready-to-use software packages.

There are remaining questions to be resolved if the presented frameworks are to be executed in fully experimental settings. Despite this, the last few years have seen immense progress when it comes to implementation of modeling, UQ, and OED frameworks in biological systems. We strongly believe that more research will lead to seamless integration of computational and experimental workflows to guarantee high-quality models and optimal experiment designs.