

ABSTRACT

WENTWORTH, MAMI TONOE. Verification Techniques for Parameter Selection and Bayesian Model Calibration Presented for an HIV Model. (Under the direction of Ralph Smith.)

Uncertainty quantification plays an important role when making predictive estimates of model responses. In this context, uncertainty quantification is defined as quantifying and reducing uncertainties, and the objective is to quantify uncertainties in parameter, model and measurements, and propagate the uncertainties through the model, so that one can make a predictive estimate with quantified uncertainties. Two of the aspects of uncertainty quantification that must be performed prior to propagating uncertainties are model calibration and parameter selection. There are several efficient techniques for these processes; however, the accuracy of these methods are often not verified. This is the motivation for our work, and in this dissertation, we present and illustrate verification frameworks for model calibration and parameter selection in the context of biological and physical models.

First, HIV models, developed and improved by [2, 3, 8], describe the viral infection dynamics of an HIV disease. These are also used to make predictive estimates of viral loads and T-cell counts and to construct an optimal control for drug therapy. Estimating input parameters is an essential step prior to uncertainty quantification. However, not all the parameters are identifiable, implying that they cannot be uniquely determined by the observations. These unidentifiable parameters can be partially removed by performing parameter selection, a process in which parameters that have minimal impacts on the model response are determined. We provide verification techniques for Bayesian model calibration and parameter selection for an HIV model.

As an example of a physical model, we employ a heat model with experimental measurements presented in [10]. A steady-state heat model represents a prototypical behavior for heat conduction and diffusion process involved in a thermal-hydraulic model, which is a part of nuclear reactor models. We employ this simple heat model to illustrate verification techniques for model calibration.

For Bayesian model calibration, we employ adaptive Metropolis algorithms to construct densities for input parameters in the heat model and the HIV model. To quantify the uncertainty in the parameters, we employ two MCMC algorithms: Delayed Rejection Adaptive Metropolis (DRAM) [33] and Differential Evolution Adaptive Metropolis (DREAM) [66, 68]. The densities obtained using these methods are compared to those obtained through the direct numerical evaluation of the Bayes' formula. We also combine uncertainties in input parameters and measurement errors to construct predictive estimates for a model response. A significant emphasis is on the development and illustration of techniques to verify the accuracy of sampling-based Metropolis algorithms. We verify the accuracy of DRAM and DREAM by comparing chains,

densities and correlations obtained using DRAM, DREAM and the direct evaluation of Bayes formula. We also perform similar analysis for credible and prediction intervals for responses. Once the parameters are estimated, we employ energy statistics test [63, 64] to compare the densities obtained by different methods for the HIV model. The energy statistics are used to test the equality of distributions.

We also consider parameter selection and verification techniques for models having one or more parameters that are noninfluential in the sense that they minimally impact model outputs. We illustrate these techniques for a dynamic HIV model but note that the parameter selection and verification framework is applicable to a wide range of biological and physical models. To accommodate the nonlinear input to output relations, which are typical for such models, we focus on global sensitivity analysis techniques, including those based on partial correlations, Sobol indices based on second-order model representations, and Morris indices, as well as a parameter selection technique based on standard errors. A significant objective is to provide verification strategies to assess the accuracy of those techniques, which we illustrate in the context of the HIV model.

Finally, we examine active subspace methods as an alternative to parameter subset selection techniques. The objective of active subspace methods is to determine the subspace of inputs that most strongly affect the model response, and to reduce the dimension of the input space. The major difference between active subspace methods and parameter selection techniques is that parameter selection identifies influential parameters whereas subspace selection identifies a linear combination of parameters that impacts the model responses significantly. We employ active subspace methods discussed in [22] for the HIV model and present a verification that the active subspace successfully reduces the input dimensions.

© Copyright 2015 by Mami Tonoe Wentworth

All Rights Reserved

Verification Techniques for Parameter Selection and Bayesian Model Calibration
Presented for an HIV Model

by
Mami Tonoe Wentworth

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Applied Mathematics

Raleigh, North Carolina

2015

APPROVED BY:

Mansoor Haider

Pierre Gremaud

Brian Reich

Ralph Smith
Chair of Advisory Committee

DEDICATION

To my husband, Thomas.

BIOGRAPHY

Mami Tonoe Wentworth was born in Saitama, Japan, in 1985. She grew up in a small town called Takanezawa, which is located about 3 hours north of Tokyo. Mami lived in Japan until she was 12, then moved to Brazil with her family. After attending a Japanese middle school for a few years, she entered an American high school in São Paulo at the age of 15. Learning in a foreign language was difficult (especially History!), but as a consequence, Mami's love for Mathematics flourished.

In 2004, she started attending Rensselaer Polytechnic Institute to pursue B.S. in Mathematics and Physics. This is where she met her husband, Thomas, who is also a mathematician.

Mami and Thomas attended North Carolina State University together and got M.S. in Applied Mathematics in 2011. They got married in 2012, and they had their first son, Koki, in 2013. Mami will graduate in 2015 with a Ph.D. in Applied Mathematics and plans to pursue her passion for teaching Mathematics after graduating.

ACKNOWLEDGEMENTS

First and foremost, I would like to extend my sincerest gratitude to my advisor, Dr. Ralph Smith, for his constant support and encouragement. His dedication to my professional and personal success has helped me become a better researcher and scholar, which I can say *with confidence*. I would also like to thank Dr. H.T. Banks for providing the HIV models, data and input that are essential parts of my research, and my committee members, Dr. Mansoor Haider, Dr. Pierre Gremaud and Dr. Brian Reich, for their time, comments and questions. A special thanks to Ms. Denise Seabrooks for all the help she provided throughout my career as a graduate student at NCSU.

I would like to thank my parents, Emiko and Nobuyuki Tonoe, for giving me the precious opportunity to see the other side of the world and for supporting my decision to pursue my passion. I would also like to thank my parents-in-law, Diana and Bruce Wentworth, for all of the help and support they provided while writing this thesis.

Finally, I would like to thank my beloved husband, Thomas. Thank you for believing in me, for staying by my side both good and difficult times, and for making me a better person each and every day. You are the best husband, friend and partner in life one could ever ask for.

TABLE OF CONTENTS

LIST OF TABLES	vii
LIST OF FIGURES	ix
Chapter 1 Introduction	1
1.1 Models	3
1.1.1 HIV Model	3
1.1.2 Heat Equation	5
1.2 Organization	9
Chapter 2 Bayesian Model Calibration	11
2.1 Introduction	11
2.2 Bayesian Inference	13
2.2.1 Bayesian Estimation Formula	14
2.3 Posterior Distributions	15
2.3.1 Direct Method	15
2.3.2 Metropolis Algorithm	16
2.3.3 Delayed Rejection Adaptive Metropolis (DRAM)	17
2.3.4 Differential Evolution Adaptive Metropolis (DREAM)	21
2.4 Predictive Estimates	23
2.4.1 Confidence Intervals and Prediction Intervals	24
2.4.2 DRAM	26
2.4.3 DREAM	26
2.5 Results and Verification	27
2.5.1 Heat Example	27
2.5.2 HIV Example	29
2.5.3 Three Parameter Case	33
2.5.4 12 Parameter Case	37
2.6 Conclusion	43
Chapter 3 Energy Statistics	48
3.1 Introduction	48
3.1.1 Energy Distance	48
3.2 Samples	49
3.2.1 Inverse Transform Sampling	50
3.2.2 Example	50
3.3 Equality of Distributions	50
3.3.1 Three Parameter Case	50
3.3.2 12 Parameter Case	54
3.4 Conclusion	58
Chapter 4 Parameter Selection	59
4.1 Introduction	59

4.1.1	HIV Model, Inputs and Responses	61
4.1.2	Previous Work and Chapter Organization	62
4.2	Global Sensitivity Motivation	63
4.3	Parameter Selection Methods	64
4.3.1	Partial Correlation	65
4.3.2	Sobol Indices	66
4.3.3	Morris Screening	74
4.3.4	Parameter Subset Selection	76
4.4	Comparison and Verification	82
4.4.1	Verification of Input Rankings	82
4.4.2	Verification of Noninfluential Inputs	89
4.5	Conclusion	92
Chapter 5	Active Subspace Methods	94
5.1	Introduction	94
5.1.1	Active Subspace Methods	95
5.1.2	Chapter Organization	95
5.2	Computing the Gradient	97
5.2.1	Finite Difference	98
5.2.2	Sensitivity Equations	98
5.2.3	Automatic Differentiation	100
5.2.4	Comparison of Gradient	100
5.3	Active Subspace Methods	101
5.3.1	Active Subspace Using SVD	101
5.3.2	Response Surface Representation	103
5.4	Example and Verification	104
5.4.1	Sufficient Summary Plot	104
5.4.2	Verification	107
5.5	Conclusion	109
BIBLIOGRAPHY	111

LIST OF TABLES

Table 1.1	Description of parameters in the model (1.1).	6
Table 1.2	Nominal values of parameters and initial conditions from [6].	7
Table 1.3	Temperature data from [10].	8
Table 1.4	Parameter values specified for the heat problem.	10
Table 2.1	Parameters in the HIV model (2.1).	12
Table 2.2	Parameter values with difference simulation numbers in DRAM.	29
Table 2.3	Parameter values with difference simulation numbers in DREAM.	31
Table 2.4	Fixed values for parameters.	32
Table 2.5	Mean and standard deviation of b_E , δ and d_1 using $n = 61, 81, 101$ in direct evaluation of Bayes' relation.	32
Table 2.6	Mean and standard deviation for b_E , δ and d_1 .	34
Table 2.7	Convergence diagnostics for DRAM.	37
Table 2.8	The lower and upper bounds for the prior distributions $q_i \sim U(a_i, b_i)$ for $i = 1, \dots, 12$.	43
Table 2.9	Mean values obtained using DRAM and DREAM for 12 parameters.	44
Table 2.10	Convergence diagnostics for DRAM with 12 parameters.	44
Table 3.1	Energy test and critical values for $\alpha = 0.1, 0.05$ for DRAM (above) and DREAM (below).	52
Table 3.2	Percentage of rejection obtained by repetition of energy test for DRAM and DREAM.	53
Table 3.3	Test and critical values for $\alpha = 0.1, 0.05$ using samples from chain and kde. R indicates the case in which the null hypothesis is rejected.	57
Table 4.1	The lower and upper bounds of parameters, $q_i \in [\ell b_i, ub_i]$, for $i = 1, \dots, 15$.	62
Table 4.2	Parameter subset selection results from [6].	81
Table 4.3	Summary of parameter selection techniques.	83
Table 4.4	Sensitivity measures provided by Partial Correlation, Sobol by the Saltelli algorithm and Morris Screening.	84
Table 4.5	Subsets of influential parameters for $n_p = 1, \dots, 4$. Here, P_{corr} , S, M and P_{ss} denote Partial Correlations, Sobol indices, Morris indices and Parameter Subset Selection.	85
Table 4.6	Subsets of influential parameters for $n_p = 5, \dots, 8$.	86
Table 4.7	Subsets of influential parameters for $n_p = 9, \dots, 12$.	87
Table 4.8	Subsets of influential parameters for $n_p = 13, \dots, 14$.	88
Table 4.9	Sensitivity measures for $y = q_1 - q_2$.	91
Table 5.1	Nominal values of q^0 , at which the gradient is computed.	97
Table 5.2	The lower and upper bounds for the distributions $q_i \sim U(a_i, b_i)$ for $i = 1, \dots, 15$.	97
Table 5.3	Gradient via finite difference using different step sizes.	98
Table 5.4	Derivatives at q^0 obtained via finite differences (FD), sensitivity equations (SE) and automatic differentiation (AD).	101

Table 5.5 Eigenvalues computed using 15 samples. 105

LIST OF FIGURES

Figure 1.1	Components of uncertainty propagation process.	2
Figure 1.2	The compartments of HIV model (1.1).	4
Figure 1.3	Plot of temperature data from [10].	9
Figure 2.1	Densities from DRAM, DREAM and direct methods for (a) Φ and (b) h	28
Figure 2.2	Contour plot from the direct method and the correlation plots from (a) DRAM and (b) DREAM.	28
Figure 2.3	Distributions of different simulation numbers for (a) DRAM and (b) DREAM.	30
Figure 2.4	Chain evolutions of b_E , δ and d_1 for DRAM.	33
Figure 2.5	Initial chain evolutions of b_E , δ and d_1 for DREAM.	34
Figure 2.6	Final chain evolutions of b_E , δ and d_1 for DREAM.	35
Figure 2.7	Marginal densities for b_E , δ and d_1 with different methods.	36
Figure 2.8	Contour plots from the direct method plotted with joint sample points from DRAM and DREAM showing the correlations between b_E , δ and d_1	36
Figure 2.9	Credible and prediction intervals for response E using (a) Direct, (b) DRAM and (c) DREAM.	37
Figure 2.10	(a) R-Statistics and (b) autocorrelation for b_E , δ and d_1 in DREAM.	38
Figure 2.11	Credible and prediction intervals for E using (a) DRAM and (b) DREAM when σ is updated.	39
Figure 2.12	Densities for 12 parameters with DRAM (solid line) and DREAM (dashed line).	40
Figure 2.13	Joint sample points obtained using DRAM.	41
Figure 2.14	Joint sample points obtained using DREAM.	42
Figure 2.15	R-statistics for 12 parameters in DREAM.	45
Figure 2.16	Autocorrelation for 12 parameters in DREAM.	46
Figure 3.1	Construction of random samples using discrete probability density functions.	51
Figure 3.2	Replicates of energy test and critical values between Direct and DRAM.	52
Figure 3.3	Replicates of energy test and critical values between Direct and DREAM.	53
Figure 3.4	Energy test for 1000 samples from DRAM and DREAM with critical values computed with $M = 199$	55
Figure 3.5	Energy test for 1000 samples from DRAM and DREAM with critical values computed with $M = 199$	56
Figure 4.1	Illustration of $y = f(q)$ for (a) identifiable, (b) unidentifiable and (c) noninfluential parameters q	60
Figure 4.2	Illustration of influential parameters where q_1 is more influential than q_2	61
Figure 4.3	Scatterplots of 1000 joint realizations of y versus (a) q_1 and (b) q_2	64
Figure 4.4	Partial correlation of the scalar response to the input parameters.	66
Figure 4.5	Response for fixed values of (a) q_1 and (b) q_2 illustrating $\mathbb{E}(Y q_i)$ and $\text{var}[\mathbb{E}(Y q_i)]$	69
Figure 4.6	Sobol indices S_{T_i} for 15 parameters.	72
Figure 4.7	Model response (solid blue), 1st order (dash-dot black) and 2nd order approximation (dash red).	73

Figure 4.8	Morris μ^* and σ computed using the Morris screening algorithm.	77
Figure 4.9	Densities obtained by fixing (a) 7, (b) 5, (c) 3 and (d) 1 least influential parameters. The density in solid is obtained by varying all the parameters.	90
Figure 5.1	Illustration of $f(q) = q_1 - q_2$ with $q_1, q_2 \sim U(0, 1)$	96
Figure 5.2	Gradient with respect to λ_T using different step sizes.	99
Figure 5.3	Log eigenvalues computed using the SVD-based active subspace method.	106
Figure 5.4	Comparison of (a) weight on parameters determined by global linear model and (b) influential parameters identified by Partial Correlation and (c) Sobol decomposition.	108
Figure 5.5	Sufficient summary plot produced using (a) the weights and (b) the first eigenvector.	109
Figure 5.6	Probability density functions of model responses using (a) 2, 5, 8, (b) 9, 10, 11 and (c) 12, 13, 14, along with probability density functions using the full dimensional input space and the active subspace based on weights.	110

Chapter 1

Introduction

Uncertainty Quantification can be defined as the process of (i) quantifying uncertainties in inputs – parameters, initial and boundary conditions; (ii) propagating the input uncertainties through the model; and (iii) making predictive estimates with associated uncertainties. As shown in Figure 1.1, model calibration and parameter selection are two of the key components of uncertainty propagation. Parameter selection allows us to identify the parameters that minimally impact the model responses, whereas model calibration allows us to quantify parameter values and their uncertainties.

In many biological and physical models, the dimensionality of input space can be as large as $\mathcal{O}(10^5)$. Some parameters are non-physical, and others can not be measured directly in experiments. Hence, these parameters must be inferred or estimated using data and their uncertainties must be quantified before one can propagate the uncertainties through the model. In these situations, estimating parameters using Bayesian inference allows us to subsequently sample from parameter densities and directly propagate the parameter uncertainties through the model. An example of techniques for efficiently estimating parameters is sampling-based methods. In particular, Delayed Rejection Adaptive Metropolis (DRAM) [33] and Differential Evolution Adaptive Metropolis (DREAM) [66, 68] are two of the methods we employ to avoid high-dimensional integration in the Bayesian model calibration formula.

At the same time, some or many of these input parameters can not be uniquely identified using the data. This necessitates the use of parameter selection techniques, in which parameters that have greater impacts on the model response are identified. The parameters whose impact on the response is minimal can be fixed at nominal values, hence reducing the number of parameters to be estimated for model calibration. For parameter selection, we study Partial Correlation, Sobol indices and Morris indices, which are based on global-sensitivity analysis, as well as a parameter subset selection algorithm based on standard errors.

The use of these techniques facilitates the model calibration and parameter selection processes;

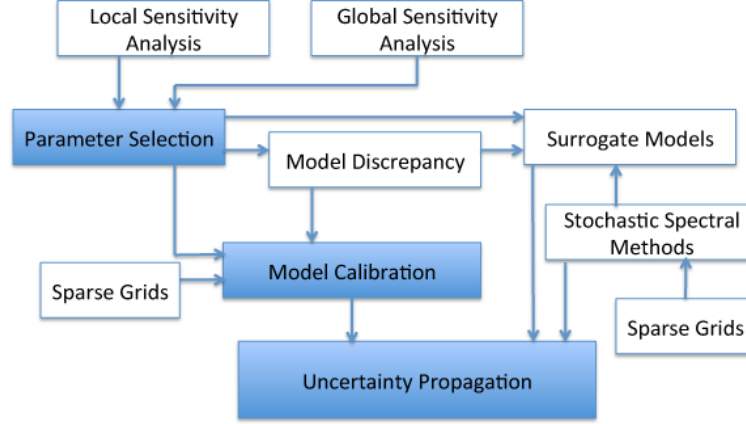


Figure 1.1 Components of uncertainty propagation process.

however, the accuracy of these methods is often not verified. This motivates us to present and demonstrate verification frameworks for model calibration and parameter selection.

For model calibration, one must carefully determine whether the chain values have converged. Even though the sampling-based methods are based on the construction of a Markov chain whose stationary distribution is the posterior, the adaptation used to improve the efficiency causes DRAM and DREAM to be no longer a Markovian process. Convergence is still guaranteed but it is difficult to establish when chains have converged. Even if a chain appears stationary for a long time based on a visual inspection, the parameter values may still jump to a region of another local minimum. For this reason, it is essential to perform a verification procedure to ensure that chains have converged and that the correct distributions are obtained for the posterior.

In parameter selection, parameters are ranked in the order of importance based on their sensitivity indices. A major issue here is that there is often no clear cut-off between important and unimportant parameters. Additionally, some parameter selection techniques are based on the linear approximation of functions and function gradients. If the function can not be accurately represented with a linear approximation, then the parameter selection techniques will fail to identify the important parameters correctly. One must perform a verification to determine that the parameters that are identified as unimportant do in fact impact the response minimally.

In this dissertation, we discuss model calibration and parameter selection techniques presented

for a biological model and a physical model. The significant emphasis is on the verification frameworks for these techniques. We present the application and verification of model calibration and parameter selection techniques in the context of an HIV model and a heat model; however, these techniques are applicable to a wide range of physical and biological models.

1.1 Models

1.1.1 HIV Model

The modeling of the HIV disease is essential for understanding the disease and developing optimal treatment regimes. A mathematical model has been developed in [8] to quantify the interactions among infected and uninfected T-cells, infectious and noninfectious viral loads and effector cells. The model's predictive capability has been verified using data from patients who underwent a clinical study involving structured treatment interruptions.

The system of ODEs modeling the HIV disease developed in [8] is given by

$$\begin{aligned}
\dot{T}_1 &= -d_1 T_1 - (1 - \xi_1(t)) k_1 V_I T_1 - \gamma_T T_1 + p_T \left(\frac{a_T V_I}{V_I + K_V} + a_A \right) T_2 \\
\dot{T}_1^* &= (1 - \xi_1(t)) k_1 V_I T_1 - \delta T_1^* - m E_1 T_1^* - \gamma_T T_1^* + p_T \left(\frac{a_T V_I}{V_I + K_V} + a_A \right) T_2^* \\
\dot{T}_2 &= \lambda_T \frac{K_s}{V_I + K_s} - \gamma_T T_1 - d_2 T_2 - (1 - f \xi_1(t)) k_2 V_I T_2 - \left(\frac{a_T V_I}{V_I + K_V} + a_A \right) T_2 \\
\dot{T}_2^* &= \gamma_T T_1^* + (1 - f \xi_1(t)) k_2 V_I T_2 - d_2 T_2^* - \left(\frac{a_T V_I}{V_I + K_V} + a_A \right) T_2^* \\
\dot{V}_I &= (1 - \xi_2(t)) 10^3 N_T \delta T_1^* - c V_I - 10^3 [(1 - \xi_1(t)) \rho_1 k_1 T_1 + (1 - f \xi_1(t)) \rho_2 k_2 T_2] V_I \\
\dot{V}_{NI} &= \xi_2(t) 10^3 N_T \delta T_1^* - c V_{NI} \\
\dot{E}_1 &= \lambda_E + \frac{b_{E1} T_1^*}{T_1^* + K_{b1}} E_1 - \frac{d_E T_1^*}{T_1^* + K_d} E_1 - \delta_{E1} E_1 - \gamma_E \frac{T_1 + T_1^*}{T_1 + T_1^* + K_\gamma} E_1 + \frac{p_E a_E V_I}{V_I + K_V} E_2 \\
\dot{E}_2 &= \gamma_E \frac{T_1 + T_1^*}{T_1 + T_1^* + K_\gamma} E_1 + \frac{b_{E2} K_{b2}}{E_2 + K_{b2}} E_2 - \delta_{E2} E_2 - \frac{a_E V_I}{V_I + K_V} E_2
\end{aligned} \tag{1.1}$$

with initial conditions $[T_1(0), T_1^*(0), T_2(0), T_2^*(0), V_I(0), V_{NI}(0), E_1(0), E_2(0)]$. Here, T_1 and T_1^* respectively denote uninfected and infected activated (antigen-specific) CD4+ T-cells. Uninfected resting, non-activated, CD4+ T-cells are denoted by T_2 and infected resting CD4+ T-cells are denoted by T_2^* . Infectious free virus is denoted by V_I ; this is the virus that is capable of infecting other cells in the plasma. On the other hand, V_{NI} denotes non-infectious free virus, which is yielded inactive by protease inhibitors. HIV-specific effector CD8+ T-cells are denoted by E_1 and HIV-specific memory CD8+ T-cells are denoted by E_2 . The compartments of the model are

depicted in Figure 1.2.

Several terms in the model (1.1) are based on the law of mass action, so that the rate of change in population size is proportional to the population size. The terms $-d_1 T_1$ and $\gamma_T T_1$ in \dot{T}_1 are examples of mass action terms. Other terms are based on Michaelis-Menten kinetics, in which the rate saturates a maximum. An example of this type is $\frac{a_T V_I}{V_I + K_V}$, which is the activation of infected HIV specific resting CD4+ T-cells with a_T being the maximum activation rate. The term $\frac{\lambda_T K_S}{V_I + K_S}$ in the differential equation for \dot{T}_2^* accounts for the source rate of naive CD4+ T-cells. In the equation for \dot{E}_1 , $\frac{b_{E1} T_1^*}{T_1^* + K_{b1}} E_1$ and $-\frac{d_E T_1^*}{T_1^* + K_d} E_1$ respectively represent dynamic effect that activated, infected CD4+ T-cells have on the effector CD8+ T-cells when $K_{b1} < K_d$ and $b_{E1} < b_E$. Also in this differential equation, $\frac{p_E a_E V_I}{V_I + K_V} E_2$ represents the activation of memory CD8+ T-cells into effector CD8+ T-cells. In the differential equation for \dot{E}_2 , $\frac{E_1 (T_1 + T_1^*)}{T_1 + T_1^* + K_\gamma}$ has an essential role that activated CD4+ T-cells play in the generation of memory CD8+ T-cells, whereas $\frac{b_{E2} K_{b2} E_2}{E_2 + K_{b2}}$ and $\delta_{E2} E_2$ are homeostatic regulation terms in E_2 .

The parameters in the HIV model (1.1) are described in Table 1.1 and nominal values for parameters and initial conditions reported in [6] are compiled in Table 1.2. The functions $\xi_1 = \epsilon_1 u(t)$ and $\xi_2 = \epsilon_2 u(t)$ represent the impact of the treatment. Here, ϵ_1 is the effectiveness of the reverse transcriptase inhibitor (RTI), whereas ϵ_2 is the effectiveness of the protease inhibitor

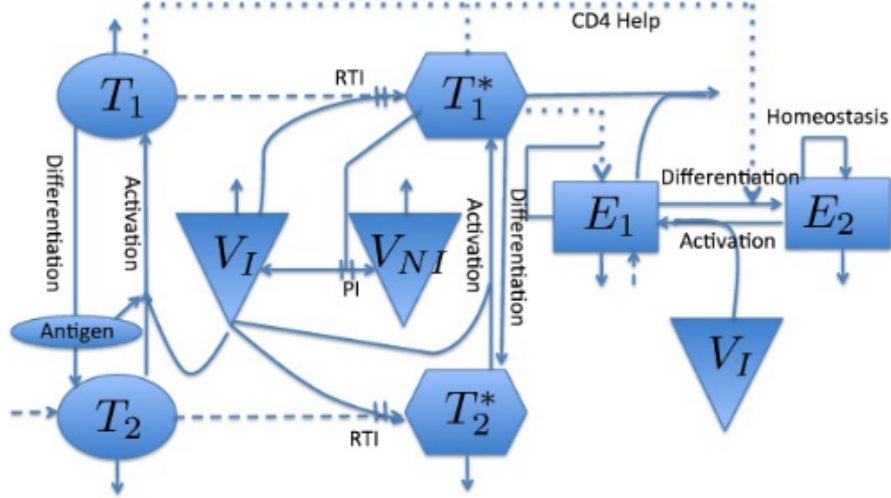


Figure 1.2 The compartments of HIV model (1.1).

(PI). Also, $u(t)$ is the HAART drug level, where $u(t) = 1$ when the patient is on treatment, and $u(t) = 0$ when the patient is off treatment. Parameters such as ϵ_1 and ϵ_2 , along with many others, can not be directly measured and hence must be estimated through a fit to data.

Among these inputs, however, some parameters do not influence model output. These parameters are identified via parameter selection prior to parameter estimation. Isolating these noninfluential parameters allows us to reduce the parameter dimensions for model calibration. In some cases, parameter selection also helps eliminate the parameter identifiability issues during model calibration since noninfluential parameters may not be uniquely estimated.

Quantifying and reducing uncertainties in the modeling of HIV disease are essential in understanding the disease better, reducing the spread and developing treatment regimes. The sources of uncertainty involved in the HIV model include parameter values that are not directly measurable as well as observed data which are either noisy or censored. We perform model calibration and parameter selection for the HIV model (1.1). Additionally, we present the illustration and verification of model calibration and parameter selection techniques in the context of the HIV model.

1.1.2 Heat Equation

Next, we discuss models for nuclear reactor design. Understanding and improvement of nuclear reactor design is dependent on developing models and simulation codes. A primary goal of Consortium for Advanced Simulation of Light water Reactors (CASL) funded by the Department of Energy is to develop simulation-based design tools and predict the interactions between the nuclear fuel, neutron transport, heat transfer and thermal-hydraulic components of a light water reactor.

A neutron transport equation, for example, models a component of a reactor core. It is essential to quantify the neutron distributions in the reactor core since neutron densities and energy levels govern various nuclear reactions. Readers are referred to [14, 26] for details on neutron transport equations. The transport equation provides a highly accurate description of neutron distributions in a reactor if a correct macroscopic cross-section information is employed. There have been developed libraries of cross-section distributions for various stable and radioactive nuclei. Since the quantification of neutron distribution is critical to nuclear reactors, numerous commercial, government laboratory and proprietary neutron transport codes have been developed.

Another component of nuclear reactor models is thermal-hydraulic models. The goal of thermal-hydraulic models is to characterize the primary coolant behavior in a nuclear reactor. However, this is a highly complex process since it includes integration of high pressure, two-phase flow dynamics, heat conduction, heat transfer and neutron interactions in complex geometries.

Table 1.1 Description of parameters in the model (1.1).

Parameter	Explanation
δ	Viral produced lysis rate of T_1^*
d_2	T_2 and T_2^* natural death rate
δ_{E2}	Death rate of E_2
m	Rate of removal by cell lysis of T_1^* from the system by E_1
γ_T	Rate at which T_1 and T_1^* differentiate into T_2 and T_2^* , respectively
c	Natural clearance rate of V_I and V_{NI}
δ_{E1}	Constant death rate of E_1
γ_E	Source term for E_1
k_2	Production rate of T_2^* due to encounters between T_2 and V_I that is less than k_1
ρ_1	Rate of removal of V_I through successful infection of T_1
ρ_2	Rate of removal of V_I through successful infection of T_2
d_1	Natural death rate of T_1
ϵ_2	Relative effectiveness of protease inhibitor (PI)
a_A	Activation rate of T_2 and T_2^* by non-HIV antigen
ϵ_1	Relative effectiveness of reverse transcriptase inhibitor (RTI)
p_T	Net proliferation of T_1 and T_1^* due to clonal expansion and programmed contraction
p_E	Net proliferation of E_1 due to clonal expansion and programmed contraction
k_1	Production rate of T_1^* from encounters between T_1 and V_I
N_T	Number of RNA copies produced during the process of T_1^* lysis
a_T	Maximum activation rate of T_2 and T_2^*
f	Efficacy of treatment $0 \leq f \leq 1$
λ_E	Constant differentiation of E_2 into E_1
K_V	Half-saturation constant of virus

Table 1.2 Nominal values of parameters and initial conditions from [6].

$\lambda_T = 3.2543$	$d_1 = 0.1317$	$\epsilon_1 = 0.5241$	$k_1 = 4.8200\text{e-}5$
$a_T = 2.3198\text{e-}4$	$\epsilon_2 = 0.7149$	$N_T = 79.26$	$b_{E2} = 0.34554$
$a_E = 1.5332\text{e-}2$	$p_E = 1.0294$	$a_A = 8.07\text{e-}5$	$p_T = 5.531$
$\gamma_T = 3.792\text{e-}4$	$d_2 = 3.096\text{e-}3$	$f = 0.5068$	$k_2 = 2.005\text{e-}9$
$\delta = 0.2095$	$m = 1.127\text{e-}3$	$c = 5.818$	$\lambda_E = 9.9930\text{e-}4$
$b_{E1} = 3.885\text{e-}2$	$K_{b1} = 2.488\text{e-}2$	$d_E = 6.278\text{e-}2$	$K_d = 0.12$
$\delta_{E1} = 5.967\text{e-}2$	$K_{b2} = 86.97$	$\gamma_E = 5.154\text{e-}4$	$K_\gamma = 1.357$
$K_V = 14.79$	$\delta_{E2} = 1.450\text{e-}3$	$K_s = 2.789\text{e+}4$	$T_1(0) = 12.135$
$T_1^*(0) = 5.8604\text{e-}4$	$T_2(0) = 823.59$	$T_2^*(0) = 7.521\text{e-}3$	$V_I(0) = 3.571\text{e+}3$
$V_{NI}(0) = 3.571\text{e+}3$	$E_1(0) = 6.821\text{e-}2$	$E_2(0) = 0.6909$	

The quantification of void fraction distributions and boiling transitions is essential for optimized performance and maintenance of safety margins. See [14] for detailed derivation and numerical analysis of thermal-hydraulic models.

The sources of uncertainties involved in nuclear reactor designs include input uncertainties, model errors, numerical errors and uncertainties in measurements. The highly complex, coupled systems mentioned above are difficult to resolve in small grid sizes, resulting in large numerical errors. Furthermore, the number of parameters in these models can be as high as $\mathcal{O}(10^5)$, several of which are nonphysical or cannot be directly measured, and must be inferred through model calibration. Finally, the numbers and types of measurements obtained for model calibration and experimental design are limited due to the harsh environment inside the reactors.

Here, we summarize a steady-state heat model for an uninsulated aluminum rod in open air with a heat source at one end and dissipation due to conduction and air-cooling along the length of the rod. The steady-state heat equation represents as a prototypical behavior of heat conduction and diffusion process involved in thermal-hydraulic models. Because of its simplicity and the availability of analytic solutions, we employ the heat equation for illustration and verification of model calibration techniques in Chapter 2.

Let $T_s(x)$ be the steady-state temperature at location x . As detailed in [10], the steady-state

heat equation is given by

$$\frac{d^2 T_s}{dx^2} = \frac{2(a+b)}{ab} \frac{h}{k} [T_s(x) - T_{amb}] \quad (1.2)$$

$$\frac{dT_s}{dx}(0) = \frac{\Phi}{k}, \quad \frac{dT_s}{dx}(L) = \frac{h}{k} [T_{amb} - T_s(L)]. \quad (1.3)$$

Here, T_{amb} denotes the ambient temperature whereas Φ denotes the source flux at $x = 0$. Also, h denotes the convective heat transfer coefficient. The parameters $q = [\Phi, h]$ are inferred using the data from [10], which is summarized in Table 1.3 and plotted in Figure 1.3. The descriptions and values of a, b, L and k are summarized in Table 1.4.

The analytic solution to (1.2) and (1.3) is

$$T_s(x, q) = c_1(q)e^{-\gamma x} + c_2(q)e^{\gamma x} + T_{amb} \quad (1.4)$$

where $\gamma = \sqrt{\frac{2(a+b)h}{abk}}$ and

$$c_1(q) = -\frac{\Phi}{k\gamma} \left[\frac{e^{\gamma L(h+k\gamma)}}{e^{-\gamma L(h-k\gamma)} + e^{\gamma L(h+k\gamma)}} \right], \quad c_2(q) = \frac{\Phi}{k\gamma} + c_1(q). \quad (1.5)$$

Table 1.3 Temperature data from [10].

Location	10	14	18	22	26
Temperature	96.1392	80.1221	67.6552	57.9609	50.9009
Location	30	34	38	42	46
Temperature	44.8437	39.7502	36.1598	33.3076	31.1506
Location	50	54	58	62	66
Temperature	29.2792	27.8847	27.1810	26.3956	25.8603

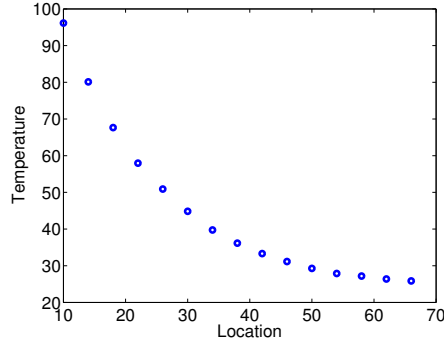


Figure 1.3 Plot of temperature data from [10].

1.2 Organization

In Chapter 2, we focus on the development and illustration of techniques to verify the accuracy of sampling-based Metropolis algorithms that are used to construct densities for input parameters. We first detail a Bayesian approach for model calibration. A Bayesian approach is more suited than a Frequentist approach since it allows us to treat inputs as random variables with associated densities. Because input parameters are assumed to be random variables, we are able to sample parameter values from densities and propagate associated uncertainties through the model. In this chapter, we discuss two Markov Chain Monte Carlo algorithms: DRAM and DREAM. We verify their accuracy by comparing chains, densities and correlations obtained using DRAM, DREAM and direct numerical evaluation of Bayes' formula. A similar analysis is performed for credible and prediction intervals for responses.

Once the parameter distributions are estimated, in Chapter 3, we apply energy statistic tests [63, 64] to verify the equality of distributions from DRAM and DREAM. In this chapter, we discuss energy distance, which is used to measure the distance between two sets of samples. We also discuss a method for constructing of random samples from a discrete probability density functions. We apply the test to the results from Chapter 2.

In Chapter 4, we illustrate the techniques for verifying the accuracy of parameter selection techniques based on global sensitivity analysis. Parameter selection is a process in which we isolate parameters that cannot be estimated using data because they have minimal impacts on

Table 1.4 Parameter values specified for the heat problem.

Parameter values	Description
$a = 0.95$	Cross-sectional dimension [cm]
$b = 0.95$	Cross-sectional dimension [cm]
$L = 70.0$	Length of the rod [cm]
$k = 2.37$	Thermal conductivity for aluminum [W/cm C]

the model responses. The results from Chapter 2 indicate that some parameters may not be uniquely estimated, and this is the motivation for performing parameter selection. A successful application of parameter selection will enable us to reduce the number of parameter dimensions during model calibration. In this chapter, we first examine four types of parameter selection techniques: Partial correlation [4], Sobol indices [55, 56, 59, 60, 61], Morris screening [18, 46, 55] and a parameter subset selection algorithm [6, 7, 9, 20]. We then apply the techniques on the HIV model using a set of 15 input parameters and initial conditions. We also present verification techniques to verify that the methods correctly identify parameters that minimally impact the model response.

Finally in Chapter 5, we discuss active subspace methods [5, 22, 52] as a part of dimension reduction technique. Active subspace methods allow us to reduce the dimensionality of the input space. We apply the technique on the HIV model (1.1), and present a verification framework to ensure that the model responses represented by reduced dimensional input subspace are accurate. Active subspace methods are similar to parameter selection in some ways. They both identify the direction in which the function varies the most and enable us to reduce the number of input parameters by isolating parameters that impact the response minimally. One key difference, however, is that active subspace methods determine a linear combination of inputs, which is the directions of input space in which the model responses exhibit the greatest change. A linear combination of parameters determined by active subspace methods makes the interpretation of results more complicated; however, active subspace methods can often reduce the dimensionality of input space further than parameter selection can do. When presenting the results of active subspace methods, we compare them to the parameter selection results from Chapter 4 to show similarities and differences.

Chapter 2

Bayesian Model Calibration

2.1 Introduction

Development of models for viral infection processes is critical to understand the disease spread of HIV, and quantified uncertainties are required when constructing optimized treatment regimes. To illustrate these issues, we employ a previous version of the HIV model (1.1)

$$\begin{aligned}\dot{T}_1 &= \lambda_1 - d_1 T_1 - (1 - \epsilon)k_1 V T_1 \\ \dot{T}_2 &= \lambda_2 - d_2 T_2 - (1 - f\epsilon)k_2 V T_2 \\ \dot{T}_1^* &= (1 - \epsilon)k_1 V T_1 - \delta T_1^* - m_1 E T_1^* \\ \dot{T}_2^* &= (1 - f\epsilon)k_2 V T_2 - \delta T_2^* - m_2 E T_2^* \\ \dot{V} &= N_T \delta (T_1^* + T_2^*) - cV - [(1 - \epsilon)\rho_1 k_1 T_1 + (1 - f\epsilon)\rho_2 k_2 T_2]V \\ \dot{E} &= \lambda_E + \frac{b_E(T_1^* + T_2^*)}{T_1^* + T_2^* + K_b} E - \frac{d_E(T_1^* + T_2^*)}{T_1^* + T_2^* + K_d} E - \delta_E E\end{aligned}\tag{2.1}$$

as discussed in [2, 3]. Here, T_1 and T_1^* represent the uninfected and infected type 1 target cells, respectively. Similarly, T_2 and T_2^* represent uninfected and infected type 2 target cells. Note that T_1 and T_2 cells could represent, for instance, CD4 T-lymphocytes and macrophages, respectively. Here V is the total number of infectious and non-infectious free viruses, and E represents the immune effector cells. Whereas this model is very simplified, it provides a framework to investigate control strategies for HIV [2]. Studying the HIV infection dynamics using mathematical models can contribute to understanding of fundamental features of infection and to improved disease monitoring and treatment. Though we perform our verification using model (2.1), there are more comprehensive models such as the one found in [8].

The description of the parameters in model (2.1) are summarized in Table 2.1. Examples of parameters that can be measured are δ , the death rate of infected cells, and c , the death rate

Table 2.1 Parameters in the HIV model (2.1).

λ_1	Target cell 1 production rate	ρ_1	Ave. virions infecting type 1 cells
λ_2	Target cell 2 production rate	ρ_2	Ave. virions infecting type 2 cells
d_1	Target cell 1 death rate	b_E	Max. birth rate immune effectors
d_2	Target cell 2 death rate	d_E	Max. death rate immune effectors
k_1	Population 1 infection rate	K_b	Birth constant, immune effectors
k_2	Population 2 infection rate	K_d	Death constant, immune effectors
c	Virus natural death rate	λ_E	Immune effector production rate
δ	Infected cell death rate	δ_E	Natural death rate, immune effectors
ε	Population 1 treatment efficacy	N_T	Virions produced per infected cell
m_1	Population 1 clearance rate	f	Treatment efficacy reduction
m_2	Population 2 clearance rate		

of free virus. The values of measurable parameters are found in [27, 45, 47, 49]. On the other hand, immune response parameters such as m_1 and m_2 are not well known, and unmeasurable parameters such as these must be estimated. The authors of [16] determine the parameter values involved in their model by applying physical constraints and using the expressions for the steady states.

In this chapter, we focus on a Bayesian approach for model calibration and the construction of predictive estimates. Bayesian techniques for parameter estimation are based on the assumption the parameters are random variables with associated densities. The goal of parameter estimation is to find the density that best reflects the distribution of parameters values so that the model response best describes the sampled observations. Bayesian inference allows the parameter densities to be updated as additional information is obtained. Additionally, parameter values can be sampled from constructed densities to propagate uncertainties through the model for uncertainty quantification.

In Section 2.2, we describe the process of model calibration using the Bayesian inference. Synthetic measurement data are used to construct densities for the parameters in the model. We then discuss different implementations that one can employ to compute the posterior distributions in Section 2.3.

Once parameter densities are constructed, the uncertainties in the input parameters and measurement errors are propagated through the model to construct uncertainty bounds. More specifically, credible and prediction intervals allow us to bound the model response distributions in terms of parameter uncertainties and the measurement noise. The process of constructing these estimates is described in Section 2.4.

Finally, in Section 2.5, we demonstrate the model calibration process using DRAM and DREAM. We first apply the model calibration techniques to the heat equation (1.2), which is a simplified application of nuclear reactor models. We then perform the model calibration on the HIV model (2.1). In both cases, we verify the accuracy of the two methods by constructing densities for three parameters and comparing the results to densities that are computed by solving the Bayes' formula directly. The accuracy of DRAM and DREAM is also verified by estimating the densities for a larger set of parameters, which vary greatly in their orders of magnitude.

We note that the verification of sampling-based Metropolis algorithms is critical and we provide a framework, which we illustrate in the context of the considered HIV model. We first verify our results by comparing DRAM, DREAM and direct numerical evaluation of Bayes' relation. Secondly, we verify by statistically validating our chain convergence. In combination, this provides a general verification framework that is applicable to a wide range of problems. Compared to an initial analysis of the model (2.1) using DRAM with six parameters presented in [59], our work extends that significantly through the development of this verification framework and estimation of a much larger set of unidentifiable parameters.

2.2 Bayesian Inference

The parameter values are estimated using Frequentist inference in [2], where the model (2.1) is developed. In this framework, the parameters are assumed to be unknown but fixed. The authors use an iterative generalized least squares method and obtain uncertainty in the estimation process. On the contrary, we perform Bayesian inference to estimate parameters in our study. In the Bayesian framework, the parameters are considered to be random variables with associated densities. This allows the densities to be updated as new information is available. Bayesian inference is also natural for uncertainty quantification since one can propagate the parameter densities through the model.

Bayesian techniques have been applied to several models arising in bioscience, including an HIV model [36]. In the field outside of HIV modeling, the importance of parameter estimation and uncertainty quantification of pharmacokinetics models is illustrated in [28]. Smoothing using Gaussian Processes on chemical reaction dynamics of producing nylon involving 6 to 24 parameters is discussed in [17], and the authors of [15] discuss accelerated sampling procedures

of Bayesian inference with applications in systems biology and nonlinear dynamic systems. Estimation of rate constants in a continuous-time Markov model of an ion channel model from a single data is discussed in [58]. Other applications of Bayesian inference include deformable neuronanatomical atlases [31] and several other models in Systems Biology Markup Language [12].

For applications such as the HIV model (2.1), where modeling and measurement errors are assumed to be unbiased and iid, we employ the statistical model

$$Y_i = f(t_i, Q) + \varepsilon_i, \quad i = 1, \dots, n \quad (2.2)$$

where Y_i, ε_i and Q are random variables representing measurements at time t_i , measurement errors and parameters, respectively. The parameters are defined as a random variable Q whose realizations are q with $\dim(Q) = p$. The response of model (2.1) at time t_i is defined to be $f(t_i, Q)$.

Let y be the realization of the random variable Y such that

$$y_i = f(t_i, q) + \epsilon_i \quad (2.3)$$

where f is the model response with realized parameters q evaluated at time $t = t_i$ for $i = 1, \dots, n$ and ϵ_i is a realization of the measurement error. We assume that models may have multiple responses. In a model with multiple responses, we let N_r be the number of responses. For example, $N_r = 6$ in model (2.1).

Given this statistical model in (2.2), it is our goal to find the posterior density that best reflects the distribution of parameter values based on the sampled observations.

2.2.1 Bayesian Estimation Formula

The Bayesian estimation formula is given by

$$\pi(q|y_{obs}) = \frac{\pi(y_{obs}|q)\pi_0(q)}{\int_{\mathbb{R}^p} \pi(y_{obs}|q)\pi_0(q)dq} \quad (2.4)$$

and it states that a posterior distribution, $\pi(q|y_{obs})$, is specified in terms of likelihood functions and prior distributions. With the assumption that the model errors are iid and $\varepsilon_i \sim N(0, \sigma_0^2)$, where σ_0^2 is fixed but unknown, the likelihood function, $\pi(y_{obs}|q)$, for a model with a single response is given by

$$\pi(y_{obs}|q) = (2\pi\sigma_0^2)^{-n/2} e^{\left(-\frac{SS_q}{2\sigma_0^2}\right)} \quad (2.5)$$

where

$$SS_q = \sum_{i=1}^n [y_i - f(t_i, q)]^2. \quad (2.6)$$

The likelihood function quantifies the probability of observing y_{obs} for given values of parameters q , whose density is specified by the prior, $\pi_0(q)$. If prior knowledge is unknown, one may employ a non-informative prior, which is often taken as an improper uniform density posed on the parameter support. We note that it is better to use non-informative prior than to use a prior of questionable accuracy. For simplicity, we let $y = y_{obs}$ in the subsequent discussion.

The likelihood function for a model with multiple responses and iid measurement errors is given by

$$\pi(y|q) = \left[\prod_{j=1}^{N_r} \frac{1}{(2\pi\sigma_j^2)^{n/2}} \right] e^{\left(-\sum_{j=1}^{N_r} \frac{SS_q^j}{2\sigma_j^2} \right)} \quad (2.7)$$

where

$$SS_q^j = \sum_{i=1}^n \left[y_i^j - f^j(t_i, q) \right]^2 \text{ for } j = 1, \dots, N_r. \quad (2.8)$$

2.3 Posterior Distributions

The goal for model calibration is to construct the posterior distributions $\pi(q|y)$ for parameters using (2.4). However, the right hand side of (2.4) requires multi-dimensional integration, where the dimension of integration depends on the number of input parameters to be estimated. When this dimension is high, it becomes impossible to compute the integral analytically, and even the integration via Monte Carlo sampling becomes prohibitively expensive. For this reason, high-dimensional integration prohibits the direct computation of marginal densities for parameters. As a solution to avoid integration, we employ sampling-based Markov Chain Monte Carlo (MCMC) methods, which avoid explicit integration. The two MCMC algorithms we consider are Delayed Response Adaptive Metropolis (DRAM) and Differential Evolution Adaptive Metropolis (DREAM). We also show a direct method, where the integral is approximated using a numerical quadrature, to compute the posterior. For low-dimensional parameter spaces, comparison with these direct solutions provides a way to verify the accuracy of densities constructed using sampling-based Metropolis methods.

2.3.1 Direct Method

We start by computing the density using the formula in (2.4) with numerical approximation. Since we evaluate the formula directly, we call this method the direct method. It is noted that the likelihood functions that appear in numerator and the denominator of (2.4) can be very

small. We avoid numerical evaluation of $\frac{0}{0}$ by reformulating the posterior so that

$$\pi(q|y) = \frac{e^{-\frac{1}{2\sigma^2}SS_q}}{\int_{\mathbb{R}^p} e^{-\frac{1}{2\sigma^2}SS_\zeta} d\zeta} \quad (2.9a)$$

$$= \frac{1}{\int_{\mathbb{R}^p} e^{-\frac{1}{2\sigma^2}(SS_\zeta - SS_q)} d\zeta}. \quad (2.9b)$$

Since the integral in (2.9b) can not be computed analytically, we use the trapezoid rule with a uniform grid to approximate the integral.

For one-dimensional integration with a uniform grid on $x \in [a, b]$ with N grid points, we let

$$x_i = a + (i-1)h \quad \text{with} \quad h = \frac{b-a}{N-1} \quad \text{for} \quad i = 1, \dots, N. \quad (2.10)$$

The 1D trapezoid rule is then

$$\int_a^b f(x) dx \approx \frac{h}{2} \sum_{i=1}^{N-1} [f(x_i) + f(x_{i+1})]. \quad (2.11)$$

The trapezoid rule is used repeatedly for higher dimensional integration. The solution obtained in the direct method is regarded as the true solution when verifying the accuracy of sampling-based methods.

2.3.2 Metropolis Algorithm

Since the evaluation of (2.4) is expensive, we seek an alternative method of model calibration that avoids high-dimensional integration. Monte Carlo sampling methods for Bayesian inference is detailed in [42]. A Markov chain Monte Carlo (MCMC) method is an algorithm in which one constructs Markov chains whose stationary distribution is the posterior density. In this section, we first describe the Metropolis algorithm. Later, we discuss variations of Metropolis algorithm that improves its efficiency by adaptation and other features.

In the Metropolis algorithm, to consider a k^{th} chain value, one first proposes a new value $q^* \sim J(q^*|q^{k-1})$, where J is called the proposal distribution. Here, the interpretation of the notation for J is that the new parameter q^* is sampled based on the previous values q^{k-1} . For noninformative prior distribution, the candidate is then accepted or rejected based on the Metropolis ratio

$$r(q^*|q^{k-1}) = \frac{\pi(q^*|y)}{\pi(q^{k-1}|y)}$$

which is a ratio of likelihood in (2.5) of the candidate parameter q^* and the previously accepted parameter q^{k-1} . The algorithm from [59] is summarized below.

1. Initialization: Choose an initial parameter value q^0 that satisfies $\pi(q^0|y) > 0$.
2. For $k = 1, \dots, M$
 - (a) Let $z \sim N(0, 1)$, and construct the candidate

$$q^* = q^{k-1} + Rz$$

where R is the Cholesky decomposition of the covariance matrix, V .

This implies that $q^* \sim J(q^*|q^{k-1}) = N(q^{k-1}, V)$, where the notation $J(q^*|q^{k-1})$ specifies the sampling of q^* based on the previous value q^{k-1} .

- (b) Compute the ratio of likelihood functions

$$r(q^*|q^{k-1}) = \frac{\pi(q^*|y)}{\pi(q^{k-1}|y)} = \frac{\pi(y|q^*)\pi_0(q^*)}{\pi(y|q^{k-1})\pi_0(q^{k-1})}.$$

- (c) Set

$$q^k = \begin{cases} q^* & \text{with probability } \alpha = \min(1, r) \\ q^{k-1} & \text{else.} \end{cases}$$

That is, accept q^* with probability 1 if $r \geq 1$ and accept it with probability r if $r < 1$.

If the proposal distribution is too wide, the candidate parameter values may be rejected too often. Alternatively, if the proposal distribution is too narrow, the parameter values may be accepted frequently but the chain does not explore the space efficiently. It is then clear that the choice of proposal distribution affects mixing of chain values. Next we discuss two variations of the Metropolis algorithm that improve its efficiency.

2.3.3 Delayed Rejection Adaptive Metropolis (DRAM)

The first MCMC variation that we use to compute posterior distributions is DRAM as discussed in [33]. DRAM is a variation of the Metropolis algorithm that is improved in two major ways.

First, Adaptive Metropolis (AM) allows the algorithm to update the chain covariance matrix in a given interval. This way, information learned about the posterior distribution through the accepted chain candidate is used to update the proposal. The covariance matrix, which determines the spread of proposal distribution, is given by

$$V_k = s_p \text{cov}(q^0, q^1, \dots, q^{k-1}) + \varepsilon I_p. \quad (2.12)$$

Here, s_p is a design parameter that depends on the parameter dimension p . A common choice is $s_p = 2.38^2/p$ as detailed in [33]. The term εI_p ensures that V_k is positive definite with $\varepsilon \geq 0$.

The covariance matrix is updated every k_0 steps, often specified as $k_0 \approx 100$ in practice. Also, the covariance is computed recursively to improve the efficiency so that

$$V_{k+1} = \frac{k-1}{k} V_k + \frac{s_p}{j} \left[k \bar{q}^{k-1} (\bar{q}^{k-1})^T - (k+1) \bar{q}^k (\bar{q}^k)^T + q^k (q^k)^T + \varepsilon I_p \right] \quad (2.13)$$

where the sample mean is also computed recursively

$$\bar{q}^k = \frac{1}{k+1} \sum_{i=0}^k q^i = q^k + \frac{k}{k+1} (\bar{q}^{k-1} - q^k). \quad (2.14)$$

Second, Delayed Rejection (DR) provides a mechanism for constructing alternative candidates when the current candidate is rejected. A second-stage candidate is chosen using the proposal function

$$q^{*2} \sim J_2(q^{*2} | q^{k-1}, q^*) = N(q^{k-1}, \gamma_2^2 V_k) \quad (2.15)$$

where $J_2(q^{*2} | q^{k-1})$ is the notation for proposing q^{*2} having started at q^{k-1} and rejected q^* , and $\gamma_2 < 1$ ensures that the second-stage proposal function is narrower than the original. The second-stage step improves the chance that the next candidate will be accepted and increases mixing. We take $\gamma_2 = \frac{1}{5}$ and employ second-stage DR in our examples in Section 2.5, though different value of γ_2 and numbers of delayed rejection stages can be employed.

The DRAM algorithm with noninformative uniform priors is summarized below.

1. Set design parameters n_s, σ_s^2, k_0
2. Determine $q^0 = \arg \min \sum_{i=1}^n [\nu_i - f_i(q)]^2$
3. Set $SS_{q^0} = \sum_{i=1}^n [\nu_i - f_i(q_0)]^2$
4. Compute the initial variance estimate $s_0^2 = \frac{SS_{q^0}}{n-p}$
5. Construct covariance estimate $V = s_0^2 [\chi^T(q^0) \chi(q^0)]^{-1}$
6. For $k = 1, \dots, M$
 - (a) Sample $z_k \sim N(0, I_p)$
 - (b) Construct candidate $q^* = q^{k-1} + R z_k$
 - (c) Compute $SS_{q^*} = \sum_{i=1}^n [\nu_i - f_i(q^*)]^2$
 - (d) Compute the ratio
$$\alpha(q^* | q^{k-1}) = \min \left(q, e^{-[SS_{q^*} - SS_{q^{k-1}}]/2s_{k-1}^2} \right)$$
 - (e) If $u_\alpha < \alpha$

Set $q^k = q^*$, $SS_{q^k} = SS_{q^*}$

Else

Enter DR Algorithm

(f) Update $s_k^2 \sim \text{Inv-gamma}(a_{val}, b_{val})$, where

$$a_{val} = 0.5(n_s, n), b_{val} = 0.5(n_s \sigma_s^2 + SS_{q^k})$$

(g) If $\text{mod}(k, k_0) = 1$

Update $V_k = s_p \text{cov}(q^0, q^1, \dots, q^k)$

Else

$V_k = V_{k-1}$

(h) Update $R_k = \text{chol}(V_k)$

The Delayed Rejection algorithm in Step 6(e) is the following.

1. Set the design parameter $\gamma_2 = \frac{1}{5}$
2. Sample $z_k \sim N(0, I_p)$
3. Construct second-state candidate $q^{*2} = q^{k-1} + \gamma_2 R_k x_k$
4. Sample $u_\alpha \sim U(0, 1)$
5. Compute $SS_{q^{*2}} = \sum_{i=1}^n [\nu_i - f_i(q^{*2})]^2$
6. Compute

$$\alpha_2(q^{*2}|q^{k-1}, q^*) = \min \left(1, \frac{\pi(q^{*2}|\nu)J(q^*|q^{*2})[1 - \alpha(q^*|q^{*2})]}{\pi(q^{k-1}|\nu)J(q^*|q^{k-1})[1 - \alpha(q^*|q^{k-1})]} \right)$$

These two mechanisms DR and AM together provide an efficient algorithm to modify and update a proposal function. The software for DRAM implementation in Matlab is available at [39, 40].

2.3.3.1 Conjugate Prior

If the prior and the posterior are in the same family, they are referred to as conjugate distributions, and the prior in this case is termed the conjugate prior. In DRAM, the user can choose to update the measurement noise as the chain evolves. If the option for updating sigma is selected, the measurement noise is updated using the conjugate prior. The computation of σ requires the specification of the prior accuracy for σ as well as the prior measurement error. In our experiment,

we assume that the measurement errors are normally distributed iid random variables with mean 0 and variance σ^2 . The likelihood of observing $y = [y_1, \dots, y_n]$ under this assumption is

$$\pi(y|\sigma^2) = \frac{1}{(2\pi\sigma^2)^{n/2}} e^{-SS_q/2\sigma^2}. \quad (2.16)$$

This likelihood function is in the inverse-gamma family and its conjugate prior is

$$\pi_0(\sigma^2) \propto (\sigma^2)^{-(\alpha+1)} e^{\beta/\sigma^2} \quad (2.17)$$

where α is the shape parameter and β is the scale parameter. It then follows that

$$\begin{aligned} \sigma|y &\sim \text{Inv-gamma}(a, b), \\ \text{with } a &= \frac{n_0 + n}{2} \quad \text{and} \quad b = \frac{n_0\sigma_0^2 + SS_q}{2}. \end{aligned} \quad (2.18)$$

Here, n_0 is specified to be a small value, such as $n_0 = 1$ or 0.1 , and σ_0^2 is the variance estimate. The computed variance is used when constructing prediction intervals. The details about this topic are discussed in [59].

2.3.3.2 Burn-in and Convergence

The period during which the chain is converging to the region of interest is termed burn-in. Depending on the initial values, it may take several chain evolutions to reach the region where the global minimum is located. The parameter values during the burn-in period are excluded from the computation of mean and variance. In DRAM, a practical way of determining the burn-in is visually inspecting the chain values. However, one must exercise care when interpreting the results even when the chains appear stationary for a very large number of simulations. The chain could transition to another region if the algorithm finds another lower minimum or parameters are unidentifiable. In an example from [35], the chain shifts after as much as 130,000 iterations.

One way to overcome the difficulty of determining burn-in and convergence is to examine the acceptance ratio, which is the percentage of accepted candidates. The acceptance ratio quantifies whether the chain is adequately sampling from the posterior. While the optimal value of acceptance ratio varies depending on the geometry of the posterior, values between 0.1 and 0.5 are often considered acceptable.

Another way of examining convergence is to check autocorrelation

$$R(k) = \frac{\sum_{i=1}^{M-k} (q_i - \bar{q})(q_{i+k} - \bar{q})}{\sum_{i=1}^M (q_i - \bar{q})^2} \quad (2.19)$$

between components that are k elements apart. Checking the autocorrelation establishes that the chain is producing iid samples from the posterior. Quantitatively, DRAM uses Geweke's convergence diagnostics based on [13]. This tests for equality of the means of the first 10% and last 50% of the chain. However, Geweke's diagnostics is a necessary, not sufficient, condition for convergence. Hence the Geweke values can only indicate when the chain is not converged, but not when it is converged.

2.3.4 Differential Evolution Adaptive Metropolis (DREAM)

The second MCMC method we employ is DREAM, which was developed in [66, 67]. DREAM is an adaptive Metropolis, global optimization algorithm, which efficiently estimates the posterior probability distributions for parameters in complex, high-dimensional problems. The candidate parameters are proposed based on differential evolution

$$z^i = x^i + (I_p + e)\gamma(\delta, p) \left[\sum_{j=1}^{\delta} x^{r_1(j)} - \sum_{n=1}^{\delta} x^{r_2(n)} \right] + \epsilon \quad (2.20)$$

where

- z^i is the candidate point in chain i
- x^i is the initial population
- γ is the value of the jump size
- δ is the number of arms used to generate the proposal
- $e \sim U_p(-b, b)$ with $|b| < 1$ and $\epsilon \sim N_p(0, I_p b^*)$ where b and b^* are small compared to the width of the target distribution.

DREAM runs multiple chains simultaneously, and the chain values are used to adapt the proposal distribution, which in turn affects the mixing. The communication between chains occur at given intervals, and the cost of communication between chains is much lower compared to the cost of model evaluations. Hence the method is inherently parallel and suitable for parameters with multi-modal, heavy-tailed densities.

The algorithm is summarized below.

1. Draw an initial population $\{x^i\}, i = 1 \dots, N$ using the prior distribution.
2. Compute $\pi(x)$.
3. Construct candidate $z^i = x^i + (I_p + e)\gamma(\delta, p) \left[\sum_{j=1}^{\delta} x^{r_1(j)} - \sum_{n=1}^{\delta} x^{r_2(n)} \right] + \epsilon$.

4. Replace each element ($j = 1, \dots, d$) of z_j^i with x_j^i using a binomial scheme with probability $1 - CR$, where CR is the crossover probability.
5. Compute $\pi(z)$ and accept the candidate point with Metropolis acceptance probability

$$\alpha(x^i, z^i) = \begin{cases} \min\left(\frac{\pi(z^i)}{\pi(x^i)}, 1\right) & \text{if } \pi(x^i) > 0 \\ 1 & \text{if } \pi(x^i) = 0 \end{cases}.$$

6. If accepted, move the chain $x^i = z^i$, otherwise remain at the old location x^i .
7. Remove potential outlier chains using the interquartile range statistics during the burn-in process.
8. Compute the Gelman and Rubin convergence diagnostic for each dimension $j = 1, \dots, p$ using the last 50% of samples in each chain.
9. If $R \leq 1.2$ for all j , stop, otherwise go to chain evolution (steps 3-6).

To compute the acceptance ratio, DREAM computes the log likelihood. For models with a single response, the log likelihood function is given by

$$\begin{aligned} \log(\pi(y|q)) &= \log\left(\frac{1}{\sqrt{2\pi\sigma^2}^n}\right) + \log\left[\exp\left(-\sum_{i=1}^n (t_i - f_i(q))^2 / 2\sigma^2\right)\right] \\ &= -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - f_i(q))^2. \end{aligned} \quad (2.21)$$

In the DREAM software, which can be requested at [65], DREAM can accommodate both homoscedastic and heteroscedastic types of measurement errors. For models with homoscedastic errors, the code can easily be modified to accommodate models with multiple responses.

2.3.4.1 Burn-in and Convergence

The burn-in period in DREAM is determined by testing the interquartile quantity $IQR = Q_3 - Q_1$ where Q_1 and Q_3 are lower and upper quartile of N different chains, respectively. Chains with $\Omega < Q_1 = 2IQR$ are considered outliers, and are moved to current best members of the population. As stated in [66], this process of outlier removal does not maintain detailed balance condition and can only be used during burn-in.

Also during the burn-in, DREAM estimates a distribution of crossover values that maximizes the quantity

$$\Delta = \sum_{i=1}^n \sum_{j=1}^p (\bar{x}_{j,t}^i - \bar{x}_{j,t-1}^i)^2 \quad (2.22)$$

between two subsequent samples. This way, DREAM adaptively chooses the CR values that yield the best mixing of chains since changing all p elements of x^i chains is often not optimal.

The convergence of sequences is monitored using R-statistics, which is a factor by which the scale of the sequences shrink to the posterior distribution as one takes more samples. As detailed in [29], R-statistics use the mean and variance between sequences as well as within sequences. R-statistics is fully quantitative, and can be used to determine convergence without visually inspecting chains.

To compute the R-statistic, take sample of the last 50% of the chain values. Let N_s be the number of samples and N_c be the number of chains. The steps for computing the factor is the following.

1. Compute the sequence means.
2. Determine the variance between sequence means, B .
3. Compute variance of various sequences.
4. Determine the average of within sequences variances, W .
5. Estimate the variance of posterior distribution by

$$\sigma^2 = \frac{N_s - 1}{N_s} W + \frac{1}{N_s} B.$$

6. Compute the R-statistics

$$R = \sqrt{\frac{N_c + 1}{N_c} \frac{\sigma^2}{W} - \frac{N_s - 1}{N_c N_s}}.$$

The scale reduction near 1 indicates that each set of the N_c chains of N_s simulated values is close to the target density. We take $R \leq 1.2$ as the cut-off for convergence. The chain evolution is repeated until $R \leq 1.2$ is achieved. The convergence is also confirmed by computing autocorrelation time.

2.4 Predictive Estimates

There are two types of predictions at a parameter value x_0 in the domain of the independent variable x but not among the data used to estimate the parameters. The first is the prediction of new observation Y_{x_0} at x_0 , and the second is the mean of the predicted response $\mu_{x_0} = E(Y_{x_0})$.

Consider first the estimation of $E(Y_{x_0})$. The intervals that bound this estimation are called confidence intervals. They quantify the accuracy of the model fit and reflect the uncertainty in parameters. On the other hand, the estimation of Y_{x_0} incorporates the result of uncertainty

in parameters as well as measurement errors. The intervals in this case are termed prediction intervals. For example, the 95% confidence interval contains 95% of the mean of the future measurement x_0 , while future measurement has 95% chance of falling inside the 95% prediction interval. We discuss the construction of confidence intervals and prediction intervals in the following subsections. Details can be found in [59] and the included references.

2.4.1 Confidence Intervals and Prediction Intervals

To motivate techniques for the nonlinear parameterized HIV model, we consider first the linearly parametrized model

$$Y = Xq + \varepsilon \quad (2.23)$$

where $q = [q_1, \dots, q_p]^T$ are the parameters and $\varepsilon = [\varepsilon_1, \dots, \varepsilon_n]^T$ are the errors, which are assumed to be independent and normally distributed with $\varepsilon_i \sim N(0, \sigma_0^2)$. An example is

$$Y_i = q_1 + \sum_{j=2}^p x_{ij}q_j + \varepsilon_i, \quad \text{for } i = 1, \dots, n \quad (2.24)$$

where q are regression coefficients and

$$X = \begin{bmatrix} 1 & x_{12} & \dots & x_{1p} \\ \vdots & \vdots & & \vdots \\ 1 & x_{n2} & \dots & x_{np} \end{bmatrix} \quad (2.25)$$

is the design matrix.

Next we summarize discussion in [30]. First we note that

$$\hat{Y}_{x_0} = x_0^T \hat{q} \quad (2.26)$$

is an unbiased estimator for $E(Y_{x_0})$. Since the unbiased error covariance estimator is

$$\hat{\sigma}^2 = \frac{1}{n-p} \hat{R}^T \hat{R} \quad (2.27)$$

where $\hat{R} = Y - Xq$ is the residual estimator, it follows that

$$\text{var}(\hat{Y}_{x_0}) = \sigma_0^2 [x_0^T (X^T X)^{-1} x_0]. \quad (2.28)$$

Here, the variance estimator $\hat{\sigma}$ is used for unknown σ and

$$\hat{\sigma}^2(\hat{Y}_{x_0}) = \hat{\sigma}^2[x_0^T(X^T X)^{-1}x_0]. \quad (2.29)$$

With the assumption that $\varepsilon_i \sim N(0, \sigma_0^2)$, the sampling distribution is a normal distribution with mean μ_{x_0} and variance in (2.28) so that

$$\frac{\hat{Y}_{x_0} - \mu_{x_0}}{\sigma_0 \sqrt{x_0^T(X^T X)^{-1}x_0}} \sim N(0, 1). \quad (2.30)$$

It then follows that

$$T = \frac{\hat{Y}_{x_0} - \mu_{x_0}}{\hat{\sigma} \sqrt{x_0^T(X^T X)^{-1}x_0}} \quad (2.31)$$

has a t-distribution with $n - p$ degrees of freedom. The $(1 - \alpha) \times 100\%$ interval estimate for μ_{x_0} is

$$\left[\hat{Y}_{x_0} \pm t_{n-p, \alpha/2} \cdot \hat{\sigma} \sqrt{x_0^T(X^T X)^{-1}x_0} \right]. \quad (2.32)$$

This is a confidence interval because it only incorporates the parameter uncertainty only. This type of interval constructed based on the Bayesian framework are called the credible intervals.

For the construction of the prediction interval, we begin by noting that the random variable $Y_{x_0} - \hat{Y}_{x_0}$ is normally distributed with mean

$$E(Y_{x_0} - \hat{Y}_{x_0}) = 0 \quad (2.33)$$

and variance

$$\begin{aligned} \text{var}(Y_{x_0} - \hat{Y}_{x_0}) &= \text{var}(Y_{x_0}) + \text{var}(\hat{Y}_{x_0}) \\ &= \sigma_0^2 [1 + x_0^T(X^T X)^{-1}x_0]. \end{aligned} \quad (2.34)$$

It follows that

$$\frac{\hat{Y}_{x_0} - \mu_{x_0}}{\sigma_0 \sqrt{1 + x_0^T(X^T X)^{-1}x_0}} \sim N(0, 1) \quad (2.35)$$

and

$$T = \frac{\hat{Y}_{x_0} - \mu_{x_0}}{\hat{\sigma} \sqrt{1 + x_0^T(X^T X)^{-1}x_0}} \quad (2.36)$$

has a t-distribution with $n - p$ degrees of freedom. Then the $(1 - \alpha) \times 100\%$ interval estimate for Y_{x_0} , i.e. the prediction interval, is

$$\left[\hat{Y}_{x_0} \pm t_{n-p, \alpha/2} \cdot \hat{\sigma} \sqrt{1 + x_0^T(X^T X)^{-1}x_0} \right]. \quad (2.37)$$

We have shown that the credible and prediction intervals are given by (2.32) and (2.37). In

case of nonlinearly parametrized problems, X is replaced by the sensitivity matrix χ where

$$\chi_{ij} = \frac{\partial f^i(t, q)}{\partial q_j} \quad (2.38)$$

as detailed in [53].

2.4.2 DRAM

To construct the intervals in DRAM, N parameter values as specified by the user are sampled and the model responses are computed at these parameter values. For example, the 95% credible interval is constructed by sorting the computed responses and finding the upper 2.5% and 97.5% values. For the 95% prediction interval, the measurement noise is added to the computed responses before finding the 95% interval. The measurement noise is generated by a normally distributed random variable with mean 0 and variance σ^2 , where σ^2 is constructed using the techniques in Section 2.3.3.1. If σ^2 is not constructed during model calibration, we use the OLS estimate of measurement noise in (2.27) to construct the prediction intervals.

2.4.3 DREAM

In DREAM, the values from all of the chains are combined in a matrix so that the i^{th} column contains all the chain values for the parameter q_i . The number of columns corresponds to the number of parameters that are estimated. The model response is computed for the last 25% of the values from each column. It is noted that 25% is chosen to ensure that the parameter values in the chains have converged. The generated responses are sorted, and the appropriate interval of the response is stored. For instance, the responses at upper 2.5% and 97.5% are stored for the 95% confidence interval. For prediction intervals, measurement noise is added to the response. The noise is generated from a distribution $\varepsilon \sim N(0, \text{RMSE}^2)$, where

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (f(t_i; q) - y_i)^2}. \quad (2.39)$$

Once the noise is added to the model responses, they are sorted and the responses at the upper 2.5% and 97.5% are stored, for example, to construct the 95% prediction interval. Since the DREAM package does not come with the feature to update measurement noise, we added this feature using the prior conjugate in the similar manner it is implemented in DRAM. If the measurement noise is updated during the chain evolutions, sampled σ^2 values are used when constructing the prediction interval.

2.5 Results and Verification

2.5.1 Heat Example

Here, we apply DRAM and DREAM as well as the Direct Method to the heat equation in (1.2). For the direct method, we compute the posterior density using the formula (2.4). Here, we use 201 uniform grid points in each dimension. First, the sum of squares in (2.6) is computed for the rectangular mesh. The likelihood is computed using the SS_q and the measurement error, $\sigma^2 = 0.0627$, for the rectangular mesh. The result is then integrated in each direction to obtain the posterior distribution for each parameter.

We use 30,000 simulations for DRAM and 30,000 maximum function evaluations for DREAM. We construct the measurement error by using the error variance obtained via OLS, $\sigma^2 = 0.0627$, as an initial measurement error. The initial parameter values are also set to be the values obtained via OLS, which are $[\Phi_{\text{initial}}, h_{\text{initial}}] = [-18.41, 0.00191]$. The minimum and maximum values for the two parameters are set to be $[\Phi_{\text{min}}, h_{\text{min}}] = [-20, 0]$ and $[\Phi_{\text{max}}, h_{\text{max}}] = [0, 1]$.

2.5.1.1 Chains and Marginal Densities

We compare the densities from the three methods by plotting them in Figure 2.1. We use 201 grid points for the direct method, 50,000 simulations for DRAM and 400,000 simulations for DREAM. The parameter correlations are compared in Figure 2.2, where the direct method correlation is represented by the contour plots, while the DRAM and DREAM correlations are represented by scatter plots.

2.5.1.2 Convergence

Here we examine the convergence for each method. We first begin with the direct method. Starting with the 51 grid points in each dimension, we observe the change when we increase the grid points to 101, 201 and to 401. In this case, we observe that the mean and the standard deviation of the parameters do not change. This implies that the computation of posterior distribution is accurate with the number of grid points as low as 51. To get a smooth curve for densities, we choose 201 grid points for the direct method as a reference.

Next we examine the convergence with DRAM and DREAM. We plot in Figure 2.3 the results with different simulations numbers along with the 201-grid points direct method as a reference. The parameter values obtained in this convergence test for DRAM and DREAM are summarized in Table 2.2 and 2.3. We see that the densities in DRAM match those of the direct well even for 25,000 simulations. We use the densities from 50,000 simulations for better accuracy. On the other hand, the densities in DREAM tend to be not so smooth at the peaks. It is noted that in some runs, the densities are very smooth with just 50,000 simulations.

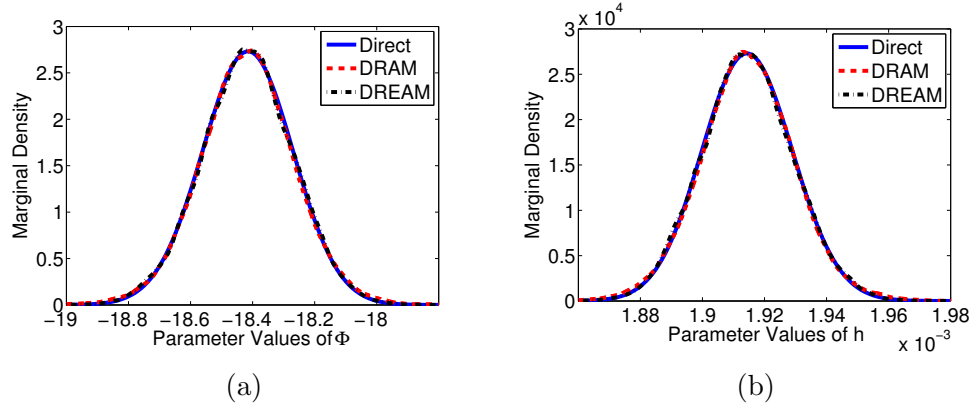


Figure 2.1 Densities from DRAM, DREAM and direct methods for (a) Φ and (b) h .

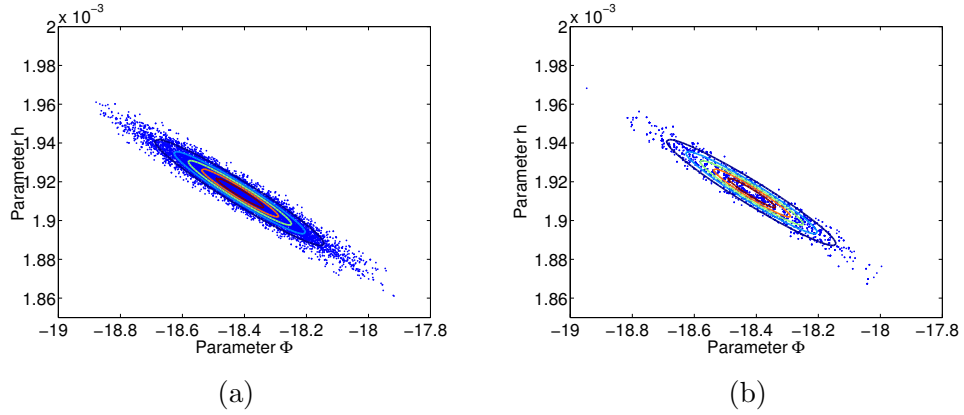


Figure 2.2 Contour plot from the direct method and the correlation plots from (a) DRAM and (b) DREAM.

Table 2.2 Parameter values with difference simulation numbers in DRAM.

Number of DRAM Simulations	μ_Φ	σ_Φ	μ_h	σ_h	σ^2
25000	-18.419	0.15146	1.9147E-3	1.5140E-5	0.06741
50000	-18.417	0.15241	1.9145E-3	1.5244E-5	0.06778
100000	-18.418	0.15148	1.9147E-3	1.5163E-5	0.06710
Direct	-18.417	0.1460	1.9145E-3	1.4611E-5	0.0627

As we can see in Table 2.3, the parameter means and variances are comparable to those from the direct method with 50,000 simulations. Visually, we need to run at least 400,000 simulations to obtain a smooth curve as illustrated in Figure 2.3(b). One gets even smoother density curves, especially for Φ , if 800000 simulations are performed. However, the number is unreasonably large for a simple example like the heat equation, especially when a good convergence is attained at 50000 simulations with DRAM.

The heat equation in (1.2) was used to develop a verification framework for sampling-based methods for the uncertainty quantification project funded by Consortium for Advanced Simulation of Light Water Reactors (CASL). The verification is presented in the milestone report for CASL [34].

2.5.2 HIV Example

In this section, we demonstrate the performance of the two MCMC methods for the HIV model (2.1) for two cases. In the first, we consider three parameters $q = [b_E, \delta, d_1]$ to permit verification of DRAM and DREAM with direct numerical evaluation of (2.4). We also illustrate the construction of credible and prediction intervals for the model response and demonstrate that DRAM and DREAM provide comparable results. Secondly, we consider the construction of posterior densities for 12 parameters to illustrate the feasibility of DRAM and DREAM for a problem with moderate dimensionality.

To generate the synthetic data, we solve the model (2.1) with the initial condition

$$y_0 = \begin{bmatrix} 9\text{e}+5, & 4000, & 1, & 1, & 1, & 12 \end{bmatrix} \quad (2.40)$$

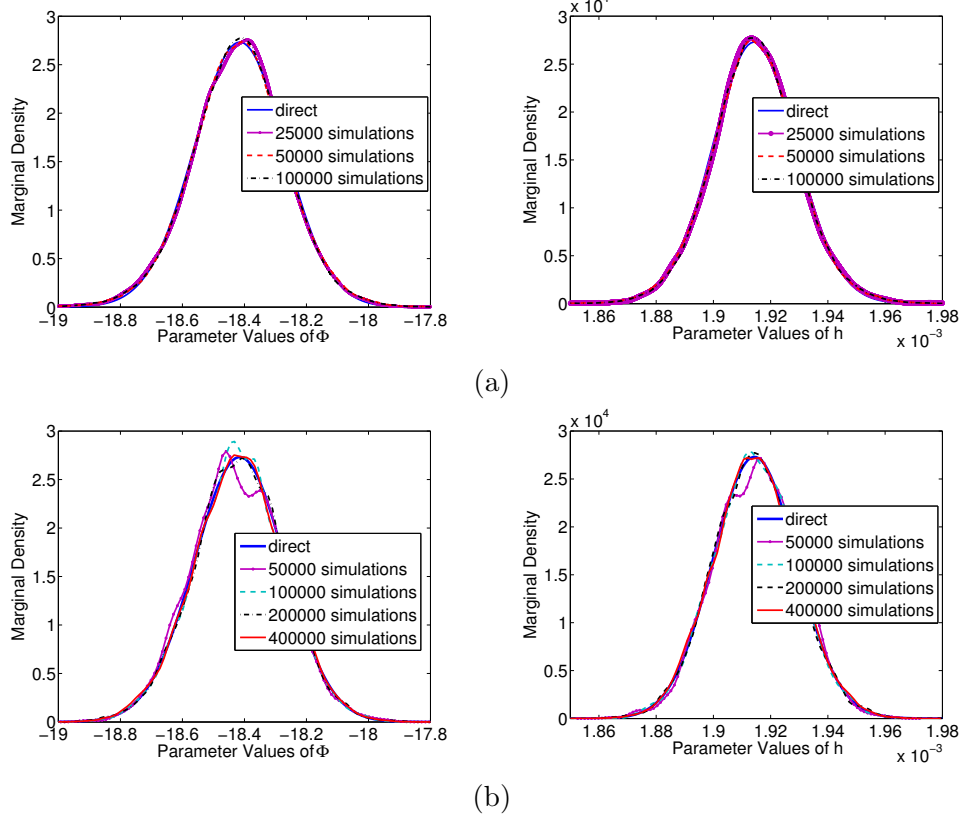


Figure 2.3 Distributions of different simulation numbers for (a) DRAM and (b) DREAM.

Table 2.3 Parameter values with difference simulation numbers in DREAM.

Number of Simulations	μ_Φ	σ_Φ	μ_h	σ_h	σ^2
50000	-18.4218	0.14588	1.9149E-3	1.4500E-5	0.06721
100000	-18.4165	0.14313	1.9145E-3	1.4380E-5	0.06697
200000	-18.4161	0.14596	1.9144E-3	1.4579E-5	0.06711
400000	-18.4167	0.14674	1.9145E-3	1.4681E-5	0.06724
800000	-18.4169	0.1458	1.9145E-3	1.4580E-5	0.06739
Direct	-18.417	0.1460	1.9145E-3	1.4611E-5	0.0627

and the parameter values summarized in Table 2.4, which are reported in [3]. Once the solution to the model is computed, we add random, normally distributed measurement errors $\varepsilon_j \sim N(0, \sigma_j^2)$ for $j = 1, \dots, 6$, where

$$\begin{aligned} \begin{bmatrix} \sigma_1, & \sigma_2, & \dots, & \sigma_6 \end{bmatrix} &= \begin{bmatrix} \sigma_{T_1}, & \sigma_{T_2}, & \sigma_{T_1^*}, & \sigma_{T_2^*}, & \sigma_V, & \sigma_E \end{bmatrix} \\ &= \begin{bmatrix} 2\text{e}+4, & 10, & 3\text{e}+3, & 10, & 10\text{e}+4, & 2.5 \end{bmatrix}. \end{aligned} \quad (2.41)$$

For comparison, we employ the direct method to compute the parameter densities. In this example, we use a uniform grid with 101 grid points in each dimension. We restricted the region of integration to $[0.296, 0.304] \times [0.69, 0.72] \times [0.007, 0.012]$ for our computations to decrease the run time since the magnitude of the likelihood function outside the specified region is negligible. The accuracy of numerical approximation of the direct method is illustrated by the convergence test with $n = 61, 81$, and 101 summarized in Table 2.5. The means and standard deviations of the three parameters remain unchanged between $n = 81$ and $n = 101$. This indicates that the numerical approximation of integrals in (2.4) has converged. Hence the direct method can be considered a baseline for the verification of the sampling methods.

For the three parameter case, we first perform numerical experiments assuming that the measurement noise is fixed. In this case, the measurement noise is taken to be the OLS estimate of measurement noise in (2.27). The chains, marginal densities, correlations and prediction intervals are compared with this setting. We then perform additional experiments, where the measurement noise is updated as the chains evolve. For this case, we compare the credible

Table 2.4 Fixed values for parameters.

$\lambda_1 = 1\text{e}+4$	$\epsilon = 0$	$k_1 = 8.0\text{e-}7$	$\lambda_2 = 31.98$
$d_2 = 0.01$	$f = 0.34$	$m_1 = 1.0\text{e-}5$	$m_2 = 1.0\text{e-}5$
$N_T = 100$	$c = 13$	$\rho_1 = 1$	$\rho_2 = 1$
$\lambda_E = 1$	$K_b = 100$	$d_E = 0.25$	$K_d = 500$
$\delta_E = 0.1$	$k_2 = 1\text{e-}4$	$\delta = 0.7$	$b_E = 0.3$
$d_1 = 0.01$			

Table 2.5 Mean and standard deviation of b_E , δ and d_1 using $n = 61, 81, 101$ in direct evaluation of Bayes' relation.

	b_E Mean	b_E Sigma	δ Mean	δ Sigma	d_1 Mean	d_1 Sigma
$n = 61$	2.9990e-1	7.5089e-4	7.0558e-1	3.0448e-3	9.5777e-3	7.5642e-4
$n = 81$	2.9990e-1	7.5090e-4	7.0558e-1	3.0448e-3	9.5777e-3	7.5644e-4
$n = 101$	2.9990e-1	7.5090e-4	7.0558e-1	3.0448e-3	9.5777e-3	7.5645e-4

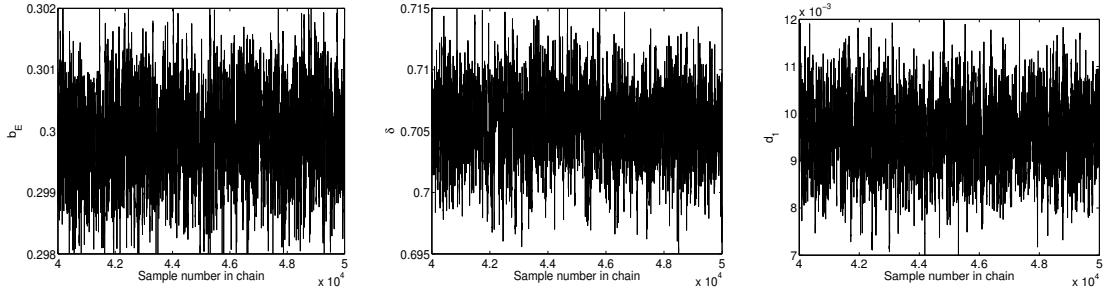


Figure 2.4 Chain evolutions of b_E , δ and d_1 for DRAM.

and prediction intervals of DRAM and DREAM. We do not include DIRECT here since the method does not update measurement noise. We point out that this is one limitation of the direct method.

2.5.3 Three Parameter Case

2.5.3.1 Chains and Marginal Densities

Whereas DRAM and DREAM each has convergence criteria to ensure the convergence of parameters, the chains can also be used to visually examine the convergence for both DRAM and DREAM. It is noted that DRAM contains a single chain, from which means and variances of the parameters are computed. In Figure 2.4, we observe that the DRAM chains have converged in probability to a stationary distribution. On the other hand, DREAM is comprised of multiple chains, which evolve to fluctuate around the mean value as the iteration number increases. The initial iterations are plotted in Figure 2.5 to illustrate the mixing between chains. The last 1000 iterations of five chains are plotted in Figure 2.6 and we see that the chains have converged.

We see from the results summarized in Table 2.6 that the means and the standard deviations for the three parameters are comparable among the three methods. In Figure 2.7, we illustrate that the marginal densities obtained with the three methods are in agreement. The densities shown in 2.7 are after applying a kernel density estimator to smooth the densities. Additional smoothing must be performed especially if the density is multi-modal. Typically, DREAM requires more smoothing than DRAM. In Figure 2.8, we show the contour plots of likelihood as a function of two parameters along with the joint sample points from DRAM and DREAM. While we can not quantify the correlation, we see qualitatively that the shapes of correlations are in agreement among the three methods. Both the marginal densities and correlations

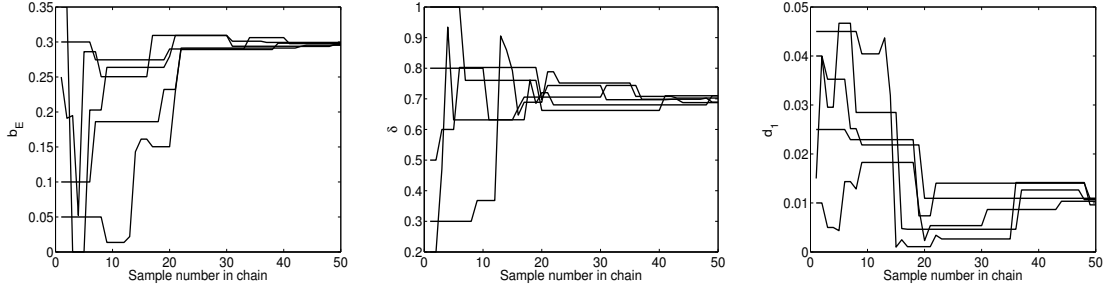


Figure 2.5 Initial chain evolutions of b_E , δ and d_1 for DREAM.

Table 2.6 Mean and standard deviation for b_E , δ and d_1 .

Method	b_E Mean	b_E Sigma	δ Mean	δ Sigma	d_1 Mean	d_1 Sigma
Direct	2.999e-1	7.509e-4	7.056e-1	3.045e-3	9.577e-3	7.564e-4
DRAM	2.999e-1	7.416e-4	7.056e-1	3.041e-3	9.565e-3	7.676e-4
DREAM	3.000e-1	7.486e-4	7.055e-1	2.989e-3	9.567e-3	7.476e-4

appear comparable, thus verifying the accuracy of the sampling-based Metropolis algorithms. Moreover, the scatter plots constructed using DRAM and DREAM coincide with the contour plots constructed using the direct method. The contour plots represent the values of joint density functions. We note that the parameters d_1 and b_E are slightly negatively correlated, while d_1 and δ , as well as δ and b_E , exhibit minimal correlation.

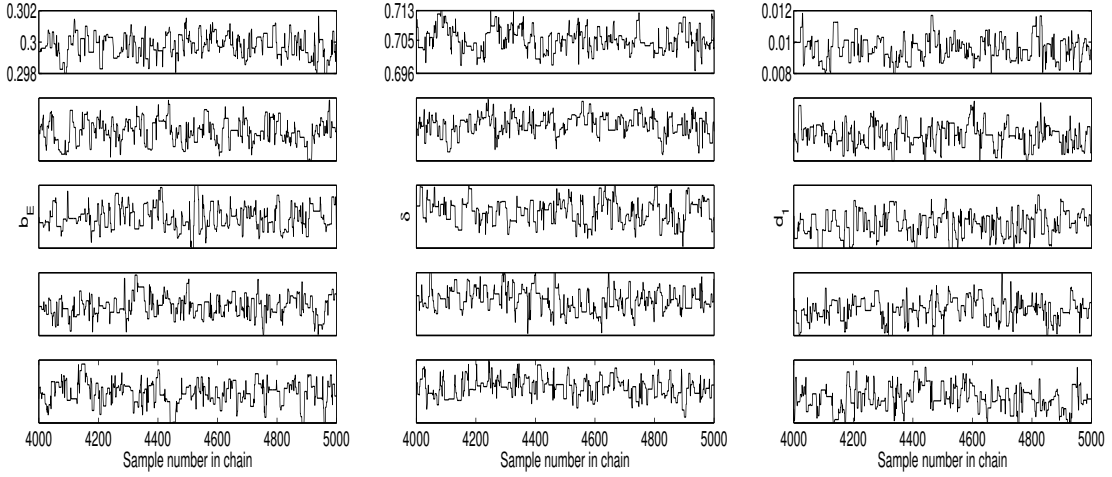


Figure 2.6 Final chain evolutions of b_E , δ and d_1 for DREAM.

2.5.3.2 Prediction Intervals

We next consider the credible and prediction intervals shown in Figure 2.9. The inner intervals represent the credible intervals, which incorporate parameter uncertainties. The outer intervals in lighter gray represent the prediction intervals. The results via Direct, DRAM and DREAM are nearly identical. One can observe that about 95 out of the 100 data points fall in the prediction intervals, which is consistent with its definition. The prediction interval can also be used to predict model response outside of the domain that is used in this experiment. In that case, the next collected measurement is expected to fall inside the prediction interval 95% of the time.

2.5.3.3 Convergence Diagnostics

Here we examine the convergence diagnostics for DRAM and DREAM. The default diagnostics used in DRAM are summarized in Table 2.7. MC error indicates the standard deviations from the batch mean, which estimates the Monte Carlo standard deviations. τ is the autocorrelation time. The Geweke diagnostic near 1 indicates that the mean from the last 10% and 50% samples are very close. The acceptance ratio without the delayed rejection stage is 0.310 and the ratio combined with 2nd stage delayed rejection is 0.527.

On the other hand, we look at R-statistics and autocorrelation time for DREAM. The R values in Figure 2.10 for all three parameters are below 1.2 after $1e+4$ iterations. The acceptance ratio for DREAM is 0.343.

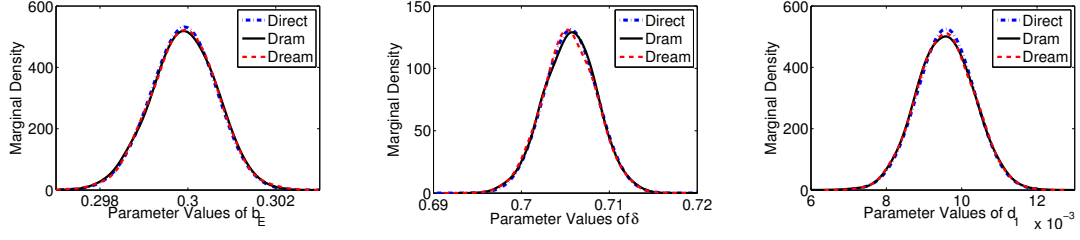


Figure 2.7 Marginal densities for b_E , δ and d_1 with different methods.

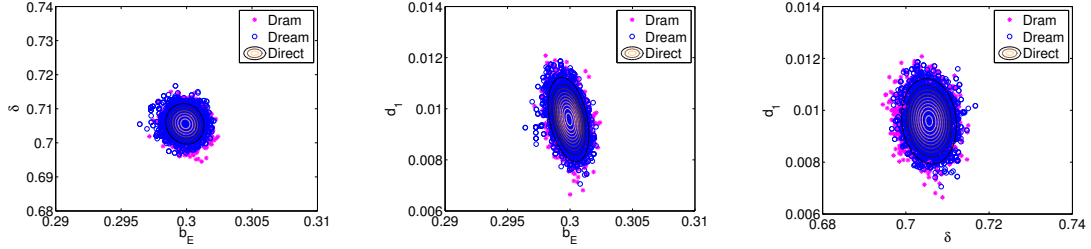


Figure 2.8 Contour plots from the direct method plotted with joint sample points from DRAM and DREAM showing the correlations between b_E , δ and d_1 .

2.5.3.4 Updating Measurement Error

We also test DRAM and DREAM to see how well they estimate the measurement noise. The OLS estimate is $\sigma_E = 2.975$, whereas the updated measurement for E using DRAM is $\sigma_E = 2.975$ and using DREAM is $\sigma_E = 2.985$. The values are comparable, though updated noise by DRAM is closer to the OLS estimate than that by DREAM. The direct method does not allow updating the measurement noise feature, hence we compute the credible and prediction intervals using the DRAM and DREAM simulations. The 95% credible and prediction intervals for the response E are plotted in Figure 2.11.

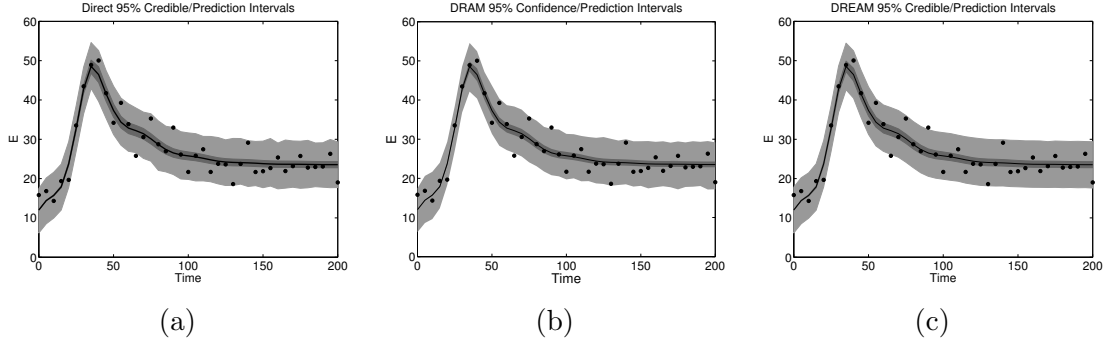


Figure 2.9 Credible and prediction intervals for response E using (a) Direct, (b) DRAM and (c) DREAM.

Table 2.7 Convergence diagnostics for DRAM.

Parameter	MC error	τ	Geweke
d_1	1.550e-05	8.9863	0.99971
δ	6.150e-05	10.055	0.99941
b_E	1.148e-05	9.447	0.98654

2.5.4 12 Parameter Case

Here, we illustrate the verification of DRAM and DREAM by comparing densities for a moderate number of parameters. We estimate the densities for

$$Q_{12} = [b_E, \delta, d_1, k_2, \lambda_1, K_b, K_d, k_1, \lambda_2, c, \rho_1, \rho_2]. \quad (2.42)$$

It is noted here that the orders of magnitudes of these parameters vary greatly. More specifically, the smallest and the largest parameter values among Q_{12} that are used to generate the synthetic data are $k_1 = 8\text{e-}7$ and $\lambda_1 = 1\text{e}+4$, respectively.

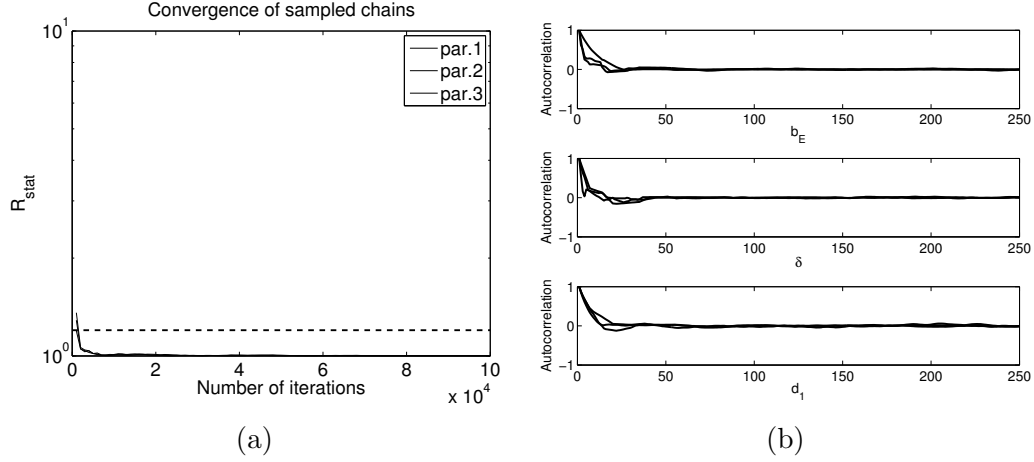


Figure 2.10 (a) R-Statistics and (b) autocorrelation for b_E , δ and d_1 in DREAM.

We use the generated data to perform model calibration using DRAM and DREAM. We take the total of 1,000,000 chain evaluations including burn-in for DRAM. In DREAM, we set the maximum number of function evaluations to be 1,000,000 and use 10 chains. The prior for both DRAM and DREAM is a uniform distribution $q_i \sim U(a_i, b_i)$ for $i = 1, \dots, 12$, where a_i and b_i are summarized in Table 5.2.

The estimated densities are plotted in Figure 2.12 and the correlation plots are shown in Figures 2.13 and 2.14. The direct method is not easily applicable in this case due to the larger dimension of input parameters. The mean values are summarized in Table 2.9, where the values listed as TRUE represent the parameter values used to generate the synthetic data.

The acceptance ratio for DRAM and DREAM are 0.201 and 0.147, respectively. The convergence diagnostics for DRAM are summarized in Table 2.10. Larger values of τ compared to the three parameter case indicate slower convergence for the 12 parameter case. The Geweke values near 1 for most parameters indicates that the means have remained the same over iterations. The R-statistics and autocorrelations are plotted in Figure 2.15 and 2.16 as convergence diagnostics for DREAM. R-statistics values are below the threshold after $1e+5$ iterations; however, the autocorrelation values are not near zero for some parameters. Again, the lack of analytic values of samples for convergence indicates the necessity for performing a verification test through comparison of DRAM and DREAM. Figure 2.12 indicates that the two methods give similar

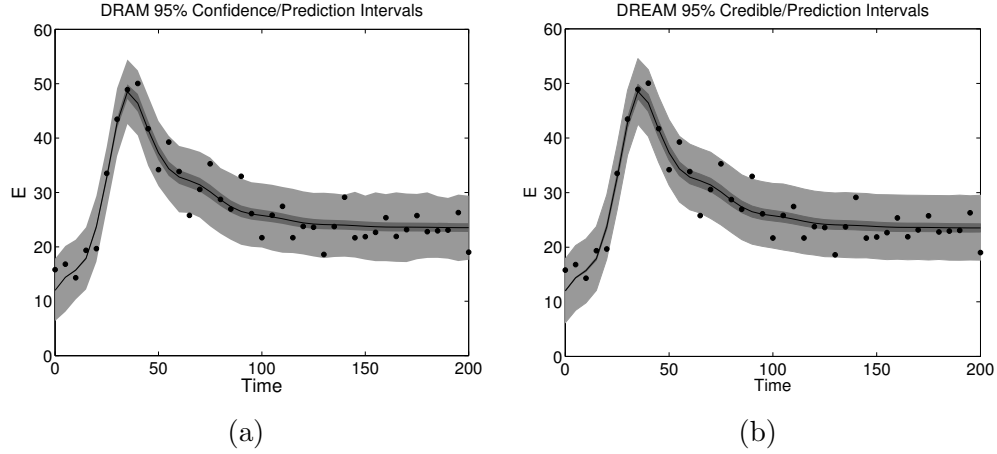


Figure 2.11 Credible and prediction intervals for E using (a) DRAM and (b) DREAM when σ is updated.

densities for all 12 parameters.

The mean values are not in agreement among TRUE, DRAM and DREAM for some parameters. In some cases, the disagreement is due to the measurement noise. In fact, in the absence of measurement error, the mean values of DRAM and DREAM are much closer to the true values. The disagreement may also be caused by parameter unidentifiability issues. For example, correlated parameters and noninfluential parameters may not be estimated uniquely. In our example, d_1 and λ_1 as well as k_1 and c appear linearly correlated, which affects the parameter identifiability. Also, K_b and K_d have wider distributions and these may indicate that the parameters can not be estimated uniquely by the observations.

There are methods to eliminate parameter unidentifiability. Though they perform Frequentist inference, the authors of [50] presents methods for detecting unidentifiable parameters prior to estimation. Also, issues with overparameterization are discussed in [58], where the ion channel model with superfluous states results in parameters that spread over a wide range of values are multi-modal or uniform distribution. Detection of unidentifiable parameters results not only in accurate parameter estimation, but also in reduction of parameter dimensions. An example is illustrated in [32], which demonstrates the reduction of 61 model parameters to 26 parameters through successful isolation of unidentifiable parameters. Alternatively, one could

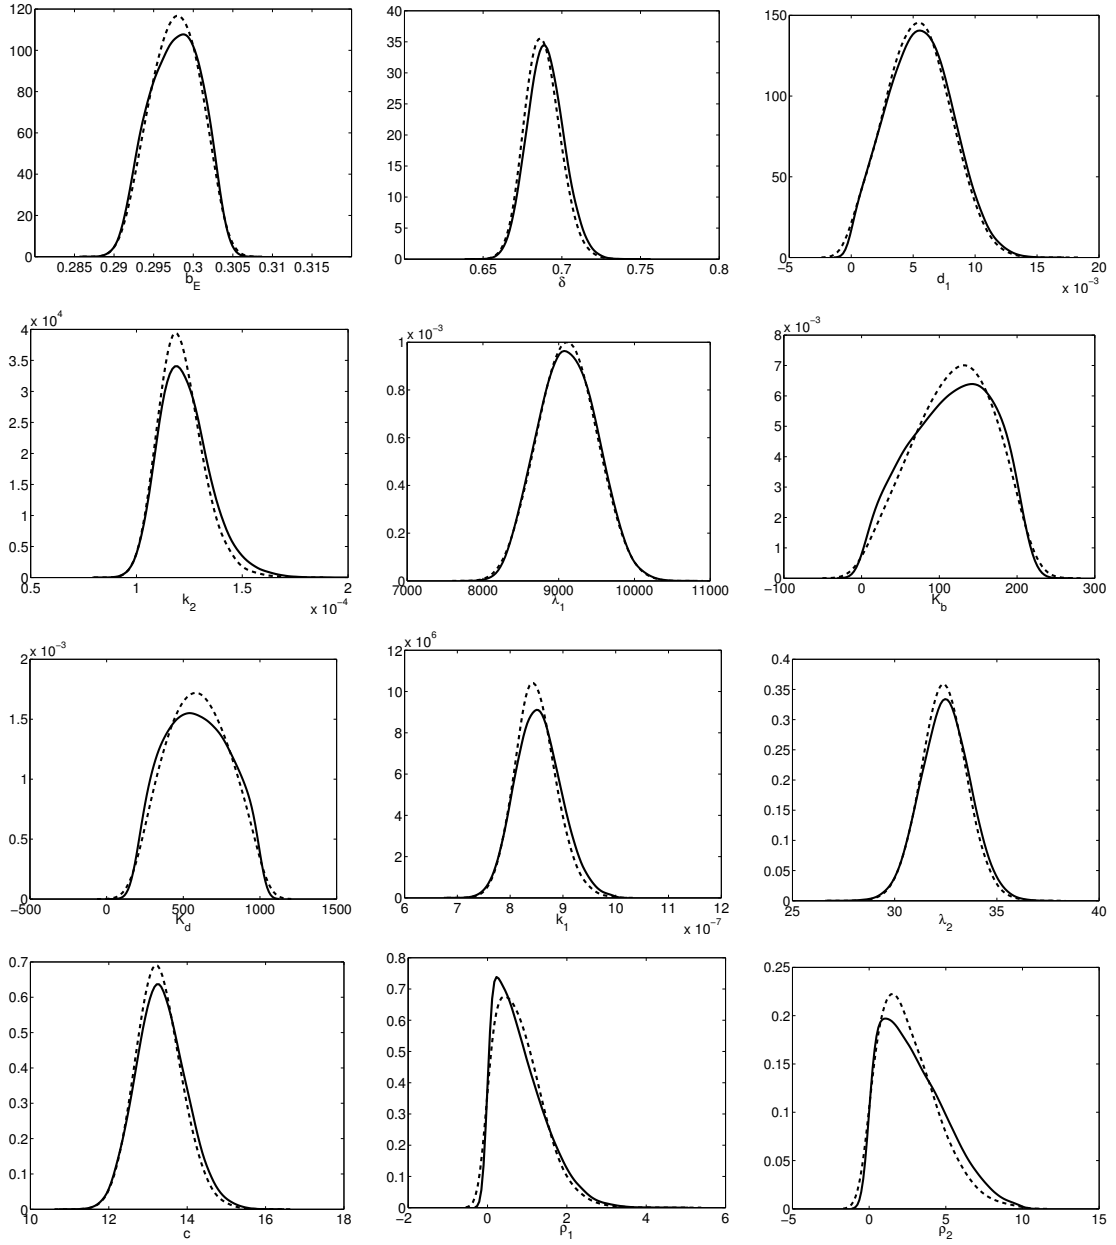


Figure 2.12 Densities for 12 parameters with DRAM (solid line) and DREAM (dashed line).

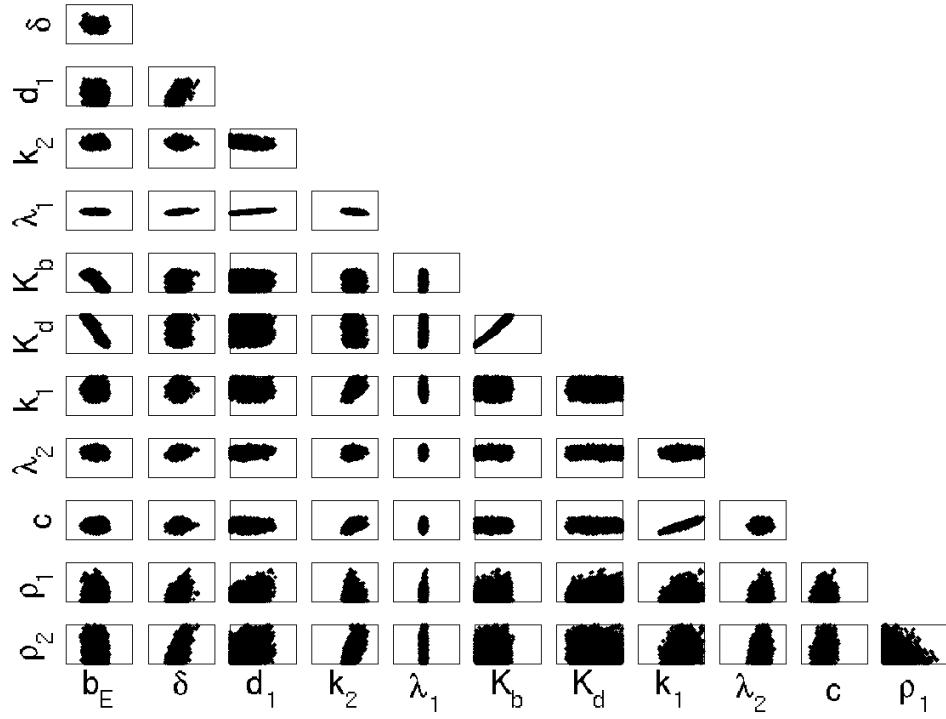


Figure 2.13 Joint sample points obtained using DRAM.

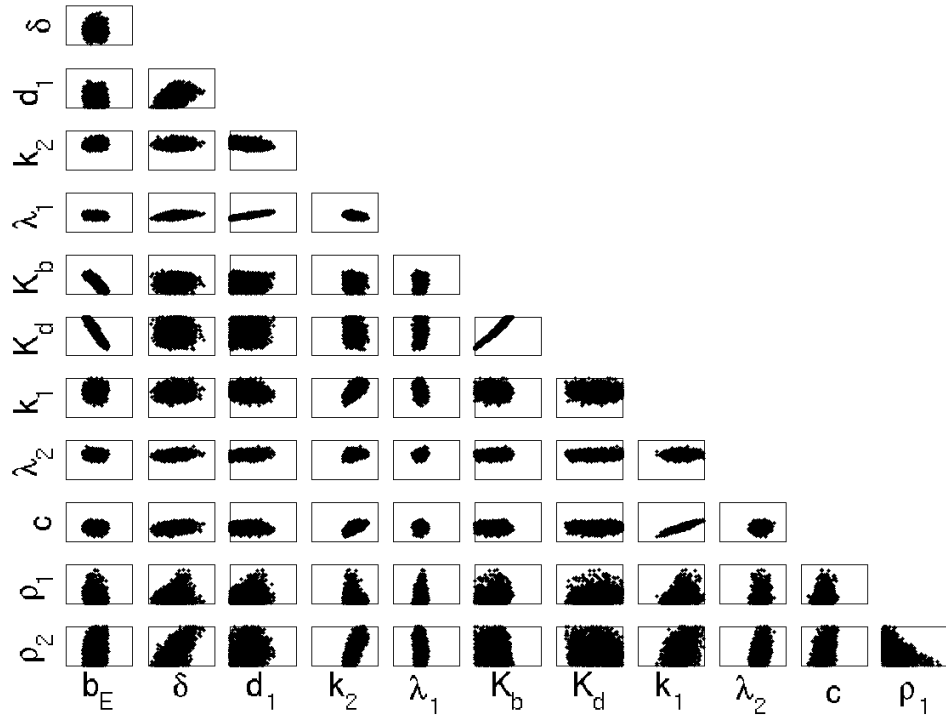


Figure 2.14 Joint sample points obtained using DREAM.

Table 2.8 The lower and upper bounds for the prior distributions $q_i \sim U(a_i, b_i)$ for $i = 1, \dots, 12$.

	b_E	δ	d_1	k_2	λ_1	K_b	K_d	k_1	λ_2	c	ρ_1	ρ_2
a_i	0	0	0	0	0	0	0	0	0	0	0	0
b_i	0.5	1	0.05	1e-3	1e+5	1e3	1e+3	1e-6	50	50	10	10

perform parameter selection to eliminate unidentifiable parameters prior to model calibration. Some parameter selection techniques include variance-based methods and screening methods based on global sensitivity analysis. The details on this topic can be found in [59].

We showed that DRAM and DREAM can be used to construct parameter densities even when individual parameters differ by many orders of magnitudes. One of the major difference between DRAM and DREAM is the single-chain in DRAM as opposed to the simultaneous multiple chains in DREAM. Running multiple chains of DRAM is also an option if a single-chain DRAM exhibits a behavior of multi-modal distribution. In the case of the HIV model (2.1), a single-chain DRAM is sufficient to efficiently sample from the posterior distribution.

2.6 Conclusion

We described Bayesian techniques for model calibration and used them to construct prediction intervals for the HIV model (2.1). We first constructed posterior densities for three parameters using DRAM and DREAM. A single chain for each parameter evolves to the posterior distribution using delayed rejection and adaptation features in DRAM, whereas multiple chains are run to increase efficiency in DREAM. Candidates from both methods are either accepted or rejected according to the Metropolis ratio, which determines whether the likelihood increased or not.

To verify the accuracy of the DRAM and DREAM methods, we demonstrated that they yield the same parameter densities and parameter correlations. We then compared the parameter densities and correlations via DRAM and DREAM to the true values computed directly from (2.4) via the direct method. The densities for three parameters using the three methods are agreeable. When we performed parameter estimation using DRAM and DREAM for 12 parameters, both methods successfully estimated parameter densities despite the large range in the parameters, which vary from $\mathcal{O}(10^{-7})$ to $\mathcal{O}(10^4)$.

Table 2.9 Mean values obtained using DRAM and DREAM for 12 parameters.

	b_E	δ	d_1	k_2	λ_1	K_b
TRUE	3.00e-1	7.00e-1	1.00e-2	1.00e-4	1.00e+4	1.00e+2
DRAM	2.98e-1	6.90e-1	5.55e-3	1.23e-4	9.12e+3	1.16e+2
DREAM	2.98e-1	6.88e-1	5.36e-3	1.21e-4	9.11e+3	1.18e+2
	K_d	k_1	λ_2	c	ρ_1	ρ_2
TRUE	5.00e+2	8.00e-7	3.20e+1	1.30e+1	1.00	1.00
DRAM	5.91e+2	8.54e-7	3.25e+1	1.33e+1	8.33e-1	2.96
DREAM	5.95e+2	8.49e-7	3.24e+1	1.33e+1	8.08e-1	2.63

Table 2.10 Convergence diagnostics for DRAM with 12 parameters.

Parameter	MC Error	τ	Geweke
b_E	5.36e-5	83.45	0.9992
δ	1.15e-4	59.92	0.9993
d_1	2.33e-5	57.80	0.9988
k_2	1.67e-7	83.78	0.9927
λ_1	3.52	57.24	0.9996
K_b	1.08	81.52	0.9678
K_d	4.13	81.91	0.9684
k_1	5.65e-10	61.40	0.9970
λ_2	1.44e-2	57.78	0.9971
c	6.06e-3	59.86	0.9984
ρ_1	8.71e-3	64.68	0.9563
ρ_2	2.96e-2	75.96	0.9937

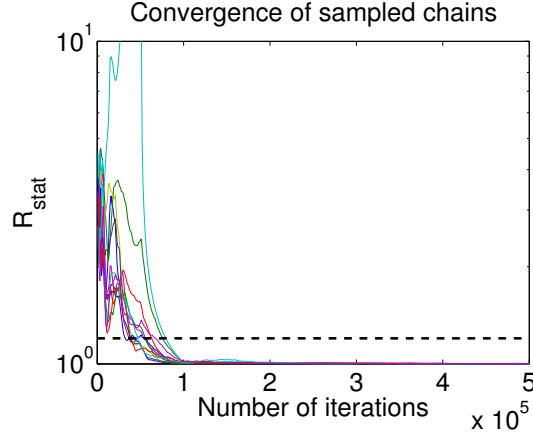


Figure 2.15 R-statistics for 12 parameters in DREAM.

We also discussed credible and prediction intervals, which are constructed by propagating the input uncertainties through the model. The credible intervals quantify uncertainties in the model responses due to the parameter uncertainties, whereas the prediction intervals quantify uncertainties in the responses due to both the parameter uncertainties and the measurement error. In the HIV example, credible and prediction intervals were constructed for the immune effector cells, E , using DRAM and DREAM as a part of verification.

The sampling methods, with verified accuracy as discussed here, can be used to perform parameter estimation and to obtain predictive estimates in other models. They can also be extended to a larger number of parameters. These methods have been successfully applied to models involving gene transcription activity [11, 51], cellular signaling pathways [38], and algae model [32]. Moreover, the application of DRAM and DREAM to models with higher parameter dimensions has also been documented. Models with input parameter dimensions, 25, 50 and 100, are examined in [66]. Also, a hydrology model in [68] has 63 to 65 parameters, where DREAM is performed for model calibration. Recently, [43] presented the use of MCMC methods, including DRAM and DREAM, as a part of uncertainty quantification for a groundwater flow problem, in addition to a 10-dimensional multi-modal mathematical function and 100-dimensional Gaussian functions. In their examples, it is reported that DREAM is more efficient than DRAM. These applications indicate the efficiency of DRAM and DREAM in model calibration for models

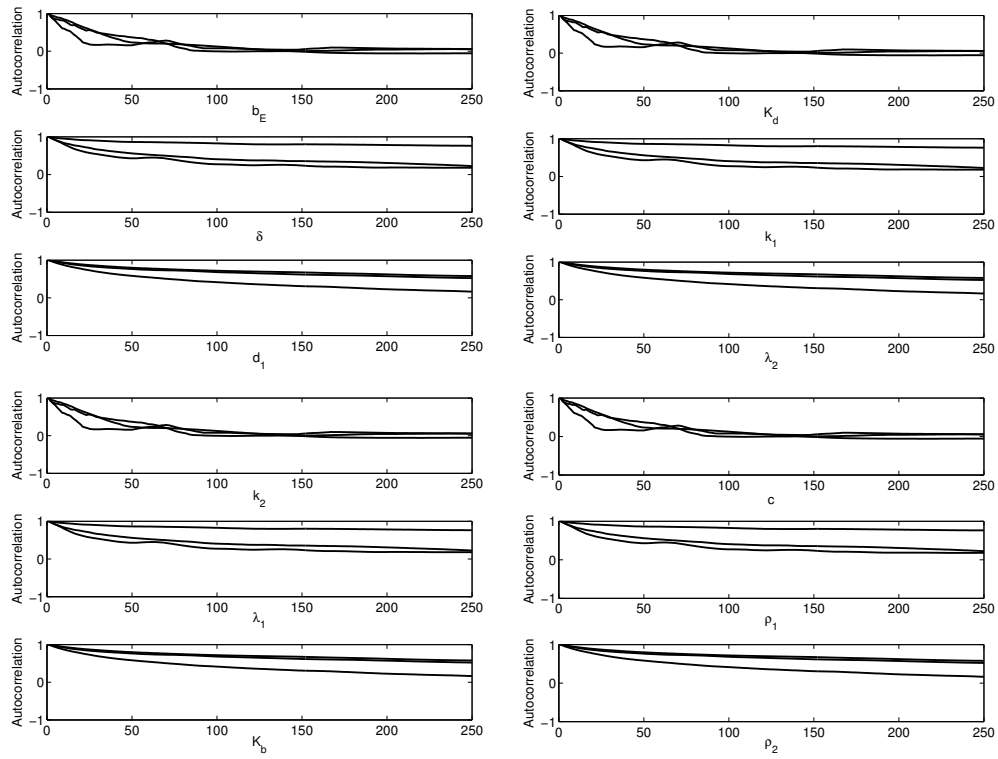


Figure 2.16 Autocorrelation for 12 parameters in DREAM.

arising in bioscience with a wide range of parameter dimensions.

Through the HIV model, we illustrated the process for model calibration and uncertainty quantification using Bayesian inference. Equally importantly, we provided a verification framework by comparing multiple methods. Since sampling methods often lack clear convergence criteria, it is essential to provide verification when discussing convergence to the posterior distribution. Numerical approximation of posterior distributions via the direct method is also possible when the parameter dimension is low.

For future work, we want to show the capabilities of DRAM and DREAM for estimating a subset or a full set of parameters using experimental data. More extensive study on parameter identifiability must also be performed to avoid unidentifiable parameters when performing parameter estimation for a large set of parameters. A final research component will focus on the use of energy statistics [64] to quantify the distance between posterior distributions generated using DRAM, DREAM and direct numerical evaluation to provide a rigorous verification metric.

Remark

M.T. Wentworth and R.C. Smith, “Bayesian model calibration, verification and uncertainty quantification for an HIV model,” submitted, 2014.

Chapter 3

Energy Statistics

3.1 Introduction

In Chapter 2, we estimated subsets of input parameters for an HIV model presented in [8] and verified the accuracy of sampling-based Bayesian model calibration methods, DRAM [33] and DREAM [66, 68]. For this work, we first constructed the posterior densities of b_E, δ and d_1 using DRAM and DREAM. The constructed densities were compared to those constructed by directly evaluating Bayes' formula, which we termed the Direct Method. We then estimated twelve parameters. In this case, due to the dimensionality of the input parameters, the Direct Method is not feasible. The verification of DRAM and DREAM was performed by analyzing convergence statistics, chain evolutions and parameter correlations. Furthermore, we concluded that the densities from DRAM and DREAM are qualitatively same by visually inspecting the densities.

Here, we use energy statistic test discussed in [63, 64] to verify that two densities are essentially equal. The energy test is a statistical hypothesis test to determine whether two sets of samples come from the same distribution by computing the distance between samples. The null hypothesis in this case is that the two sets of samples, whether from the Direct Method, DRAM or DREAM, come from the same distribution.

3.1.1 Energy Distance

To test that two distributions are equal, let X_1, \dots, X_{n_1} and Y_1, \dots, Y_{n_2} be independent random variables. First, we compute the energy distance derived in [63]

$$\epsilon_{n_1, n_2}(X, Y) = \frac{2}{n_1 n_2} \sum_{i=1}^{n_1} \sum_{m=1}^{n_2} |X_i - Y_m| - \frac{1}{n_1^2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_1} |X_i - X_j| - \frac{2}{n_2^2} \sum_{k=1}^{n_2} \sum_{m=1}^{n_2} |Y_k - Y_m| \quad (3.1)$$

and the test statistic

$$T_{n_1, n_2} = \frac{n_1 n_2}{n_1 + n_2} \epsilon_{n_1, n_2}. \quad (3.2)$$

By summing the absolute difference between samples from two distributions, the test statistics indicate the distance between two sets of samples. Smaller test values indicate that samples from two distributions are close and hence it is more likely that the two sets of samples come from the same distribution.

For the hypothesis testing, we must compute critical values. The algorithm for computing the critical values detailed in [63] is the following.

1. Combine the samples from the two sets, X and Y , into one set $W = [XY]$.
2. Randomly choose, without replacement, n_1 entries from W , call it W_1 , and the rest of n_2 entries are in the set W_2 .
3. Compute the test statistic for W_1 and W_2 .
4. Repeat Steps 2 and 3 for a total of M times to obtain M replicates.

The value M is chosen so that $\alpha(1 + M)$ is an integer, where α is the significance level for the test. In our examples in Section 3.3, the value of M is chosen based on the computational capability.

Once M test values are computed, the critical value is chosen to be the test statistic at the $(1 - \alpha) \times 100$ percentile. For example, for $\alpha = 0.1$, the critical value is the test statistic at the 90th percentile. If the test statistic from the original two sets X and Y is smaller than the critical value, we do not reject the null hypothesis. Otherwise, we reject the null hypothesis.

The idea behind this hypothesis testing is the following: if two sets of samples are indeed coming from the same distribution, then the test statistic for the two sets should be comparable to the test statistics computed using randomly separated sets of samples. Assuming that the null hypothesis is true, if we repeat the hypothesis test multiple times, each time with different samples in X and Y , the test statistic would fall below the critical value $(1 - \alpha) \times 100\%$ of the time.

3.2 Samples

The samples from DRAM and DREAM are obtained by running 1,000,000 simulations for DRAM and 1,000,000 total function evaluations with 10 chains for DREAM. To exclude samples during the burn-in period, we select the last 10,000 parameter values as samples to be used in the test. For the Direct Method, we compute the probability density function values at uniformly

distributed parameter values. We use the inverse transform sampling to obtain samples for the Direct Method.

3.2.1 Inverse Transform Sampling

We use the following algorithm to obtain random samples given a set of parameter values and corresponding probability density function.

1. Start with a set of values, $x \in [x_0, x_1]$, and corresponding pdf, $f(x)$, for a distribution D .
2. Compute the cumulative density function $F(x)$ at the discrete values, x .
3. Interpolate the inverse of cdf, F^{-1} , at appropriately distributed points in $[F(x_0), F(x_1)]$.
4. Sample points from a uniform distribution, $x_{new} \sim U(0, 1)$.
5. Compute $y_{new} = F^{-1}(x_{new})$. The points y_{new} are equivalent to samples from the distribution D .

It is noted that in Step 1, one must choose an appropriate range of $[x_0, x_1]$ so that $f(x_0) < \epsilon$ and $|f(x_1) - 1| < \epsilon$ for some ϵ small. Also, in Step 3, one must adjust F^{-1} so that $F^{-1}(0) = 0$ and $F^{-1}(1) = 1$.

3.2.2 Example

As an example, we take the known distribution to be $N(0, 4)$, and compute the pdf at points $[-10 : .1 : 10]$. We construct 10,000 random samples using the inverse transform sampling. In Figure 3.1, we plot the points at which pdf is available, kernel density estimator (kde) of samples from the true distribution $N(0, 4)$ and kde of samples constructed using the pdf. The pdf values, the samples from the distribution and the samples constructed via the cdf all match well.

3.3 Equality of Distributions

3.3.1 Three Parameter Case

We first focus on the three parameter case. Since the density computed using the Direct Method is considered the “true solution,” we perform an energy test to see if the samples from DRAM and DREAM come from distribution constructed via the Direct Method. Here, the null hypothesis H_0 is that the two sets of samples come from the same distribution. For this test, we use 100 samples from each method with $M = 999$ replicates. We use a small number of samples and a comparatively large number of replicates to keep the computational cost reasonable, while

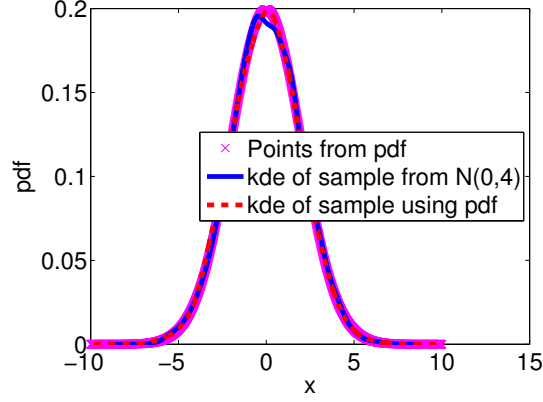


Figure 3.1 Construction of random samples using discrete probability density functions.

approximating the test value as accurately as possible. Moreover, 100 samples are taken randomly from respective chains after burn-in.

The results are summarized in Table 3.1. The replicates used to approximate the critical value and the test statistic are plotted in Figures 3.2 and 3.3. For this experiment, the null hypothesis was not rejected for all cases except b_E with $\alpha = 0.1$. Not rejecting the null hypothesis implies that the samples from the Direct Method and DRAM, as well as from the Direct Method and DREAM, may come from the same distribution with $\alpha = 0.05$ and 0.1 . For δ and d_1 , the test value (solid line) is to the left of the critical values at $\alpha = 0.1$ (dotted line) and at $\alpha = 0.05$ (dashed line). Hence, the distance between the two sets of samples from the Direct Method and DRAM, as well as the Direct Method and DREAM, was smaller than the distance between randomly divided two sets of samples.

By definition, we expect that 95% of the energy test will not reject the hypothesis for $\alpha = 0.05$. Similarly, for $\alpha = 0.1$, we expect that 90% of the energy test will not reject the hypothesis. To show this, we repeated energy tests 950 times for $\alpha = 0.1$ and $\alpha = 0.05$. We summarize the results in Table 3.2. For example, the top left value 8.94 indicates that 8.94% of the energy test rejected the null hypothesis with $\alpha = 0.1$ for the parameter b_E when samples from the Direct Method and DRAM were compared. For the same parameter b_E , 17.47% of 950 energy tests rejected the hypothesis with $\alpha = 0.1$ for the comparison of samples from the Direct Method and DREAM. The percentage of rejection is much more consistent with the definition for DRAM than for DREAM. Also, for most other cases, the percentage of rejection is much higher for DREAM than that for DRAM. This implies that the samples from DRAM and Direct

Table 3.1 Energy test and critical values for $\alpha = 0.1, 0.05$ for DRAM (above) and DREAM (below).

DRAM	Test Value	$\alpha = 0.1$	$\alpha = 0.05$
b_E	1.64e-3	1.55e-3	2.03e-3
δ	1.91e-3	6.52e-3	7.99e-3
d_1	1.21e-3	1.76e-3	2.27e-3

DREAM	Test Value	$\alpha = 0.1$	$\alpha = 0.05$
b_E	1.19e-3	1.71e-3	2.38e-3
δ	1.38e-3	7.16e-3	8.84e-3
d_1	4.68e-4	1.90e-3	2.39e-3

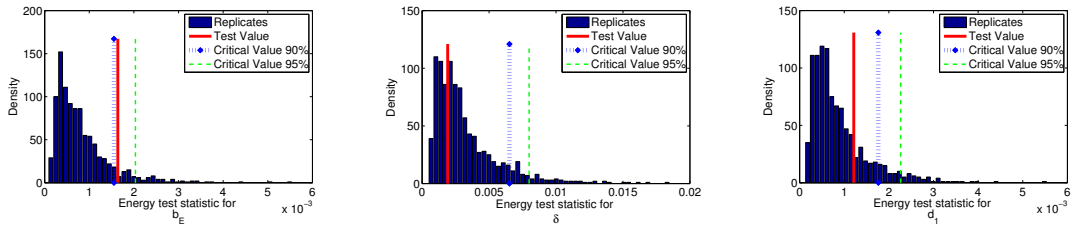


Figure 3.2 Replicates of energy test and critical values between Direct and DRAM.

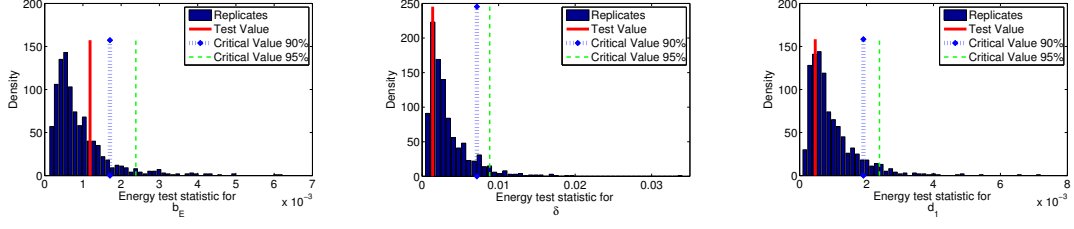


Figure 3.3 Replicates of energy test and critical values between Direct and DREAM.

Table 3.2 Percentage of rejection obtained by repetition of energy test for DRAM and DREAM.

Parameter	$\alpha = 0.1$		$\alpha = 0.05$	
	DRAM	DREAM	DRAM	DREAM
b_E	8.94	17.47	4.32	10.74
δ	8.95	12.53	4.21	6.32
d_1	10.74	9.26	5.47	4.53

are closer than the samples from DREAM and Direct.

Using the hypothesis test, we showed that the samples from the Direct Method and DRAM are very likely coming from the same distribution for the three parameter case. We showed the similar results for the samples from the Direct Method and DREAM. When we repeatedly performed the energy tests, the results using DRAM samples were more consistent with the definition of hypothesis tests. For DRAM, the energy tests rejected the null hypothesis about 5% of the time for $\alpha = 0.05$ and approximately 10% of the time for $\alpha = 0.1$. On the other hand, for DREAM, the test rejected as high as 11% and 17% for $\alpha = 0.05$ and $\alpha = 0.1$, respectively. The difference in the DRAM and DREAM results arises from the fact that the DREAM chains

tend to yield densities that are less smooth. The rough chains in DREAM are apparent when constructing densities from the chain values using a minimum smoothing. Because of this roughness in DREAM chains, the energy tests are likely to determine that the samples from the Direct Method and DREAM are not coming from the same distribution.

3.3.2 12 Parameter Case

Next, we perform the energy tests for the 12 parameters from (2.42). There are no samples from the Direct Method in this case due to the size of the input parameter space. For this test, we take 1000 samples from each DRAM and DREAM and $M = 199$ replicates. The samples are chosen randomly from respective chain evaluations. The results of one energy test are plotted in Figure 3.4. We see that in many cases, we rejected the null hypothesis. For example, the null hypothesis was rejected for all parameters except λ_1 with $\alpha = 0.1$. For $\alpha = 0.05$, we did not reject the hypothesis for the three parameters d_1, λ_1, k_1 .

Based on the energy test using the samples from the chain evolutions, we are unable to show that samples from DRAM and samples from DREAM come from the same distribution. This may be due to the fact that the densities constructed using DREAM tend to be rougher than those using DRAM. During the model calibration process in Chapter 2, more smoothing was required for DREAM than for DRAM to construct parameter densities from chain values using kde for visualization.

For this reason, we next perform the energy test using samples from the constructed densities. After densities are obtained using a kernel density estimator, samples are obtained using the method described in Section 3.2. The results are plotted in Figure 3.5.

The comparison of test values, along with the corresponding critical values, are summarized in Table 3.3. Here, chain and kde respectively represent tests performed using samples from the chains and from the constructed posterior density. The letter R next to the critical values indicates that we rejected the null hypothesis.

The test values are, in general, much smaller for the samples from the posterior density obtained using kde. Furthermore, we rejected the null hypothesis for more parameters using the samples from chain than using the samples from posterior density with kde. For example, at a $\alpha = 0.05$ significance level, the null hypothesis is rejected for all but one parameter using the chain samples, whereas the hypothesis is rejected for seven parameters using the posterior density. However, testing the hypothesis using samples from constructed densities is not statistically rigorous.

For the 12 parameter case, we first performed the energy test using 1000 samples from DRAM chains and 1000 samples from DREAM chains. The null hypothesis was rejected for 9 parameters even at $\alpha = 0.05$. We then performed the test using samples from posterior densities

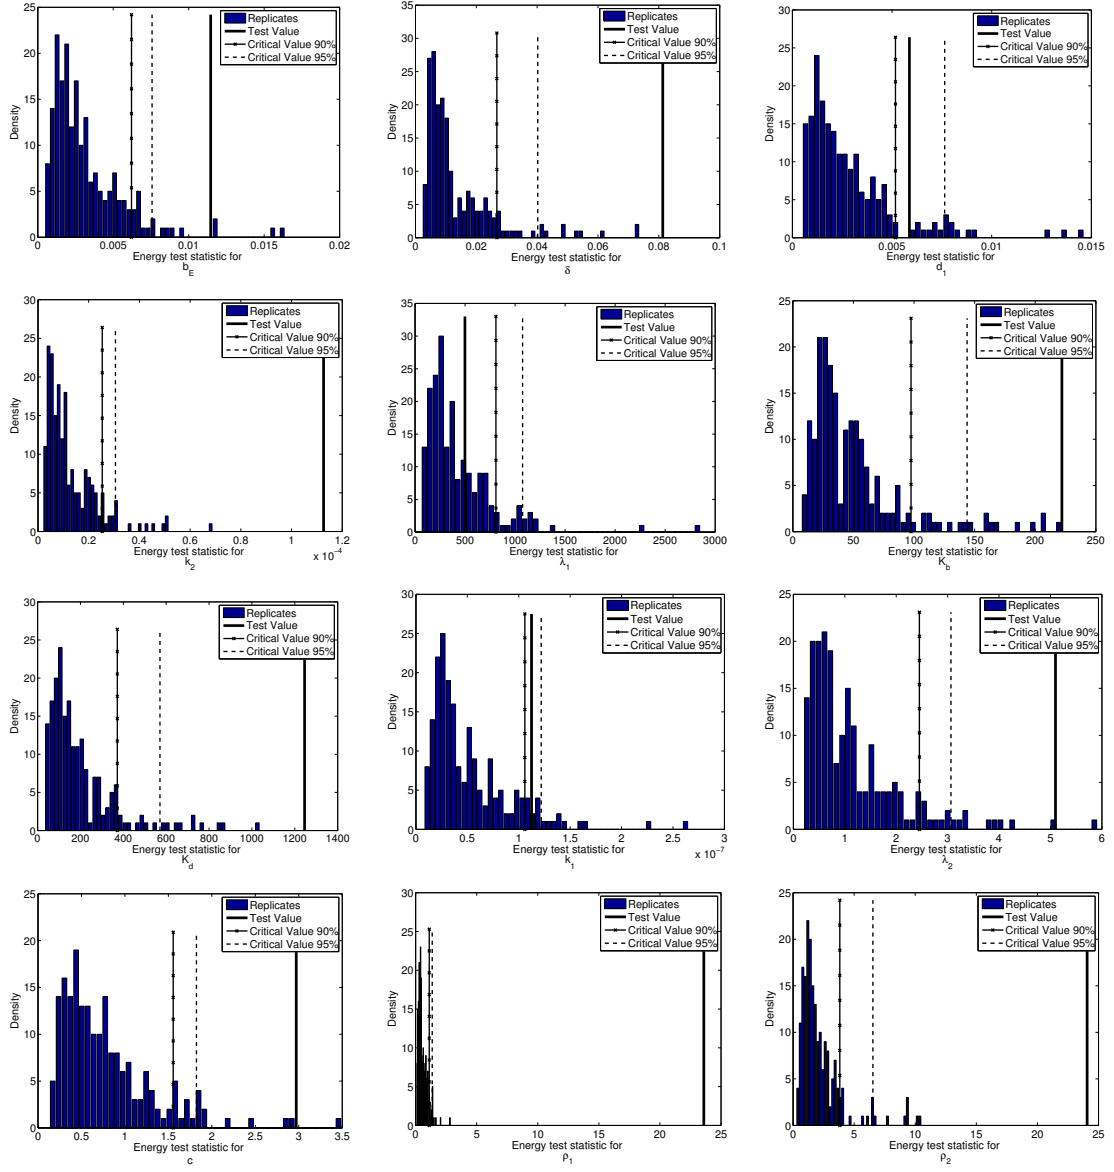


Figure 3.4 Energy test for 1000 samples from DRAM and DREAM with critical values computed with $M = 199$.

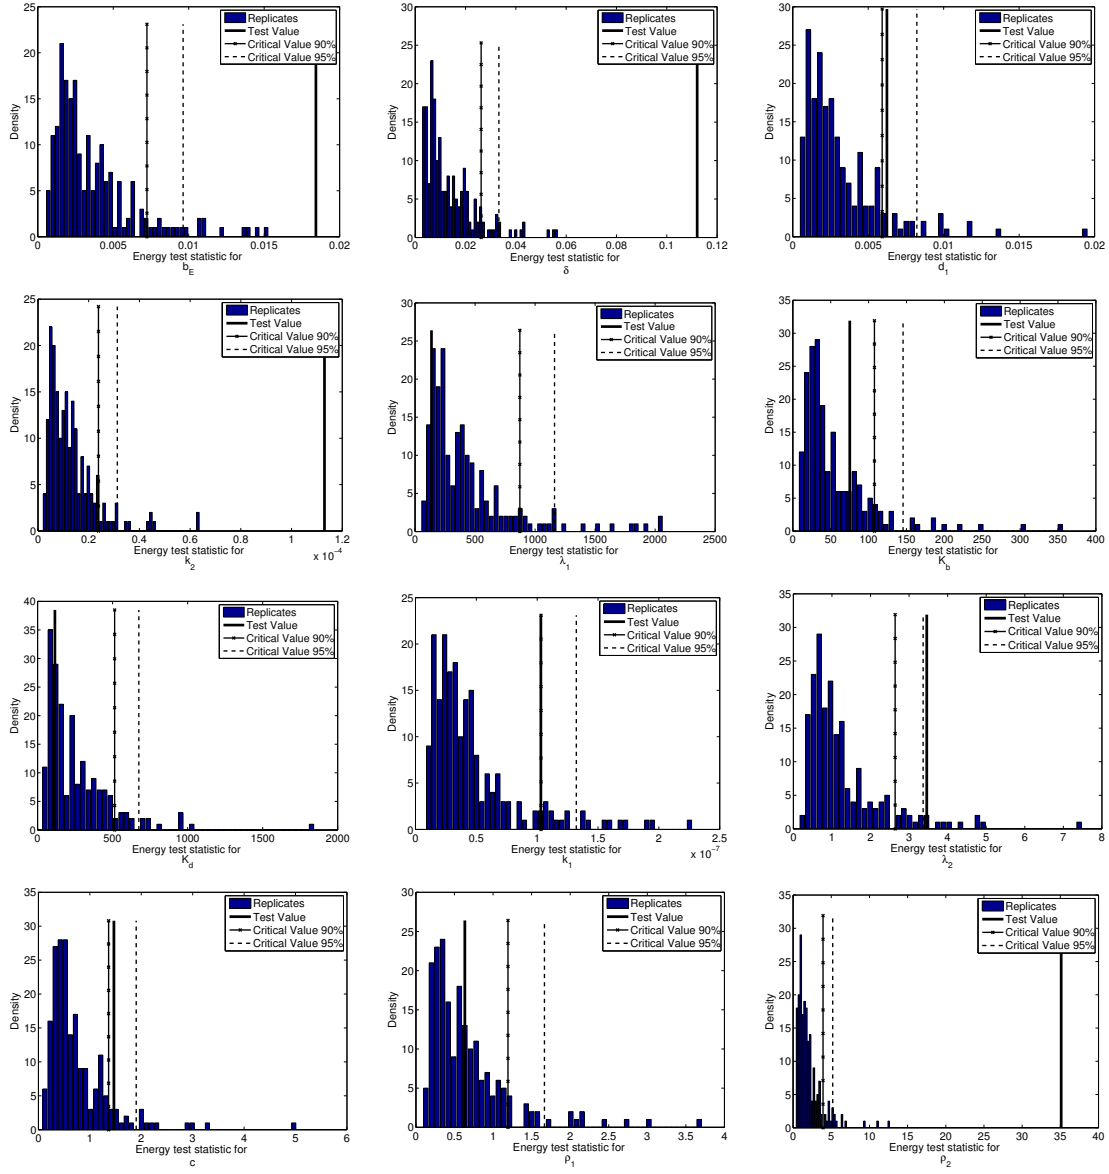


Figure 3.5 Energy test for 1000 samples from DRAM and DREAM with critical values computed with $M = 199$.

Table 3.3 Test and critical values for $\alpha = 0.1, 0.05$ using samples from chain and kde. R indicates the case in which the null hypothesis is rejected.

	Test Value		Critical Value ($\alpha = 0.10$)				Critical Value ($\alpha = 0.05$)			
	chain	kde	chain		kde		chain		kde	
b_E	1.15e-2	1.84e-2	6.21e-3	R	7.24e-3	R	7.58e-3	R	9.63e-3	R
δ	8.13e-2	1.12e-1	2.67e-2	R	2.62e-2	R	4.02e-2	R	3.33e-2	R
d_1	5.87e-3	6.24e-3	5.17e-3	R	5.93e-3	R	7.64e-3		8.22e-3	
k_2	1.13e-4	1.13e-4	2.53e-5	R	2.38e-5	R	3.05e-5	R	3.13e-5	R
λ_1	4.99e+2	1.37e+2	8.08e+2		8.73e+2		1.08e+3		1.16e+3	
K_b	2.22e+2	7.51e+1	9.75e+1	R	1.08e+2		1.44e+2	R	1.45e+2	
K_d	1.25e+3	1.16e+2	3.72e+2	R	5.14e+2		5.71e+2	R	6.75e+2	
k_1	1.12e-7	1.03e-7	1.06e-7	R	1.03e-7		1.22e-7		1.32e-7	
λ_2	5.10e+0	3.46e+0	2.45e+0	R	2.64e+0	R	3.06e+0	R	3.36e+0	R
c	2.97e+0	1.47e+0	1.56e+0	R	1.37e+0	R	1.82e+0	R	1.90e+0	
ρ_1	2.36e+1	6.35e-1	1.12e+0	R	1.19e+0		1.36e+0	R	1.67e+0	
ρ_2	2.41e+1	3.51e+1	3.83e+0	R	3.93e+0	R	6.53e+0	R	5.19e+0	R

constructed with kde. Because of our observation that the parameter values tend to be rougher in DREAM than those in DRAM during the model calibration, sampling from the posterior densities is intended to help mitigate the difference in the nature of methods. Among the 12 parameters, λ_1 was the only parameter for which the null hypothesis was not rejected for all four cases, chain and kde, $\alpha = 0.1$ and $\alpha = 0.05$. This corroborates our observations during model calibration in Chapter 2 since λ_1 was one of the parameters for which the densities from DRAM and DREAM visually matched well. Other parameters did not result in the way we expected based on our observation from the model calibration process.

3.4 Conclusion

Using the energy test, we were not able to show with sufficient confidence that the samples from DRAM and DREAM come from the same distribution for many parameters. Although the use of samples from posterior density improved the chance at which the null hypothesis is not rejected, there were still parameters for which the null hypothesis is rejected. Furthermore, a major issue that remains is the amount of smoothing we performed to construct the posterior densities. While smoothing helps convert the values from chain evolutions to a density, the amount of smoothing we can apply without changing the nature of density is not quantified. More importantly, smoothing is for visualization purpose only, and the energy test using the samples from constructed densities is not statistically rigorous. We are interested in going back to the methods of model calibration and examine the possibility of adjusting the parameter settings, especially for DREAM, so that the densities are smoother without using kde.

Chapter 4

Parameter Selection

4.1 Introduction

Biological and physical models commonly have tens to hundreds of inputs – comprised of parameters, discretized spatially-varying coefficients, initial or boundary conditions, or exogenous forces – many of which have minimal influence on model responses. This necessitates the development of robust analysis techniques to establish subsets or subspaces of influential parameters or inputs. This challenge is exacerbated for models such as neutronics equations, which can have 10^6 inputs, of which only 50-100 are considered influential. The need for robust parameter selection techniques is further motivated by the following objectives: (i) determine those inputs that can be uniquely estimated from measured data; (ii) establish the robustness or fragility of models with respect to certain parameter sets; (iii) simplify models by fixing insensitive inputs; and (iv) guide experimental design by ascertaining parameter subsets or subspaces that have the greatest impact on parameter or response sensitivity.

To establish notation and terminology, we consider the nonlinear input-output relation

$$y = f(q) \tag{4.1}$$

where $q = [q_1, \dots, q_p]$ denotes the model inputs – e.g., parameters, initial or boundary conditions – and f denotes the mathematical model. For this discussion, we consider real-valued responses $y \in \mathbb{R}^1$.

A significant goal of input or parameter selection techniques is to establish subsets or subspaces of inputs or parameters that can be uniquely identified from data or that strongly influence model responses. Such subsets can be characterized by the concepts of identifiable and influential parameter sets.

The concept of *identifiability* is classical and can be defined as follows. The parameters

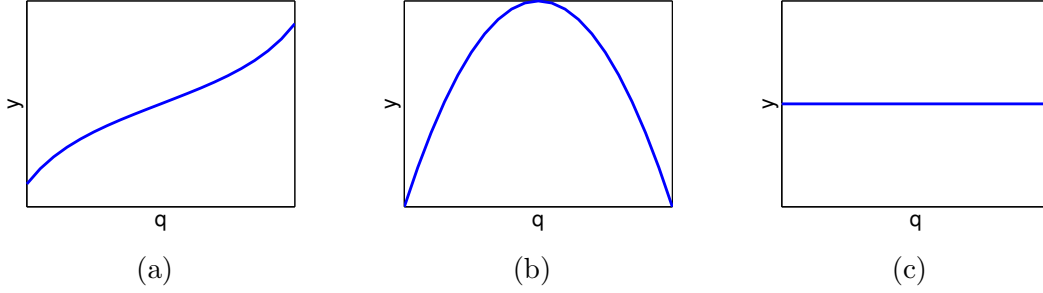


Figure 4.1 Illustration of $y = f(q)$ for (a) identifiable, (b) unidentifiable and (c) noninfluential parameters q .

$q = [q_1, \dots, q_p]$ are identifiable at q^* if $f(q) = f(q^*)$ implies that $q = q^*$ for all admissible $q \in \mathbb{Q}$. The parameters q are identifiable with respect to a space $\mathcal{I}(q)$, termed the identifiable subspace, if this holds for all $q^* \in \mathcal{I}(q)$. Hence identifiable parameters can be uniquely determined from observations. An example of identifiable and unidentifiable parameters are illustrated in Figure 4.1 (a) and (b).

Influential parameter spaces are sometimes defined differently in various disciplines. We define the parameter set $q = [q_1, \dots, q_p]$ to be *noninfluential* on the space $\mathcal{NI}(q)$ if $f(q) - f(q^*) < \varepsilon$ for all q and $q^* \in \mathcal{NI}(q)$. The space $\mathcal{I}(q)$ of influential parameters is defined to be the orthogonal complement of $\mathcal{NI}(q)$. Noninfluential parameters, like unidentifiable parameters, can be fixed for model calibration and uncertainty propagation. Hence, the space of noninfluential parameters is a subspace of the space of unidentifiable parameters. An example of a noninfluential parameter is illustrated in Figure 4.1 (c). Furthermore, parameter q_1 is more influential than parameter q_2 if changes in q_1 affect greater changes in y than changes in q_2 do. See Figure 4.2 for an example of highly and minimally influential parameters. We will quantify the degree of influence using global sensitivity analysis.

For linearly parameterized problems $y = Aq$, it is shown in Chapter 6 of [59] that deterministic and parametrized QR or SVD algorithms can be used to determine subspaces of influential parameters. For the nonlinearly parametrized problems, one typically resorts to global sensitivity analysis or active subspace techniques.

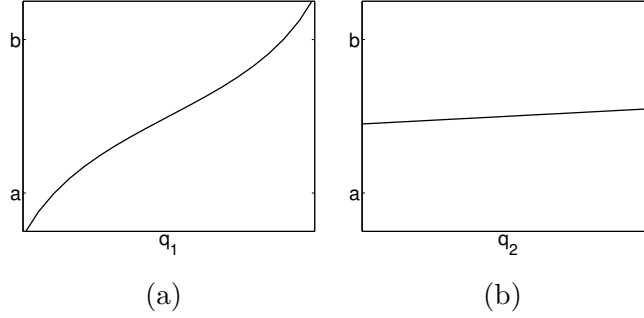


Figure 4.2 Illustration of influential parameters where q_1 is more influential than q_2 .

In this chapter, we focus on global sensitivity analysis and subset selection based on standard errors to determine subsets $q^s = \{q_1^s, \dots, q_p^s\} \subset q = \{q_1, \dots, q_p\}$ of influential parameters. This differs from subspace selection techniques – typically based on QR or SVD algorithms with inputs randomly selected from the admissible input space – which can include linear combinations of inputs [5, 22, 59]. The comparison of active input subspaces with the subset established here for the HIV model constitutes an area of future research.

4.1.1 HIV Model, Inputs and Responses

For the application of parameter selection techniques, we employ the HIV model (1.1). Based on results from [6], we focus on the 15 parameters and initial conditions

$$q = [\lambda_T, d_1, \epsilon_1, k_1, a_T, \epsilon_2, N_T, b_{E2}, a_E, p_E, a_A, p_T, T_1(0), T_1^*(0), T_2(0)] \quad (4.2)$$

whose values tend to be patient specific. Here, the input dimensions is $p = 15$. The associated random variable, considered for global sensitivity analysis, is denoted by Q . Also, we denoted the admissible input space of biologically feasible parameters and initial conditions by \mathbb{Q} . The lower and upper bounds for each parameter, where $q_i \in [\ell b_i, ub_i]$, is summarized in Table 4.1. For more details of the terms and parameters, see [6, 8].

As detailed in [8], data was collected from patients, in a clinical trial, who underwent anti-retroviral therapy (ART) and had at least one ART interruption. The total CD4+ T-cell count/micro L-blood ($T_1 + T_1^* + T_2 + T_2^*$) as well as total RNA copies/mL-plasma ($V_I + V_{NI}$)

Table 4.1 The lower and upper bounds of parameters, $q_i \in [\ell b_i, ub_i]$, for $i = 1, \dots, 15$.

	λ_T	d_1	ϵ_1	k_1	a_T	ϵ_2	N_T	b_{E2}
ℓb_i	3.1	0.11	0.43	4.0e-5	2.0e-4	0.63	65	0.28
ub_i	3.5	0.15	0.6	5.5e-5	2.7e-4	0.78	85	0.45
	a_E	p_E	a_A	p_T	$T_1(0)$	$T_1^*(0)$	$T_2(0)$	
ℓb_i	1.40e-3	0.85	6.5e-5	5	10.5	5.0e-4	720	
ub_i	1.75e-3	1.3	9.0e-5	6.5	13.5	7.0e-4	950	

were recorded during this process.

For global sensitivity analysis, we require a scalar response. At the same time, we are interested in how parameters affect the model output for the feasible input as well for the entire duration of therapy. For these reasons, we choose our scalar model response to be

$$f(q) = \int_0^{1500} T_1(t; q) + T_1^*(t; q) + T_2(t; q) + T_2^*(t; q) dt + \int_0^{1500} V_I(t; q) + V_{NI}(t; q) dt. \quad (4.3)$$

To test the parameter selection techniques, we generate synthetic data using the mean values from the model calibration performed in [8], which are summarized in Table 1.2. The model is solved numerically using `ode15s` in MATLAB.

4.1.2 Previous Work and Chapter Organization

Whereas global sensitivity analysis techniques for parameter selection have not previously been investigated for this dynamic HIV model, certain techniques have been used to analyze other biological models.

Readers are referred to [21] for a case study illustrating the use of sensitivity analysis for a rice model, and [41, 62] for examples of parameter selection in computational and systems biology. The subset selection developed in [7, 9, 20] is applied to the HIV model (1.1) in [6] and we compare our sensitivity-based parameter subsets to those of [6] in Section 4.4.

In Section 4.2, we illustrate the difference between local and global sensitivity analysis using a simple portfolio model. In Section 4.3, we discuss four different techniques for parameter selection. We start with Partial Correlation [4], which quantifies the linear effects of parameters on the

model response. Secondly, we discuss Sobol indices, which are variance-based methods based on a second-order Sobol decomposition. For the HIV example, we discuss the limited accuracy of this decomposition and its affect on parameter selection. Thirdly, we summarize Morris indices using a screening method that ranks parameters in the order of importance. Finally, we discuss the parameter subset selection algorithm discussed in [6]. In Section 4.4, we present our results of applying parameter selection techniques to the HIV model. We interpret the sensitivity indices from each method and provide a comparison for identifying influential parameters. We present verification techniques to illustrate that non-influential parameters should not affect the model output when fixed at nominal values. Finally, we provide comprehensive implications of parameter selection techniques on the HIV model.

4.2 Global Sensitivity Motivation

There are two types of sensitivity analysis: local and global. In literature, sensitivity analysis often refers to local sensitivity analysis, which is typically performed by evaluating the derivative of the response with respect to inputs at nominal input values. On the other hand, global sensitivity analysis considers the effect of parameters over the entire range of input values. Global sensitivity analysis is also used to ascertain how uncertainty in model outputs is apportioned to uncertainties in model inputs; see [55, 57, 59, 60] for details.

We note that global sensitivity techniques rank the relative impact of influential inputs or parameters. Further tests are required to establish that minimally influential parameters are non-influential in the sense defined in Section 2.1.

To illustrate the difference between local sensitivity analysis and global sensitivity analysis, we begin by considering the linear portfolio model

$$Y = c_1 Q_1 + c_2 Q_2 \tag{4.4}$$

considered in [57, 59]. Here, the random variable Y is the return for the investment and $Q_1 \sim N(0, \sigma_1^2)$ and $Q_2 \sim N(0, \sigma_2^2)$ represent hedged portfolios, where c_1 and c_2 are the amounts invested in each portfolio. In this example, we take $c_1 = 2$, $c_2 = 1$, $\sigma_1 = 1$ and $\sigma_2 = 3$. The fact that $\sigma_2 > \sigma_1$ implies that the second portfolio is more volatile than the first. The scatterplots of 1000 joint realizations of q_1, q_2 and y in Figure 4.3 indicate that Q_2 has more influence on Y than Q_1 . Hence, globally, Y is more sensitive to Q_2 than Q_1 .

However, the local sensitivity $s_i = \frac{\partial Y}{\partial Q_i}$ for $i = 1, 2$ yields $s_1 = 2$ and $s_2 = 1$, indicating that q_1 is more sensitive. This reflects the amounts invested in the two portfolios rather than the effects of their volatility of the return. Hence the local sensitivity does not incorporate the nonlinear uncertainty structure over the global admissible parameter space nor the effect of

parameter variability on the response.

In our HIV example, we are interested in how parameters affect the model response in the entire parameter space, rather than at some nominal parameter values. For this reason, we use global sensitivity analysis as a parameter selection technique and isolate influential parameters from noninfluential parameters. In the next section, we discuss three methods of parameter selection based on global sensitivity analysis and one method based on standard errors.

4.3 Parameter Selection Methods

The first of the four parameter selection methods that we discuss is termed the Partial Correlation, or Pearson's Correlations. This method quantifies the linear effect of parameters on the model response. Secondly, we detail the use of Sobol indices based on a variance-based, second-order Sobol decomposition. As an initial step, we examine and verify the accuracy of the second-order expansions. Thirdly, we consider the Morris screening method. We note that this method provides a mechanism of ranking parameters but does not necessarily quantify their relative importance. Finally, we summarize the parameter subset selection detailed in [6]. This method quantifies the importance of parameters by comparing a dimensionless ratio of standard error and mean for each parameter.

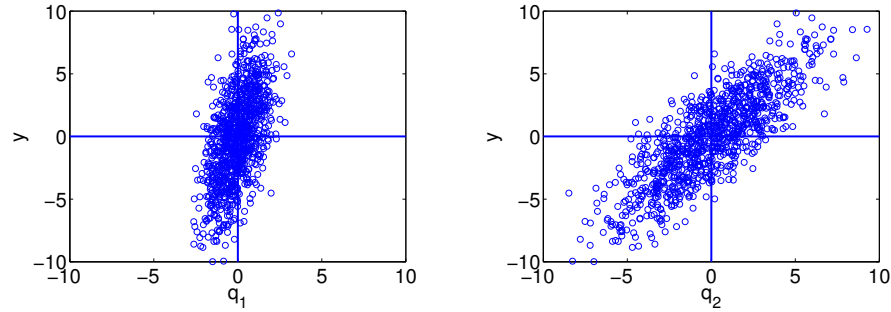


Figure 4.3 Scatterplots of 1000 joint realizations of y versus (a) q_1 and (b) q_2 .

4.3.1 Partial Correlation

We begin by computing partial correlations as detailed in [4]. For two random variables X and Y , the covariance is given by

$$\text{cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}(X))(Y - \mathbb{E}(Y))] = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y). \quad (4.5)$$

The partial correlation is then given by

$$\rho_{XY} = \frac{\text{cov}(X, Y)}{\sigma_X \sigma_Y}. \quad (4.6)$$

The partial correlation quantifies the degree to which two random variables are correlated. For example, $\rho_{XY} = 0$ indicates that X and Y are not correlated. We note that $\rho_{XY} = 0$ does not imply that the two random variables are independent since (4.6) only quantifies linear dependencies between parameters. On the other hand, $\rho_{XY} = \pm 1$ indicates a linear algebraic relation between the variables, in which case they are not jointly identifiable. Values greater than 0.5 generally indicate significant correlations. However, one must study the parameters with partial correlation values less than 0.5 for possible confounding factors or nonlinearities before determining that they are insignificant.

For the HIV example, $X = Q_i$ denotes the random variable for the i^{th} parameter, and Y is the random variable representing the model response. The partial correlation then quantifies the degree of linear correlation between a parameter Q_i and model response Y . We approximate the correlation

$$\rho_{Q_i Y} \approx \frac{\sum_j ((q_i)_j - \bar{q}_i)(y_j - \bar{y})}{\sqrt{\sum_j ((q_i)_j - \bar{q}_i)^2 \sum_k (y_k - \bar{y})^2}}, \quad (4.7)$$

where q_i and y are realizations of Q_i and Y , respectively, and \bar{q}_i and \bar{y} respectively are the sample means of Q_i and Y . The number of function evaluations required to compute the partial correlation for p parameters using M Monte Carlo evaluations is M .

For this method, variables with large partial correlations are considered more influential on the response than those yielding small values of $\rho_{Q_i Y}$. For the portfolio model (4.4), this would reflect the results shown in Figure 4.3, which indicate that Q_2 is more influential than Q_1 .

4.3.1.1 Partial Correlation Results

The partial correlation values are computed for the model (1.1) using $M = 2000$ function evaluations per parameter. The result is plotted in Figure 4.4 to provide visual comparison for

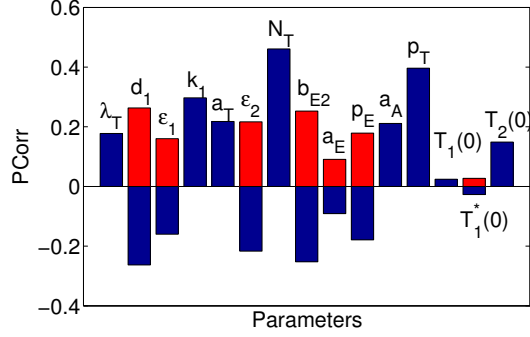


Figure 4.4 Partial correlation of the scalar response to the input parameters.

the overall input-output correlation. Since we are interested in the magnitude of correlation values, the negative correlation values are also shown in the positive direction using red bars.

The result indicates that N_T is most correlated to the model response. Also, p_T and k_1 are more correlated to the model response than other variables. On the other hand, two of the initial conditions, $T_1(0)$ and $T_1^*(0)$ are not correlated to the model response, implying that they have minimum influence.

4.3.2 Sobol Indices

To construct Sobol indices, we assume that parameters have been mapped to $[0, 1]$ and that $q \sim \mathcal{U}[0, 1]^p$. Details regarding the construction of Sobol indices for general densities are provided in [59].

4.3.2.1 Sobol Decomposition

Sobol indices are based on a second-order High Dimensional Model Representation (HDMR) or Sobol representation

$$f(q) \approx f_0 + \sum_{i=1}^p f_i(q_i) + \sum_{1 \leq i < j \leq p} f_{ij}(q_i, q_j). \quad (4.8)$$

Since the representation (4.8) is not unique, additional conditions are imposed to ensure the uniqueness of component functions f_i and f_{ij} . As detailed in [44, 48, 59, 60], each component

function is uniquely specified by minimizing the functional

$$\mathcal{J} = \int_{\Gamma^p} \left[f(q) - \left(f_0 + \sum_i^p f_i(q_i) + \cdots + \sum_{i_1 < \cdots < i_s} f_{i_1, \dots, i_s}(q_{i_1}, \dots, q_{i_s}) \right) \right]^2 dq \quad (4.9)$$

subject to

$$\int_{\Gamma} f_{i_1, \dots, i_s}(q_{i_1}, \dots, q_{i_s}) dq_{i_k} = 0 \quad (4.10)$$

for $k = 1, \dots, s$ and $s = 1, \dots, p$, where $\Gamma^n = [0, 1]^n$ for a positive integer n .

The component functions are given by

$$f_i = \int_{\Gamma^{p-1}} f(q) dq_{\sim i} \quad (4.11)$$

$$f_{ij} = \int_{\Gamma^{p-2}} f(q) dq_{\sim i, j} \quad (4.12)$$

where the notation $dq_{\sim i}$ denotes $dq_1, \dots, dq_{i-1}, dq_{i+1}, \dots, dq_p$.

The variance-based method employs the expansion (4.8) to quantify the contribution of each parameter to the variance of response. As detailed in [59], the total variance of response Y is given by

$$D = \text{var}(Y) = \int_{\Gamma} f^2(q) dq - f_0^2 \quad (4.13)$$

where f_0 is the mean response given by

$$f_0 = \int_{\Gamma} f(q) dq. \quad (4.14)$$

The total variance can be expressed as a sum of variances due to first-order and second-order parameter interactions by expressing D as

$$D = \sum_{i=1}^p D_i + \sum_{1 \leq i < j \leq p} D_{ij} \quad (4.15)$$

where

$$\begin{aligned} D_i &= \int_{\Gamma} f_i^2(q_i) dq_i \\ D_{ij} &= \int_{\Gamma^2} f_{ij}^2(q_i, q_j) dq_i dq_j. \end{aligned} \quad (4.16)$$

The Sobol indices are then defined to be

$$S_i = \frac{D_i}{D}, \quad S_{ij} = \frac{D_{ij}}{D}, \quad i, j = 1, \dots, p. \quad (4.17)$$

Here S_i are often called the importance measures or first-order sensitivity indices, and they measure the contribution of the parameter q_i on the response variance. A large value of S_i implies stronger influence of parameter q_i on the response variance. Similarly, S_{ij} measures the contribution of parameter interactions between q_i and q_j on the response variance. Since the computation of first- and second- order sensitivity indices requires $p + \frac{p(p-1)}{2}$ model responses, we instead consider the total sensitivity indices

$$S_{T_i} = S_i + \sum_{j=1}^p S_{ij} \quad (4.18)$$

which quantify the total effect of the parameter q_i on the response [59].

4.3.2.2 Statistical Interpretation

The Sobol indices, along with the expansion terms and partial variances, have expectation or variance interpretations. Let

$$\mathbb{E}(Y|q_i) = \int_{\Gamma^{p-1}} f(q) dq_{\sim i} \quad (4.19)$$

$$\mathbb{E}(Y|q_i, q_j) = \int_{\Gamma^{p-2}} f(q) dq_{\sim \{ij\}} \quad (4.20)$$

$$(4.21)$$

denote the expected response when q_i and q_i, q_j are fixed. The component functions are

$$\begin{aligned} f_0 &= \mathbb{E}(Y) \\ f_i(q_i) &= \mathbb{E}(Y|q_i) - f_0 \\ f_{ij}(q_i, q_j) &= \mathbb{E}(Y|q_i, q_j) - f_i(q_i) - f_j(q_j) - f_0. \end{aligned} \quad (4.22)$$

As detailed in [59],

$$D_i = \text{var}[\mathbb{E}(Y|q_i)] \quad (4.23)$$

and hence

$$S_i = \frac{\text{var}[\mathbb{E}(Y|q_i)]}{\text{var}(Y)}. \quad (4.24)$$

Similarly, using the equality

$$D_{ij} = \text{var}[\mathbb{E}(Y|q_i, q_j)] - \text{var}[\mathbb{E}(Y|q_i)] - \text{var}[\mathbb{E}(Y|q_j)], \quad (4.25)$$

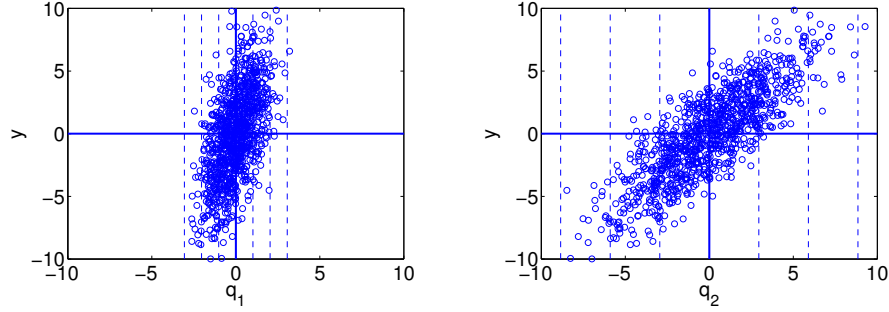


Figure 4.5 Response for fixed values of (a) q_1 and (b) q_2 illustrating $\mathbb{E}(Y|q_i)$ and $\text{var}[\mathbb{E}(Y|q_i)]$.

the total sensitivity index has the variance interpretation

$$S_{T_i} = 1 - \frac{\text{var}[\mathbb{E}(Y|q_{\sim i})]}{\text{var}(Y)} = \frac{\mathbb{E}[\text{var}(Y|q_{\sim i})]}{\text{var}(Y)}. \quad (4.26)$$

The interpretation of $\mathbb{E}(Y|q_i)$ and $\text{var}[\mathbb{E}(Y|q_i)]$ is illustrated in Figure 4.5 from the portfolio example in Section 4.2; see also Chapter 15 of [59]. The conditional expectations for fixed q_1 and q_2 are the average values of Y along vertical slices. Again, we see that mean of response for fixed values of q_2 has more variance than that for fixed values of q_1 .

4.3.2.3 Sobol Indices Algorithm

Since the computation of the Sobol indices requires high-dimensional integration, the indices are approximated numerically. If one uses M Monte Carlo evaluations to approximate the mean $\mathbb{E}(Y|q_i)$ and repeats the procedure M times to approximate the variance $\text{var}[\mathbb{E}(Y|q_i)]$, a total of M^2 evaluations will be required to evaluate a single index. The total number of function evaluations required is M^2p , which is computationally prohibitive for a large parameter dimensions p . This motivated the author of [55] to provide a more efficient algorithm to compute Sobol indices that reduces the required evaluations to $M(p+2)$, based on Sobol's original approach in [60]. The algorithm was further improved by the authors of [56, 61] and is summarized here.

Algorithm

1. Create two sample matrices A and B

$$A = \begin{bmatrix} q_1^1 & \dots & q_i^1 & \dots & q_p^1 \\ \vdots & & & & \vdots \\ q_1^M & \dots & q_i^M & \dots & q_p^M \end{bmatrix}, \text{ and } B = \begin{bmatrix} \hat{q}_1^1 & \dots & \hat{q}_i^1 & \dots & \hat{q}_p^1 \\ \vdots & & & & \vdots \\ \hat{q}_1^M & \dots & \hat{q}_i^M & \dots & \hat{q}_p^M \end{bmatrix}. \quad (4.27)$$

The entries q_i^j and \hat{q}_i^j are quasi-random numbers drawn from the respective densities.

2. Create $A_B^{(i)}$

$$A_B^{(i)} = \begin{bmatrix} q_1^1 & \dots & \hat{q}_i^1 & \dots & q_p^1 \\ \vdots & & & & \vdots \\ q_1^M & \dots & \hat{q}_i^M & \dots & q_p^M \end{bmatrix} \quad (4.28)$$

which is the matrix A except that i^{th} column is taken from B . Similarly, create $B_A^{(i)}$.

3. Create C which is the matrix B appended to matrix A such that

$$C = \begin{bmatrix} A \\ - \\ B \end{bmatrix}. \quad (4.29)$$

The rows of C are linearly independent, and this matrix C is used when estimating the total variance.

4. Compute column vectors $f(A)$, $f(B)$, $f(A_B^{(i)})$ and $f(B_A^{(i)})$ by evaluating the model at input values from the rows of matrices A , B , $A_B^{(i)}$ and $B_A^{(i)}$. Let $f(A)_j$ denote the output computed from the j^{th} row of A . The computation of $f(A)$ and $f(B)$ requires $2M$ model evaluations, whereas the evaluation of $f(A_B^{(i)})$ and $f(B_A^{(i)})$ for $i = 1, \dots, p$ requires $2Mp$ evaluations. The total number of model evaluations is $2M(1 + p)$.
5. Estimate the Sobol indices. The first-order Sobol indices are approximated by

$$S_i = \frac{\frac{1}{M} \sum_{j=1}^M [f(A)_j f(B_A^{(i)})_j - f(A)_j f(B)_j]}{\frac{1}{2M} \sum_{j=1}^{2M} f(C)_j f(C)_j - \langle f(C) \rangle^2} \quad (4.30)$$

and the total Sobol indices are approximated by

$$S_{T_i} = \frac{\frac{1}{2M} \sum_{j=1}^M [f(A)_j - f(A_B^{(i)})_j]^2}{\frac{1}{2M} \sum_{j=1}^{2M} f(C)_j f(C)_j - \langle f(C) \rangle^2} \quad (4.31)$$

where $\langle \cdot \rangle$ denotes the mean.

In the last step, variances are approximated using Monte Carlo approximation. The denominator in (4.30) and (4.31) is the approximation for the total variance with $\mathbb{E}(Y^2) \approx \frac{1}{2M} \sum_{j=1}^{2M} f(C)_j f(C)_j$ and $(\mathbb{E}(Y))^2 = \langle f(C) \rangle^2$. In (4.30), the term $\frac{1}{M} \sum_{j=1}^M f(A)_j f(B_A^{(i)})_j$ approximates $\mathbb{E}(\mathbb{E}(Y|q_i))^2$. In essence, we are taking the mean of responses when all input parameters are varied except q_i . The effect of q_i is fixed since the i^{th} column is the same in both A and $B_A^{(i)}$.

The second term in (4.30),

$$\frac{1}{M} \sum_{j=1}^M f(A)_j f(B)_j, \quad (4.32)$$

represents the squared mean, f_0^2 , using the identity

$$f_0^2 = \int_{\Gamma^2} f(x) f(x') dx dx'. \quad (4.33)$$

This approximation is shown in [61] to reduce the loss of accuracy when computing D , compared to

$$f_0^2 \approx \left(\frac{1}{M} \sum_{j=1}^M f(A)_j \right) \left(\frac{1}{M} \sum_{j=1}^M f(B)_j \right), \quad (4.34)$$

which is used in the previous versions of the algorithm.

The computation of S_{T_i} follows from the derivations in [37], which uses the approximation

$$\mathbb{E}[\text{var}(Y|q_{\sim i})] \approx \frac{1}{2M} \sum_{j=1}^M [f(A)_j - f(A_B^{(i)})_j]^2 \quad (4.35)$$

instead of the approximation

$$\text{var}[\mathbb{E}(Y|q_{\sim i})] \approx \frac{1}{M} \sum_{j=1}^M f(A)_j f(A_B^{(i)})_j - f_0^2 \quad (4.36)$$

in (4.26). The comparison of different versions of the algorithm can be found in [56].

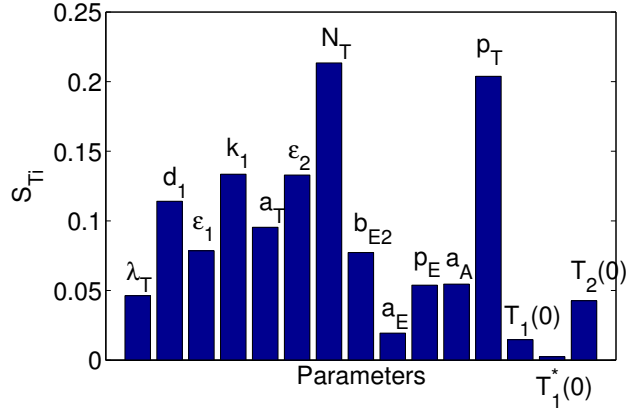


Figure 4.6 Sobol indices S_{T_i} for 15 parameters.

4.3.2.4 Sobol Indices Results

The Sobol indices for the 15 parameters in (4.2) are plotted in Figure 4.6. It is clear that N_T has the largest S_{T_i} value, indicating that N_T affects the model output the most. We note that p_T is also almost as significant. On the other hand, $T_1^*(0)$ affects the output the least. The parameters a_E and $T_1(0)$ are also very insignificant.

4.3.2.5 Verification of the Sobol Decomposition

Since the accuracy of the Sobol indices depends on the accuracy of the approximated second-order Sobol representation, we test whether the function is accurately approximated by the second-order Sobol decomposition.

To ensure that we can adequately approximate the integrals, we consider four parameters $q = [\lambda_T, d_1, \epsilon_1, k_1]$ with values in the 4-D hypercube $[3.1, 3.5] \times [0.11, 0.15] \times [0.43, 0.60] \times [4e-5, 5.5e-5]$. We compute the model response using $n = 41$ equally-spaced quadrature points in each dimension to evaluate the integrals (4.11) and (4.12). The function is expanded with a zero-th, first, second, third and fourth order component functions so that

$$f = f_0 + \sum_i f_i(q_i) + \sum_{i < j} f_{ij}(q_i, q_j) + \sum_{i < j < k} f_{ijk}(q_i, q_j, q_k) + \sum_{i < j < k < h} f_{ijkh}(q_i, q_j, q_k, q_h). \quad (4.37)$$

In Figure 4.7, we plot the model response along with first- and second-order approximations, where the fixed parameter values are taken to be $\lambda_T = 3.19$, $\epsilon_1 = 0.119$, $d_1 = 0.46825$, $k_1 = 4.3375\text{e-}5$. The model response is represented by the blue solid line, while the first-order approximation and the second-order approximation are represented by dashed-dot black and by dashed red, respectively.

We note that both the first- and second-order approximations smooth out the jumps in the model response, and they do not accurately represent the model response. There is little difference between the first-order and second-order approximations, which explains the similarity between the reported values of S_i and S_{T_i} . In the HIV model (1.1), the higher order interactions are non-negligible, and the second-order approximation is not sufficiently accurate to completely represent the model response. This may introduce some inaccuracy when determining the relative influence of parameters using the Sobol indices.

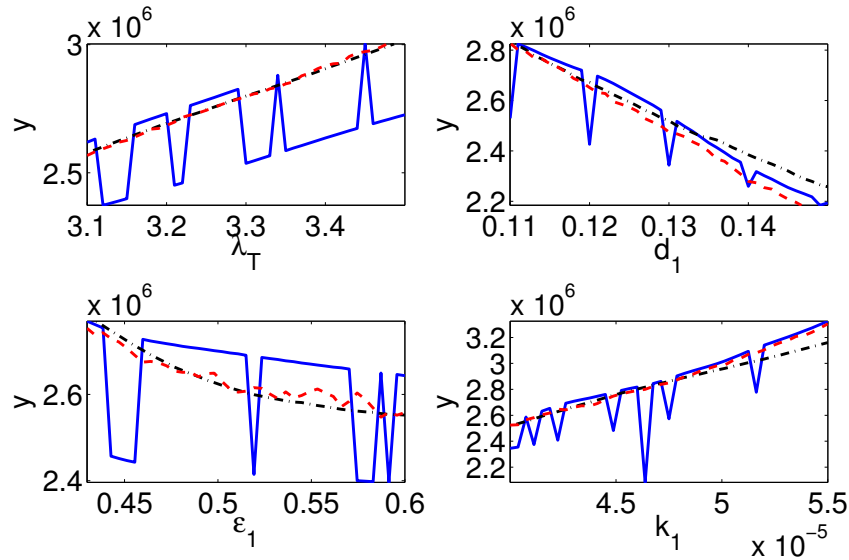


Figure 4.7 Model response (solid blue), 1st order (dash-dot black) and 2nd order approximation (dash red).

4.3.3 Morris Screening

The third method we consider is Morris screening [46, 55]. Screening methods rank the importance of parameters by averaging coarse difference relations termed elementary effects. The elementary effects are then used to compute sensitivity measures. The mean and variance of elementary effects represent the linear effect of parameters and the effect of interaction terms on the model response, respectively. Morris Screening employs neighbors to compute elementary effects, which reduces the total model evaluations by approximately a half. Whereas Morris Screening can only rank the parameter importance, and does not quantify the relative importance of each parameter, this method is significantly more efficient than computing Sobol indices. More details regarding the method can be found in [18, 59].

As with Sobol indices, we first map parameters to $[0, 1]$. We also assume no prior information about parameters and hence take them to be uniformly distributed. This latter assumption can be modified if prior parameter information is available. The elementary effect is given by

$$d_i(q) = \frac{f(q_1, \dots, q_{i-1}, q_i + \Delta, q_{i+1}, \dots, q_p) - f(q)}{\Delta} = \frac{f(q + e_i) - f(q)}{\Delta}, \quad (4.38)$$

where Δ is the step size chosen from the set $\Delta \in \left\{ \frac{1}{\ell-1}, \dots, 1 - \frac{1}{\ell-1} \right\}$. Constructed in this way, d_i quantifies the approximate, large scale, local sensitivity of the model response relative to the i^{th} parameter. We note that the step size is taken large to cover the entire parameter space. As detailed in [18, 46, 59], taking ℓ to be even and choosing $\Delta = \frac{\ell}{2(\ell-1)}$ has the advantage that it guarantees equal probability sampling from the distribution.

Let

$$d_i^k = \frac{f(q^k + \Delta e_i) - f(q^k)}{\Delta} \quad (4.39)$$

be the elementary effect associated with the i^{th} parameter and k^{th} sample. For r sample points, the Morris indices for the parameter q_i are

$$\begin{aligned} \mu^* &= \frac{1}{r} \sum_{k=1}^r |d_s^k| \\ \sigma^2 &= \frac{1}{r-1} \sum_{k=1}^r (d_s^k - \mu)^2, \text{ where } \mu = \frac{1}{r} \sum_{k=1}^r d_s^k. \end{aligned} \quad (4.40)$$

The mean quantifies the individual effect of the input on output, whereas the variance incorporates the influence of parameter interactions. Since we must consider both the mean and the variance, we rank the parameter using the quantity $\sqrt{\mu^{*2} + \sigma^2}$ when ordering the importance of parameters. Computing (4.39) requires two model evaluations per parameter per sample. Hence, a total of $2pr$ model evaluations is required to compute the Morris indices, μ^*

and σ^2 . As detailed in Algorithm 3.3.1 below, taken from [18], one employs neighbors to reduce the number of total model evaluations to $(p+1)r$.

4.3.3.1 Morris Screening Algorithm

1. Create a $(p+1) \times p$ matrix A with ones in the lower triangle such that

$$A = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ \vdots & & \ddots & \\ 1 & 1 & \dots & 1 \end{bmatrix}. \quad (4.41)$$

2. Choose the step size Δ . Unless specified by the user, take $\Delta = \frac{\ell}{2(\ell-1)}$.
3. Select a starting vector q^* .
4. Construct a diagonal matrix D^* , whose entries are randomly chosen from $\{-1, 1\}$.
5. Calculate the sampling matrix A_s as the following

$$A_s = J_{p+1,p}q^* + \frac{\Delta}{2} [(2A - J_{p+1,p})D^* + J_{p+1,p}]P^*, \quad (4.42)$$

where $J_{i,j}$ is a $i \times j$ matrix with all ones and P^* is a $p \times p$ permutation of the identity matrix.

6. If the parameters are not defined in the hypercube $[0, 1]^p$ and instead $q \in [\ell b_i, ub_i]$ for $i = 1, \dots, p$, take $\ell b = [\ell b_1, \ell b_2, \dots, \ell b_p]$ and $ub = [ub_1, ub_2, \dots, ub_p]$. The sampling matrix is then scaled to match the range of parameters

$$C = J_{p+1,1}\ell b + A_s(D(ub - \ell b)) \quad (4.43)$$

where $D(ub - \ell b)$ is a diagonal matrix with entries $ub - \ell b$.

7. Compute the elementary effect for $s = 1, \dots, p$. We let C_k denote the k^{th} row of C . Then

$$d_s = \frac{f(C_i) - f(C_j)}{\Delta}, \quad (4.44)$$

where i and j denote the indices such that i^{th} row and j^{th} row differ in the s^{th} entry.

8. Repeat the steps 1 – 7 for r samples. The Morris mean μ^* and σ^2 are computed by taking the average of the local elementary effect

$$\begin{aligned}\mu^* &= \frac{1}{r} \sum_{k=1}^r |d_s^k| \\ \sigma^2 &= \frac{1}{r-1} \sum_{k=1}^r (d_s^k - \mu)^2, \text{ where } \mu = \frac{1}{r} \sum_{k=1}^r d_s^k.\end{aligned}\tag{4.45}$$

We note that the denominator of (4.44) in Step 7 is Δ for all $q_i, i = 1, \dots, p$. The elementary effects must be computed using the scaled step size, even though model responses are computed at the parameter values, which are mapped using (4.43).

4.3.3.2 Morris Indices Results

We use $\ell = 20$, $r = 50$ and the default step size $\Delta = \ell/2(\ell - 1)$. We plot the elementary effects μ^* and σ^2 in Figure 4.8 to visualize those parameters that are more influential. The most influential parameter is again N_T followed by p_T and ϵ_2 . The results also coincide with those from Partial Correlation and Sobol for the least influential parameters, which are $T_1^*(0)$ and $T_1(0)$. The parameter a_E , which is one of the least influential parameters in Partial Correlation and Sobol after $T_1^*(0)$ and $T_1(0)$, is still ranked low. One difference with Morris screening, however, is that all three initial conditions are identified as three least influential parameters.

4.3.4 Parameter Subset Selection

Finally, we discuss the parameter subset selection method presented in [6, 7, 9, 20]. This method can be used to determine a subset of n_p parameters, $n_p \leq p$, that are identifiable with the smallest uncertainty measure. The subset selection algorithm uses the optimal parameter estimates as well as standard errors associated with the parameters in the estimation process. We consider a ratio of standard errors and parameter estimates to rank the set of parameters that are most influential for a given n_p . The parameter subset selection can be used as a parameter selection technique since identifiable parameters are influential, and the n_p parameters isolated in this algorithm correspond to the n_p most influential parameters.

For a vector of parameters $q = [q_1, \dots, q_p]$, we first require the optimal parameter estimates of q , denoted by $\hat{q} = [\hat{q}_1, \dots, \hat{q}_p]$, and the corresponding standard errors, $SE = [SE_1, \dots, SE_p]$. Then, given the parameter vector q of size p and a number $n_p \leq p$, the subset selection algorithm returns a set of parameters of size n_p that minimizes the selection score,

$$\alpha(\hat{q}) = |\nu(\hat{q})|.\tag{4.46}$$

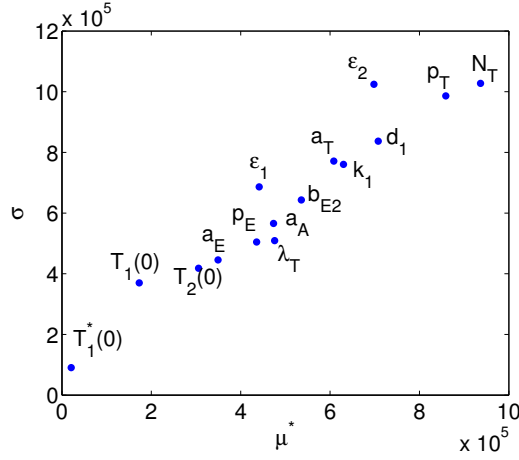


Figure 4.8 Morris μ^* and σ computed using the Morris screening algorithm.

Here, $\nu(\hat{q}) = [\nu(\hat{q}_1), \dots, \nu(\hat{q}_{n_p})]^T$, and $\nu(\hat{q}_i)$ is the coefficient of variation for \hat{q}_i defined by

$$\nu(\hat{q}_i) = \frac{SE_i}{\hat{q}_i}, \quad i = 1, \dots, n_p. \quad (4.47)$$

The set of parameters with the smallest selection score gives n_p most influential parameters.

4.3.4.1 Optimal Parameter Estimates

This technique utilizes time-dependent responses. Following the strategy in [6], we employ the responses

$$\begin{aligned} z_1 &= T_1 + T_1^* + T_2 + T_2^* \\ z_2 &= V_I + V_{NI}, \end{aligned} \quad (4.48)$$

which are the total CD4+ T-cells and the total RNA copies, respectively. We assume a statistical model of the form

$$\begin{aligned} Y_1^i &= z_1(t_1^i; q_0) + e_1^i, \quad i = 1, 2, \dots, N_1 \\ Y_2^j &= z_2(t_2^j; q_0) + e_2^j, \quad j = 1, 2, \dots, N_2, \end{aligned} \quad (4.49)$$

where y_1^i and y_2^j are realizations of the random variables Y_1^i and Y_2^j , respectively, and e_1^i and e_2^j are independently identically distributed such that $\mathbb{E}[e_1^i] = \mathbb{E}[e_2^j] = 0$ with $\text{Var}(e_1^i) = \sigma_1^2$ and $\text{Var}(e_2^j) = \sigma_2^2$ for $i = 1, \dots, N_1$, $j = 1, \dots, N_2$. Also, q_0 represents the hypothesized true parameter values.

The weighted least squares estimator is given by

$$\hat{q} = \arg \min_{q \in Q} \left(\frac{1}{N_1} \sum_{i=1}^{N_1} \frac{(y_1^i - z_1(t_1^i; q_0))^2}{\sigma_1^2} + \frac{1}{N_2} \sum_{j=1}^{N_2} \frac{(y_2^j - z_2(t_2^j; q_0))^2}{\sigma_2^2 z_2^{2\gamma}(t_2^j; q_0)} \right) \quad (4.50)$$

where the variance components are given by

$$\begin{aligned} \sigma_1^2(q_0) &= \frac{1}{N_1 - \dim(q_0)} \sum_{i=1}^{N_1} (y_1^i - z_1(t_1^i; q_0))^2 \\ \sigma_2^2(q_0) &= \frac{1}{N_2 - \dim(q_0)} \sum_{j=1}^{N_2} \frac{(y_2^j - z_2(t_2^j; q_0))^2}{z_2^{2\gamma}(t_2^j; q_0)}. \end{aligned} \quad (4.51)$$

The value of γ is determined based on the underlying assumption for the statistical models (4.49). More specifically, it was determined in [6, 10] that choosing $\gamma = 1.2$ results in the residuals being approximately iid, which is an assumption for the model (4.49). For this reason, the parameter estimation was performed with $\gamma = 1.2$.

Since the estimates in (4.50) and (4.51) involve an unknown, to-be-estimated parameter vector q_0 , the optimal parameter is estimated iteratively with the initial variance $\sigma_k^2 = 1$ for $k = 1, 2$ and the weights $z_2^{2\gamma}(t_2^j; q_0) = 1$ for $j = 1, \dots, N_2$. We summarize the parameter estimation algorithm from [6].

Parameter Estimation Procedure Algorithm

1. Obtain initial estimate $\hat{q}^{(0)}$ using (4.50) with $\sigma_k^2 = 1$ for $k = 1, 2$ and the weights $z_2^{2\gamma}(t_2^j; q_0) = 1$ for $j = 1, \dots, N_2$.
2. Compute the variances σ_k^2 using (4.51), and the weights $z_2^{2\gamma}(t_2^j; q_0)$ with q_0 replaced by $\hat{q}^{(0)}$.
3. Initialize the iteration counter ℓ with the value 1.
4. Do each of the following:
 - Compute $\hat{q}^{(\ell)}$ using (4.50) with current variances σ_k^2 and weights $z_2^{2\gamma}(t_2^j; \hat{q}^{(\ell-1)})$.
 - Update the variances σ_k^2 using (4.51) and the weights $z_2^{2\gamma}(t_2^j; \hat{q}^{(\ell-1)})$ with $q_0, \hat{q}^{(\ell-1)}$ replaced by $\hat{q}^{(\ell)}$.

- Compute $\Delta_\varepsilon = ||[\hat{q}^{(\ell)} - \hat{q}^{(\ell-1)}] ./ [\hat{q}^{(\ell-1)}]||$.
- Increment ℓ by 1.

5. If $\Delta_\varepsilon > \varepsilon$, go back to Step 4. Otherwise, terminate the algorithm.

In this algorithm, ε is a user-defined threshold tolerance for a termination criterion, and $./$ denotes element-by-element division.

4.3.4.2 Computing Standard Errors

The parameter subset selection algorithm also requires the computation of standard errors for the parameters. The standard errors are computed using standard asymptotic theory for generalized least squares (GLS) estimators q_{GLS}^n following the procedure discussed in [6]. The $p \times p$ Fisher Information Matrix (FIM) corresponding to z_1 and z_2 in (4.48) is approximated by

$$\Sigma_0^{N_1+N_2} \approx \left[\left(\sum_{i=1}^{N_1} \frac{1}{\sigma_1^2(\hat{q}^n)} \frac{\partial z_1(t_1^i; \hat{q}^n)}{\partial q_k} \frac{\partial z_1(t_1^i; \hat{q}^n)}{\partial q_\ell} + \sum_{j=1}^{N_2} \frac{1}{\sigma_2^2(\hat{q}^n) z_2^{2\gamma}(t_2^j; \hat{q}^n)} \frac{\partial z_2(t_2^j; \hat{q}^n)}{\partial q_k} \frac{\partial z_2(t_2^j; \hat{q}^n)}{\partial q_\ell} \right) \right]_{k,\ell} \quad (4.52)$$

where σ_1^2 and σ_2^2 are defined in (4.51) with q_0 approximated by \hat{q}^n .

To approximate q_0 , we first let $z_1 = T_1 + T_1^* + T_2 + T_2^*$ and $z_2 = V_I + V_{NI}$. The sensitivities are computed by solving the system of equations

$$\frac{d}{dt} \left(\frac{\partial z_m}{\partial q} \right) = \frac{\partial g_m}{\partial x} \left(\frac{\partial x}{\partial q} \right) + \frac{dg_m}{dq}, \quad m = 1, 2. \quad (4.53)$$

Here, x and q respectively denote the state variables and the parameters being estimated. Define the $2 \times p$ matrices

$$\begin{aligned} D_1^i(q_0) &= \begin{bmatrix} \frac{\partial z_1}{\partial q_1}(t_1^i; q_0) & \cdots & \frac{\partial z_1}{\partial q_p}(t_1^i; q_0) \\ 0 & \cdots & 0 \end{bmatrix} \quad \text{for } i = 1, \dots, N_1 \\ D_2^i(q_0) &= \begin{bmatrix} 0 & \cdots & 0 \\ \frac{\partial z_1}{\partial q_1}(t_1^i; q_0) & \cdots & \frac{\partial z_1}{\partial q_p}(t_1^i; q_0) \end{bmatrix} \quad \text{for } i = 1, \dots, N_2. \end{aligned} \quad (4.54)$$

and define the 2×2 matrix

$$V_0(t; q_0) = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 z_2^{2\gamma}(t; q_0) \end{bmatrix}. \quad (4.55)$$

The matrices $D_1^i V_0^{-1}(t_1^i) D_1^i$ and $D_2^j V_0^{-1}(t_2^j) D_2^j$ respectively have entries

$$\begin{aligned} F_{k,\ell}^{1,i}(q_0) &= \sigma_1^{-2} \frac{\partial z_1}{\partial q_k}(t_1^i; q_0) \frac{\partial z_1}{\partial q_\ell}(t_1^i; q_0), \quad k, \ell = 1, \dots, p, \quad i = 1, \dots, N_1 \\ F_{k,\ell}^{2,j}(q_0) &= \sigma_2^{-2} z_2^{-2\gamma}(t_2^j; q_0) \frac{\partial z_2}{\partial q_k}(t_2^j; q_0) \frac{\partial z_2}{\partial q_\ell}(t_2^j; q_0), \quad k, \ell = 1, \dots, p, \quad i = 1, \dots, N_2 \end{aligned} \quad (4.56)$$

Then, we define the $p \times p$ Fisher matrix $F(q_0) = F_{k,\ell}(q_0)$ with entries

$$F_{k,\ell}(q_0) = \sum_{i=1}^{N_1} F_{k,\ell}^{1,i}(q_0) + \sum_{j=1}^{N_2} F_{k,\ell}^{2,j}(q_0). \quad (4.57)$$

The approximate Fisher matrix (4.52) is obtained by evaluating (4.57) at $\hat{q}^n \approx q_0$. Using the Fisher matrix approximations, F , the standard errors for \hat{q}_k^n , $k = 1, \dots, p$, are given by

$$SE_k = SE(\hat{q}_k^n) = \sqrt{(F^{-1}(\hat{q}^n))_{k,k}}. \quad (4.58)$$

It is illustrated in [59] that the standard errors are related to the variance of parameter estimates so they quantify the uncertainty of each parameter. Parameters with small standard errors are estimated with a high degree of certainty, so one can conclude that their impact on the response is influential. On the other hand, parameters that are noninfluential have minimal impact on responses, which yields more uncertainty and larger standard error when estimating optimal parameter values.

4.3.4.3 Parameter Subset Selection Results

As presented in [6], we compile the parameters that give the smallest selection score for a given number of parameters in the set, n_p , in Table 4.2. We note that these results are patient-dependent and N_T was not in the top three for the considered patient. For other patients, though not shown here, N_T is in the top 3. See [6] for more details on patient-dependent parameters.

For example, if we want a subset of three parameters that are most influential, we select λ_T, ϵ_2 and p_T . In this way, the parameter subset selection algorithm selects a set of parameters for a given value of n_p ; however, it does not specify which parameter is the most influential among the selected parameters. We see that the set for $n_p = k$ is a subset for $n_p = k + 1$ for all k , except $k = 2$. Unlike Partial Correlation, Sobol indices and Morris indices, Parameter Subset Selection relies on a local sensitivity approach since the sensitivity matrices are computed around the mean values. Nevertheless, we can use the parameter subset selection result to provide a comparison regarding which parameters to include when we specify a number of parameters to choose from the entire set.

Table 4.2 Parameter subset selection results from [6].

n_p	N_T	λ_T	ϵ_2	p_T	p_E	$T_2(0)$	$T_1(0)$	ϵ_1	d_1	b_{E2}	a_E	a_T	k_1	a_A	$T_1^*(0)$
1	x														
2	x	x													
3		x	x	x											
4	x	x	x	x											
5	x	x	x	x	x										
6	x	x	x	x	x	x									
7	x	x	x	x	x	x	x								
8	x	x	x	x	x	x	x	x							
9	x	x	x	x	x	x	x	x	x						
10	x	x	x	x	x	x	x	x	x	x					
11	x	x	x	x	x	x	x	x	x	x	x				
12	x	x	x	x	x	x	x	x	x	x	x	x			
13	x	x	x	x	x	x	x	x	x	x	x	x	x		
14	x	x	x	x	x	x	x	x	x	x	x	x	x	x	
15	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

4.4 Comparison and Verification

In this section, we illustrate two techniques for verifying the accuracy of the parameter selection techniques. We first verify the results provided by the global sensitivity techniques, which rank the impact or influence of the inputs, and the parameter subset selection. We do this in Section 4.4.1 by comparing the input rankings provided by the four methods. In Section 4.4.2, we verify the noninfluential inputs by comparing responses obtained with various input combinations.

4.4.1 Verification of Input Rankings

Here, we provide comparisons of the four methods. First, we summarize in Table 4.3 sensitivity measures, a description and the computational cost of each method. The Sobol indices, Morris indices and Partial Correlation indices are summarized in Table 4.4 in the order of importance. For Partial Correlation, we rank the importance by the absolute values of the partial correlation. For the Sobol indices, the parameters are ranked by the magnitude of S_{T_i} . For the Morris indices, we consider the quantity $\sqrt{\mu^{*2} + \sigma^2}$ to rank the parameters.

To provide a comparison among the four methods, we summarize in Tables 4.5, 4.6, 4.7 and 4.8, the parameters to be selected for a given number of parameters. Here, P_{corr} , S , M and P_{SS} respectively denote Partial Correlation, Sobol indices via Saltelli algorithm, Morris indices and Parameter Subset Selection. For Partial Correlation, Sobol indices and Morris indices, n_p influential parameters correspond to the top n_p parameters from Table 4.4.

Overall, Partial Correlation is the cheapest method to measure linear correlation between parameters and response. This often corresponds to the first order Sobol indices. Computing Sobol indices is expensive and it becomes prohibitively slow as the number of input parameters increases. For a model with a moderate number of input parameters, we can apply Morris screening. This employs neighbors to compute statistically averaged local, very coarse approximations to derivatives. Morris indices are a good measure to isolate influential parameters from noninfluential parameters with much fewer evaluations than Sobol indices. Finally, the parameter subset selection algorithm provides sensitivity in terms of uncertainties involved in the estimation process. The noninfluential parameters determined by this method did not match the results from the other three.

In terms of accuracy, Sobol indices measure the first- and second-order interaction effects of parameters most accurately. However, we showed that second-order Sobol decomposition may not be sufficiently accurate depending on the model. Even though the Sobol indices are widely used for global sensitivity analysis, one must always consider the accuracy of Sobol decomposition as an approximation to the model before applying the results of Sobol indices.

When the parameter selection techniques are applied to the HIV model, we found that

Table 4.3 Summary of parameter selection techniques.

Method	Sensitivity Measure	Description	Cost
Partial Correlation	Degree of linear correlation between parameters and response	Ranks the parameters in the order of strong linear correlation to response. Considers linearity only.	M model evaluations for M Monte Carlo samples.
Sobol by Saltelli	First order Sobol S_i and total sensitivity indices S_{Ti}	A type of variance-based method. Uses 2nd order Sobol decomposition. Ranks the parameters and quantifies relative importance. Measures the effects of individual parameters as well as interaction terms.	$2M(p + 1)$ model evaluations for M Monte Carlo samples and p parameters.
Morris Screening	Mean μ_i^* and variance σ^2 of elementary effects	Averages coarse local derivative approximations. Only ranks parameters. Employs neighbors to reduce the cost.	$(p + 1)r$ model evaluations for r sample points for averaging and p parameters.
Parameter Subset Selection	n_p parameters with minimum uncertainty	Provides identifiable subset of n_p parameters. Requires the optimal parameter estimate \hat{q} and standard errors SE .	$C(p, n_p)$ subsets to check for minimum uncertainty for subset of n_p parameters among p parameters.

Table 4.4 Sensitivity measures provided by Partial Correlation, Sobol by the Saltelli algorithm and Morris Screening.

Rank	Partial Correlation		Sobol by Saltelli			Morris Indices	
	q	$\text{Corr}(q, y)$	q	S_i	S_{T_i}	q	$\sqrt{\mu^{*2} + \sigma^2}$
1	N_T	4.608e-1	N_T	1.455e-1	2.134e-1	N_T	1.390e+6
2	p_T	3.964e-1	p_T	1.046e-1	2.038e-1	p_T	1.308e+6
3	k_1	2.970e-1	k_1	1.066e-1	1.335e-1	ϵ_2	1.240e+6
4	d_1	-2.630e-1	ϵ_2	7.879e-3	1.329e-1	d_1	1.096e+6
5	b_{E2}	-2.526e-1	d_1	6.276e-2	1.141e-1	k_1	9.876e+5
6	a_T	2.178e-1	a_T	5.568e-2	9.541e-2	a_T	9.824e+5
7	ϵ_2	-2.167e-1	ϵ_1	5.426e-2	7.849e-2	b_{E2}	8.371e+5
8	a_A	2.111e-1	b_{E2}	8.948e-2	7.723e-2	ϵ_1	8.161e+5
9	p_E	-1.789e-1	a_A	3.727e-2	5.458e-2	a_A	7.378e+5
10	λ_T	1.774e-1	p_E	2.013e-2	5.384e-2	λ_T	6.971e+5
11	ϵ_1	-1.601e-1	λ_T	4.682e-2	4.626e-2	p_E	6.666e+5
12	$T_2(0)$	1.487e-1	$T_2(0)$	-1.032e-2	4.266e-2	a_E	5.662e+5
13	a_E	-9.089e-2	a_E	1.057e-2	1.943e-2	$T_2(0)$	5.180e+5
14	$T_1^*(0)$	-2.721e-2	$T_1(0)$	-3.974e-3	1.464e-2	$T_1(0)$	4.084e+5
15	$T_1(0)$	2.410e-2	$T_1^*(0)$	-8.715e-3	2.423e-3	$T_1^*(0)$	9.285e+4

Table 4.5 Subsets of influential parameters for $n_p = 1, \dots, 4$. Here, P_{corr} , S, M and P_{ss} denote Partial Correlations, Sobol indices, Morris indices and Parameter Subset Selection.

	1				2				3				4			
	P_{corr}	S	M	P_{ss}	P_{corr}	S	M	P_{ss}	P_{corr}	S	M	P_{ss}	P_{corr}	S	M	P_{ss}
λ_T								x				x				x
d_1													x		x	
ϵ_1																
k_1									x	x			x	x		
a_T																
ϵ_2											x	x		x	x	x
N_T	x	x	x	x	x	x	x	x	x	x	x		x	x	x	x
b_{E2}																
a_E																
p_E																
a_A																
p_T					x	x	x		x	x	x	x	x	x	x	x
$T_1(0)$																
$T_1^*(0)$																
$T_2(0)$																

Table 4.6 Subsets of influential parameters for $n_p = 5, \dots, 8$.

	5				6				7				8			
	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}
λ_T				x				x				x				x
d_1	x	x	x		x	x	x		x	x	x		x	x	x	
ϵ_1										x				x	x	x
k_1	x	x	x		x	x	x		x	x	x		x	x	x	
a_T					x	x	x		x	x	x		x	x	x	
ϵ_2			x	x		x	x	x	x	x	x	x	x	x	x	x
N_T	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
b_{E2}	x				x				x		x		x	x	x	
a_E																
p_E				x				x				x				x
a_A													x			
p_T	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
$T_1(0)$												x				x
$T_1^*(0)$																
$T_2(0)$							x				x					x

Table 4.7 Subsets of influential parameters for $n_p = 9, \dots, 12$.

	9				10				11				12			
	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}
λ_T				x	x		x	x	x	x	x	x	x	x	x	x
d_1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
ϵ_1		x	x	x		x	x	x	x	x	x	x	x	x	x	x
k_1	x	x	x		x	x	x		x	x	x		x	x	x	
a_T	x	x	x		x	x	x		x	x	x		x	x	x	x
ϵ_2	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
N_T	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
b_{E2}	x	x	x		x	x	x	x	x	x	x	x	x	x	x	x
a_E												x			x	x
p_E	x			x	x	x		x	x	x	x	x	x	x	x	x
a_A	x	x	x		x	x	x		x	x	x		x	x	x	
p_T	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
$T_1(0)$				x				x				x				x
$T_1^*(0)$																
$T_2(0)$				x				x				x	x	x		x

Table 4.8 Subsets of influential parameters for $n_p = 13, \dots, 14$.

	13				14			
	P _{corr}	S	M	P _{ss}	P _{corr}	S	M	P _{ss}
λ_T	x	x	x	x	x	x	x	x
d_1	x	x	x	x	x	x	x	x
ϵ_1	x	x	x	x	x	x	x	x
k_1	x	x	x	x	x	x	x	x
a_T	x	x	x	x	x	x	x	x
ϵ_2	x	x	x	x	x	x	x	x
N_T	x	x	x	x	x	x	x	x
b_{E2}	x	x	x	x	x	x	x	x
a_E	x	x	x	x	x	x	x	x
p_E	x	x	x	x	x	x	x	x
a_A	x	x	x		x	x	x	x
p_T	x	x	x	x	x	x	x	x
$T_1(0)$				x		x	x	x
$T_1^*(0)$					x			
$T_2(0)$	x	x	x	x	x	x	x	x

certain parameters are determined to be highly influential by all four methods. An example of highly influential parameters are N_T and p_T . These parameters respectively represent the number of RNA copies during the process of T_1^* lysis and net proliferation of T_1 and T_1 due to clonal expansion and programmed contraction. We also observed that both ϵ_1 and ϵ_2 were ranked above average in their importance. This is essential in designing the optimal control for drug therapy. We see from our global sensitivity analysis that the relative effectiveness of protease inhibitor, ϵ_2 , has more affect on the model response than that of reverse transcriptase inhibitor, ϵ_1 . On the other hand, for our specific response, it was shown that initial conditions $T_1(0)$, $T_1^*(0)$ and $T_2(0)$ do not play a strong role in determining the response.

4.4.2 Verification of Noninfluential Inputs

To verify that the influential parameters are correctly identified, we compute the probability density functions of model responses by fixing noninfluential parameters.

Verification Procedure

1. For a set of n_p influential parameters, sample $n = 1000$ parameter values from their respective distributions. For the results reported here, we took the distributions to be uniform with lower and upper bounds summarized in Table 4.1.
2. Fix $p - n_p$ noninfluential parameters at pre-specified values, which we take to be the lower bounds of the parameters.
3. Compute the model response with parameter values from Steps 1 and 2.
4. Construct probability density function using a kernel density estimation.

We then compare the densities for the model responses where all parameters are sampled randomly. We construct densities for $n_p = 8, 10, 12, 14$ influential parameters. That is, we examine four cases where the numbers of fixed parameters are 7, 5, 3 and 1, and plot the densities along with the density obtained by varying all the parameters. In Figure 4.9 (a) and (b), we see that we have fixed too many parameters. In Figure 4.9 (c), densities using Partial Correlation, Sobol and Morris match the sample density, whereas the density from parameter subset selection does not match the rest. This is reasonable since the influential parameters determined via Partial Correlation, Sobol and Morris indices are very similar as shown in Tables 4.5-4.8. Finally in Figure 4.9 (d), we see that all four methods give comparable densities. The agreement of densities indicates that the parameter $T^*(0)$ was determined to be noninfluential and it did not affect the output significantly.

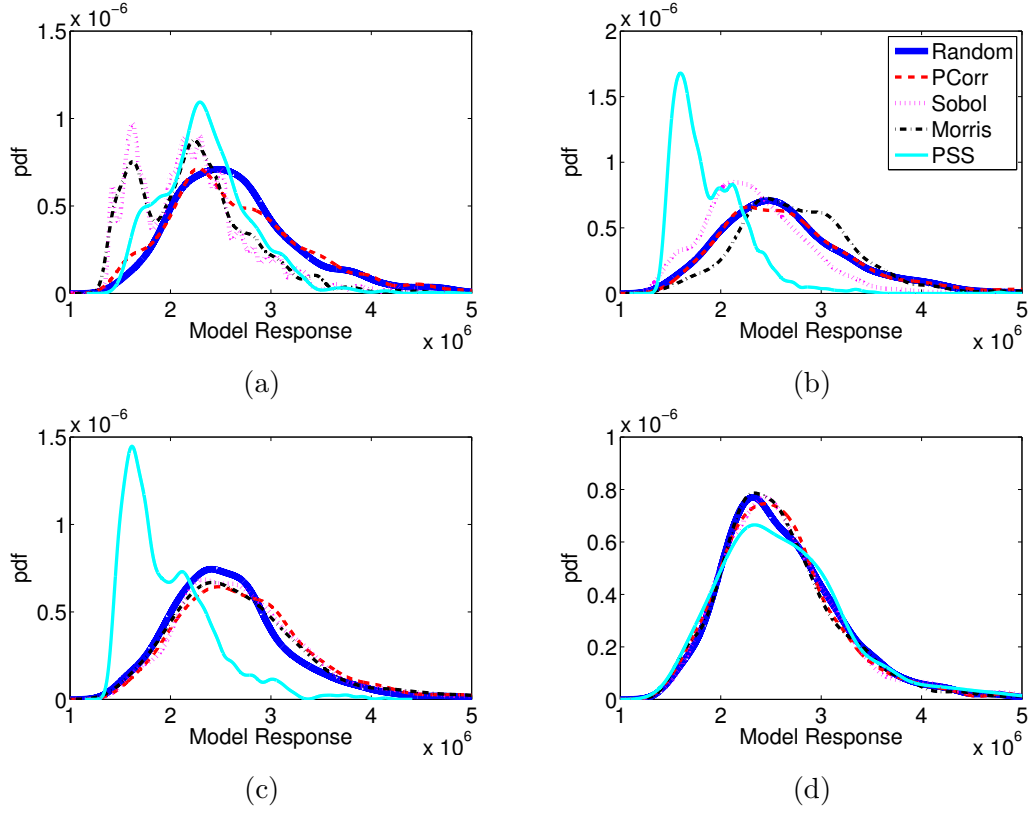


Figure 4.9 Densities obtained by fixing (a) 7, (b) 5, (c) 3 and (d) 1 least influential parameters. The density in solid is obtained by varying all the parameters.

Table 4.9 Sensitivity measures for $y = q_1 - q_2$.

Parameter	Partial Correlation	S	S_T	μ^*	σ^2
q_1	0.690	0.498	0.516	1	0.980
q_2	-0.700	0.481	0.498	1	1.020

In terms of parameter estimation to be followed after parameter selection, there are two inherent difficulties. First, the cut-off between influential and noninfluential parameters is not always clear. Depending on the model, one might observe a cluster of parameters with high sensitivity measures and another cluster with lower sensitivity measures. For the HIV example, it was only after we performed a verification test that we learned that fixing one parameter resulted in insignificant variability of sample densities. This verification requires additional model evaluations and there is not a simple way to check which parameters are influential just by observing sensitivity measures.

Second, even if we are successful at isolating influential parameters, the parameter identifiability issues may still remain. Consider a simple example

$$y = q_1 - q_2 \tag{4.59}$$

with $q_1, q_2 \sim U(0, 1)$. As we can see from the sensitivity measures summarized in Table 4.9, q_1 and q_2 are equally influential. Suppose that we have the observation $y = 0$. It is easy to see that parameter estimation using this observation will fail to estimate the densities of q_1 and q_2 correctly since there are several values of q_1 and q_2 that match the observation. Therefore, unless some prior knowledge is specified, q_1 and q_2 are unidentifiable.

This simple example illustrates that determining influential parameters may not eliminate parameter identifiability issues completely. In this regard, the parameter subset selection algorithm has the advantage that the selected subset is identifiable. Since Partial Correlation, Sobol and Morris methods only determine influential parameters, care must be exercised if these parameter selection techniques are used to isolate identifiable parameters for model calibration.

4.5 Conclusion

In this chapter, we examined parameter selection techniques based on global sensitivity analysis and compared the results to a local sensitivity-based method originally performed on the model (1.1). Four parameter selection techniques were applied to the HIV model (1.1) to determine the set of influential parameters. This process enables us to fix the noninfluential parameters and hence reduce the parameter dimensions for subsequent uncertainty quantification. We also showed that the accuracy of Sobol indices depends greatly on the model. In our HIV model, the second-order decomposition was not sufficiently accurate to represent the response. If one requires an insight into determining parameter identifiability issues, then the parameter subset selection algorithm is recommended since it returns a set of identifiable parameters with minimized uncertainty.

It is important to note that there are several alternate choices for the model response. Our choice of the model response was motivated by the types of data that are available to us. However, one must carefully examine the cases when different model responses are chosen. It is important to remember that the parameters were determined influential in our analysis for our specific choice of model response. One idea for future work is to examine global sensitivity analysis using solely the T-cell counts as a response. Since the treatment attempts to increase the T-cell counts in patients, it is reasonable to focus on the T-cell counts alone. Similarly, one could focus on the viral loads $V_I + V_{NI}$ in an attempt to keep the viral loads low. Another aspect of analysis that we did not cover in this chapter is to consider model response as a function of time. Recall that in our analysis we integrated the response in time to take into consideration of response at several time steps. In reality, we see that the states T_1, T_1^*, \dots, E_2 can be mostly flat except for some jumps. Considering time-dependent model response will enable us to incorporate jumps that occur at certain times.

Finally, there are other methods of parameter selection that eliminate parameter identifiability issues. In particular, Active Subspace Methods detailed in [5, 22] determine a subspace of input parameter space which affects the response the most. This method does not isolate influential parameters from noninfluential parameters and the interpretation of the results may be more complicated. However, finding a linear combination of parameters that affects the response will resolve the parameter identifiability issues. Moreover, responses can be approximated based on the reduced parameter space, which is useful in subsequent model calibration and uncertainty quantification. Examining active subspace methods more closely as a part of parameter selection techniques will likely add more complete analysis on parameter selection techniques.

Remark

M.T. Wentworth, R.C. Smith and H.T. Banks, “Parameter Selection and Verification Techniques Based on Global Sensitivity Analysis Illustrated for an HIV Model,” submitted, 2015.

Chapter 5

Active Subspace Methods

5.1 Introduction

To reduce the dimensions of input parameters, parameter selection techniques have been applied to the model (1.1) in Chapter 4. Parameter selection is a subset selection method, in which a subset of parameters that strongly influence the model response is identified. Such parameters are termed *influential*, and identifying influential parameters allows us to fix noninfluential parameters at nominal values without affecting the response.

In Chapter 4, we employed parameter selection techniques based on global sensitivity analysis and a method based on standard errors. More specifically, we applied Partial Correlation [4], Sobol indices [55, 60, 61], Morris indices [18, 46, 55, 59] and Parameter Subset Selection [6, 7, 9, 20]. The sensitivity indices were computed for the 15 parameters listed in (4.2). In this work, N_T and p_T are identified influential with large sensitivity indices, indicating that the changes in these parameters highly influence the changes in the model response. On the other hand, $T_1^*(0)$ is ranked the least influential among the 15 parameters that were considered.

For the verification of these techniques, probability density functions were constructed by computing 1000 model responses, in which highly influential parameters were sampled randomly from their corresponding distributions and minimally influential parameters were fixed at nominal values. In this verification, we found that we must keep most of the parameters to represent the model response accurately. In fact, $T_1^*(0)$ was the only parameter whose value could be fixed for the resulting density to match the response constructed using the full dimensions for all four methods considered. For Partial Correlation, Sobol indices and Morris indices, three parameters could be fixed without affecting the output.

Parameter selection techniques can sometimes be used to determine identifiable parameters, which are the parameters that can be estimated from observations. However, when two parameters are correlated, those two parameters may not be jointly identifiable even if they are both

influential. As illustrated in Chapter 4, an example of a response with two correlated parameters is given by

$$f(q) = q_1 - q_2 \quad (5.1)$$

where $q_1, q_2 \sim U(0, 1)$. The two parameters q_1 and q_2 are equally influential, as Partial Correlation, Sobol indices and Morris indices indicate, however in this example, the two parameters can not be uniquely identified from an observation without further information.

5.1.1 Active Subspace Methods

As an alternative to parameter subset selection techniques, we examine parameter subspace selection techniques to reduce the input parameter dimensions. Unlike parameter selection methods, subspace selection methods identify a subspace of the input space that strongly affects the model response. Such subspaces are often termed *Active Subspaces*, which was introduced by Russi in [52]. Russi proposed a method called *Active Subspace Method* to construct active subspaces using function gradients. Active subspace methods determine the directions of input parameter space in which the model responses exhibit the greatest change. If the active subspace aligns with coordinate axes, then the result is identical to parameter subset selection.

An active subspace method has been applied to an elliptic PDE model with a 100-parameter model for the coefficients in [22]. Other applications of active subspace methods include design optimization [19, 25], inverse analysis [23] and spatial sensitivity [24]. Additionally, an example based on Morris' experiments [46] is detailed in [1], and a radiation transport model employed in nuclear reactor design calculations is detailed in [5].

In the simple example (5.1), the two parameters are equally influential; however, the model response is most sensitive to the linear combination of the parameters in the form $q_1 - q_2$. As illustrated in Figure 5.1, the active subspace is $[1, -1]^T$, which is the direction in which the response vary the greatest. On the other hand, the model output is constant in the direction $[1, 1]^T$.

To construct the active subspace for a nontrivial function, the authors of [22, 52] use SVD decompositions of the gradient of functions. The theoretical foundation of gradient-based dimension reduction techniques is detailed in [22]. In [5], the authors employ QR decompositions for dimension reduction. In both cases, one uses the gradients of functions and determines the directions of the greatest response change. The two methods are discussed in Section 5.3.

5.1.2 Chapter Organization

In Section 5.2, we discuss three ways to compute gradients that are required in active subspace methods. First, the finite difference method is commonly used to compute gradient, but we show

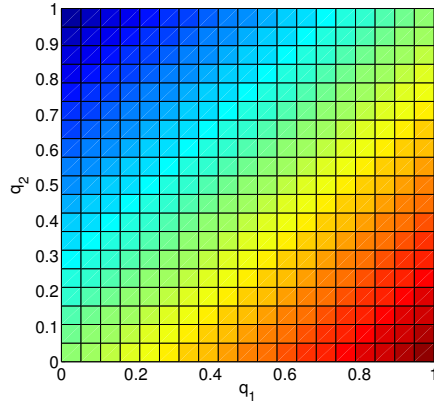


Figure 5.1 Illustration of $f(q) = q_1 - q_2$ with $q_1, q_2 \sim U(0, 1)$.

using our HIV example that the finite differences may not accurately approximate the gradient. Secondly, we discuss the method of sensitivity equations, in which we explicitly derive sensitivity equations of the ODE's. The system of ODE's, augmented with the sensitivity equations, is then solved using an ODE solver. While this method does not have any assumptions, generating sensitivity equations is time-consuming and is prone to errors. Thirdly, we employ automatic differentiation, in which the gradient is computed using sensitivity equations in a automated manner.

Once we obtain the gradients, we detail two active subspace method based on SVD in Section 5.3. One is based on the singular value decomposition detailed in [22] and the other is based on the QR decomposition as detailed in [5]. The active subspace is constructed using the eigenvalue decomposition of gradients. Once the active subspace is found, one can use interpolation to construct a response surface.

In Section 5.4, we apply the active subspace methods to the HIV model (1.1). We show sufficient summary plots that allow us to determine whether dimension reduction is possible. We then perform a verification test to verify the accuracy of response surface that is constructed using a subspace of input parameter space.

Table 5.1 Nominal values of q^0 , at which the gradient is computed.

$\lambda_T = 3.2669$	$d_1 = 1.1199\text{e-}1$	$\epsilon_1 = 5.8346\text{e-}1$	$k_1 = 5.4172\text{e-}1$
$a_T = 2.3436\text{e-}4$	$\epsilon_2 = 7.0339\text{e-}1$	$N_T = 7.1754\text{e+}1$	$b_{E2} = 4.3301\text{e-}1$
$a_E = 1.5292\text{e-}3$	$p_E = 9.0040\text{e-}1$	$a_A = 8.4506\text{e-}5$	$p_T = 5.5846$

Table 5.2 The lower and upper bounds for the distributions $q_i \sim U(a_i, b_i)$ for $i = 1, \dots, 15$.

	λ_T	d_1	ϵ_1	k_1	a_T	ϵ_2	N_T	b_{E2}
a_i	3.1	0.11	0.43	4.0e-5	2.0e-4	0.63	65	0.28
b_i	3.5	0.15	0.6	5.5e-5	2.7e-4	0.78	85	0.45
	a_E	p_E	a_A	p_T	$T_1(0)$	$T_1^*(0)$	$T_2(0)$	
a_i	1.40e-3	0.85	6.5e-5	5	10.5	5.0e-4	720	
b_i	1.75e-3	1.3	9.0e-5	6.5	13.5	7.0e-4	950	

5.2 Computing the Gradient

Before discussing active subspace methods, we discuss different methods to compute the gradient. We compute the gradient via three different methods: finite difference, sensitivity equations, and automatic differentiation. The gradient is computed using all three methods at a randomly chosen set of parameter values, q^0 given in Table 5.1, which are sampled uniformly with appropriate lower and upper bound obtained from [6]. The values of lower and upper bounds are summarized in Table 5.2.

Table 5.3 Gradient via finite difference using different step sizes.

h	λ_T	d_1	ϵ_1	k_1	a_T	ϵ_2
1.00e-1	1.01e+6	-9.32e+5	-1.70e+5	7.72e+5	9.85e+5	-1.95e+5
1.00e-2	1.01e+6	-9.54e+5	-1.78e+5	7.89e+5	9.85e+5	-2.06e+5
1.00e-3	1.02e+6	-9.52e+5	-1.69e+5	7.91e+5	9.94e+5	-1.98e+5
1.00e-4	1.01e+6	-8.62e+5	-8.37e+4	7.73e+5	1.01e+6	-1.13e+5
1.00e-5	9.29e+5	-7.16e+5	-1.90e+5	1.02e+6	1.00e+6	-1.98e+5
1.00e-6	1.41e+6	2.01e+5	8.21e+5	3.04e+6	2.16e+6	3.49e+5
1.00e-7	1.92e+7	8.31e+6	1.25e+6	8.76e+6	1.70e+7	2.94e+6
1.00e-8	8.83e+7	1.02e+8	1.13e+8	4.89e+7	7.06e+7	6.39e+7

5.2.1 Finite Difference

We initially computed the gradient using the finite difference

$$\frac{\partial f}{\partial q_i} \approx \frac{f(q_1, \dots, q_i + h_i, \dots, q_p) - f(q_1, \dots, q_i, \dots, q_p)}{h_i} \quad (5.2)$$

for $i = 1, \dots, p$. To test the convergence, the finite difference is computed at different step sizes $h = [1e-1, \dots, 1e-8]$ for the parameters (4.2) without the initial conditions. The result for six parameters is summarized in Table 5.3. The gradient seems to converge as the step size decreases to $1e-3$. However, for most parameters, the gradient then diverges when we further decrease the step size. As an example, the result for λ_T with different step sizes is plotted in Figure 5.2.

5.2.2 Sensitivity Equations

Next, we compute the gradient using analytically-derived sensitivity equations. The sensitivity equations are derived by taking the derivative of the model (1.1) with respect to each parameter. For example, the sensitivity equations for λ_T are

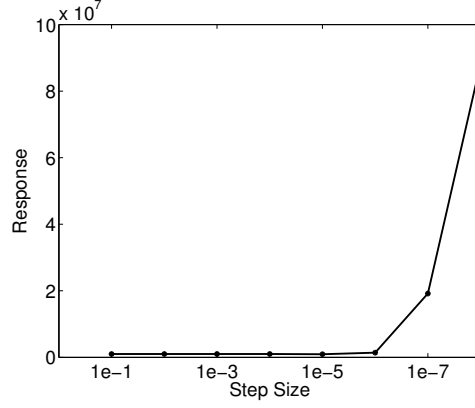


Figure 5.2 Gradient with respect to λ_T using different step sizes.

$$\begin{aligned}
\frac{d}{dt} \frac{\partial T_1}{\partial \lambda_T} &= -d_1 \frac{\partial T_1}{\partial \lambda_T} - (1 - \xi_1(t))k_1 \left(\frac{\partial V_I}{\partial \lambda_T} T_1 + V_I \frac{\partial T_1}{\partial \lambda_T} \right) - \gamma_T \frac{\partial T_1}{\partial \lambda_T} + p_T D_f(T_2) \\
\frac{d}{dt} \frac{\partial T_1^*}{\partial \lambda_T} &= (1 - \xi_1(t))k_1 \left(\frac{\partial V_I}{\partial \lambda_T} + V_I \frac{\partial T_1}{\partial \lambda_T} \right) - \delta \frac{\partial T_1^*}{\partial \lambda_T} - m \left(\frac{\partial E_1}{\partial \lambda_T} T_1^* + E_1 \frac{\partial T_1^*}{\partial \lambda_T} \right) \\
&\quad - \gamma_T \frac{\partial T_1^*}{\partial \lambda_T} + p_T D_f(T_2^*) \\
\frac{d}{dt} \frac{\partial T_2}{\partial \lambda_T} &= \frac{K_s}{V_I + K_s} - \frac{\lambda_T K_s}{(V_I + K_s)^2} \frac{\partial V_I}{\partial \lambda_T} - \gamma_T \frac{\partial T_1}{\partial \lambda_T} - d_2 \frac{\partial T_2}{\partial \lambda_T} \\
&\quad - (1 - f\xi_1(t))k_2 \left(\frac{\partial V_I}{\partial \lambda_T} T_2 + V_I \frac{\partial T_2}{\partial \lambda_T} \right) - D_f(T_2) \\
\frac{d}{dt} \frac{\partial T_2^*}{\partial \lambda_T} &= \gamma_T \frac{\partial T_1^*}{\partial \lambda_T} + (1 - f\xi_1(t))k_2 \left(\frac{\partial V_I}{\partial \lambda_T} T_2 + V_I \frac{\partial T_2}{\partial \lambda_T} \right) - d_2 \frac{\partial T_2^*}{\partial \lambda_T} - D_f(T_2^*) \\
\frac{d}{dt} \frac{\partial V_I}{\partial \lambda_T} &= (1 - \xi_2(t))10^3 \delta T_1^* + (1 - \xi_2(t))10^3 N_T \delta \frac{\partial T_1^*}{\partial \lambda_T} - c \frac{\partial V_I}{\partial \lambda_T} \\
&\quad - 10^3 (1 - \xi_1(t))\rho_1 k_1 \left(\frac{\partial T_1}{\partial \lambda_T} V_I + T_1 \frac{\partial V_I}{\partial \lambda_T} \right) \\
&\quad - 10^3 (1 - f\xi_1(t))\rho_2 k_2 \left(\frac{\partial T_2}{\partial \lambda_T} V_I + T_2 \frac{\partial V_I}{\partial \lambda_T} \right) \\
\frac{d}{dt} \frac{\partial V_{NI}}{\partial \lambda_T} &= \xi_2(t)10^3 T_1^* + \xi_2(t)10^3 N_T \delta \frac{\partial T_1^*}{\partial \lambda_T} - c \frac{\partial V_{NI}}{\partial \lambda_T} \\
\frac{d}{dt} \frac{\partial E_1}{\partial \lambda_T} &= \frac{b_{E1}}{T_1^* + K_{b1}} \left(\frac{\partial T_1^*}{\partial \lambda_T} E_1 + T_1^* \frac{\partial E_1}{\partial \lambda_T} \right) - \frac{b_{E1}}{(T_1^* + K_{b1})^2} \left(\frac{\partial T_1^*}{\partial \lambda_T} \right) T_1^* E_1 \\
&\quad - \frac{d_E}{T_1^* + K_d} \left(\frac{\partial T_1^*}{\partial \lambda_T} E_1 + \frac{\partial E_1}{\partial \lambda_T} T_1^* \right) + \frac{d_E}{(T_1^* + K_d)^2} \frac{\partial T_1^*}{\partial \lambda_T} T_1^* E_1 - \delta_{E1} \frac{\partial E_1}{\partial \lambda_T} \\
&\quad - \frac{\gamma_E}{T_1 + T_1^* + K_\gamma} \left(\frac{\partial T_1}{\partial \lambda_T} E_1 + T_1 \frac{\partial E_1}{\partial \lambda_T} \right) + \frac{\gamma_E}{(T_1 + T_1^* + K_\gamma)^2} \left(\frac{\partial T_1}{\partial \lambda_T} + \frac{\partial T_1^*}{\partial \lambda_T} \right) T_1 E_1 \\
&\quad - \frac{\gamma_E}{T_1 + T_1^* + K_\gamma} \left(\frac{\partial T_1^*}{\partial \lambda_T} E_1 + T_1^* \frac{\partial E_1}{\partial \lambda_T} \right) + \frac{\gamma_E}{(T_1 + T_1^* + K_\gamma)^2} \left(\frac{\partial T_1}{\partial \lambda_T} + \frac{\partial T_1^*}{\partial \lambda_T} \right) T_1^* E_1 \\
&\quad + \frac{p_E a_E}{V_I + K_E} \left(\frac{\partial V_I}{\partial \lambda_T} E_2 + V_I \frac{\partial E_2}{\partial \lambda_T} \right) - \frac{p_E a_E}{(V_I + K_E)^2} \left(\frac{\partial V_I}{\partial \lambda_T} \right) V_I E_2
\end{aligned}$$

where

$$\begin{aligned} D_f(T) &\equiv \frac{\partial}{\partial \lambda_T} \left[\left(\frac{a_T V_I}{V_I + K_V} + a_A \right) T \right] \\ &= \frac{a_T}{V_I + K_V} \left(\frac{\partial V_I}{\partial \lambda_T} T + V_I \frac{\partial T}{\partial \lambda_T} \right) - a_T V_I T \frac{1}{(V_I + K_V)^2} \frac{\partial V_I}{\partial \lambda_T} + a_A \frac{\partial T}{\partial \lambda_T}. \end{aligned}$$

The sensitivity equations (5.3) are then solved simultaneously with the original model (1.1) using `ode45` with the initial condition

$$Y_0 = [T_1(0), T_1^*(0), T_2(0), T_2^*(0), V_I(0), V_{NI}(0), E_1(0), E_2(0), 0, 0, 0, 0, 0, 0, 0]. \quad (5.4)$$

The response is again taken as the sum of T_1 , T_1^* , T_2 , T_2^* , V_I and V_{NI} integrated in time; that is,

$$\frac{\partial f}{\partial \lambda_T} = \int_0^{1500} \left[\frac{\partial T_1}{\partial \lambda_T}(t) + \frac{\partial T_1^*}{\partial \lambda_T}(t) + \frac{\partial T_2}{\partial \lambda_T}(t) + \frac{\partial T_2^*}{\partial \lambda_T}(t) + \frac{\partial V_I}{\partial \lambda_T}(t) + \frac{\partial V_{NI}}{\partial \lambda_T}(t) \right] dt. \quad (5.5)$$

5.2.3 Automatic Differentiation

Thirdly, we compute the automatic differentiation using Automatic Differentiation. This method is based on the fact that all computer programs can be decomposed into a combination of elementary arithmetic operations, which are additions, subtractions, multiplications and divisions, along with the composition of the functions, and function evaluations, which are exponentials, logarithms, cosines and sines. The derivatives are then computed by applying chain rules to these operations. It is noted here that automatic differentiation should not be confused with symbolic differentiation. Given the ODE system

$$\frac{dy}{dt} = f(t, y(q); q), \quad (5.6)$$

differentiation yields

$$\frac{d}{dt} \frac{\partial y}{\partial q} = \frac{\partial f}{\partial y} \frac{\partial y}{\partial q} + \frac{\partial f}{\partial q}. \quad (5.7)$$

Here, $\frac{\partial f}{\partial y}$ and $\frac{\partial f}{\partial q}$ are computed using AD while $\frac{\partial y}{\partial q}$ is obtained using `ode15s` since it is a stiff system. The augmented system has states, sensitivities for parameters, and sensitivity for IC's. One can obtain the gradients by inputting a system of ODE $\frac{dy}{dt} = f(t, y(q); q)$ without further breaking it into pieces using the `myAD` package available at the Matlab Central.

5.2.4 Comparison of Gradient

We now compare the gradient obtained by the three methods. We use the result for $h=1e-3$ for the finite difference. The derivatives at q^0 with respect to 12 parameters from (4.2) without the

Table 5.4 Derivatives at q^0 obtained via finite differences (FD), sensitivity equations (SE) and automatic differentiation (AD).

	FD	SE	AD
λ_T	1.0169e+6	1.0086e+6	1.0360e+6
d_1	-9.5177e+5	-9.5171e+6	-1.0538e+7
ϵ_1	-1.6900e+5	-4.2746e+5	-2.3624e+5
k_1	7.9136e+5	9.5006e+11	2.2509e+10
a_T	9.9433e+5	5.6167e+9	4.9934e+9
ϵ_2	-1.9811e+5	-5.5008e+5	-2.9964e+5
N_T	1.7486e+6	3.4898e+4	3.5015e+4
b_{E2}	-9.7448e+5	-2.3067e+6	-2.0806e+6
a_E	-4.8006e+5	-5.5569e+8	-5.1713e+8
p_E	-1.2164e+6	-1.0882e+6	-9.8223e+5
a_A	6.9051e+5	1.1099e+10	1.3500e+10
p_T	1.6110e+6	4.2887e+5	4.3642e+5

initial conditions are summarized in Table 5.4.

We see that the values obtained by solving the sensitivity equations and automatic differentiation compare well, whereas the gradient via finite difference yields different values for certain parameters. This indicates the limitation of finite difference. Since we verified that automatic differentiation is accurate based on the comparison to the analytic sensitivity equations, we use automatic differentiation for the computation of gradient in the subsequent analysis.

5.3 Active Subspace Methods

5.3.1 Active Subspace Using SVD

We first illustrate the construction of active subspaces based on the SVD decomposition. As detailed in [22], for a function f with p continuous inputs, we first consider

$$f = f(q), q \in \mathcal{Q} \subset \mathbb{R}^p. \quad (5.8)$$

The gradient of f is denoted by

$$\nabla f(q) = \left[\frac{\partial f}{\partial q_1} \cdots \frac{\partial f}{\partial q_p} \right]^T \quad (5.9)$$

and define the $m \times m$ matrix C

$$C = \mathbb{E}[(\nabla f)(\nabla f)^T]. \quad (5.10)$$

Here, C is the uncentered covariance of the gradient vector. The eigenvalue decomposition of C is

$$C = W\Lambda W^T, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_p), \quad \lambda_1 \geq \dots \geq \lambda_p \geq 0. \quad (5.11)$$

The eigenvectors W define the domain of f . The eigenvalues and eigenvectors can be partitioned such that

$$\Lambda = \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix}, \quad W = [W_1 \ W_2] \quad (5.12)$$

where $\Lambda_1 = \text{diag}(\lambda_1, \dots, \lambda_n)$ with $n < p$ and W_1 is $p \times n$. The rotated coordinates $y \in \mathbb{R}^n$ and $z \in \mathbb{R}^{p-n}$ are defined by

$$y = W_1^T q, \quad z = W_2^T q. \quad (5.13)$$

We are interested in finding the subspace $\mathcal{Y} \subset \mathbb{R}^n$ with $n \leq p$ such that $f(q) \approx \hat{G}(y) = \hat{G}(W_1^T q)$ for $y = W_1^T q \in \mathcal{Y}$. Here, \hat{G} is approximated using the Monte Carlo approximation

$$\hat{G}(y) = \frac{1}{N} \sum_{j=1}^N f(W_1 y + W_2 z^j). \quad (5.14)$$

The choice of n is determined by inspecting the eigenvalues. Ideally, when $\lambda_{n+1} = \dots = \lambda_m = 0$, the choice is clear. In other cases, the eigenvalues often have a gap in their magnitude, and we choose n so that $\lambda_n - \lambda_{n+1}$ is large. See [22] for more detail.

We employ the algorithm detailed in [22] to find the active subspaces. Here, we let q , y and f denote the full dimensional parameter vector, reduced dimensional parameter vector, and model response, respectively. We let the superscript denote the sample number, and let the subscript denote the entry number. For example, q_i^j indicates the i^{th} entry of j^{th} sample vector $q^j \in \mathbb{R}^p$. We also use the notation q^{jn} to indicate the j^{th} sample of the parameter vector that is constructed using n -dimensional active subspace.

Algorithm 1

1. Begin by choosing a set of N random points $q^j \in \mathbb{R}^p$ for $j = 1, \dots, N$, where $q_i^j \sim U(-1, 1)$ for $i = 1, \dots, p$, and compute $f^j = f(q^j)$ and $\nabla f^j = \nabla f(q^j)$ for $j = 1, \dots, N$.

2. Compute the SVD of the matrix

$$G = \frac{1}{\sqrt{N}}[\nabla f^1 \dots \nabla f^N] = W\Sigma V^T \quad (5.15)$$

and set $\Lambda = \Sigma^2$.

3. Considering the decay of eigenvalues, partition $W = [W_1 W_2]$, where $W_1 \in \mathbb{R}^{p \times n}$ and $W_2 \in \mathbb{R}^{p \times (p-n)}$ for $n \leq p$.
4. Choose a set of P points y^j such that $y_i^j \sim U(-1, 1)$ for $i = 1, \dots, n$ and construct the full dimensional parameter vectors $q^{jn} = W_1 y^j$ for $j = 1, \dots, P$.

We note here that in Step 1, the input parameters q^j are normalized so that $q^j \in [-1, 1]^p$. When computing the responses $f^j = f(q^j)$, the inputs are mapped to their appropriate ranges using the values in Table 5.2. The normalization for q^j is essential here to allow us to sample $y_i^j \sim U(-1, 1)$ for $i = 1, \dots, n$ in Step 4.

5.3.2 Response Surface Representation

Once the active subspaces are identified, the input subspace is no longer physical. This requires us to represent the response surface with a surrogate model with minimum cost and whose input may be non-physical. In [22], a standard Kriging is employed to construct the response surface with the active subspaces.

As detailed in [54, 59], the Kriging emulator is given by $\tilde{f}(q, \beta) = g^T(q)\beta + Z(q)$, where $g^T(q)\beta$ and $Z(q)$ respectively denote a deterministic trend function and a Gaussian process error model. In ordinary Kriging, the trend function is constant so that $g^T(q)\beta = \beta_0$. In the absence of measurement noise, Z is assumed to be a stationary random process with zero mean, variance σ^2 and nonzero covariance

$$\text{cov}[Z(q^{j1}), Z(q^{j2})] = \sigma^2 R(q^{j1}, q^{j2}) + \sigma_0^2 \delta(q^{j1} - q^{j2}) \quad (5.16)$$

where

$$\delta(q^{j1} - q^{j2}) = \begin{cases} 1, & q^{j1} - q^{j2} = 0, \\ 0, & \text{else.} \end{cases} \quad (5.17)$$

The correlation function is given by

$$R(q^{j1}, q^{j2}) = \exp\left(-\sum_{i=1}^p \theta_i |q_i^{j1} - q_i^{j2}|^{\gamma_i}\right), \quad 0 < \gamma_i \leq 2, \quad \theta_i > 0. \quad (5.18)$$

Here, σ_0 and σ are the unadjusted variance and variance of measurement noise, respectively. Also, the hyperparameters θ_i and γ_i can be tuned to achieve varying degrees of correlation. Based on [4, 22], we choose $\theta_i = 0.1$ and $\gamma_i = 2$.

Using the set of data $\{q^j, \hat{G}^j\}_{j=1}^P$, the Kriging prediction $\tilde{f}(q, \beta_0)$ for new values of q is given by

$$\tilde{f}(q, \beta_0) = \beta_0 + r^T(q) \mathcal{R}^{-1} [\hat{G} - \beta_0 \mathbf{1}]. \quad (5.19)$$

Here,

$$\beta_0(\theta, \gamma) = [\mathbf{1}^T \mathcal{R}^{-1} \mathbf{1}]^{-1} \mathbf{1}^T \mathcal{R}^{-1} \hat{G} \quad (5.20)$$

is the least squares estimate for β_0 , where $\mathbf{1} = [1, \dots, 1]^T \in \mathbb{R}^P$ and \mathcal{R} is an $P \times P$ correlation matrix defined by $\mathcal{R}_{j_1 j_2} = R(q^{j_1}, q^{j_2})$. Also, $r_j(q) = R(q^j, q)$ for $j = 1, \dots, P$ is a $P \times 1$ vector that quantifies the Gaussian process correlations between values at which trained data is computed and the new input.

For our HIV example, we are interested in the parameter selection aspect of active subspace methods. We focus on finding active subspaces and verifying their accuracy in the next section.

5.4 Example and Verification

In this section, we apply the SVD-based active subspace method detailed in Section 5.3.1 to the HIV model (1.1). We obtain 15 sample points for each parameter. Here, the parameters are sampled from uniform distributions with appropriate lower and upper limits summarized in Table 5.2. The response gradients are computed at these 15 points.

We then compute active subspaces using the SVD. The computed eigenvalues are summarized in Table 5.5. We note that the order of magnitudes of eigenvalues is large, so to avoid numerical error, we take the log transform of gradient before computing the SVD. The log transformed eigenvalues are plotted in Figure 5.3. We see that the trends with and without log transform are similar. Aside from the gap between the first two eigenvalues, there is no clear gap in eigenvalues. From this, we expect that we may find a one dimensional active subspace. For an active subspace with dimension larger than one, it would be difficult to find a cut-off for the number of dimensions to include in the active subspace.

5.4.1 Sufficient Summary Plot

The purpose of sufficient summary plots is to visually inspect whether or not the function can indeed be expressed as a function of low dimensional active variables. We produce sufficient summary plots in two ways. First, we use the eigenvectors associated with the SVD of the gradient, as computed in the previous section. The active variables are $y = W_1 q$, where W_1 is

Table 5.5 Eigenvalues computed using 15 samples.

	Eigenvalues	Eigenvalues of Log Transform
λ_1	2.86e+21	5.94e+4
λ_2	3.64e+19	6.16e+1
λ_3	4.55e+18	1.65e+1
λ_4	1.17e+16	4.48e+0
λ_5	2.15e+12	2.05e+0
λ_6	1.79e+12	3.28e-1
λ_7	4.16e+11	1.77e-1
λ_8	6.58e+10	1.04e-1
λ_9	2.44e+10	5.81e-2
λ_{10}	1.12e+10	4.99e-2
λ_{11}	6.51e+9	2.91e-2
λ_{12}	1.45e+9	1.79e-2
λ_{13}	1.14e+8	7.19e-3
λ_{14}	7.61e+6	1.35e-3
λ_{15}	1.86e+5	3.11e-5

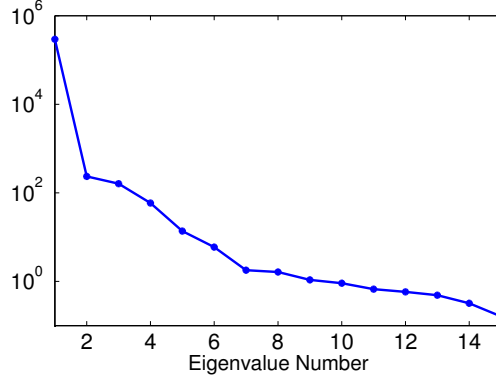


Figure 5.3 Log eigenvalues computed using the SVD-based active subspace method.

the first column of W and q is the full dimensional parameter.

Secondly, we approximate the function by a global linear model, and use the coefficients of gradient to represent the active subspace as detailed in [22]. To compute the coefficients of gradients, we employ the model

$$f^j = f(q^j) \approx c + b^T q^j \quad (5.21)$$

for $j = 1, \dots, N$, where N is the number of samples. The constant term c and the vector of coefficients b are computed using least squares. The weight on parameters is given by

$$w = b / \|b\| \quad (5.22)$$

which is the normalized gradient of coefficients.

For this problem, the weight is computed using $N = 150$ samples yielding the result plotted in Figure 5.4(a). The negative weights are plotted in red in the positive direction to compare the magnitudes of weights. The active variable in this case is $y = w^T q$. We note here that the parameters with largest and smallest weights correspond to the influential and noninfluential parameters identified by global sensitivity analysis. Here, N_T and p_T are the parameters with the greatest weights, whereas $T_1^*(0)$ and $T_1(0)$ receive the smallest weights. The result is reasonable since the global sensitivity analysis methods we considered in Chapter 4 incorporated linear and quadratic effects of parameters. Since the global linear model linearizes the effect of parameters

onto the model, we expect that influential parameters have larger weights than noninfluential parameters. For example, influential parameters, as determined by Partial Correlation and the Sobol decomposition are plotted in Figure 5.4(b) and (c).

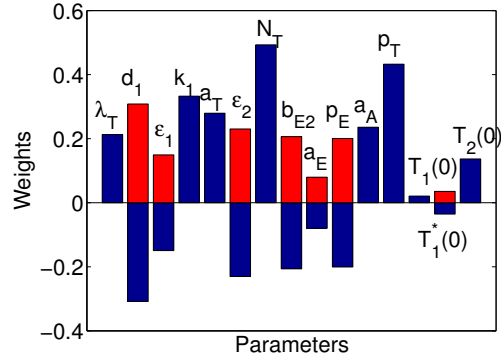
We plot the function evaluations in terms of active variables in Figure 5.5. In Figure 5.5(a), we employ $f(wq)$, where w is the weight computed in (5.22). Similarly, in Figure 5.5(b), we employ $f(W_1q)$, where W_1 is the eigenvector corresponding to the largest eigenvalue from (5.12). In both cases, $q \sim U(-1, 1)$ and we sample 100 points. We see that the model response as a function of the active variance using weight is approximately univariate in Figure 5.5(a). This indicates the possibility of approximating the response surface with one dimensional active variable. As illustrated in the next section, however, this one dimensional representation is only moderately accurate. On the other hand, in Figure 5.5(b), the function in terms of the principal active variable is not nearly univariate, indicating that we will require more than one dimensional active subspace to represent the input space.

5.4.2 Verification

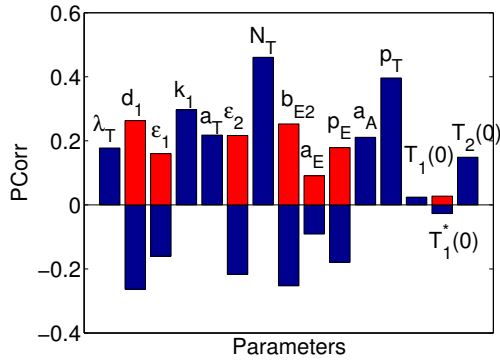
To verify the accuracy of model responses computed using active subspaces based on the eigenvectors from (5.11), we perform the following verification procedure.

1. Sample $N = 1000$ random points in the full dimensional space $\{q^j\}_{j=1}^N$ where $q_i^j \sim U(-1, 1)$ for $i = 1, \dots, p$.
2. For $n = 1 : p$
 - (a) Let W_n be the first n columns of W , where W comes from SVD in (5.12).
 - (b) Transform the reduced domain $q_{1:n}^j$ into the full domain $q^{jn} \equiv W_n q_{1:n}^j$ for $j = 1, \dots, N$.
 - (c) Compute the model response at $f^j(n) \equiv f(q^{jn})$ for $j = 1, \dots, N$.
 - (d) Using a kde, compute the probability density function of $[f^1(n), f^2(n), \dots, f^N(n)]$.

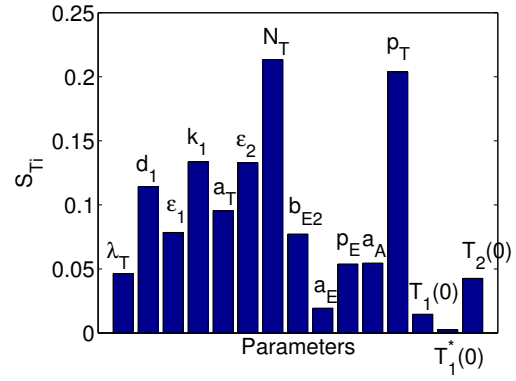
To verify the active subspace using the weights, we perform a similar procedure with $n = 1$, where W in Step 2(a) is w from (5.22). In both cases, n represents the dimension of the active subspace. In Figure 5.6, we plot the resulting densities for (a) $n = 1, 5, 8$, (b) $n = 9, 10, 11$ and (c) $n = 12, 13, 14$ and compare them to the density using the full dimensional input space with $n = 15$, as well as the weights in (5.22). In Figure 5.6(a), we see that the densities with $n \leq 8$ have completely different shapes. We note that the active subspace using the weights is not sufficiently accurate, but its performance is remarkable considering that its dimension is one. In Figure 5.6(b) and (c), the densities for $n \geq 10$ closely approximate the density with $n = 15$. For $n \geq 12$, the densities are almost identical, indicating that the input dimension was definitely reduced to $n = 12$ using the active subspace.



(a)



(b)



(c)

Figure 5.4 Comparison of (a) weight on parameters determined by global linear model and (b) influential parameters identified by Partial Correlation and (c) Sobol decomposition.

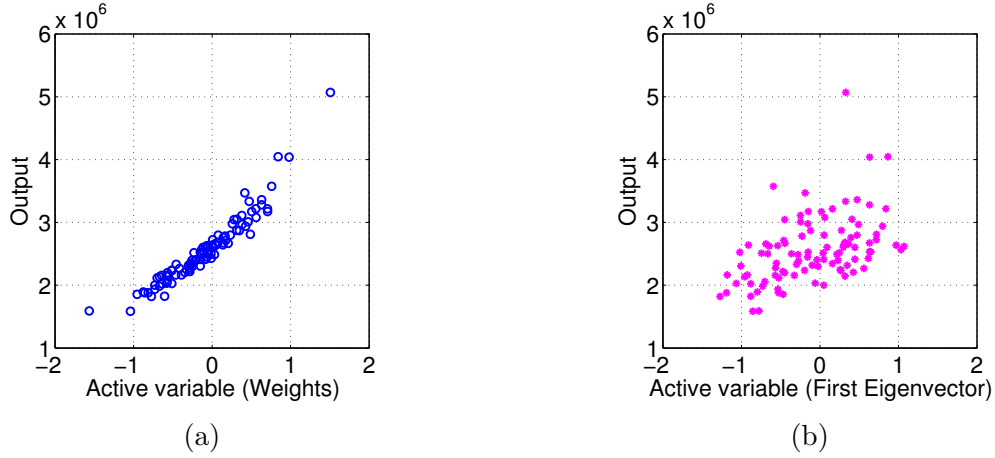


Figure 5.5 Sufficient summary plot produced using (a) the weights and (b) the first eigenvector.

A similar verification test performed in Chapter 4 indicated that one could reduce the input dimension by three using the parameter selection based on the global sensitivity analysis. We see here that by determining the parameter subspace using a linear combination of parameters, we can further reduce the input parameter dimensions.

5.5 Conclusion

We first discussed the methods for computing the gradients. Finite differences are frequently employed because of their ease of implementation. However, we observed that finite differences may not approximate gradients accurately. Computing gradients using sensitivity equations and automatic differentiation is more accurate. To avoid mistakes deriving sensitivity equations, we used automatic differentiation for the computation of gradients in this chapter.

The sufficient summary plots indicated that the model response can be approximately expressed in terms of one dimensional reduced input parameter subspace. The weights on parameters for the global linear model were almost identical to the sensitivity indices obtained using the global sensitivity analysis reported in Chapter 4. The density constructed using the weights performed reasonably well considering that the number of dimensions is one. However, comparison to the density using the full dimension indicated that one dimensional active subspace

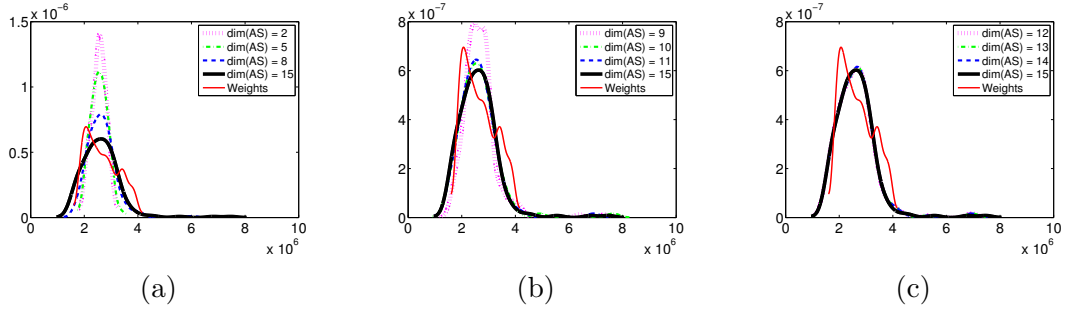


Figure 5.6 Probability density functions of model responses using (a) 2, 5, 8, (b) 9, 10, 11 and (c) 12, 13, 14, along with probability density functions using the full dimensional input space and the active subspace based on weights.

based on the weights is not sufficiently accurate to represent the response.

The active subspace constructed using eigenvectors of the gradients successfully reduced the input parameter dimension to $n = 12$. The active subspace with $n \geq 10$ also did reasonably well. The global sensitivity analysis indicated that 12 out of 15 parameters are influential, but the active subspace methods reduced the number of dimensions more than the global sensitivity did by finding a linear combination of parameters that has large impacts on the response.

BIBLIOGRAPHY

- [1] H.S. Abdel-Khalik. Hybrid uncertainty and sensitivity algorithms for high dimensional nonlinear models. *Transactions of the American Nuclear Society*, 103, 2010.
- [2] B.M. Adams, H.T. Banks, M. Davidian, Hee-Dae Kwon, H.T. Tran, S.N. Wynne, and E.S. Rosenberg. HIV dynamics: modeling, data analysis, and optimal treatment protocols. *J. Comput. Appl. Math.*, 184(1):10–49, 2005.
- [3] B.M. Adams, H.T. Banks, M. Davidian, and E.S. Rosenberg. Estimation and Prediction With HIV-Treatment Interruption Data. *Bulletin of Mathematical Biology*, 69(2):563–584, 2007.
- [4] B.M. Adams, M.S. Ebeida, M.S. Eldred, J.D. Jakeman, L.P. Swiler, W.J. Bohnhoff, K.R. Dalbey, J.P. Eddy, K.T. Hu, D.M. Vigil, L.E. Bauman, and P.D. Hough. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 5.3.1 user’s manual. *Sandia National Labs Report*, No. SAND2010-2183, 2013.
- [5] Y. Bang, H.S. Abdel-Khalik, and J.M. Hite. Hybrid reduced order modeling applied to nonlinear models. *Internat. J. Numer. Methods Engrg.*, 91(9):929–949, 2012.
- [6] H.T. Banks, R. Baraldi, K. Cross, K. Flores, C. McChesney, L. Poag, and E. Thorpe. Uncertainty Quantification in Modeling HIV Viral Mechanics. Technical report, CRSC-TR13-16, NC State University, Raleigh, NC, 2013. *Journal of Theoretical Biology*, submitted.
- [7] H.T. Banks, A. Cintrón-Arias, and F. Kappel. Parameter selection methods in inverse problem formulation. In *Mathematical modeling and validation in physiology*, volume 2064 of *Lecture Notes in Math.*, pages 43–73. Springer, Heidelberg, 2013.
- [8] H.T. Banks, M. Davidian, S. Hu, G.M. Kepler, and E.S. Rosenberg. Modelling HIV immune response and validation with clinical data. *J. Biol. Dyn.*, 2(4):357–385, 2008.
- [9] H.T. Banks, S.L. Ernstberger, and S.L. Grove. Standard errors and confidence intervals in inverse problems: sensitivity and associated pitfalls. *J. Inverse Ill-Posed Probl.*, 15(1):1–18, 2007.
- [10] H.T. Banks and H.T. Tran. *Mathematical and experimental modeling of physical and biological processes*. Textbooks in Mathematics. CRC Press, Boca Raton, FL, 2009.
- [11] M. Barenco, D. Tomescu, D. Brewer, R. Callard, J. Stark, and M. Hubank. Ranked prediction of p53 targets using hidden variable dynamic modeling. *Genome Biology*, 7(3):R25, 2006.
- [12] F.Y. Bois. GNU MCSim: Bayesian statistical inference for SBML-coded systems biology models. *Bioinformatics*, 25(11):1453–1454, 2009.
- [13] S.P. Brooks and G.O. Roberts. Convergence assessment techniques for Markov chain Monte Carlo. *Statistics and Computing*, 8(4):319–335, 1998.

- [14] D.G. Cacuci. *Handbook of Nuclear Engineering*. Springer-Verlag, New York, 2010.
- [15] B. Calderhead and M. Girolami. Estimating Bayes factors via thermodynamic integration and population MCMC. *Comput. Statist. Data Anal.*, 53(12):4028–4045, 2009.
- [16] D.S. Callaway and A.S. Perelson. HIV-1 infection and low steady state viral loads. *Bulletin of Mathematical Biology*, 64(1):29–64, 2002.
- [17] D. Campbell and R.J. Steele. Smooth functional tempering for nonlinear differential equation models. *Stat. Comput.*, 22(2):429–443, 2012.
- [18] F. Campolongo and A. Saltelli. Sensitivity analysis of an environmental model: an application of different analysis methods. *Reliability Engineering & System Safety*, 57(1):49–69, 1997.
- [19] H. Chen, Q. Wang, R. Hu, and P. Constantine. Conditional sampling and experiment design for quantifying manufacturing error of a transonic airfoil. *AIAA Aerospace Sciences Meeting*, 2011.
- [20] A. Cintrón-Arias, H. T. Banks, A. Capaldi, and A. L. Lloyd. A sensitivity matrix based methodology for inverse problem formulation. *J. Inverse Ill-Posed Probl.*, 17(6):545–564, 2009.
- [21] R. Confalonieri, G. Bellocchi, S. Bregaglio, M. Donatelli, and M. Acutis. Comparison of sensitivity analysis techniques: A case study with the rice model WARM. *Ecological Modelling*, 221(16):1897–1906, 2010.
- [22] P.G. Constantine, E. Dow, and Q. Wang. Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM J. Sci. Comput.*, 36(4):A1500–A1524, 2014.
- [23] P.G. Constantine, Q. Wang, A. Doostan, and G. Iaccarino. A surrogate-accelerated bayesian inverse analysis of the HyShot II flight data. 2011.
- [24] P.G. Constantine, Q. Wang, and G. Iaccarino. A method for spatial sensitivity analysis. pages 2011–2037, 2011.
- [25] E. Dow and Q. Wang. Output based dimensionality reduction of geometric variability in compressor blades. 2013.
- [26] J.J. Dunderstadt and L.J. Hamilton. *Nuclear Reactor Analysis*. John Wiley and Sons, New York, 1976.
- [27] N.M. Ferguson, F. deWolf, A.C. Ghani, C. Fraser, C.A. Donnelly, P. Reiss, J.M.A. Lange, S.A. Danner, G.P. Garnett, J. Goudsmit, and R.M. Anderson. Antigen-driven CD4+ T cell and HIV-1 dynamics: Residual viral replication under highly active antiretroviral therapy. *Proceedings of the National Academy of Sciences*, 96(26):15167–15172, 1999.
- [28] A. Gelman, F. Bois, and J. Jiang. Physiological pharmacokinetic analysis using population modeling and informative prior distributions. *J. Am. Stat. Assoc*, 91(436):1400–1412, 12 1996.

- [29] A. Gelman and D.B. Rubin. Inference from iterative simulation using multiple sequences. *Statist. Sci.*, 7(4):457–472, 1992.
- [30] M.A. Golberg and H.A. Cho. *Introduction to regression analysis*. WIT Press, Southampton, 2004.
- [31] U. Grenander and M.I. Miller. Representations of knowledge in complex systems. *J. Roy. Statist. Soc. Ser. B*, 56(4):549–603, 1994.
- [32] H. Haario, L. Kalachev, and M. Laine. Reduced models of algae growth. *Bull. Math. Biol.*, 71(7):1626–1648, 2009.
- [33] H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM: efficient adaptive MCMC. *Stat. Comput.*, 16(4):339–354, 2006.
- [34] J.M. Hite, H.S. Abdel-Khalik, R.C. Smith, M. Wentworth, E. Prudencio, and B. Williams. Uncertainty Quantification and Data Assimilation (UQ/DA) Study on a VERA Core Simulator Component for CRUD analysis. Technical report, CASL-I-2013-0184-000, Milestone Report for L2:VUQ:P7.02, 2013.
- [35] Z. Hu, R.C. Smith, N. Burch, M. Hays, and W.S. Oates. A modeling and uncertainty quantification framework for a flexible structure with macrofiber composite actuators operating in hysteretic regimes. *Journal of Intelligent Material Systems and Structures*, 25(2):204–228, 2014.
- [36] Y. Huang, D. Liu, and H. Wu. Hierarchical Bayesian methods for estimation of parameters in a longitudinal HIV dynamic system. *Biometrics*, 62(2):413–423, 2006.
- [37] M.J.W. Jansen. Analysis of variance designs for model output. *Computer Physics Communications*, 117(1–2):35–43, 1999.
- [38] D. Klinken II. An empirical Bayesian approach for model-based inference of cellular signaling networks. *BMC Bioinformatics*, 10(1), 2009.
- [39] M. Laine. MCMC toolbox for Matlab. <http://helios.fmi.fi/~lainema/mcmc/>, 2013.
- [40] M. Laine. Adaptive markov chain monte carlo methods. <https://wiki.helsinki.fi/display/inverse/Adaptive+MCMC>, 2014.
- [41] G. Lillacci and M. Khammash. Parameter estimation and model selection in computational biology. *PLoS Comput. Biol.*, 6(3):1–17, 2010.
- [42] J.S. Liu. *Monte Carlo strategies in scientific computing*. Springer Series in Statistics. Springer, New York, 2008.
- [43] D. Lu, M. Ye, M.C. Hill, E.P. Poeter, and G.P. Curtis. A computer program for uncertainty analysis integrating regression and Bayesian methods. *Environmental Modelling & Software*, 60(0):45–56, 2014.

- [44] X. Ma and N. Zabaras. An adaptive high-dimensional stochastic model representation technique for the solution of stochastic partial differential equations. *J. Comput. Phys.*, 229(10):3884–3915, 2010.
- [45] J.E. Mittler, M. Markowitz, D.D. Ho, and A.S. Perelson. Improved estimates for HIV-1 clearance rate and intracellular delay. *AIDS*, 13(11), 1999.
- [46] M.D. Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991.
- [47] A.S. Perelson, P. Essunger, Y. Cao, M. Vesanen, A. Hurley, K. Saksela, M. Markowitz, and D.D. Ho. Decay characteristics of HIV-1-infected compartments during combination therapy. *Nature*, 387(6629):188–191, 1997.
- [48] H. Rabitz and Ö.F. Alış. General foundations of high-dimensional model representations. *J. Math. Chem.*, 25(2-3):197–233, 1999.
- [49] B. Ramratnam, S. Bonhoeffer, J. Binley, A. Hurley, L. Zhang, J.E. Mittler, M. Markowitz, J.P. Moore, A.S. Perelson, and D.D. Ho. Rapid production and clearance of HIV-1 and hepatitis C virus assessed by large volume plasma apheresis. *The Lancet*, 354(9192):1782–1785, 1999.
- [50] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009.
- [51] S. Rogers, R. Khanin, and M. Girolami. Bayesian model-based inference of transcription factor activity. *BMC Bioinformatics*, 8:S2, 2007.
- [52] T.M. Russi. *Uncertainty quantification with experimental data and complex system models*. ProQuest LLC, Ann Arbor, MI, 2010. Thesis (Ph.D.)—University of California, Berkeley.
- [53] G.A. Saber and C.J. Wild. *Nonlinear Regression*. Wiley-Interscience, NY, 2003.
- [54] J. Sacks, W.J. Welch, T.J. Mitchell, and H.P. Wynn. Design and analysis of computer experiments. *Statist. Sci.*, 4(4):409–435, 1989.
- [55] A. Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297, 2002.
- [56] A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Comput. Phys. Comm.*, 181(2):259–270, 2010.
- [57] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global sensitivity analysis. The primer*. John Wiley & Sons, Ltd., Chichester, 2008.

- [58] I. Siekmann, L.E. Wagner II, D. Yule, C. Fox, D. Bryant, E.J. Crampin, and J. Sneyd. MCMC estimation of Markov models for ion channels. *Biophysical Journal*, 100(8):1919–1929, 2011.
- [59] R.C. Smith. *Uncertainty Quantification: Theory, Implementation and Applications*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2014.
- [60] I.M. Sobol’. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math. Comput. Simulation*, 55(1-3):271–280, 2001.
- [61] I.M. Sobol’, S. Tarantola, D. Gatelli, S.S. Kucherenko, and W. Mauntz. Estimating the approximation error when fixing unessential factors in global sensitivity analysis. *Reliability Engineering & System Safety*, 92(7):957–960, 2007.
- [62] T. Sumner, E. Shephard, and I.D.L. Bogle. A methodology for global-sensitivity analysis of time-dependent outputs in systems biology modelling. *Journal of The Royal Society Interface*, 9(74):2156–2166, 2012.
- [63] G.J. Székely and M.L. Rizzo. Testing for equal distributions in high dimensions. *InterStat*, 143(5), 2004.
- [64] G.J. Székely and M.L. Rizzo. Energy statistics: a class of statistics based on distances. *J. Statist. Plann. Inference*, 143(8):1249–1272, 2013.
- [65] J.A. Vrugt. Software: Matlab packages. <http://faculty.sites.uci.edu/jasper/sample/>, 2014.
- [66] J.A. Vrugt, C.J.F. ter Braak, M.P. Clark, J.M. Hyman, and B.A. Robinson. Treatment of input uncertainty in hydrologic modeling: Doing hydrology backward with Markov chain Monte Carlo simulation. *Water Resources Research*, 44(12), 2008.
- [67] J.A. Vrugt, C.J.F. ter Braak, C.G.H. Diks, B.A. Robinson, J.M. Hyman, and D. Higdon. Accelerating Markov chain Monte Carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *International Journal of Nonlinear Sciences and Numerical Simulation*, 10(3):273–290, 2009.
- [68] J.A. Vrugt, C.J.F. ter Braak, H.V. Gupta, and B.A. Robinson. Equifinality of formal (DREAM) and informal (GLUE) Bayesian approaches in hydrologic modeling? *Stochastic Environmental Research and Risk Assessment*, 23(7):1011–1026, 2009.