Optimal Sensor Placement for Bayesian Parametric Identification of Structures

Thesis by
Pinaky Bhattacharyya

In Partial Fulfillment of the Requirements for the
degree of
Doctor of Philosophy in Civil Engineering

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2017
Defended December 5, 2016

Pinaky Bhattacharyya
ORCID: [0000-0003-3773-0392]

ACKNOWLEDGEMENTS

This thesis would not have been possible without the mentorship and guidance of my advisor, Prof. James (Jim) Beck, who has helped me learn and investigate the subject matter, not in small part through long afternoon discussions at Steele House. I am also very grateful for his help through tough spots in my career. Working with him has been a great privilege and also very enjoyable.

I would also like to thank Professors Joel Burdick, Steven Low and Thomas Heaton for their inputs as my thesis committee members.

The administrative staff at Gates-Thomas and Annenberg, the International Students Program, Caltech Security and the rest of the Caltech support system, including Ernie, have worked hard to keep my stay at Caltech smooth and for that I am very thankful.

I would like to thank Professors Ravichandran and Bhattacharya for the teaching assistantship opportunities and guidance; Dr. Michael Mello for several fun and informative hours in the mechanical engineering lab. Thanks also to Prof.Swaminathan Krishnan for helping me get started at Caltech.

Among many others, I'd like to thank my friends Keng-Wit, Jonathan, Subrahmanyam, Priya, Hemanth, Ramses, Ramya, Scott, Eric, Lucy and Utkarsh as well as Sujeet, Vikas and Nisha for late night discussions over donuts at Winchell's.

Finally, I would like to thank my family for their unwavering support and for putting up with my odd hours and erratic contact through these six years.

ABSTRACT

There exists a choice in where to place sensors to collect data for Bayesian model updating and system identification of structures. It is desirable to use an available deterministic predictive model, such as a finite-element model, along with prior information on the uncertain model parameters and the uncertain accuracy of the predictive model, to determine which optimal sensor locations should be instrumented in the structure. In this thesis, an information-theoretic framework for optimality is considered.

The mutual information between the uncertain model predictions for the data and the uncertain model parameters is presented as a natural measure of reduction in uncertainty to maximize over sensor configurations. A combinatorial search over all sensor configurations is usually prohibitively expensive. A convex optimization method is developed to provide a fast sub-optimal, but possibly optimal, sensor configuration when certain simplifying assumptions can be made about the chosen stochastic model class for the structure. The optimization method is demonstrated to work for a 50-story uniform shear building, with 20 sensors to be installed.

The stability of optimal sensor configurations under refinement of the mesh of the underlying finite-element model is investigated and related to the choice of prediction-error correlations in the model. An example problem of placement of a single sensor on the continuum of an elastic axial bar is solved analytically.

In order to solve the optimal sensor placement problem in the more general case, numerical estimation of mutual information between the model predictions for the data and the model parameters becomes necessary. To this end, a thermodynamic integration scheme based on path sampling is developed with the aim of estimating the entropy of the data prediction distribution. The scheme is demonstrated to work for an example that uses synthetic data for model class comparison between linear and Duffing oscillator model classes. The thermodynamic integration method is then used to determine the optimal location of a single sensor for a two degree-of-freedom oscillator model.

TABLE OF CONTENTS

Acknowledgements . . . iii
Abstract . . . iv
Table of Contents . . . v
List of Illustrations . . . vii
List of Tables . . . ix
Chapter I: Introduction . . . 1
1.1 General Overview . . . 3
1.2 Contributions . . . 3
1.3 Thesis outline . . . 4
1.4 A note on terminology . . . 5
Chapter II: Mutual Information-based Optimal Sensor Placement . . . 7
2.1 Introduction . . . 7
2.2 Bayesian model updating . . . 8
2.3 Mutual information . . . 10
2.4 Statement of the general problem . . . 13
2.5 Differences in mutual information . . . 15
2.6 Computing the mutual information . . . 15
2.7 Sensor rearrangement problem . . . 23
Chapter III: Efficient Solution Using Convex Relaxation . . . 26
3.1 Log-determinant formulation . . . 26
3.2 Entropy-based optimal sensor location . . . 29
3.3 Convex relaxation of the combinatorial optimization problem . . . 31
3.4 Application to multistory structure . . . 36
Chapter IV: Modeling allowable sensor locations . . . 46
4.1 Sensor placement on a continuum . . . 46
4.2 Finite-element mesh refinement . . . 48
4.3 Comparison with the finite selection-space case . . . 52
4.4 Example: free-vibration of a cantilever beam with prediction-error correlations . . . 57
Chapter V: Bayesian model class selection using thermodynamic integration . 63
5.1 Introduction . . . 63
5.2 Numerical Evaluation of Model Evidence . . . 64
5.3 Duffing oscillator . . . 67
5.4 Numerical example . . . 70
5.5 Simulation Results . . . 76
5.6 Conclusion . . . 83
Chapter VI: Conclusions and Future Work . . . 85
6.1 Future Work . . . 86
Appendix A: Matrix Calculus Identities . . . 87

A.1 Useful matrix calculus identities and properties 87
Appendix B: Derivatives for Convex Problem 89
B.1 Calculation of parametric-gradient for modal response 89

LIST OF ILLUSTRATIONS

Number *Page*

1.1 Organization of the thesis 4
3.1 The resulting sensor locations after solving the convex optimization problem for a 50-story structure 41
3.2 Another solution for the 50-story building but with different prior samples 42
4.1 Optimal location of the second sensor as a function of the prediction-error correlation length scale (note that the first sensor is always at the free end of the cantilever beam) 61
4.2 Additional bits of information from the optimally placed second sensor plotted as a function of the correlation length scale 62
5.1 The closed state-space trajectory of a double-well Duffing oscillator . 68
5.2 Period doubling bifurcation for forced, damped Duffing oscillator . . 69
5.3 Chaotic trajectory for a forced, damped Duffing oscillator 69
5.4 Displacement trajectory plots of the data data from a linear oscillator (blue) and a Duffing oscillator (red) 73
5.5 Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^L, \mathcal{M}^L)$ . . 76
5.6 Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^L, \mathcal{M}^D)$ . . 77
5.7 Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^D, \mathcal{M}^L)$ . . 77
5.8 Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^D, \mathcal{M}^D)$ . . 78
5.9 Displacement trajectory plots for data from a linear oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a linear oscillator model class 79
5.10 Displacement trajectory plots for data from a linear oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a Duffing oscillator model class 79
5.11 Probability density function for the log-normal prior marginal for the non-linear coefficient, α 80
5.12 Displacement trajectory plots for data from a Duffing oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a linear oscillator model class 80

5.13 Displacement trajectory plots for data from a Duffing oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a Duffing oscillator model class 81
5.14 Thermodynamic integration curve for the $(\mathcal{D}^L, \mathcal{M}^L)$ case 83

LIST OF TABLES

Number *Page*

2.1 Simulation results for entropy-of-evidence 20
2.2 Conditional entropy 20
2.3 Mutual information estimate for a single DOF system 21
2.4 Single degree-of-freedom example: Simulation parameters 21
2.5 Two degree-of-freedom example: Simulation parameters 22
2.6 Mutual information estimates for 2-DOF system 22
3.1 Simulation results 40
3.2 Description of various sensor configurations for comparative example 43
3.3 Comparison of information gain relative to optimal configuration . . 43
5.1 Specification of the data-generating parameters for $\mathcal{M}^L$ and $\mathcal{M}^D$ . . . 71
5.2 Specification of priors for $\mathcal{M}^L$ and $\mathcal{M}^D$ 72
5.3 Data generating samples and MAP estimates for each scenario 81
5.4 AIMS parameters 82
5.5 Log-evidence values from thermodynamic integration 82

Chapter 1

INTRODUCTION

An important goal of the field of structural monitoring is to use dynamic response data from an instrumented structure to perform system identification to quantify engineering quantities of interest such as substructure stiffness, as well as to make predictions about the future behavior of the structure. Response data is typically available in the form of acceleration records from ambient vibrations, vibration tests or seismic events. The parameters of a mechanical model that explains the behavior of a structure are typically sought. To this end, parametric system identification is used to quantify the values of parameters of the model that is assumed to govern the structure.

It is often not possible to postulate a true and exact model that explains the observations and makes good predictions. The primary reason for this is that the mechanical model used to describe the structure, however complex, will usually fail to account for every mechanical phenomenon, geometric and material detail correctly. In addition, it is not clear how complex and detailed a model is appropriate, since a sufficiently adjustable model would always be able to explain the given observations at the cost of being able to make accurate predictions about future observations.

Given these limitations, in order to formulate a more meaningful system identification problem, it is important to account for modeling uncertainties related to the prediction accuracy and the parameters associated with the mechanical model being considered. This is done by specifying a probability model for predicting the response based on the deterministic mechanical model of the structure and a probability distribution over the uncertain parameters of the mechanical model, together known as a stochastic embedding of the parameterized deterministic model class, to form a stochastic model class [1].

The axioms of Bayesian probability form a foundation for performing plausible inference in a principled manner. Therefore, the data, together with the stochastic embedding, allow for the update of the distribution over the uncertain parameters of the model class.

The structural response data obtained is from locations in the structure that have been chosen to be instrumented. There is an element of choice as to the nature of

instrumentation to be performed. Given that Bayesian inference will be done with the data, there is a problem to be addressed: the choice of optimal data-set to be collected for this inference.

This thesis will address the problem of optimal sensor location for Bayesian parameteric system identification. For the purpose of illustration, consider a simple version of the sensor-selection problem of selecting a certain number of locations to be instrumented for structural monitoring, from a finite number of possibilities. For instance, one could imagine being budgeted with 20 sensors to be used to instrument a 50-story structure. In this situation, the number of possible configurations is roughly 47 trillion! The evaluation of any objective function for each such configuration would be infeasible for even state-of-the-art computers available as of this writing.

There are several important factors to be considered when formulating an optimal sensor location problem:

- The measure of optimality used
- The type of data to be collected
- The choice of locations available for instrumentation
- Prior information about the structure to be instrumented

Each of these factors can have a significant impact on the solution to the problem and its solvability and will be addressed in this thesis.

The primary objective of this research is to introduce new methods for more efficiently solving existing formulations of the optimal sensor placement problem in the context of structural engineering. In addition, the assumptions made by contemporary formulations are relaxed to determine if the resulting more general framework is solvable.

Optimality in the current research is based on objective functions that quantify an information-theoretic measure of the anticipated uncertainty of the parameters to be inferred, or quantities derived from them. This is in contrast to other objectives such as those based on financial cost, which have been considered in the literature [2].

1.1 General Overview

In the domain of earthquake engineering and structural dynamics, the optimal sensor location problem can be cast into a determinant maximization problem for an entropy-based objective function, if certain assumptions can be made about the stochastic model class. The first is that it should be globally identifiable under the selected data. In addition, it is assumed that a sufficiently large number of data points would be available so as to justify the use of the Laplace approximation to the posterior that would be obtained. Finally, a small parametric uncertainty is assumed. The optimization problem here is still one over all the possible sensor configurations available The use of genetic algorithms was proposed in [3] to obtain a sub-optimal and possibly near-optimal solution to the problem.

When such assumptions cannot be made for the problem under consideration, then the computational complexity of and theoretical insight into the problem deteriorates significantly. More advanced theoretical and computational formulations are needed to tackle such a problem.

So far, nothing has been said about what the possible sensor locations are or can be. In some problems, the possible sensor locations can be a pre-ordained constraint. For instance, consider the problem of selecting which stories of a building are to be instrumented based on a shear-building model. If the possible location in the floor for sensor placement in each story is fixed, then the possible sensor locations are already determined by the constraint of being only on floors.

In some problems, however, the possible sensor locations is a modeling choice. There are certain issues that can arise in this scenario. For instance, if there exists a highly informative "hotspot", then the algorithm may have a tendency to cluster all available sensors near that location. This in turn is a result of a flaw in the modeling of the mutual informativeness of the sensors in such a configuration. This issue will be addressed here.

1.2 Contributions

A contribution of this work has been to improve the existing approximate sensor location problem by providing a sub-optimal solution much more efficiently than combinatorial algorithms by using existing techniques from the theory of convex optimization. This method of solution may produce an optimal solution to the problem, although there are no guarantees as to when this happens. Otherwise, the method is able to provide a bound on how sub-optimal the resulting sensor

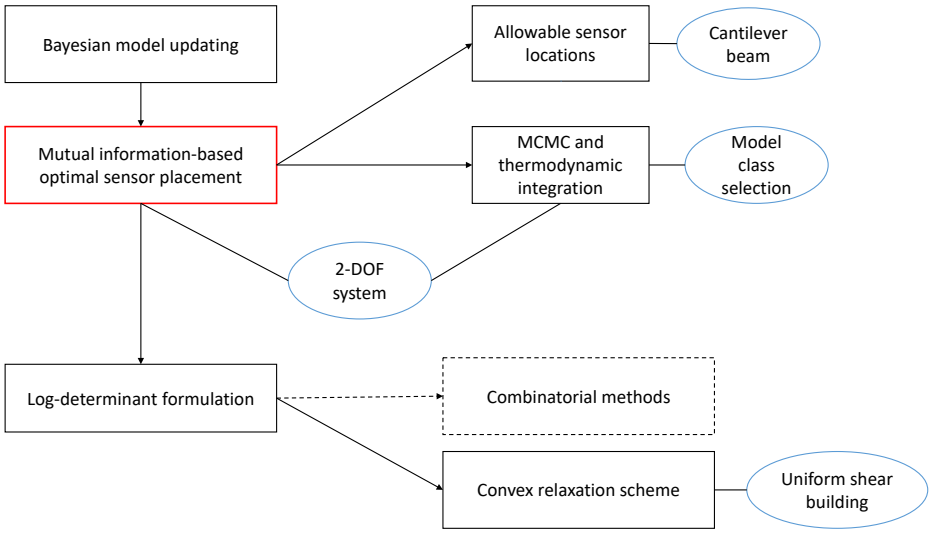

Figure 1.1: Organization of the thesis

configuration is.

In the process of attempting to computationally solve the fully general information-theoretic formulation, the evaluation of the Bayesian probabilistic quantity known as the evidence for the stochastic model class is necessary. A framework known as path sampling is shown to naturally arise out of a specific sampling algorithm. The model evidence thus calculated is used to perform model class selection using synthetic data from a linear oscillator versus one from a cubic non-linear, or Duffing, oscillator.

Finally, the current research draws on existing information-theoretic concepts to theoretically analyze and model certain specific situations that are encountered in the sensor selection problem, such as the effect of grid refinement on sensor selection, as well as sensor relocation.

1.3 Thesis outline

Figure 1.1 presents a flowchart depicting the thesis outline.

Chapter 2 briefly recaps the problem of Bayesian inference in a structural dynamics context and lays out the sensor placement problem in its full generality as a mutual information maximization problem while discussing theoretical and computational aspects of the problem. Since the estimation of mutual information is very expensive, the sensor selection problem is tackled for only a two degree-of-freedom problem.

In Chapter 3 the problem is tackled for the situation in which simplifying assumptions about model class identifiablity and the suitability of applying Laplace's approximation to the posterior that is to be obtained, can be made. Methods developed in the field of convex optimization are incorporated into the solution method for the problem, in order to avoid a combinatorially large search for the optimal sensor configuration.

Chapter 4 discusses the effect of choices regarding the allowable sensor locations. Optimal sensor location on a continuous elastic bar is used as an example to illustrate the sensor placement on a continuum and to illustrate the relationship between the objective function and the mesh size. Finally, the effect of prediction-error correlation on sensor locations is discussed in the context of a cantilever beam.

Chapter 5 is a bit of a detour from the optimal sensor selection problem. The focus here is on demonstrating how a path sampling scheme arises naturally out of the AIMS Markov sampling algorithm for drawing posterior samples and can be used to estimate model evidence. The problem of model class selection between a linear and non-linear Duffing oscillator is tackled for data from either type of oscillator, to demonstrate consistency with existing interpretations of model evidence.

1.4 A note on terminology

The nomenclature for probabilistic quantities in Bayesian inference is based on the presumption that the data is already available. In the optimal sensor location problem, the data is not available yet. Thus, expressions for already named quantities but with the term data considered variable are called by a different name. The term, "stochastic forward model" is used for the probability density of the stochastic prediction of the data given the parameters. Note that the prior distribution is still referred to as the prior.

References

[1] J. L. Beck, L. S. Katafygiotis, "Updating Models and Their Uncertainties. I: Bayesian Statistical Framework," *Journal of Engineering Mechanics*, vol. 124, no. 4, pp. 455–461, Apr. 1998.

[2] J. Mitrani-Resier, S. Wu, J. L. Beck, "Virtual Inspector and its application to immediate pre-event and post-event earthquake loss and safety assessment of buildings," *Natural Hazards*, vol. 81, no. 3, pp. 1861–1878, Apr. 2016.

[3] C. Papadimitriou, J. L. Beck, S.-K. Au, "Entropy-Based Optimal Sensor Location for Structural Model Updating," *Journal of Vibration and Control*, vol. 6, no. 5, pp. 781–800, Jan. 2000.

Chapter 2

MUTUAL INFORMATION-BASED OPTIMAL SENSOR PLACEMENT

The Bayesian optimal sensor placement problem, in its full generality, seeks to maximize the mutual information between the uncertain parameters and the data to be collected for the purpose of performing Bayesian inference. This maximization is done over all possible sensor configurations for a given sensor budget.

Thus, in this chapter, a recap of the Bayesian model updating framework is first presented. The general expression for mutual information between prior and data for typical stochastic dynamical systems encountered in structural dynamics is presented next. Different interpretations of the mutual information objective are considered.

It is established that the computation of mutual information in the context of dynamical systems is computationally expensive. In order to make the problem tractable, certain simplifying assumptions are made to yield a more tractable objective function. Later, Chapter 3 discusses an efficient convex relaxation scheme to determine possibly-optimal solution to the problem.

2.1 Introduction

The mutual information between two stochastic variables is an important quantity in experimental design for Bayesian system identification of stochastic dynamical systems. In this context, several possible interpretations of mutual information exist.

Prior information on the parameter values is expressed in terms of a prior distribution, $p(\theta|\mathcal{M})$. The stochastic forward model, typically denoted by $p(y|\theta,\mathcal{M})$, prescribes the probability of observing a particular data set given the parameters. The probability model class, $\mathcal{M}$, is used to indicate a specific pairing of stochastic forward model and prior. The stochastic forward model with the actual data $\mathcal{D}$ substituted for y, when viewed simply as a function of the uncertain parameters, θ, is called the likelihood function.

Bayes' rule implies that the posterior probability density, $p(\theta|y,\mathcal{M})$, is proportional to the product of the likelihood and prior probability densities. The reciprocal of the normalizing factor is called the evidence, which when viewed as a function of

the stochastic predictions can be denoted by $p(y|\mathcal{M})$. This can be determined using the Total Probability Theorem, as $p(y|\mathcal{M}) = \int p(y|\theta, \mathcal{M})p(\theta|\mathcal{M})d\theta$

For a given data set, the parameter vector that maximizes the likelihood function is known as the maximum likelihood estimate (MLE), usually denoted by $\hat{\theta}_{MLE}$. Similarly, the parameter vector that maximizes the posterior probability is known as the maximum *a posteriori* (MAP) estimate of the parameters, denoted usually by $\hat{\theta}_{\text{MAP}}$.

2.2 Bayesian model updating

This section lays out the framework for Bayesian model updating for structural dynamics which forms the basis for the optimal sensor placement problem. The measured data on which inference is to be performed is assumed to be available in the form of dynamic test data in the time domain. It is assumed, here, that the complete system input can be determined given the parameters.

In a typical structural dynamics problem, one would have observed the acceleration, velocity or displacement time history at a certain number of locations. We shall assume, for the time being, that these locations are selected from among the degrees-of-freedom of a finite-element model that governs the structure.

Denote the uncertain system parameters by θ_s. These could correspond, for instance, to sub-structure stiffness, Rayleigh damping parameters or input ground acceleration amplitude. In general, they would depend on the mechanical model of the structure and the finite element approximation used to predict its response.

The prediction-error parameters are denoted by θ_e. When the prediction errors are modeled as Gaussian, with a stationary and isotropic co-variance matrix with diagnonal entries σ^2, θ_e could just be the scalar, σ. For an anisotropic model for the co-variance matrix without prediction error correlation between different locations, this could be a vector of variances corresponding to the diagonal entries of the co-variance matrix.

Denote by y_n the stochastic prediction of the quantity to be observed at the time-point t_n. Note that $y_n \in \mathbb{R}^{N_o}$, where N_o is the number of observed degrees-of-freedom out of a total of N_d degrees-of-freedom. The set of stochastic predictions for the observations over the whole duration of the measurement is denoted by $y_{1:N}$.

The vector of deterministic predictions at each time-step of interest is denoted by $x_n(\theta_s)$, where $x_n(\theta_s) \in \mathbb{R}^{N_d}$. The prediction-error is additive and is denoted by

ϵ_n. $\delta \in \{0,1\}^{N_d}$ is a Boolean vector that indicates the presence or absence of a measurement sensor, so that the sum of entries of δ equals the number of observed DOFs, N_o.

A sensor selection matrix, $S_o(\delta) \in \mathbb{R}^{N_o \times N_d}$ selects the components that are to be observed from the full stochastic prediction given by the sum of the deterministic prediction and noise terms.

The equation relating stochastic predictions to be observed, to the deterministic predictions, is, then:

$$y_n = S_o(\delta)\left(x_n(\theta_s) + \epsilon_n(\theta_e)\right) \tag{2.1}$$

In Equation (2.1), δ refers to a Boolean vector whose entries equal 1 when the response from a particular degree-of-freedom is observed, and 0 when not. The matrix $S_o(\delta)$ is a rectangular selection matrix and picks out those entries corresponding to observed degrees of freedom.

Let $\mathcal{M}$ denote the probability model class. This includes the deterministic forward model with a stochastic embedding. The stochastic embedding, in turn, is specified by the stochastic forward model together with a prior distribution over the uncertain parameters. Our statistical modeling assumptions lead us to the following stochastic forward model:

$$p(y_{1:N}|\theta, \mathcal{M}) = \prod_{n=1}^{N} p\left(y_n|\theta, \mathcal{M}\right) \tag{2.2}$$

Bayes' theorem can be used to perform the update as in Equation (2.3):

$$p(\theta|y_{1:N}, \mathcal{M}) = \frac{p(y_{1:N}|\theta, \mathcal{M})p(\theta|\mathcal{M})}{p(y_{1:N}|\mathcal{M})} \tag{2.3}$$

The term in the denominator is called the evidence for the model class, $\mathcal{M}$, under the data, $\mathcal{D}_N$ and is given by the Total Probability Theorem in Equation (2.4):

$$p(y_{1:N}|\mathcal{M}) = \int p(y_{1:N}|\theta, \mathcal{M})p(\theta|\mathcal{M})d\theta \tag{2.4}$$

This quantity will be dealt with in greater detail in Chapter 5 for the purposes of model selection.

The Prediction-Error Model

The examples here will deal with the case where we model the prediction errors as Gaussian with equal variance at every degree of freedom, and uncorrelated across

time and location, so that:

$$\mathbb{E}\left[\epsilon_{i,n}\epsilon_{j,m}\right] = \sigma^2\delta_{ij}\delta_{mn} \text{ where } \epsilon_{i,n} \sim \mathcal{N}(\cdot|0,\sigma^2) \text{ i.i.d.} \tag{2.5}$$

We then have an expression for the stochastic forward model:

$$p(y_{1:N}|\theta,\mathcal{M}) = \prod_{n=1}^{N} \mathcal{N}\left(y_n|S_o(\delta)x_n(\theta),\sigma^2\right) \tag{2.6}$$

Using Bayes' theorem, the distribution for the posterior, given the data is:

$$\begin{aligned} p(\theta|y_{1:N},\mathcal{M}) &= \frac{p(y_{1:N}|\theta,\mathcal{M})p(\theta|\mathcal{M})}{p(y_{1:N}|\mathcal{M})} \\ &= c^{-1}p(y_{1:N}|\theta,\mathcal{M})p(\theta|\mathcal{M}) \end{aligned} \tag{2.7}$$

The evidence, c, for the model class in Equation (2.7) is treated simply as a normalizing constant for the purpose of this section. In later sections such as in Chapters 5, it will be viewed as a function of the stochastic predictions as well as the model class.

The logarithm of the posterior is sometimes easier to work with. In the case of this example, we have:

$$\log p(\theta|y_{1:N},\mathcal{M}) = -\frac{NN_o}{\log 2\pi\sigma^2} - \frac{1}{2\sigma^2}\sum_{n=1}^{N}\|y_n - S_o(\delta)x_n(\theta_s)\|^2 + \log p(\theta|\mathcal{M}) \tag{2.8}$$

Note that when observations are actually made, then the collection of data at hand is denoted by $\mathcal{D}_N = \{\hat{y}_n\}_{n=1}^N$ or $\hat{y}_{1:N}$.

2.3 Mutual information

The mutual information between two stochastic variables is a degree of measure about how much they imply about each others' possible values. Formally, the mutual information between two continuous stochastic variables is defined as follows:

$$I(X;Z) = \mathbb{E}_{X,Z}\log\left[\frac{p_{X,Z}(x,z)}{p_X(x)p_Z(z)}\right] \tag{2.9}$$

The mutual information between two sets of stochastic variables is zero if and only if they are independent, since the joint density factors into the marginals. Additionally, if two stochastic variables are linearly correlated as in the case of a bivariate Gaussian with a non-zero correlation coefficient, then the mutual information controls the correlation coefficient, ρ:

$$I(X;Z) = -\frac{1}{2}\log(1-\rho^2) \tag{2.10}$$

The mutual information can be expressed as a difference of two entropies:

$$I(X;Z) = H(X) - H(X|Z) \tag{2.11}$$

Equation (2.11) states that the mutual information is the difference in entropy of the marginal distribution for one variable and the conditional entropy. The expression is symmetric in either variable because of Bayes' rule.

The conditional entropy is the entropy of one variable, conditioned on a particular value of the other and averaged over that value according to its probability distribution. That is,

$$H(X|Z) = \int P_Z(z)[H(X|Z=z)]\,dz \tag{2.12}$$

The mutual information is always non-negative, since conditioning always reduces entropy or leaves it unchanged, on average.

In the Bayesian system identification setting, the mutual information quantity has several interpretations, that are discussed in the following sections.

Mutual information and Bayesian system identification

With the usual notation, the mutual information between the stochastic predictions for the observed data, $y_{1:N}$, and the parameter vector, θ, is:

$$I(y_{1:N}(\delta);\theta) = \mathbb{E}_{y_{1:N}(\delta),\theta|\mathcal{M}}\left[\ln \frac{p(y_{1:N}(\delta),\theta|\mathcal{M})}{p(y_{1:N}(\delta)|\mathcal{M})p(\theta|\mathcal{M})}\right] \tag{2.13}$$

Equation (2.13) emphasizes the dependence of the stochastic prediction for the data, $y_{1:N}(\delta)$, on the sensor configuration, δ. This explicit dependence will be omitted for convenience until a discussion of the optimal sensor location problem in Section 2.4.

Divergence from posterior to prior

The equation (2.13) can be rewritten after applying Bayes' theorem to express mutual information as the expected Kullback-Leibler (KL) divergence, or expected relative entropy, from the posterior to the prior. That is:

$$I(y_{1:N}(\delta), \theta) = \mathbb{E}_{y_{1:N}(\delta),\theta|\mathcal{M}} \left[\ln \frac{p(y_{1:N}(\delta), \theta|\mathcal{M})}{p(y_{1:N}(\delta)|\mathcal{M})p(\theta|\mathcal{M})} \right] \tag{2.14}$$

$$= \int \int p(y_{1:N}(\delta), \theta|\mathcal{M}) \left[\ln \frac{p(\theta|y_{1:N}(\delta), \mathcal{M})}{p(\theta|\mathcal{M})} \right] dY d\theta \tag{2.15}$$

$$= \int \int p(\theta|y_{1:N}(\delta), \mathcal{M}) p(y_{1:N}(\delta)|\mathcal{M}) \left[\ln \frac{p(\theta|y_{1:N}(\delta), \mathcal{M})}{p(\theta|\mathcal{M})} \right] dY d\theta \tag{2.16}$$

$$= \mathbb{E}_{y_{1:N}(\delta)|\mathcal{M}} D_{KL}[p(\theta|y_{1:N}(\delta), \mathcal{M})||p(\theta|\mathcal{M})] \tag{2.17}$$

Here, the Kullback-Leibler divergence from one continuous probability distribution, P, to another, Q, specified by their probability density functions $p(x)$ and $q(x)$, is:

$$D_{KL}(p(x)||q(x)) = \int_{-\infty}^{\infty} p(x) \log \frac{p(x)}{q(x)} \, dx \tag{2.18}$$

The KL-divergence in this case can also be referred to as the relative entropy of P with respect to Q.

If one thinks of information maximization in the sensor placement framework, then maximizing this expected KL divergence would be equivalent to choosing the sensor configuration that has the largest expected reduction in entropy of the posterior distribution from the prior. Of all sensor configurations, the optimal one would result in the maximum reduction in entropy, on average, upon measuring the data.

Expression in terms of known quantities

The mutual information in a Bayesian system identification framework, usually involves only two specified distributions: the prior and the stochastic forward model. We can use Bayes' theorem to express Equation (2.13) in terms of these distributions as:

$$I(y_{1:N}(\delta); \theta) = \int \int p(y_{1:N}(\delta)|\theta, \mathcal{M}) p(\theta|\mathcal{M}) \left[\log \frac{p(y_{1:N}(\delta)|\theta, \mathcal{M})}{p(y_{1:N}(\delta)|\mathcal{M})} \right] dy_{1:N}(\delta) d\theta \tag{2.19}$$

The denominator in the logarithm term is what one would associate with the evidence term, based on given data. Here, the data is not available and this term is simply a function obtained upon marginalizing out the parameters from the joint:

$$p(y_{1:N}(\delta)|\mathcal{M}) = \int p(y_{1:N}(\delta)|\theta, \mathcal{M})p(\theta|\mathcal{M})d\theta \tag{2.20}$$

In principle, all these quantities should be calculable, given $p(\theta)$ and $p(y_{1:N}(\delta)|\theta, \mathcal{M})$. In practice, calculating these quantities turns out to be quite challenging. Section 2.6 addresses this issue.

2.4 Statement of the general problem

We propose to use the mutual information between the uncertain parameters and the data variable as the objective function to be maximized over sensor configuration. Consider the objective function:

$$\begin{aligned} &\underset{\delta}{\text{maximize}} && I(y_{1:N}(\delta); \theta) \\ &\text{subject to} && \delta_i \in \{0, 1\}, \qquad i = 1, \ldots, N_d \\ &\text{and} && \mathbb{1}^T\delta = N_o. \end{aligned} \tag{2.21}$$

Equation (2.19) can be used for the objective function for the problem in (2.21). Recall that δ is a binary N_d-dimensional vector specifying the sensor configuration over the degrees of freedom in the system.

Log-determinant approximation

The mutual information criterion can be simplified to a more manageable log-determinant criterion when the following assumptions can be made about the probability model class:

- The posterior distribution, upon collection of the data, is globally identifiable
- The Laplace approximation holds; that is, a sufficient of data points are measured so that the posterior is approximately Gaussian

Under these assumptions, the expression for mutual information may be simplified as follows. First, recognize the expression for mutual information as the KL-divergence from the posterior to the prior:

$$I(y_{1:N}(\delta); \theta) = \mathbb{E}_{y_{1:N}(\delta),\theta}\left[\log \frac{p(\theta|y_{1:N}(\delta))}{p(\theta)}\right] \tag{2.22}$$

Begin by noting that the prior term is irrelevant in the sensor placement problem and may be discarded to yield the utility function, $U(\delta)$, as:

$$U(\delta) = \mathbb{E}_{y_{1:N}(\delta)}\left[p(\theta|y_{1:N}(\delta)) \log p(\theta|y_{1:N}(\delta))\right] \tag{2.23}$$

Next, apply the Gaussian approximation to the posterior. That is,

$$p(\theta|y_{1:N}(\delta)) \approx \mathcal{N}(\theta|\hat{\theta}(y_{1:N}), A^{-1}(\hat{\theta}(y_{1:N}), \delta)) \tag{2.24}$$

Note that the precision matrix here is given by the Hessian of the logarithm of the stochastic forward model:

$$[A(\hat{\theta}(y_{1:N}), \delta)]_{pq} = -\frac{\partial^2 \log p(y_{1:N}(\delta)|\theta)}{\partial\theta_p\theta_q} \tag{2.25}$$

The utility function becomes is now simply related to the entropy of this approximately Gaussian posterior which may be expressed in terms of the log-determinant of its precision matrix as:

$$\Rightarrow U(\delta) = -\mathbb{E}_{y_{1:N}(\delta)}\left[\text{constant} - \frac{1}{2}\log\det A(\hat{\theta}(y_{1:N}), \delta)\right] \tag{2.26}$$

Finally, we assume that the expectation under the marginal distribution for the data equals that under the prior or nominal parameters, so that:

$$U(\delta) \approx \text{constant} + \frac{1}{2}\mathbb{E}_\theta\left[\log\det A(\theta, \delta)\right] \tag{2.27}$$

Hence, the mutual information utility function can be approximated by a Gaussian-like posterior entropy. Yet, there is still the matter of searching over possible sensor configurations. The convex optimization technique developed in Chapter 3 reduces the computational complexity of this search and provides expressions for the Hessian mentioned in Equation (2.26).

Comparison of two sensor configurations

Consider the difference in mutual information between two sensor configurations:

$$\Delta I_{12} \triangleq I(y_{1:N}(\delta_2); \theta) - I(y_{1:N}(\delta_1)|\theta) = H(y_{1:N}(\delta_1)|\theta) - H(y_{1:N}(\delta_2)|\theta) \tag{2.28}$$

If the same approximations hold in either case, then we have:

$$\Delta I_{12} = \frac{1}{2}\mathbb{E}_\theta\left[\log\det A(\theta, \delta_2)\right] - \frac{1}{2}\mathbb{E}_\theta\left[\log\det A(\theta, \delta_1)\right] \tag{2.29}$$

Now, note that the Hessian $A(\theta, \delta)$ consists of two block diagonal parts corresponding to the system parameters, θ_s and the single prediction error parameter, σ^2. We may write:

$$\mathbb{E}_\theta\left[\log\det A(\theta, \delta)\right] = \mathbb{E}_{\theta_s}\left[\log\det B(\theta_s, \delta)\right] + \mathbb{E}_{\sigma^2}\left[\log C(\sigma^2)\right] \tag{2.30}$$

For the same number of sensors, N_o, the scalar factor, $C(\sigma^2) = \frac{NN_o}{2\sigma^2}$, cancels out in the difference.

We are left with the result:

$$\Delta I_{12} = H(y_{1:N}(\delta_1)|\theta) - H(y_{1:N}(\delta_2)|\theta) \tag{2.31}$$

$$= \frac{1}{2}\mathbb{E}_{\theta_s}\left[\log\det B(\theta_s, \delta_2)\right] - \frac{1}{2}\mathbb{E}_{\theta_s}\left[\log\det B(\theta_s, \delta_1)\right] \tag{2.32}$$

We can therefore compare the effect of sensor configuration on the difference in differential entropies of the posteriors that would result in the mean in each case.

2.5 Differences in mutual information

The mutual information between two stochastic variables is the mean increase in information of one variable upon learning the value of the other. This can be thought of as the mean gain in the number of bits of Shannon information, if the logarithm to the base 2 is used in the definition of entropy. (**ref. CoverThomas**)

In the expression above, the difference in values of the mutual information for two sensor configurations, therefore corresponds to the mean increase in bits of information gained relative to the prior for one configuration, δ_1 versus the other, δ_2.

2.6 Computing the mutual information

Recall that Equation (2.19) defines the mutual information that can in principle be solved. A piece of that equation is the evidence term which is given in equation (2.20).

To evaluate the mutual information, we split the expression from Eqn. (2.19) into two terms using the difference of the logarithms:

$$I(y_{1:N}(\delta), \theta) = \int\int p(y_{1:N}(\delta)|\theta, \mathcal{M})p(\theta|\mathcal{M})[\log p(y_{1:N}(\delta)|\theta, \mathcal{M}) - \log p(y_{1:N}(\delta)|\mathcal{M})]dy_{1:N}(\delta)d\theta \tag{2.33}$$

$$= \int p(\theta|\mathcal{M}) \int p(y_{1:N}(\delta)|\theta, \mathcal{M}) \log p(y_{1:N}(\delta)|\theta, \mathcal{M})dy_{1:N}(\delta)d\theta - \int p(y_{1:N}(\delta)|\mathcal{M}) \log p(y_{1:N}(\delta)|\mathcal{M})\, dy_{1:N}(\delta) \tag{2.34}$$

$$\Rightarrow I(y_{1:N}(\delta), \theta) = H(y_{1:N}(\delta)) - H(y_{1:N}(\delta)|\theta) \tag{2.35}$$

The conditional entropy is usually easy to compute: it is the expected conditional entropy of the stochastic forward model over the parameters, which, in the case

of a conditionally independent Gaussian stochastic forward model reduces to the expectation of a simple expression over the prediction-error uncertainty parameters only as will be seen in the examples that follow.

The entropy of the marginal distribution for the data prediction, on the other hand, is much harder to calculate in general, since an analytic expression for this marginal is usually unavailable. In addition, a single evaluation of the integrand involves one evaluation of evidence for the model class under that particular sample from the prediction for the data. As will be explored in Chapter 5, this is not a trivial task.

Methods to compute mutual information between low-dimensional variables exist, using samples from the joint distribution.

Monte Carlo approximation

Recall that mutual information can be decomposed into the difference of two entropies:

$$I(y_{1:N}(\delta);\theta) = H(y_{1:N}(\delta)) - H(y_{1:N}(\delta)|\theta) \tag{2.36}$$

Evaluation of the conditional entropy (the second term) is usually trivial in many situations such as when the prediction-errors are modeled as additive and independently Gaussian. This would typically involve calculating the expectation of the log-determinant of the co-variance matrix over the prediction-error uncertainty parameters.

The challenging computation is that of the entropy of evidence. A naive Monte Carlo approximation to this quantity is presented:

$$H(y_{1:N}(\delta)) = -\int p(y_{1:N}(\delta)) \log p(y_{1:N}(\delta)) y_{1:N}(\delta) \tag{2.37}$$

$$\approx \frac{1}{N_s} \sum_{k=1}^{N_s} \log p(y_{1:N}(\delta)^{(k)}) \tag{2.38}$$

Here, $y_{1:N}(\delta)^{(k)}$ is a sample from the joint distribution, $p(y_{1:N}(\delta), \theta)$. The variance of a naive estimate of this might be quite high, since the data variable can be of a high dimension. However, this example is illustrative for the purpose of discussing the complexity of any sample-based estimate that uses roughly the same number of samples.

For each sample of the data variable drawn from the joint distribution, one instance of model evidence needs to be evaluated. This is an expensive computation that is

not straightforward, since again naive estimators tend to have high variance. To this end, a method of evaluating log-evidence is discussed in Chapter 5.

Example: Gaussian data with uncertain mean

The purpose of this example is to give insight into what the expression for mutual information may looks like. Consider a model class $\mathcal{M}_G$ that specifies a N_o-variate Gaussian stochastic forward model:

$$p(y_{1:N}(\delta)|\mu(\theta), \Sigma(\theta), \mathcal{M}_G) = \prod_{n=1}^{N} \mathcal{N}(y_n|\mu(\theta), \Sigma(\theta)) \tag{2.39}$$

$$= \frac{1}{(2\pi)^{NN_o/2}|\Sigma(\theta)|^{N/2}} \times \exp\left(-\frac{1}{2}\sum_{n=1}^{N}(y_n - \mu(\theta))^T\Sigma^{-1}(\theta)(y_n - \mu(\theta))\right) \tag{2.40}$$

It is also assumed that a prior model, $p(\theta|\mathcal{M}_G)$ is specified. In this case, the first term of Equation (2.35) is given by the expected entropy of a multivariate Gaussian distribution:

$$H(y_{1:N}|\theta) = -\mathbb{E}_\theta\left[\frac{N}{2}\log(2\pi e) + \frac{1}{2}\log\det\Sigma(\theta)\right] \tag{2.41}$$

$$= -\frac{NN_o}{2}\log(2\pi e) - \frac{N}{2}\mathbb{E}_\theta[\log\det\Sigma(\theta)] \tag{2.42}$$

On the other hand, the entropy of the distribution over the prediction of the data, or the entropy of the evidence density, is hard to calculate. This is because the marginal likelihood, $p(y_{1:N}|\mathcal{M})$, and its logarithm needs to be evaluated for every integration point over $y_{1:N}$. There is usually no closed-form expression to $p(y_{1:N}|\mathcal{M})$, except perhaps when the predictions are linear in the parameters which in turn have a Gaussian prior, or when conjugate priors are employed. Again, these introduce restrictions on the class of models that may be employed and are usually not appropriate modeling choices for dynamical systems. However, this example will look at the case where the prior for the mean parameter is chosen to be conjugate to the stochastic forward model.

To simplify this example further, assume that $N_o = 1$ so that the observations are modeled as univariate Gaussian data. Let only the mean be uncertain and have a Gaussian distribution, that is $\theta = \mu \sim \mathcal{N}(\cdot|\mu_0, \sigma_0)$. For simplicity, the prediction-error standard deviation, σ, is taken as known.

The resulting marginal likelihood is Gaussian:

$$p(y_{1:N}|\mathcal{M}_G) = \int p(y_{1:N}|\mu)p(\mu|\mu_0, \sigma_0)\, d\mu \tag{2.43}$$

$$= \int \frac{1}{(2\pi\sigma^2)^{N/2}\sqrt{2\pi\sigma_0^2}} \exp\left[-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(y_n - \mu)^2 - \frac{(\mu - \mu_0)^2}{2\sigma_0^2}\right] d\mu \tag{2.44}$$

$$= \mathcal{N}\left(y_{1:N}|\mu_0\mathbf{1}, \tilde{\Sigma}(\sigma_0, \sigma)\right) \tag{2.45}$$

where

$$|\tilde{\Sigma}(\sigma_0, \sigma)| = \sigma^{2N}\left(1 + \frac{N\sigma_0^2}{\sigma^2}\right) \tag{2.46}$$

The mutual information between the prediction of the data and the uncertain mean parameter is therefore:

$$I(y_{1:N}; \mu) = -\frac{N}{2}\log(2\pi e\sigma^2) + \frac{N}{2}\log(2\pi e) + \frac{1}{2}\log|\tilde{\Sigma}(\sigma_0, \sigma)| \tag{2.47}$$

$$= \frac{N}{2}\log\left(1 + \frac{N\sigma_0^2}{\sigma^2}\right) \tag{2.48}$$

By inspection, Equation (2.48) tells us the following:

1. More observations, N, result in a higher mutual information
2. The data to be observed, $y_{1:N}$, is more informative about the uncertain mean parameter, μ, for a larger prior variance, σ_0^2
3. For this contrived example, when the known prediction-error variance, σ^2, is small, then this is an assertion that the observations are close to the predictions and therefore more informative

To conclude this example, it is interesting to note that for $N = 1$, the expression for mutual information in Equation (2.48) is the same as that for a Gaussian communication channel with signal power σ_0^2 and noise variance σ^2, as described in [1].

Example: Spring-mass oscillator

The simple case of a single degree-of-freedom dynamical system is considered next, to provide insight into mutual information calculations in more complicated cases. There is no analytical expression for the mutual information in this case.

Consider an undamped mechanical oscillator with uncertain natural frequency ω, with unit initial displacement and starting from rest. The governing equation is:

$$\ddot{x} + \omega^2 x = 0, \text{ with } x(0) = 1, \dot{x}(0) = 0 \tag{2.49}$$

The solution to the displacement time history of this oscillator is:

$$x(t) = \cos(\omega t) \tag{2.50}$$

We assume that data, $y_{1:N} = \{y_n\}_{n=1}^N$, is obtained from points in time equally spaced at Δt. The equation for the stochastic prediction is:

$$y_n = \cos(\omega n \Delta t) + \epsilon_n \tag{2.51}$$

As in the previous example, we assume the prediction-error to be stationary and Gaussian with uncertain variance σ^2, so that the stochastic forward model may be specified as:

$$p(y_{1:N}|\omega, \sigma^2, \mathcal{M}_O) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left[-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(y_n - \cos(\omega n \Delta t))^2\right] \tag{2.52}$$

Let the uncertain parameters be distributed according to a prior, as $\omega \sim p(\omega)$, $\sigma^2 \sim p(\sigma^2)$ and $p(\omega, \sigma^2) = p(\omega)p(\sigma^2)$. Then, the model evidence is given by:

$$p(y_{1:N}|\mathcal{M}_O) = \int\int \frac{p(\omega)p(\sigma^2)}{(2\pi\sigma^2)^{N/2}} \exp\left[-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(y_n - \cos(\omega n \Delta t))^2\right] d\omega d\sigma^2 \tag{2.53}$$

Chapter 5 discusses the application of Asymptotically Independent Markov Sampling (AIMS) and path sampling for the the purpose of evaluating model evidence, or integrals of the form:

$$p(y_{1:N}|\mathcal{M}) = \int p(y_{1:N}|\theta, \mathcal{M})p(\theta|\mathcal{M})d\theta \tag{2.54}$$

As mentioned in Section 2.6, the computation of mutual information requires the evaluation of several model evidence-like terms, each of which is typically quite expensive to evaluate and requires a full run of AIMS.

In this example, the entropy-of-evidence term is determined as follows:

$$H(\mathcal{D}) \approx \frac{1}{N_s} \sum_{k=1}^{N_s} \log p(y_{1:N}^{(k)}|\mathcal{M}) \tag{2.55}$$

where

$$(y_{1:N}^{(k)}, \theta^{(k)}) \sim p(y_{1:N}, \theta|\mathcal{M}) = p(y_{1:N}|\theta, \mathcal{M})p(\theta|\mathcal{M}) \tag{2.56}$$

Equations (2.55) and (2.56) specify that the entropy of evidence according to its Monte-Carlo approximations using samples from the predicted data variable. These samples come from the joint density. One sample from the joint is obtained by drawing a prior sample and running the stochastic forward model on it.

Each logarithmic term in the sum of Equation (2.55) is calculated using AIMS and path sampling. The sample average is an estimator for the mean. The standard deviation of the estimator of the mean is itself estimated using the standard relationship between the sample and population variances.

Table 2.1: Simulation results for entropy-of-evidence

Quantity	**Trapezoid**	**Gauss 7-point**	**Kronrod 15-point**
Mean estimate for entropy of evidence	-22.5746	-23.9345	-23.2962

The conditional entropy is evaluated much more easily, since it requires the expectation of the log variance under the prior:

$$H(y_{1:N}|\theta) = \frac{N}{2} \log 2\pi + \frac{N}{2} \mathbb{E}_{\sigma^2}[\log \sigma^2] \tag{2.57}$$

In this case, we don't even need use a Monte-Carlo estimate, since we have already specified that our prior standard deviation follows a log-uniform.

Table 2.2: Conditional entropy

Quantity	**Analytic Value**
Conditional entropy	-30.81061

Using the values for the trapezoidal estimate, the estimated value for mutual information for this example is given in Table 2.3.

Table 2.3: Mutual information estimate for a single DOF system

Quantity	Mean
Mutual Information	8.2360

It should be noted that this evaluation is quite expensive. The simulation parameters for this example are listed in Table 2.4.

Table 2.4: Single degree-of-freedom example: Simulation parameters

Quantity	Value
Number of log-evidence samples	5000
Number of AIMS samples	300×30 or 9000
Total forward evaluations	45 million
Number of time-steps	10
Time-step	0.2
Prior frequency	$\omega \sim U[\frac{1}{2}, \frac{3}{2}]$
Prior standard deviation	$\log \sigma \sim U[-5, -3]$

Example: Two-DOF unidentifiable system

Consider a 2-DOF system, consisting of two masses connected using two springs, with the end of one spring fixed. The equations of motion for this system are given by:

$$m\ddot{x}(t) + \begin{bmatrix} k_1 + k_2 & -k_2 \\ -k_2 & k_2 \end{bmatrix} x(t) = 0 \tag{2.58}$$

with

$$x(0) = \begin{Bmatrix} 1 \\ 1 \end{Bmatrix} \tag{2.59}$$

We consider a unit mass, $m = 1$, for simplicity. This system can be diagonalized using a transformation matrix Φ that relates the modal co-ordinates q_j to the original ones x_i. The resulting decoupled equations of motion have solutions:

$$q_j(t) = (\Phi_{1j} + \Phi_{2j}) \cos(\omega_j t) \tag{2.60}$$

The solution to the original system, is therefore:

$$x_i(t) = \sum_j \Phi_{ij} (\Phi_{1j} + \Phi_{2j}) \cos(\omega_j t) \tag{2.61}$$

This system has the peculiar property that the forward model output at the second degree-of-freedom can be matched for two very different stiffness configurations. The configurations $[k_1 = k_1^*, k_2 = k_2^*]$ and $[k_1 = 2k_2^*, k_2 = k_1^*/2]$ result in identical predictions for $x_2(t)$. On the other hand, the output $x_1(t)$ is defined uniquely.

For a specific prior, therefore, we expect to see a higher mutual information between the displacement at the first sensor location as opposed to the second. As in the previous example, the prediction-error parameter, σ, is uncertain. For this example, however, there are two uncertain system stiffness parameters, k_1 and k_2. The stochastic forward model in either case is specified by:

$$p(y_{1:N}(\delta = e_i)|k_1, k_2, \sigma) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left[-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(y_n - x_{i,n}(k_1, k_2))^2\right] \quad (2.62)$$

The parameters used for this simulation are listed in Table 2.5.

Table 2.5: Two degree-of-freedom example: Simulation parameters

Quantity	**Value**
Number of log-evidence samples	5000
Number of AIMS samples	300×30 or 9000
Total forward evaluations	45 million
Number of time-steps	10
Time-step	0.2
Prior k_1	$k_1 \sim U[2, 6]$
Prior k_2	$k_2 \sim U[2, 6]$
Prior standard deviation	$\log \sigma \sim U[-5, -3]$

The mutual information is evaluated at both locations for each case and is presented in Table 2.6.

Table 2.6: Mutual information estimates for 2-DOF system

Location	**Mean**
DOF1	4.7524
DOF2	5.0922

The lack of identifiability of the system with data from the second degree-of-freedom surprisingly did not result in a lower information gain compared to the first location.

This warrants further analysis that is relegated to future work at this time, to verify that this is not an artefact of numerical computations.

Note that the number of time-steps was kept small and the measurement time-step fairly quite in order to keep in check the dimensionality of the data variable, while simultaneously being able to capture one whole period of oscillation for every frequency in the range of the prior.

2.7 Sensor rearrangement problem

The optimal sensor placement problem, in the absence of previous instrumentation of a structure, relies on a prior distribution over the parameters that has the maximum entropic uncertainty for its model class, possibly with location and variance constraints.

When a structure has already been instrumented, however, a posterior distribution over the parameters is usually available after performing a Bayesian update. Here, we investigate whether or not a posterior could simply be used as the prior for a new optimal sensor placement problem.

Assumptions

To begin with, we assume that the same model class is chosen for the structure to be re-instrumented. Denote this model class by $\mathcal{M}$.

Denote the data to be obtained from the new sensor configuration by $Y'_{N'}(\delta')$. The data from the old configuration is, before performing the update, denoted by $Y_N(\delta)$. Then, we make the assumption that:

$$p(Y'_{N'}(\delta'), Y_N(\delta)|\theta, \mathcal{M}) = p(Y'_{N'}(\delta')|\theta, \mathcal{M})p(Y_N(\delta)|\theta, \mathcal{M}) \tag{2.63}$$

This states that the new data is conditionally independent of the old data, given the parameter vector.

For convenience of notation, denote $Y_N(\delta)$ by D_1 and $Y'_{N'}(\delta')$ by D_2. For this problem, we are given a realization $d_1 \sim D_1$. Also, conditioning on $\mathcal{M}$ is omitted.

In this notation, Equation (2.63) becomes:

$$p(D_2, D_1 = d_1|\theta) = p(D_2|\theta)p(D_1 = d_1|\theta) \tag{2.64}$$

Optimal placement for sensor reconfiguration

For the new sensor placement problem, we would like to use an information-theoretic measure as the utility function to be maximized over then new sensor configuration,

δ'. We use the expected value of the KL-divergence from the updated prior to the pre-posterior given the as-yet-unmeasured data. That is:

$$U(\delta') = \mathbb{E}_{D_2|D_1=d_1}\left[\mathrm{KL}\left\{p(\theta|D_2, D_1 = d_1)||p(\theta|D_1 = d_1)\right\}\right] \tag{2.65}$$

This may be further expressed as follows:

$$U(\delta') = \mathbb{E}_{D_2|D_1=d_1}\left[\int p(\theta|D_2, D_1 = d_1)\log\left\{\frac{p(\theta|D_2, D_1 = d_1)}{p(\theta|D_1 = d_1)}\right\}\, d\theta\right] \tag{2.66}$$

Conditional probability and Bayes' theorem allows us to express this in terms of quantities that can be estimated numerically. Denote the logarithm term in Eqn. (2.66) by $L(\theta, D_2, d_1)$. Then, we have:

$$U(\delta') = \int p(D_2|D_1 = d_1)p(\theta|D_2, D_1 = d_1)L(\theta, D_2, d_1)\, d\theta\, dD_2 \tag{2.67}$$

Using the conditional independence assumption of Eqn. (2.63), the logarithm term can be simplified:

$$\begin{aligned} L(\theta, D_2, d_1) &= \log\left\{\frac{p(D_2|\theta)p(D_1 = d_1|\theta)p(\theta)p(D_1 = d_1)}{p(D_2|D_1 = d_1)p(D_1 = d_1)p(D_1 = d_1|\theta)p(\theta)}\right\} \\ \Rightarrow L(\theta, D_2, d_1) &= \log\left\{\frac{p(D_2|\theta)}{p(D_2|D_1 = d_1)}\right\} \end{aligned} \tag{2.68}$$

Thus, upon simplifying the expression for the remaining terms, we get:

$$U(\delta') = \int p(D_2|\theta)p(\theta|D_1 = d_1)\log\left\{\frac{p(D_2|\theta)}{p(D_2|D_1 = d_1)}\right\}\, d\theta\, dD_2 \tag{2.69}$$

The new problem, therefore requires evaluating the following entropy-like expression:

$$H_{21} = -\int p(D_2|D_1 = d_1)\log p(D_2|D_1 = d_1)\, dD_2 \tag{2.70}$$

Further,

$$p(D_2|D_1 = d_1) = \int p(D_2|\theta)p(\theta|D_1 = d_1)\, d\theta \tag{2.71}$$

This implies that:

$$H_{21} = \log p(D_1 = d_1) - \int \frac{p(D_2, D_1 = d_1)}{p(D_1 = d_1)}\log\left\{p(D_2, D_1 = d_1)\right\}\, dD_2 \tag{2.72}$$

Then, the optimal sensor rearrangement reconfiguration problem can proceed as in the original optimal sensor location problem. Note, however, that integrals

against θ in Equations (2.69) and (2.72) are against the posterior distribution over the parameters $p(\theta|D_1 = d_1)$, that have been updated using data from the first deployment. Thus, the complexity of this problem is compounded by the fact that unlike the sensor location problem for which the prior distribution is easy to sample from, the sensor reconfiguration problem required drawing several sets of samples from the posterior distribution under the first set of data. An example problem is not attempted therefore, since reconfiguration problem, in the absence of any simplifying assumptions, is considerably more computationally expensive than the already costly sensor location problem.

References

[1] T. M. Cover, J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: John Wiley & Sons, Inc., Sep. 2005, ISBN: 9780471748823.

Chapter 3

EFFICIENT SOLUTION USING CONVEX RELAXATION

In the previous chapter, the Bayesian information-theoretic foundations of the optimal sensor placement problem were laid out. Here, The entropy-based optimal sensor placement problem is formulated as in [1] and an efficient solution using convex optimization techniques is presented. Model identifiability is discussed and the Laplace approximation to the posterior in a Bayesian system identification problem is described.

3.1 Log-determinant formulation

This section expands on the derivation of the log-determinant formulation set up in Section 2.4. The Laplace approximation to the posterior together with additional assumptions about the prediction-errors, allows for the development of a log-determinant entropy-based objective.

System Identifiability

The notion of identifiability is important to characterize the topography of the posterior distribution. The definition used here is in line with [2]. Essentially, the posterior distribution can either be globally identifiable, locally identifiable or unidentifiable as follows.

Denote by $S(\theta)$, the set of all admissible values of the parameter vector, θ. Also, let $\hat{\theta}$ correspond to the optimal parameters that maximize the likelihood function in Equation (2.6).

If $S_{\text{opt}}(\theta; \mathcal{D}_N)$, a subset of $S(\theta)$, is the set of all optimal models in the model class $\mathcal{M}$ given data $\mathcal{D}_N$, then a parameter θ_j is said to be *globally system identifiable* if $S_{\text{opt}}(\theta; \mathcal{D}_N)$ contains only one optimal parameter, or else if for two different optimal parameter vectors $\hat{\alpha}^{(1)}$ and $\hat{\theta}^{(2)}$, the corresponding components of the j'th parameter are equal as $\hat{\theta}_j^{(1)} = \hat{\theta}_j^{(2)}$.

In general, a parameter is *system identifiable* if no two optimal parameters are arbitrarily close together. In other words, for $\hat{\theta}_j^{(1)}, \hat{\theta}_j^{(2)} \in S_{\text{opt}}(\theta; \mathcal{D}_N)$, there exists a positive number, ϵ_j such that exactly one of $|\hat{\theta}_j^{(1)} - \hat{\theta}_j^{(2)}| > \epsilon_j$ and $\hat{\theta}_j^{(1)} = \hat{\theta}_j^{(2)}$ holds true.

Finally, not that if a model is system identifiable but not globally so, then it is said to be *locally system identifiable*.

Laplace's asymptotic approximation

When the system is globally identifiable, then we may expand the log-posterior about the most probable value to second order using a Taylor series to give us a Gaussian approximation to the posterior density:

$$p(\theta|y_{1:N}, \mathcal{M}) \approx p(\hat{\theta}|y_{1:N}, \mathcal{M}) \exp\left[-\frac{1}{2}(\theta - \hat{\theta})^T A_N(\hat{\theta})(\theta - \hat{\theta})\right] \tag{3.1}$$

The normalizing constant in front of the exponential can be determined by inspection:

$$p(\theta|y_{1:N}, \mathcal{M}) \approx \frac{\sqrt{\det A_N(\hat{\theta})}}{(2\pi)^{NN_o/2}} \exp\left[-\frac{1}{2}(\theta - \hat{\theta})^T A_N(\hat{\theta})(\theta - \hat{\theta})\right] \tag{3.2}$$

The precision matrix of the multivariate Gaussian in Equation (3.1) is given by the Hessian of the logarithm of the posterior density, evaluated at the unique most-probable parameter vector. That is,

$$A_{N,ij}(\theta) = -\frac{\partial^2 \log p(\theta|y_{1:N}, \mathcal{M})}{\partial\theta_i \partial\theta_j} \tag{3.3}$$

It is to be borne in mind that the parameter vector and its associated derivatives can be partitioned into two parts - the system parameters and the prediction-error parameters:

$$A_N(\theta_s) = \begin{bmatrix} B_N(\theta_s) & 0 \\ 0 & C_N(\hat{\sigma}) \end{bmatrix} \tag{3.4}$$

where

$$[B_N(\theta_s)]_{ij} = \left[\frac{\partial^2 \log p(y_{1:N}|[\theta_s, \hat{\sigma}], \mathcal{M})}{\partial\theta_{s,i}\partial\theta_{s,j}}\right] \tag{3.5}$$

and

$$C_N(\hat{\sigma}) = \left.\frac{\partial^2 \log p(y_{1:N}|[\theta_s, \sigma], \mathcal{M})}{\partial\sigma\partial\sigma}\right|_{\sigma=\hat{\sigma}} \tag{3.6}$$

The off-diagonal terms in the Hessian of Equation (3.4) are zero because of our choice of stochastic forward model and prior.

The posterior distribution of the system parameters only, with the prediction-error parameter marginalized out, can therefore be expressed independently as:

$$p(\theta_s|y_{1:N}, \mathcal{M}) = \frac{\sqrt{\det B_N(\hat{\theta}_s)}}{(2\pi)^{NN_o/2}} \exp\left[-\frac{1}{2}(\theta_s - \hat{\theta}_s)^T B_N(\hat{\theta}_s)(\theta_s - \hat{\theta}_s)\right] \tag{3.7}$$

We now seek to provide an approximation to the diagonal block of the Hessian that corresponds to the system parameters, $B_N(\hat{\theta}_s)$, in terms of the sensitivity coefficients of the predictions with respect to the system parameters.

For fixed system parameters θ_s, the optimal (in the maximum-likelihood sense) prediction-error variance in terms of the data as:

$$\hat{\sigma}^2(\theta_s) = \frac{1}{NN_o} \sum_{n=1}^{N} \|y_n - S_o(\delta)x_n(\theta_s)\|^2 \triangleq J(\theta_s) \tag{3.8}$$

Equation (3.8) can be used to re-express the stochastic forward model at the optimal parameter:

$$p(y_{1:N}|\theta = [\theta_s, \hat{\sigma}], \mathcal{M}) = [2\pi e J(\theta_s)]^{-NN_o/2} \tag{3.9}$$

We can evaluate the block diagonal terms of the Hessian corresponding to the system parameters by:

$$B_N(\hat{\theta}_s) = \frac{NN_o}{2}\left[\frac{\partial^2 J(\theta_s)}{\partial\theta_s\partial\theta_s}\right]_{\theta_s=\hat{\theta}_s} \tag{3.10}$$

We expand the second-order derivative in Equation (3.10) using the expression for $J(\theta_s)$ from Equation (3.8):

$$[B_N(\hat{\theta}_s)]_{pq} = \frac{1}{\hat{\sigma}^2}\sum_{n=1}^{N}\left[\frac{\partial x_n}{\partial\theta_{s,p}} S_o S_o^T \frac{\partial x_n^T}{\partial\theta_{s,q}} + \epsilon_n\left(S_o\frac{\partial^2 x_n}{\partial\theta_{s,p}\partial\theta_{s,q}}\right)^T\right] \tag{3.11}$$

This is where an additional assumption is needed in order to determine this second-derivative approximately: it is assumed that the term involving the second derivative in Equation 3.11 can be neglected. This can be justified if at least one of the following is true: the prediction errors, ϵ_n are small in magnitude; or the deterministic predictions, x_n, vary slowly with respect to the system parameters and therefore have a small second derivative. The curvature assumption can be assessed before instrumenting the structure, while the prediction-error assumption can only be validated after checking the agreement between the deterministic predictions and the data.

With this assumption, the relevant portion of the Hessian matrix can be expressed in terms of the prediction sensitivity coefficients at the observed degrees of freedom. The resulting approximation for the second derivative is:

$$[B_N(\hat{\theta}_s)]_{pq} \approx \frac{1}{\hat{\sigma}^2} \sum_{n=1}^{N} \left[\frac{\partial x_n}{\partial \theta_{s,p}} S_o S_o^T \frac{\partial x_n^T}{\partial \theta_{s,q}} \right] \tag{3.12}$$

In addition, notice the dependence on the sensor configuration as $S_o^T S_o = \text{diag}(\delta)$. The expression in Equation (3.12) can therefore be transformed to a double-sum:

$$[B_N(\hat{\theta}_s)]_{pq} = \frac{1}{\hat{\sigma}^2} \sum_{i=1}^{N_d} \delta_i \sum_{n=1}^{N} \frac{\partial x_{n,i}}{\partial \theta_{s,p}} \frac{\partial x_{n,i}^T}{\partial \theta_{s,q}} \tag{3.13}$$

Equation (3.13) is a linear sum of contributions of terms from each sensor location. For convenience of notation in the formulation of the optimal sensor placement problem, we use the following notation while discussing the matrix of the equation:

$$[B_N(\hat{\theta}_s)]_{pq} = \frac{1}{\hat{\sigma}^2} Q_{pq}(\delta) = \frac{1}{\hat{\sigma}^2} \sum_{i=1}^{N_d} \delta_i Q_{pq}(e_i) \tag{3.14}$$

In Equation (3.14),

$$Q_{pq}(e_i) = \sum_{n=1}^{N} \frac{\partial x_{n,i}}{\partial \theta_{s,p}} \frac{\partial x_{n,i}^T}{\partial \theta_{s,q}} \tag{3.15}$$

Having developed an approximate expression for the posterior distribution over the system parameters, we are now in a position to discuss the optimal sensor location problem.

3.2 Entropy-based optimal sensor location

So far, we have discussed how data can be used to update uncertainties about our model. We now address the optimal sensor location problem. We begin by defining an objective function to be optimized with respect to sensor locations, that quantifies the posterior uncertainty that we expect in the parameters after measuring the data.

Note that the optimal sensor location problem is to be solved typically *before* any response data is available from the real structure to be instrumented. What is available is information about the structure either from its design or from preliminary tests, in the form of a structural model and a prior distribution over the system parameters.

What is very convenient in the formulation of Section 2.2, is that the dependence on the data of the approximate posterior distribution over the uncertain parameters is only through the optimal parameters, $\hat{\theta} = \hat{\theta}(y_{1:N})$.

That is,

$$p(\theta_s|y_{1:N}, \mathcal{M}) = p(\theta_s|\hat{\theta}(y_{1:N}), \mathcal{M}) \tag{3.16}$$

The nominal values for the system parameters are used in place of the optimal parameters:

$$p(\theta_s|\theta_0, \mathcal{M}) = \frac{\sqrt{\det B_N(\theta_{s,0})}}{(2\pi)^{NN_o/2}} \exp\left[-\frac{1}{2}(\theta_s - \theta_{s,0})^T B_N(\theta_0)(\theta_s - \theta_{s,0})\right] \tag{3.17}$$

Since the designer is uncertain about the nominal values, a prior distribution, $p(\theta_0)$, is specified on the nominal parameters.

Equation (3.17) replaces the dependence of the posterior on the unavailable data with prior information that is already available in terms of the hyper-parameters, $\theta_0 = [\theta_{s,0}, \sigma_0]$. These hyper-parameters are either available in the form of structural design information, or in the form of information obtained from preliminary tests.

Entropy-based objective

We are now in a position to specify an objective function to be optimized with respect to sensor location. Over all the possible sensor configurations, we would like to select the configuration that results in the lowest possible expected entropy of the pre-posterior, relative to the prior, upon arrival of the data.

$$\Delta H(\delta) = H(\theta_s|\theta_0) - H(\theta_0) \tag{3.18}$$

$$= \mathbb{E}_{(\theta_s,\theta_0)}[-\log p(\theta_s|\theta_0)] - \mathbb{E}_{\theta_0}[-\log p(\theta_0)] \tag{3.19}$$

Only the conditional entropy term in Equation (3.18) depends on the value of the sensor configuration, δ. That is, we want to minimize w.r.t δ:

$$\Delta H(\delta) = -\frac{1}{2}\int \log\det Q(\delta) p(\theta_{s,0})\, d\theta_{s,0} + \text{constant} \tag{3.20}$$

We now have an objective function to be optimized. The expression for $Q(\delta)$ comes from Equation (3.14). Denote the optimal sensor configuration by δ^*. Then, ignoring the irrelevant terms in Equation (3.20) we have:

$$\delta^* = \arg\max_{\delta} \int \log\det Q(\delta, \theta_{s,0}) p(\theta_{s,0})\, d\theta_{s,0} \tag{3.21}$$

The entropy-based optimal sensor location problem is now fully specified in Equation (3.21).

Complexity of the problem

Equation (3.21) specifies a combinatorial optimization problem. That is, the problem is of choosing N_o sensor locations among N_d ones. Depending upon the problem, this could become prohibitively expensive. Each sensor configuration requires one evaluation of the expectation of the log-determinant of an expensive prediction. An example of a prohibitively expensive sensor placement problem is say one of selecting 20 locations from among 50 locations, to instrument: the number of possible combinations is about 47 trillion.

At this point, it is clear that a brute-force search for the optimal solution over all possible sensor configurations is not a wise approach. Instead, information obtained by evaluating different configurations should be used to guide the optimization process towards the optimal configuration.

It turns out that it is, in general, not possible to guarantee an optimal solution in a manner that does not require an exhaustive search. Heuristic methods such as genetic algorithms can produce an acceptable sub-optimal value if run for long enough. Incremental greedy algorithms are also guaranteed to produce a suboptimal value within $(1 - 1/e)$ of the optimal [3], relative to the range from the least to the most optimal configuration.

The method developed in Section 3.3 along the lines of [4], applies a convex relaxation technique to provide a sub-optimal solution to the problem. An upper bound on the optimal value is automatically provided by the solution.

3.3 Convex relaxation of the combinatorial optimization problem

Here, a relaxed version of the optimization problem is set up and argued to be convex. The relevant partial derivatives of the objective function, that is, its gradient and its Hessian matrix are derived for the purpose of computing the solution. The final step involves replacing the expectation integral over the prior distribution on the uncertain parameters by its Monte Carlo approximation so that the problem may be solved by a generic convex solver. This is discussed in Section 3.3.

In Section 3.4, the relevant quantities are calculated for a multistory structure subject to sinusoidal ground motion. Assumptions about the parameterization, made for tractability of the problem, are stated. Results for optimal sensor locations are

presented for specific scenarios.

Appendix A lists some of the useful matrix calculus identities used in evaluating the gradient of the objective. Appendix B contains detailed calculations relevant to the parametric gradient of the displacement predictions for the multistory structure example considered.

Combinatorial optimization problem

The original problem reflects the reality that a sensor must be present or absent at a degree of freedom. Thus, the optimization problem is over a boolean vector:

$$\begin{aligned} &\underset{\delta}{\text{maximize}} && h(\delta) \\ &\text{subject to} && \delta_i \in \{0,1\}, \quad i = 1,\ldots,N_d \\ &\text{and} && \mathbb{1}^T\delta = N_o. \end{aligned} \tag{3.22}$$

The objective function $h(\delta)$ corresponds to the gain in Shannon information associated with a given sensor configuration, δ. It is given ([5]) by:

$$h(\delta) = \mathbb{E}_{\theta_s}\left[\log\det Q(\delta)\right] \tag{3.23}$$

The quantities in Equation 3.23 are summarized in Chapter 3. The negative sign simply converts the maximization problem to a minimization one in order to be consistent with the standard description of an optimization problem.

The practical difficulty in solving this problem is the fact that a naive exploration of all possibilities quickly becomes intractable over even moderate values of N_o and N_d. For instance, the placement of 20 sensors in a 50 DOF structure would involve $\binom{50}{20} \approx 4.71 \times 10^{13}$ evaluations of the objective function, $h(\delta)$. In order to overcome this difficulty, the original problem in Problem 3.22 is relaxed to allow the boolean vector δ to take values continuously between zero and one. The problem thus specified can be solved using convex optimization techniques that guarantee a relatively cheap and unique solution, as shown next.

It is to be noted that the solution to this new optimization problem need not be a Boolean vector, but it could contain entries between 0 and 1. The solution is still very meaningful, however. If the solution is indeed a Boolean vector, then it is *the* optimal solution to the original combinatorial problem. If not, then it is an upper bound to the value of the objective function

Relaxed Problem statement

The original Boolean sensor placement vector, δ, is relaxed into a vector, z, in the hypercube $[0,1]^{N_d}$, resulting in the relaxed optimization problem:

$$\begin{aligned} \underset{z}{\text{minimize}} \quad & f(z) \\ \text{subject to} \quad & -z_i \leq 0 \\ & z_i \leq 1, \qquad i = 1, \ldots, N_d \\ \text{and} \quad & \mathbb{1}^T z = N_o. \end{aligned} \tag{3.24}$$

Here, the objective is the continuous extension of the previous, converted to a minimization problem using a minus sign:

$$f(z) = -h(z) = -\mathbb{E}_{\theta_s}\left[\log\det Q(z)\right] \tag{3.25}$$

Note that this allows one to use continuous, rather than discrete, optimization packages to solve the problem.

Convex nature of the relaxed problem

Problem 3.24 describes a convex optimization problem. This is because the objective function is a convex in z, and the equality and inequality constraints are affine, the log-determinant function in the objective is convex in z [4] and the expectation operator preserves the convexity of its argument. Thus, the problem has a unique global optimum. This can be determined computationally, as long as the objective can be computed at every z.

For larger problems, it becomes necessary to supply the gradient and hessian of the objective with respect to z. This avoids expensive computations of their finite-difference counterparts. Fortunately, their expressions are tractable and, along with their Monte Carlo approximations, are described next.

Derivatives of the objective

The gradient and hessian of the objective will be evaluated here. Identities in Appendix A are used. Consider the gradient of the objective,

$$\frac{\partial f}{\partial z_i}(z) = -\frac{\partial}{\partial z_i}\int \log\det Q(z,\theta_s)p(\theta_s)d\theta_s \tag{3.26}$$

$$= -\int \frac{\partial \log\det Q}{\partial z_i}p(\theta_s)d\theta_s \qquad \text{(interchange)} \tag{3.27}$$

$$= -\int \operatorname{tr}\left(Q^{-1}\frac{\partial Q}{\partial z_i}\right)p(\theta_s)d\theta_s \qquad \text{(derivative of log-det)} \tag{3.28}$$

$$= -\int \operatorname{tr}\left(Q^{-1}(z,\theta_s)Q(e^i,\theta_s)\right)p(\theta_s)d\theta_s \qquad \text{(evaluation of partial)} \tag{3.29}$$

$$= -\int \left(\sum_{j=1}^{N_a}\sum_{k=1}^{N_a} Q^{-1}_{jk}(z,\theta_s)Q_{kj}(e^i,\theta_s)\right)p(\theta_s)d\theta_s \qquad \text{(trace of product)} \tag{3.30}$$

$$= -\int \left(\sum_{j=1}^{N_a}\sum_{k=1}^{N_a} Q^{-1}_{jk}(z,\theta_s)Q_{jk}(e^i,\theta_s)\right)p(\theta_s)d\theta_s \qquad \text{(symmetry)} \tag{3.31}$$

$$= -\mathbb{E}_{\theta_s}\left[Q^{-1}(z,\theta_s){:}Q(e^i,\theta_s)\right] \qquad \text{(notation)} \tag{3.32}$$

The Hessian of the objective,

$$\frac{\partial^2 f}{\partial z_p \partial z_q}(z) = -\frac{\partial}{\partial z_q}\mathbb{E}_{\theta_s}\left[\sum_{j=1}^{N_a}\sum_{k=1}^{N_a} Q^{-1}_{jk}(z,\theta_s)Q_{kj}(e^p,\theta_s)\right] \tag{3.33}$$

$$= -\mathbb{E}_{\theta_s}\left[\sum_{j=1}^{N_a}\sum_{k=1}^{N_a} \frac{\partial Q^{-1}_{jk}(z,\theta_s)}{\partial z_q}Q_{kj}(e^p,\theta_s)\right] \qquad \text{(interchanges)} \tag{3.34}$$

$$= \mathbb{E}_{\theta_s}\left[\sum_{j=1}^{N_a}\sum_{k=1}^{N_a} \left(Q^{-1}\frac{\partial Q}{\partial z_q}Q^{-1}\right)_{jk}(z,\theta_s)Q_{jk}(e^p,\theta_s)\right] \qquad \text{(derivative of inverse)} \tag{3.35}$$

$$= \mathbb{E}_{\theta_s}\left[\operatorname{tr}\left(Q^{-1}(z)Q(e^q)Q^{-1}(z)Q(e^p)\right)(\theta_s)\right] \qquad \text{(derivative, trace)} \tag{3.36}$$

$$= \mathbb{E}_{\theta_s}\left[((Q^{-1}(z)Q(e^q)){:}((Q^{-1}(z)Q(e^p))^T)(\theta_s)\right] \qquad \text{(associativity, trace)} \tag{3.37}$$

The expectation integral in the objective and its derivatives can be numerically approximated using stochastic simulations. These numerical approximations would be required for each constrained Newton step of the optimization scheme.

Numerical approximations

Given N_k samples $\theta_s^{(k)}$ distributed according to the prior $p(\theta_s)$ specified by the designer, the integrals in question may be approximated by their corresponding Monte Carlo estimates.

For the objective function,

$$f(z) \approx \frac{1}{N_k} \sum_{k=1}^{N_k} \log\det Q(z, \theta_s^{(k)}) \tag{3.38}$$

For the gradient of the objective,

$$\frac{\partial f}{\partial z_i}(z) \approx -\frac{1}{N_k} \sum_{k=1}^{N_k} \sum_{j=1}^{N_a} \sum_{k=1}^{N_a} Q_{jk}^{-1}(z, \theta_s^{(k)}) Q_{jk}(e^i, \theta_s^{(k)}) \tag{3.39}$$

Finally, for the hessian of the objective,

$$\frac{\partial^2 f}{\partial z_p \partial z_q}(z) \approx \frac{1}{N_k} \sum_{k=1}^{N_k} \sum_{j=1}^{N_a} \sum_{k=1}^{N_a} \left(Q^{-1} Q(e^q) Q^{-1} \right)_{jk} (z, \theta_s^{(k)}) Q(e^p, \theta_s^{(k)})_{kj} \tag{3.40}$$

Some computational effort may be spared by noting that,

$$Q(z, \theta_s) = \sum_{l=1}^{N_d} z_l \sum_{n=1}^{N} \frac{\partial x_l}{\partial \theta_s} \frac{\partial x_l}{\partial \theta_s}^T (t_n, \theta_s) \tag{3.41}$$

$$= \sum_{l=1}^{N_d} z_l Q(e^l, \theta_s) \tag{3.42}$$

Hence, stored values of $Q(e^i, \theta_s)$ may be used to determine $Q(z, \theta_s)$.

Solver for optimization

Since Problem 3.24 is convex, it may be solved using Newton's method. In order to apply this in practice, numerical approximations to the objective and its derivatives, Equations 3.38 - 3.40, would need to be used.

The optimization problem has equality and inequality constraints. Hence, a constrained convex optimization solver needs to be employed. The expression for the Newton step is not as straightforward as in the unconstrained case [6].

3.4 Application to multistory structure

The problem of optimally placing a fixed number of sensors over a structure with several degrees of freedom is considered. The structure is assumed to behave linearly, with classical damping. The structure is subject to a sinusoidal ground acceleration at the base specified by an amplitude and frequency.

In this problem, the nominal natural frequency, input frequency, input amplitude and the Rayleigh damping parameters are taken to be uncertain, along with the prediction-error variance for a stationary, isotropic prediction error co-variance matrix. The designer specifies a prior distribution on these uncertain parameters.

Problem description

Consider a linear structure subject to sinusoidal input ground motion. The governing ordinary matrix differential equation for its displacements $x(t)$ as a function of time is given by,

$$M\ddot{x}(t) + C\dot{x}(t) + Kx(t) = -M\mathbb{1}a_0 \sin(\omega t), \text{ with } x(0) = 0, \dot{x}(0) = 0 \tag{3.43}$$

Before converting this system into its modal co-ordinates, $q_j(t)$ with corresponding natural frequency and damping ratio ω_j and ζ_j respectively, the problem is simplified to take into account uncertainties in the system. This is done largely through the substructuring scheme. All the assumptions have consequences that result in a simple expression, relative to the uncertain parameters, for each mode.

The following simplifications are assumed:

1. $M = mM^*$ and $K = kK^*$, implying that the mass and stiffness matrices can be factored into an uncertain scalar and a deterministic matrix part. k and m will be called the nominal stiffness and mass of the system respectively.
2. $\omega_0 = \sqrt{\dfrac{k}{m}}$ is the uncertain nominal natural frequency.
3. $C = \alpha M + \beta K = \alpha m M^* + \beta k K^*$, is the Rayleigh damping assumption with uncertain parameters α and β.

Here, although k and m are uncertain, their uncertainty is irrelevant in the sensor placement problem. Only their square-root ratio ω_0 is relevant and may be treated explicitly.

As a consequence of these simplifying assumptions, the following conditions hold while one attempts to determine the modal solution of the problem:

1. The eigenvalues of the (K^*, M^*) system are constants $c_j^2 = \dfrac{\omega_j^2}{\omega_0^2}$.
2. The eigenvector matrix Φ that diagonalizes the (K^*, M^*) system, also diagonalizes the (K, M) system and can be used to solve the original system.
3. $\Phi^T M^* \Phi = \text{diag}\{\mu_i\}$ and $\Phi^T K^* \Phi = \text{diag}\{\kappa_i\}$, such that $\dfrac{\kappa_i}{\mu_i} = c_i^2$.

Given these simplifications, a modal solution to the problem may be calculated analytically.

Modal solution

We let $x(t) = \Phi q(t)$, where $q(t)$ are the modal displacement co-ordinates. Then, the original equation may be simplified to,

$$\Phi^T M \Phi \ddot{q}(t) + \Phi^T C \Phi \dot{q}(t) + \Phi^T K \Phi q(t) = \Phi^T M \mathbb{1} a_0 \sin(\omega t) \tag{3.44}$$

Using the simplifying assumptions in Section 3.4, this simplifies to a decoupled set of modal ordinary differential equations as,

$$\ddot{q}_j(t) + 2\zeta_j \omega_j \dot{q}_j(t) + \omega_j^2 q_j(t) = a_j \sin(\omega t) \tag{3.45}$$

Here,

$$\omega_j^2 = c_j^2 \omega_0^2 \tag{3.46}$$

$$2\zeta_j \omega_j = \alpha + \beta \omega_j^2 \text{ and} \tag{3.47}$$

$$a_j = -\frac{a_0}{\mu_j} \Phi^T \mathbb{1} \tag{3.48}$$

Each equation is the governing equation to a driven, damped oscillator. The solution to each equation is given by the expression,

$$q_j(t) = a_j \left[\left(\frac{\omega^3 + \omega_j^2 \omega (2\zeta_j^2 - 1)}{\omega_{d_j}} \right) \exp(-\zeta_j \omega_j t) \sin(\omega_{d_j} t) \right. \tag{3.49}$$

$$+ (2\zeta_j \omega_j \omega) \exp(-\zeta_j \omega_j t) \cos(\omega_{d_j} t) + (\omega_j^2 - \omega^2) \sin(\omega t) \tag{3.50}$$

$$\left. -(2\zeta_j \omega_j \omega) \cos(\omega t) \right] / \left[(\omega_j^2 - \omega^2)^2 + (2\zeta_j \omega_j \omega)^2 \right] \tag{3.51}$$

The prediction equation for the displacements then becomes,

$$x_i(t) = \sum_{j=1}^{N_d} \Phi_{ij} q_j(t) \tag{3.52}$$

Sensor placement algorithm

In order to solve the relaxed optimal sensor placement problem in this case, the following operations need to be performed:

1. Determine the eigenvalues c_j^2 and eigenvector matrix Φ of the (K^*, M^*) system
2. Begin with an initial guess z_0 for the sensor positions that satisfies the constraints
3. Generate N_k samples of $\theta_s = [\omega_0, \alpha, \beta, a_0, \omega]$ from the designer-specified distribution $p(\theta_s)$
4. For each $\theta_s^{(k)}$, calculate the gradient of $q_j(t)$ with respect to θ_s for several values $t_n = n\Delta t$ for $n = 1, \ldots, N$
5. Calculate the gradient $\nabla_{\theta_s} x_i(t_n, \theta_s^{(k)})$
6. Calculate the elementary matrices $Q(e^i, \theta_s^{(k)})$ for every i and k. Store them.
7. Compute the objective, its gradient and hessian, at the current vector z_m
8. Update z_m to $z_{(m+1)}$ and repeat steps 7 and 8 until convergence

Example: Uniform building

Consider the problem of placing a single displacement sensor on a uniform N_d-DOF structure.

Substructuring scheme

In this case, the mass and stiffness matrices are given by:

$$M = mM^* = mI_{N_d} \tag{3.53}$$

$$K = kK^* = k\begin{bmatrix} 2 & -1 & 0 & \dots & 0 & 0 \\ -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & \dots & 0 \\ & \ddots & \ddots & \ddots & \ddots & \\ 0 & \dots & 0 & -1 & 2 & -1 \\ 0 & 0 & \dots & 0 & -1 & 1 \end{bmatrix} \tag{3.54}$$

Prior parameters

The prior parameters are taken to be independent. They are distributed as:

$$\omega_0 \sim \ln \mathcal{N}(\cdot|\mu = 2\pi, \sigma = 0.25) \tag{3.55}$$

$$\alpha \sim \ln \mathcal{N}(\cdot|\mu = 0.1, \sigma = 0.01) \tag{3.56}$$

$$\beta \sim \ln \mathcal{N}(\cdot|\mu = 10^{-4}, \sigma = 10^{-5}) \tag{3.57}$$

$$a_0 \sim \mathcal{N}(\cdot|\mu = 0, \sigma = 40\%g) \tag{3.58}$$

$$\omega \sim \ln \mathcal{N}(\cdot|\mu = 2\pi, \sigma = 0.25) \tag{3.59}$$

These numbers are all physically reasonable and encountered in typical situations. However, they have been not chosen to model any specific physical system.

Resulting sensor locations

The resulting locations for several values of N_d are determined. While the stochastic forward model is different for structures with a different number of stories, the same prior is used for each structure considered in this example. The results are tabulated in Table 3.1.

For larger values of N_d, there are sometimes difficulties encountered in interpreting the results, since a relaxed representation of the presence or absence of the sensor is being used.

Once the $Q(e^i, \theta_s^{(k)})$ matrices have been computed and stored, the optimization algorithm always converges to the same solution irrespective of the initial starting

Table 3.1: Simulation results

Case				**Sensor DOF #**
N_d	N_o	N	N_k	(1 ≡ Base, N_d ≡ Roof)
2	1	1000	1000	2
2	1	1000	2000	2
2	1	2000	1000	2
2	1	2000	2000	2
4	1	1000	1000	4
4	1	1000	2000	4
4	1	2000	1000	4
4	1	2000	2000	4
4	2	1000	1000	2, 4
4	2	1000	2000	2, 4
4	2	2000	1000	2, 4
4	2	2000	2000	2, 4
8	2	1000	1000	6, 8
8	2	2000	2000	6, 8

point, z_0. This is expected, since the problem is convex and is guaranteed to have a global minimum.

Table 3.1 contains the results of solving the relaxed, convex optimization problem for structures with up to 8 degrees-of-freedom on which 2 sensors are to be installed. For each structure or value of N_d, the sensor configuration remains stable to changes in the number of time-steps or prior samples used. There does seem to be a preference for a sensor on the roof in every case. However, one should not read into this result too much since these positions depend a lot on the parameters that are to be learned as well as the prior uncertainties associated with them.

Next, the results for the sensor location scheme for a 50-story shear building, to be instrumented with 20 sensors, are discussed. Figures 3.1 and 3.2 present typical solutions to the convex problem. Since this problem involves a larger number of degrees-of-freedom, there is a higher variance associated with the computation of expectation of the log-determinant of the sensitivity matrix using prior samples. For this example, this shows up as a difference in the value of the non-Boolean sensor configuration at stories 26 through 28. It is to be noted, however, that there does not appear to be much of an ambiguity upon rounding the numbers to their Boolean values. For example, no sensors should be placed at stories 26 and 27 in Figure 3.1, or at story 27 in Figure 3.2.

The results suggest that all floors above the 27th, except floors 40, 43, 46 and 49,

and no floor below the 28th except floor 2 should be instrumented.

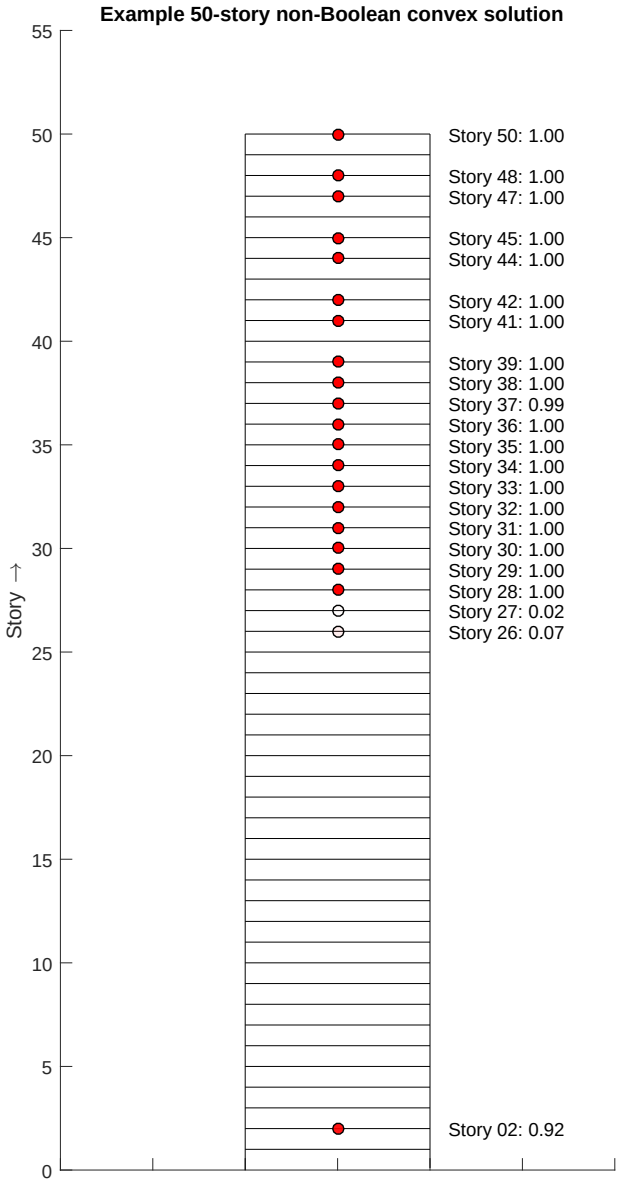

Figure 3.1: The resulting sensor locations after solving the convex optimization problem for a 50-story structure

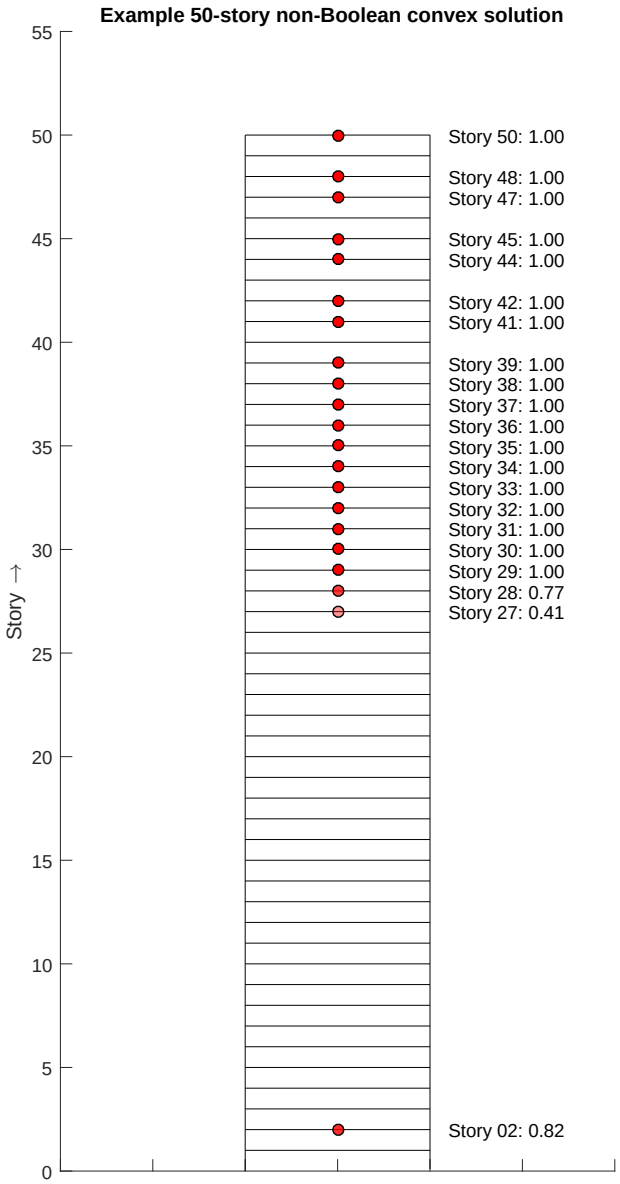

Figure 3.2: Another solution for the 50-story building but with different prior samples

It is to be remarked, as in the introduction, that an exhaustive search for the optimal configuration would have required over 47 trillion evaluations. The gradient-based method requires fewer than 100 evaluations! Of course, heuristic methods could have provided a good sub-optimal solution. In our case, however, we appear to have converged to the optimal solution itself.

Comparison of different sensor configurations

The number of bits of information gain of the optimal configuration, z^*, obtained using the convex technique over some other configurations is considered here.

The following configurations are considered, as described in Table 3.2:

Table 3.2: Description of various sensor configurations for comparative example

Case	Description
z^*	Optimal configuration using convex relaxation
z_{low}	Stories 1 through 20 instrumented
z_{high}	Stories 31 through 50 instrumented
z_{common}	Sensors evenly spaced (Stories 1 through 50 in steps of 2.5, rounded up)
z_{greedy}	Solution using greedy sequential placement

The configuration z_{greedy} is the sensor configuration obtained using a greedy sequential placement algorithm that picks the next best sensor at each stage.

Here, z_{common} is a commonly-used scheme of distributing the sensors evenly in the structure.

Table 3.3: Comparison of information gain relative to optimal configuration

Case	Objective Value	# bits gain using z*
z^*	6.14E+01	
z_{low}	5.75E+01	5.7
z_{high}	6.12E+01	0.3
z_{common}	6.05E+01	1.3
z_{greedy}	6.14E+01	0.0

In this case, it appears to be very inefficient to place sensors at at the lower stories relative to the higher ones.

It turns out in this case that $z_{\text{greedy}} = z^*$. With the convex relaxation scheme, however, it is immediately understood that the solution is optimal, if Boolean.

The commonly-used sensor configuration appears to be significantly sub-optimal for the given model class. This configuration is sometimes popular, however, for

considerations beyond the model class being considered; for instance, when it is expected that the behavior of the structure is not well understood.

Variation of optimal configuration with sample set

For a Monte Carlo sample size, N_k, of 1000 or 2000, the optimal configuration is consistent for smaller structures ($N_d < 10$).

For larger structures ($N_d \approx 50$), however, the optimal configuration may differ from sample to sample. The severity of the problem increases with the size of the structure and the number of parameters used. It turns out that for this example, the majority of the sensor locations agree. Those in agreement are typically found in the upper half of the structure. This may change with the specific model choices made in formulating the problem, however.

Possible solutions to this problem:

- Select a sufficiently large sample size to avoid variation between sample sets
- Avoid naive Monte Carlo integration and use an integral that provides an estimate with smaller variance. The method would depend on the nature of the integrand. An algorithm similar to simulated annealing could be used to compute the expectation value [7]
- Obtain better design information about the uncertain parameters to obtain a more peaked prior with less variation between samples within a sample set

Variation of optimal configuration with number of time-steps

For a fixed sample set, there is a variation in the optimal configuration when the number of observation time steps is changed from 1000 to 2000 for structures with $N_d > 10$. Sensor locations near the roof typically remain the same. However, below some story, the sensor locations can be quite different.

Conclusion

The original combinatorial determinant maximization problem for optimal sensor placement was cast into a relaxed convex optimization problem and solved efficiently to yield sub-optimal solutions. Locations for which a sensor was unambiguously present or absent were typically obtained, across differing Monte Carlo samples and differing number of time steps, as well as ambiguous locations where sample variability led to varying results.

Note that in the absence of design uncertainties, no variation would be present and the algorithm would always converge to a stable, sub-optimal, and possibly optimal, solution. Ambiguous locations could be resolved using a brute-force combinatorial approach since the problem size would be reduced and manageable.

References

[1] C. Papadimitriou, J. L. Beck, S.-K. Au, "Entropy-Based Optimal Sensor Location for Structural Model Updating," *Journal of Vibration and Control*, vol. 6, no. 5, pp. 781–800, Jan. 2000.

[2] J. L. Beck, L. S. Katafygiotis, "Updating Models and Their Uncertainties. I: Bayesian Statistical Framework," *Journal of Engineering Mechanics*, vol. 124, no. 4, pp. 455–461, Apr. 1998.

[3] C. Guestrin, A. Krause, A. P. Singh, "Near-optimal sensor placements in Gaussian processes," *Proceedings of the 22nd international conference on Machine learning - ICML '05*, New York, New York, USA: ACM Press, 2005, pp. 265–272, ISBN: 1595931805.

[4] S. Joshi, S. Boyd, "Sensor Selection via Convex Optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, Feb. 2009.

[5] J. L. Beck, C. Papadimitriou, S.-K. Au, M. W. Vanik, "Entropy-based optimal sensor location for structural damage detection," S.-C. Liu, Ed., Jun. 1998, pp. 161–172.

[6] S. Boyd, L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004, ISBN: 0521833787.

[7] J. L. Beck, K. M. Zuev, "ASYMPTOTICALLY INDEPENDENT MARKOV SAMPLING: A NEW MARKOV CHAIN MONTE CARLO SCHEME FOR BAYESIAN INFERENCE," *International Journal for Uncertainty Quantification*, vol. 3, no. 5, pp. 445–474, 2013.

Chapter 4

MODELING ALLOWABLE SENSOR LOCATIONS

So far, we have considered problems where the possible sensor locations correspond to finite element nodes of the model at hand. This choice was made partially due to the fact that the nodes may actually correspond to sensor locations in a real problem, for example in a coarse model with nodes corresponding to stories of a building; and also partially due to the face that it is convenient to do so. It is to be recognized that the set of possible sensor locations is a design choice in some problems.

4.1 Sensor placement on a continuum

Here, we consider the case where the entire continuum of an elastic bar is available for the placement of a single sensor, and apply the formulation developed in Chapter 3.

Standing wave on an elastic bar

Consider a bar of unit length defined on the domain $x \in [0, 1]$. The bar is elastic and is governed by the wave equation with wave speed α, as:

$$u_{,tt} = \alpha^2 u_{,xx} \tag{4.1}$$

with fixed, zero displacement boundary conditions:

$$u(0, t) = 0 \text{ and } u(1, t) = 0 \tag{4.2}$$

and a sinusoidal initial displacement:

$$u(x, 0) = \sin(\pi x) \text{ and } u_{,t}(x, 0) = 0 \tag{4.3}$$

The solution to this problem is given by:

$$u(x, t) = \sin(\pi x) \sin(\pi \alpha t) \tag{4.4}$$

Given the wave speed, α, the predictions at equally spaced intervals of time, $t_n = n\Delta t$, is given by:

$$u(x, n\Delta t) = u(x, t) = \sin(\pi x) \sin(\pi \alpha n \Delta t) \tag{4.5}$$

We are now in a position to evaluate the matrix of sensitivity coefficients, $Q(x, \alpha)$. Note that it is a function of every possible sensor location, $x \in [0, 1]$.

$$Q(x, \alpha) = \sum_{n=1}^{N} \left(\frac{\partial u(x, n\Delta t)}{\partial \alpha} \right)^2 \tag{4.6}$$

$$= (\pi n \Delta t)^2 \sin^2(\pi x) \left(\sum_{n=1}^{N} \cos(n\pi\alpha\Delta t) \right)^2 \tag{4.7}$$

The uncertain parameter, α, is distributed according to a prior, $p(\alpha)$. We wish find the optimal sensor location, x^*, as the solution to the following optimization problem:

$$x^* = \arg\max_{x \in [0,1]} h(x) \tag{4.8}$$

Here, according to the theory developed in Chapter 3, the objective function, $h(x)$ is given by:

$$h(x) = \mathbb{E}_\alpha[\log \det Q(x, \alpha)] \tag{4.9}$$

The previous matrix of sensitivity coefficients, $Q(x, \alpha)$, is simply a scalar here and is a differentiable function of the position, x. We can set the first derivative of the objective function with respect to x in Equation (4.8) to determine an extremum as:

$$0 = \left. \frac{dh(x)}{dx} \right|_{x=x^*} \tag{4.10}$$

$$\Rightarrow 0 = \left. \frac{d}{dx} \mathbb{E}_\alpha \log \left[(\pi n \Delta t)^2 \sin^2(\pi x) \left(\sum_{n=1}^{N} \cos(n\pi\alpha\Delta t) \right)^2 \right] \right|_{x=x^*} \tag{4.11}$$

$$\Rightarrow 0 = 2\frac{d}{dx} \mathbb{E}_\alpha \log(\pi n \Delta t) + \left. 2\frac{d}{dx} \mathbb{E}_\alpha \log \sin(\pi x) \right|_{x=x^*} + 2\,\mathbb{E}_\alpha \log \left(\sum_{n=1}^{N} \cos(n\pi\alpha\Delta t) \right) \tag{4.12}$$

$$\Rightarrow 0 = \cot(\pi x^*) \tag{4.13}$$

$$\Rightarrow x^* = \frac{1}{2} \tag{4.14}$$

As expected by an argument from symmetry, the location of the optimal sensor is at the mid-point of the bar. Notice in Equations (4.10) through (4.13) that the characterization of the prior uncertainty of α does not affect the sensor location. This is a property of this example.

4.2 Finite-element mesh refinement

Typically, unlike in the previous example of Section 4.1, it is not possible to solve the governing differential equation analytically. In practice, the finite element method is used to determine an approximate solution to the field of interest. The solution to a finite element problem begins with obtaining the value of the approximating field at a finite number of nodes. For locations between nodes, shape functions are used to determine the solution there. The number of nodes in the model depends on the mesh size, which could be uniform or non-uniform.

In the original log-determinant sensor-placement formulation, the finite-element nodes were considered to be the only possible sensor locations. Under mesh-refinement, these possible locations change and increase in number. While the effect of refining the mesh of the underlying finite-element is to increase the accuracy of the approximate field, it also increases the number of possible sensor locations due to the increased number of nodes.

Example: Finite-element formulation of an axial bar

Consider an elastic axial rod with wave speed v of length L. One end at $x = 0$ is fixed and the other end is subject to a sinusoidal strain, $\varepsilon_f = \varepsilon_0 \sin \omega t$. This can be thought of as a forcing stress if the elastic modulus of the bar is known.

The governing PDE is:

$$\begin{aligned} \frac{\partial^2 u}{\partial t^2} &= v^2 \frac{\partial^2 u}{\partial x^2} \\ \text{with BCs: } u(0,t) &= 0, \\ \partial u/\partial x(L,t) &= \varepsilon_0 \sin \omega t, \\ \text{and ICs: } u(x,0) &= 0, \\ \partial u/\partial t(x,0) &= 0. \end{aligned} \tag{4.15}$$

A finite element solution is introduced with the help of shape-functions:

$$\tilde{u}(x,t) = \sum_{i=1}^{N_d} u_i(t)\phi_i(x) \tag{4.16}$$

The weak form of the wave equation is now set up:

$$\begin{aligned}
&\int_{\mathcal{V}} \ddot{\tilde{u}} w dx && -v^2 \int_{\mathcal{V}} \frac{\partial^2 \tilde{u}}{\partial x^2} w dx && = 0 \\
\Rightarrow &\int_{\mathcal{V}} \sum_{i=1}^{N_d} \left(\ddot{u}_i(t)\phi_i(x)\right) u_j(t)\phi_j(x) dx && +v^2 \int_{\mathcal{V}} \frac{\partial u}{\partial x}\frac{\partial w}{\partial x} dx && = 0 \qquad \text{(int by parts)} \\
\Rightarrow &\sum_{i=1}^{N_d} \ddot{u}_i(t) \int_{\mathcal{V}} \phi_i(x)\phi_j(x) dx && +v^2 \sum_{i=1}^{N_d} u_i(t) \int_{\mathcal{V}} \phi_i'(x)\phi_j'(x) dx && = 0 \ \forall\ j
\end{aligned} \tag{4.17}$$

In matrix form, this may be written as:

$$M\ddot{U}(t) + KU(t) = f(t) \tag{4.18}$$

The forcing term on the right is from enforcing the Neumann BC at the last node.

A Rayleigh damping term may also be introduced in the final matrix equation. The damping matrix would be a linear combination of the mass and stiffness matrices.

$$M\ddot{U}(t) + C\dot{U}(t) + KU(t) = f(t) \tag{4.19}$$

The matrix equation 4.19 can now be solved using Newmark integration, given the parameters L, v, ε_0 and ω. For a full simulation, additional parameters such as the Rayleigh parameters α and β, as well as the time step Δt, would need to be specified.

The time-integrated solution yields the nodal displacements:

$$U(t_n) = [u_1(t_n), \ldots, u_{N_d}(t_n)]^T \tag{4.20}$$

These can be used to construct a spatially continuous displacement field, at discrete times, according to equation 4.16:

$$\tilde{u}(x, t_n) = \sum_{i=1}^{N_d} u_i(t_n)\phi_i(x) \tag{4.21}$$

Shape functions

A linear, 2-node element is chosen, so that the shape functions are defined as follows:

$$\begin{aligned}
\phi_i(x) &= \begin{cases} \dfrac{x - x_{i-1}}{h}, & x \in [x_{i-1}, x_i] \\ 1 - \dfrac{x - x_i}{h}, & x \in [x_i, x_{i+1}] \\ 0, & x \notin [x_{i-1}, x_{i+1}] \end{cases}, i \in \{2, \ldots, N_d - 1\} \\
\phi_1(x) &= \begin{cases} 1 - \dfrac{x - x_1}{h}, & x \in [x_1, x_2] \\ 0, & x \notin [x_1, x_2] \end{cases} \\
\phi_{N_d}(x) &= \begin{cases} \dfrac{x - x_{N_d-1}}{h}, & x \in [x_{N_d-1}, x_{N_d}] \\ 0, & x \notin [x_{N_d-1}, x_{N_d}] \end{cases}
\end{aligned} \tag{4.22}$$

The shape function derivatives are simply piece-wise constant corresponding to the slopes of the lines in the non-zero regions.

Uncertainty quantification

Consider a specific case, where three uncertain system parameters are considered:

- Bar wave speed, v
- Forcing amplitude, ε_0
- Forcing frequency, ω

In addition, a prediction error uncertainty parameter, σ^2, is considered.

The parameter vector is therefore represented by $\theta = [\theta_s, \sigma^2]$, where the system parameters are given by $\theta_s = [v, \varepsilon_0, \omega]$.

A bar of unit length is considered and the damping parameters can be chosen so that, at the prior modal values of the parameters in the list above, a damping ratio of 5% is achieved in the first two modes.

The finite-element solution notation is now augmented to include the uncertain parameters. That is, the FEM solution is now denoted by:

$$\tilde{u}(x, t, \theta_s) = \sum_{i=1}^{N_d} u_i(x, \theta_s)\phi_i(x) \tag{4.23}$$

In order to have the continuous treatment, we assume that observed displacement $Y_x(t)$ is governed by a Gaussian process (GP) over spatial coordinate index, x, with mean function $\mathcal{M}_t(x|\theta)$ and co-variance kernel $\mathcal{K}_t(x, y|\theta)$. The mean prediction is given by $\mathcal{M}_t(x|\theta_s) = \tilde{u}(x, t, \theta_s)$. The co-variance kernel is assumed to be stationary and isotropic, so that $\mathcal{K}_t(x, y|\sigma^2) = \sigma^2\delta_{xy}$.

Formulation of the optimization problem

Points in the domain, $[0, L]$ are indexed by the set $\mathcal{V}$ and the sensor location set is denoted by $\mathcal{A}$. Thus, $\mathcal{A} \subset \mathcal{V}$. Note that $\mathcal{A}$ is finite with $|\mathcal{A}| = N_o$. The observed displacements at the locations indexed by $\mathcal{A}$ are distributed according to a multivariate Gaussian as per the GP as:

$$y_{\mathcal{A}n}|\theta \sim \mathcal{N}\left(\cdot|\tilde{u}(\mathcal{A}, t_n, \theta_s), \sigma^2 I\right) \tag{4.24}$$

(Notation: $\tilde{u}(\mathcal{A}, t_n, \theta_s)$ is a vector of size N_o, corresponding to displacement predictions at each location in $\mathcal{A}$ at time t_n)

Modeling the prediction of the observations at different times as probabilistically independent, the likelihood of a full observation record over N observations is:

$$p(Y_{\mathcal{A}N}|\theta) = \prod_{n=1}^{N} p(y_{\mathcal{A}n}|\theta) \tag{4.25}$$

Now, it is assumed that the model is globally identifiable, so that there exists a unique global maximum likelihood estimate:

$$\hat{\theta} = [\hat{\theta_s}, \hat{\sigma}^2] = \arg\max_{\theta} p(Y_{\mathcal{A}N}|\theta) \tag{4.26}$$

The partial MLE estimate of σ^2, viz. $\hat{\sigma}^2(\theta_s)$, follows through as per the original formulation:

$$\hat{\sigma}^2(\theta_s) = J(\theta_s) = \frac{1}{NN_o}\sum_{n=1}^{N} \|y_{\mathcal{A}n} - \tilde{u}(\mathcal{A}, t_n, \theta_s)\|^2 \tag{4.27}$$

The Laplace approximation can now be applied to the posterior on the system parameters as in the original formulation, to obtain:

$$p(\theta_s|Y_{\mathcal{A}N}) \approx \frac{[\det Q(\mathcal{A}, \hat{\theta}_s)]^{1/2}}{(2\pi\hat{\sigma}^2)^{N_a/2}} \exp\left(-\frac{1}{2\hat{\sigma}^2}(\theta_s - \hat{\theta}_s)^T Q(\mathcal{A}, \hat{\theta}_s)(\theta_s - \hat{\theta}_s)\right) \tag{4.28}$$

The elements of the matrix, Q, have their origins in the Hessian matrix of the Laplace approximation:

$$Q(\mathcal{A}, \theta_s) = \sum_{n=1}^{N} (\nabla_{\theta_s} \tilde{u})(\nabla_{\theta_s} \tilde{u})^T(\mathcal{A}, t_n, \theta_s) \tag{4.29}$$

The optimization problem is therefore:

$$\mathcal{A}^* = \underset{\substack{\mathcal{A} \subset \mathcal{V} \\ |\mathcal{A}| = N_o}}{\arg\max}\ \mathbb{E}_{\theta_s}[\log\det Q(\mathcal{A}, \theta_s)] \tag{4.30}$$

Here, the expectation is taken over a nominal distribution on the system parameters specified by the designer.

For convenience, denote the objective by $f(\mathcal{A}) = \mathbb{E}[\log\det Q(\mathcal{A}, \theta_s)]$. Note that the expectation variable, θ_s, is implicit in the notation from now on.

4.3 Comparison with the finite selection-space case

When the list of possible sensor locations was finite, then a selection matrix, $S_o(\delta)$, based on the boolean sensor-selection vector, δ, could be employed to split the matrix $Q(\delta, \theta_s)$ into N_d elementary matrices $Q_i(\theta_s)$. The optimization was then over a finite set of points of cardinality N_d. That is, the optimization problem was a combinatorial one of selecting N_o of the N_d elementary sensing matrices $Q_i(\theta_s)$ whose sum maximized the expectation-log-determinant objective function.

In the continuous case, we have an continuum of points to select from. This hints towards a gradient-based optimization as opposed to a combinatorial one.

Simple case: single sensor placement

Consider the simple problem of selecting just one sensor location optimally. In this case, $|\mathcal{A}| = 1$. We would like to determine:

$$\begin{aligned} x^* &= \underset{x \in \mathcal{V}}{\arg\max}\, f(x) \\ &= \underset{x \in \mathcal{V}}{\arg\max}\ \mathbb{E}[\log\det Q(x, \theta_s)] \end{aligned} \tag{4.31}$$

With a slight abuse of notation, we consider the derivative of the objective with respect to location, x, assuming that it exists. If the optimal location is in the interior, then this derivative should be zero. Else, the optimal location is at the boundary according to the intermediate value theorem.

Specifically,

$$f'(x^*) = 0 \text{ or } x^* = 0 \text{ or } x^* = L \tag{4.32}$$

This calls for an investigation of the derivative:

$$\begin{aligned} f'(x) &= \frac{d}{dx}\mathbb{E}\log\det Q(x,\theta_s) \\ &= \mathbb{E}\frac{d}{dx}\log\det Q(x,\theta_s) && \text{(assuming allowed)} \\ &= \mathbb{E}\left[\text{tr}\left(Q^{-1}(x,\theta_s)\frac{d}{dx}Q(x,\theta_s)\right)\right] && \text{(derivative of log-det)} \end{aligned} \tag{4.33}$$

Consider the second term inside the trace involving the derivative of the sensing matrix:

$$\frac{d}{dx}Q(x,\theta_s) = \frac{d}{dx}\sum_{n=1}^{N}\nabla\tilde{u}\,\nabla\tilde{u}^T(x,t_n,\theta_s) \tag{4.34}$$

(Note: The implicit vector gradient, ∇, is with respect to the system parameters, θ_s)

The derivative needs to be performed entry-wise and we therefore need to look at:

$$\begin{aligned} \frac{d}{dx}Q_{pq}(x,\theta_s) &= \frac{d}{dx}\sum_{n=1}^{N}\frac{\partial\tilde{u}}{\partial\theta_{sp}}\frac{\partial\tilde{u}}{\partial\theta_{sq}}(x,t_n,\theta_s) \\ &= \sum_{n=1}^{N}\left(\frac{\partial}{\partial\theta_{sp}}\frac{\partial\tilde{u}}{\partial x}\right)\frac{\partial\tilde{u}}{\partial\theta_{sq}} + \left(\frac{\partial}{\partial\theta_{sq}}\frac{\partial\tilde{u}}{\partial x}\right)\frac{\partial\tilde{u}}{\partial\theta_{sp}} \end{aligned} \tag{4.35}$$

(assuming all interchanges allowed)

At this point, it is convenient to denote the x derivative of the displacement by strain, $\tilde{\varepsilon}(x,t_n,\theta_s)$. We then have,

$$\frac{d}{dx}Q_{pq}(x,\theta_s) = \sum_{n=1}^{N}\left(\frac{\partial\tilde{\varepsilon}}{\partial\theta_{sp}}\frac{\partial\tilde{u}}{\partial\theta_{sq}} + \frac{\partial\tilde{u}}{\partial\theta_{sp}}\frac{\partial\tilde{\varepsilon}}{\partial\theta_{sq}}\right)(x,t_n,\theta_s) \tag{4.36}$$

Finally, we consider the derivative of the finite-element displacement field with respect to location. One can see that this derivative involves, ultimately, the derivative of the finite element shape functions.

The localized nature of the shape functions comes into play here, assuming that we have used **linear, 2-node elements** to discretize the bar. For a given location, x,

we consider the contribution from only its immediately adjacent shape functions associated with nodes j and $j+1$, so that $x \in [x_j, x_{j+1}]$. Accordingly,

$$\begin{aligned}\tilde{\varepsilon}(x, t_n, \theta_s) &= \frac{\partial \tilde{u}}{\partial x}(x, t_n, \theta_s) \\ &= \sum_{i=1}^{N_d} u_i(t_n, \theta_s)\phi_i'(x) \\ &= u_j(t_n, \theta_s)\phi_j'(x) + u_{j+1}(t_n, \theta_s)\phi_{j+1}'(x)\end{aligned} \tag{4.37}$$

If x happens to be one of the nodes, then only the shape function associated with that node is relevant:

$$\begin{aligned}\tilde{\varepsilon}(x_j, t_n, \theta_s) &= \frac{\partial \tilde{u}}{\partial x}(x_j, t_n, \theta_s) \\ &= u_j(t_n, \theta_s)\phi_j'(x_j)\end{aligned} \tag{4.38}$$

The 2-node, linear elements are constant-strain elements. That is, the derivatives in question have a constant value in the interior of elements. The strain field determined by the finite element solution with this linear interpolation is ill-defined at the finite-element nodes, since the displacement function has a kink there.

This means that the $\phi_j'(x_j)$ term in the resulting expression of Equation (4.38) is undefined. So for now, the possibility of a finite-element nodal point being the optimal solution is considered as a case to be tackled separately.

For a point x not at a node, we have the resulting forward difference equation:

$$\begin{aligned}\tilde{\varepsilon}(x, t_n, \theta_s) &= \frac{\partial \tilde{u}}{\partial x}(x, t_n, \theta_s) \\ &= u_j(t_n, \theta_s)\phi_j'(x) + u_{j+1}(t_n, \theta_s)\phi_{j+1}'(x) \\ &= \frac{(u_{j+1} - u_j)}{h}(t_n, \theta_s)\end{aligned} \tag{4.39}$$

Here, $h = L/(N_d - 1)$, is the mesh width.

We are now equipped to solve the optimization problem. Note that the objective is not convex, since the matrix, $Q(\mathcal{A}, \theta_s)$, is not an affine function of the set of locations in $\mathcal{A}$.

General case: multiple sensors

In this case, we attempt to set the whole gradient, with respect to N_o possible locations in the set $\mathcal{A}$, to zero. Begin by evaluating the gradient:

$$\begin{aligned}\frac{\partial f(\mathcal{A})}{\partial x_a} &= \frac{\partial}{\partial x_a}\mathbb{E}\left[\log\det Q(\mathcal{A},\theta_s)\right] \\ &= \mathbb{E}\left[\operatorname{tr}\left\{Q^{-1}(\mathcal{A},\theta_s)\frac{\partial}{\partial x_a}Q(\mathcal{A},\theta_s)\right\}\right]\end{aligned} \tag{4.40}$$

Now, for the matrix derivative term, we have:

$$\begin{aligned}\frac{\partial}{\partial x_a}Q(\mathcal{A},\theta_s) &= \frac{\partial}{\partial x_a}\sum_{n=1}^{N}\sum_{b\in\mathcal{A}}\left[\nabla\tilde{u}\,\nabla\tilde{u}^T(x_b,t_n,\theta_s)\right] \\ &= \sum_{n=1}^{N}\frac{\partial}{\partial x_a}\sum_{b\in\mathcal{A}}\left[\nabla\tilde{u}\,\nabla\tilde{u}^T(x_b,t_n,\theta_s)\right] \\ &= \sum_{n=1}^{N}\frac{\partial}{\partial x_a}\left[\nabla\tilde{u}\,\nabla\tilde{u}^T(x_a,t_n,\theta_s)\right]\end{aligned} \tag{4.41}$$

Note that, in general,

$$\begin{aligned}\tilde{u}(x_a,t_n,\theta_s) &= u_j(t_n,\theta_s)\phi_j(x_a)+u_{j+1}(t_n,\theta_s)\phi_{j+1}(x_a) \\ \Rightarrow\nabla\tilde{u}(x_a,t_n,\theta_s) &= \nabla u_j(t_n,\theta_s)\phi_j(x_a)+\nabla u_{j+1}(t_n,\theta_s)\phi_{j+1}(x_a) \\ \Rightarrow\nabla\tilde{u}\nabla\tilde{u}^T(x_a,t_n,\theta_s) &= \nabla u_j\nabla u_j^T(t_n,\theta_s)\phi_j(x_a)\phi_j(x_a)+ \\ &\quad \nabla u_{j+1}\nabla u_{j+1}^T(t_n,\theta_s)\phi_{j+1}(x_a)\phi_{j+1}(x_a)+ \\ &\quad 2\nabla\tilde{u}_j\nabla\tilde{u}_{j+1}^T(t_n,\theta_s)\phi_j(x_a)\phi_{j+1}(x_a)\end{aligned} \tag{4.42}$$

The spatial partial derivative may now be evaluated:

$$\begin{aligned}\frac{\partial}{\partial x_a}\nabla\tilde{u}\nabla\tilde{u}^T(x_a,t_n,\theta_s) &= [2\nabla u_j\nabla u_j^T\phi_j(x_a)\phi_j'(x_a)+ \\ &\quad 2\nabla u_{j+1}\nabla u_{j+1}^T\phi_{j+1}(x_a)\phi_{j+1}'(x_a)+ \\ &\quad 2\nabla\tilde{u}_j\nabla\tilde{u}_{j+1}^T(\phi_j'(x_a)\phi_{j+1}(x_a)+\phi_j(x_a)\phi_{j+1}'(x_a))]\end{aligned} \tag{4.43}$$

Solution and sensor clumping

Refer back to the equation 4.40 for the location gradient of the objective. At the optimal sensor locations, this gradient should be zero as long as none of the sensors are at the boundaries of the domain, which in our case is the ends of the bar.

$$\begin{aligned}\frac{\partial f(\mathcal{A})}{\partial x_a} &= 0 \;\forall\; a \in \mathcal{A} \subset \mathcal{V} \\ &= \frac{\partial}{\partial x_a}\mathbb{E}\left[\log\det Q(\mathcal{A},\theta_s)\right] \\ &= \mathbb{E}\left[\operatorname{tr}\left\{Q^{-1}(\mathcal{A},\theta_s)\frac{\partial}{\partial x_a}Q(\mathcal{A},\theta_s)\right\}\right]\end{aligned} \tag{4.44}$$

We therefore have N_o equations in N_o unknowns which can presumably solved for the sensor locations.

There is a problem with this approach and it is this: the gradient vector is symmetric with respect to every sensor location co-ordinate, x_a. What this means is that if one component of the gradient is determined to be zero at a particular location, then the exclusion of that point results in an ill-posed problem, since the point immediately 'adjacent' to the excluded point also result in a zero gradient. We therefore have the problem of all sensor locations concentrating at a single location.

Resolution

The root of the problem is the fact that the formulation developed in Chapter 3 does not model an important aspect of the sensor selection problem: prediction-error correlations at small length scales.

The isotropic co-variance kernel fails to capture that fact that the prediction errors at points arbitrarily close to each other are correlated in practice. In order to capture this correlation, one could select a co-variance kernel that is not diagonal as:

$$\mathcal{K}_t(x, y|\theta) = \sigma^2 \exp(-c|x - y|) \tag{4.45}$$

The exponential factor in Equation (4.45) takes on a non-zero value for points that are not coincident. This correlation decays exponentially as the distance between the points considered. A multiplicative factor controls the desired length scale.

The problem with this approach is that the kernel violates the underlying assumptions of the framework, and the new resulting uncertainty objective cannot be expressed neatly in terms of a sum of positive definite matrices associated with each possible sensor location. A preliminary sensor deployment might provide some insight into determining the structure of the co-variance kernel [1], although this is ultimately a modeling choice. It should be noted, however, that in the absence of prior information about this correlation structure, the maximum entropy principle tells us

to choose an isotropic co-variance matrix for the discretized version of the problem. This, as we know, results in sensor clumping. In this situation, the best resolution is to specify a pre-defined set of possible sensor locations that are sufficiently far apart or appropriately positioned such that no correlation between the prediction errors at those locations is suspected *a priori*. This is the approach that was used in Chapter 3.

An example examining the effect of prediction error correlation on sensor placement on a cantilever beam is discussed next.

4.4 Example: free-vibration of a cantilever beam with prediction-error correlations

Here, we look at the optimal sensor placement problem with prediction-error correlations included as considered in [2] and [3]. Here, we consider the case of a free-vibrating, uniform Euler-Bernoulli cantilever beam . A cantilever beam is fixed at one end and free at the other for displacements and rotations. The governing equation for the beam is:

$$EI\frac{\partial^4 v}{\partial x^4} + \mu\frac{\partial^2 v}{\partial t^2} = 0 \tag{4.46}$$

with:

$$\text{Initial conditions: } v(x,0) = f(x) \text{ and } \dot{v}(x,0) = 0 \tag{4.47}$$

$$\text{Boundary conditions: } v(0,t) = 0, \frac{\partial v}{\partial x}(0,t) = 0 \tag{4.48}$$

$$\frac{\partial v^2}{\partial x^2}(L,t) = 0 \text{ and } \frac{\partial v^2}{\partial x^3}(L,t) = 0 \tag{4.49}$$

This problem can be solved using separation of variables as:

$$v(x,t) = \sum_{m=1}^{\infty} X_m(x)T_m(t) \tag{4.50}$$

where

$$X_m(x) = F_m\left[(\cosh\beta_m x - \cos\beta_m x) + \frac{\cosh\beta_m L + \cos\beta_m L}{\sinh\beta_m L + \sin\beta_m L}(\sin\beta_m x - \sinh\beta_m x)\right] \tag{4.51}$$

and

$$T_m(t) = \cos(\omega_m t) \tag{4.52}$$

Equation (4.51) is for the mode shapes of vibration. Here, m is used to index the mode corresponding to the natural frequency of vibration, ω_m.

The coefficients F_m are the Fourier coefficients under each mode shape, of the function $f(x)$ that describes the initial displacement condition.

The coefficients, β_m are obtained as a result of looking for valid solutions to the separation-of-variables method of solution. They are solutions to the equation:

$$\cosh \beta_m L \cos \beta_m L + 1 = 0 \tag{4.53}$$

The natural frequencies are given by:

$$\omega_m = \beta_m^2 \sqrt{\frac{EI}{\mu}} \tag{4.54}$$

Finally, $T_m(t)$ is the time-dependent part of the contribution from the m'th mode.

Correlated sensor placement for first modal frequency parameter

Here, we wish to compare the difference in mutual information, in bits, of the optimal single sensor configuration versus the optimal two-sensor configuration for different values of the prediction-error correlation parameter for the cantilever beam, for the purpose of updating uncertainties associated with the first modal frequency parameter, ω_1.

Entropy of the posterior

For two sensors located at co-ordinates p and q with correlated prediction-errors as:

$$\Sigma = \sigma^2 \begin{bmatrix} 1 & c(p,q) \\ c(p,q) & 1 \end{bmatrix} \tag{4.55}$$

Here, the chosen model for the correlation coefficient decays exponentially as the square of the distance between the two locations:

$$c(p,q) = \exp\left[-\left(\frac{p-q}{\lambda}\right)^2\right] \tag{4.56}$$

The quantity, λ, can be interpreted as a correlation length. In this analysis, variation of the entropic quantities of interest as a function of λ are investigated. The determinant of the covariance matrix is:

$$|\Sigma| = \sigma^4(1 - c^2(p,q)) \tag{4.57}$$

The inverse of the covariance matrix is:

$$\Sigma^{-1} = \frac{1}{\sigma^2(1 - c^2(p,q))} \begin{bmatrix} 1 & -c(p,q) \\ -c(p,q) & 1 \end{bmatrix} \tag{4.58}$$

Laplace approximation

We begin by writing out the likelihood function and then expressing the posterior over the parameters as a Gaussian distribution at the MLE parameters.

$$p(y_{1:N}|\theta_s, \sigma^2) = \frac{1}{(2\pi)^{NN_o/2}|\Sigma|^{N/2}} \exp\left[-\frac{1}{2}\sum_{n=1}^{N}(y_n - x_n(\theta_s))^T\Sigma^{-1}(y_n - x_n(\theta_s))\right] \tag{4.59}$$

Here, $N_o = 2$ and we may substitute the expressions for the covariance matrix determinant and inverse:

$$f(y_{1:N}, \theta_s, \sigma^2) = p(y_{1:N}|\theta_s, \sigma^2) \tag{4.60}$$

$$= \frac{1}{(2\pi\sigma^2\sqrt{1 - c^2(p,q)})^{NN_o/2}} \exp\left[-\frac{\sum_{n=1}^{N}(e_p^2 + e_q^2 - 2c(p,q)e_p e_q)}{2\sigma^2(1 - c^2(p,q))}\right] \tag{4.61}$$

For the convenience of notation, the following quantities have been introduced:

$$e_p \equiv e_p(\theta_s, p, q) = (y_{np} - x_{np}(\theta_s)) \text{ and } e_q \equiv e_q(\theta_s, p, q) = (y_{nq} - x_{nq}(\theta_s)) \tag{4.62}$$

The optimal prediction error parameter, $\hat{\sigma}^2$, as a function of the system parameters is:

$$\hat{\sigma}^2(\theta_s) = \frac{\sum_{n=1}^{N}(e_p^2 + e_q^2 - 2c(p,q)e_p e_q)(\theta_s)}{NN_o(1 - c^2(p,q))} \tag{4.63}$$

We have, therefore:

$$f(y_{1:N}, \theta_s, \hat{\sigma}^2(\theta_s)) = (2\pi e\hat{\sigma}^2(\theta_s)\sqrt{1 - c^2(p,q)})^{NN_o/2} \tag{4.64}$$

The relevant part of the inverse covariance matrix, in the Laplace approximation to the posterior, is therefore:

$$B_{ij} = \frac{\sum_{n=1}^{N}\left(\frac{\partial x_n(p)}{\partial\theta_{si}}\frac{\partial x_n(p)}{\partial\theta_{sj}} + \frac{\partial x_n(q)}{\partial\theta_{si}}\frac{\partial x_n(q)}{\partial\theta_{sj}} - 2c(p,q)\frac{\partial x_n(p)}{\partial\theta_{si}}\frac{\partial x_n(q)}{\partial\theta_{sj}}\right)\Big|_{\hat{\theta}_s}}{\hat{\sigma}^2} \tag{4.65}$$

We would like to compare the the information gain in bits by placing two sensors versus placing a single sensor. In the case of the cantilever beam, in order to learn the first mode of a free-vibration-from-initial-displacement response, we have the deterministic prediction:

$$x_n(p) = \sum_{m=1}^{\infty} X_m(p)\cos(\omega_m t_n) \tag{4.66}$$

The sensitivity coefficient for the first modal frequency is therefore:

$$\frac{\partial x_n(p)}{\partial \omega_1} = -t_n X_1(p) \sin(\omega_1 t_n) \tag{4.67}$$

The objective of interest is the entropy of the posterior density, expressed as:

$$H(p,q) = \frac{1}{2}\log(X_1^2(p)+X_1^2(q)-2c(p,q)X_1(p)X_1(q))-\frac{1}{2}\log(1-c^2(p,q))+K \tag{4.68}$$

Here, K includes constant and irrelevant time-dependent terms:

$$K = -\frac{1}{2}\log 2\pi e - \frac{1}{2}\log\hat{\sigma}^2 + \frac{1}{2}\mathbb{E}\left[\log\left(\sum_{n=1}^{N} t_n^2 \sin^2(\omega_1 t_n)\right)\right] \tag{4.69}$$

Let the reference be the placement of a single sensor on the cantilever beam for the same problem. We would like to know the difference in the mutual information in terms of bits as before. This can be calculated by simply taking the difference in the log-determinants of the corresponding block diagonals parts of the covariance matrix in the Laplace approximation to the posterior. The expected information gain is therefore:

$$\Delta I = H_0(r^*) - H(p^*, q^*) \tag{4.70}$$

In Equation (4.70), the quantity $H_0(r^*)$ refers to the objective evaluated at the optimal sensor location for a single sensor only:

$$H_0(r) = \log X_1(r) + K \tag{4.71}$$

The logarithm is expressed in base 2 to express the difference in entropies in bits. Let B_0 denote the covariance matrix of the posterior for a single sensor's placement:

$$B_0 = \frac{1}{\hat{\sigma}^2}\sum_{n=1}^{N} X_1^2(r) t_n^2 \sin^2(\omega_1 t_n) \tag{4.72}$$

The determinant is not needed since we are working with only the single scalar fundamental frequency parameter. The term, K, is common to both entropies and cancels out in the subtraction. The resulting expression in:

$$\Delta I = \frac{1}{2}\log_2(X_1^2(p^*) + X_1^2(q^*) - 2c(p^*, q^*)X_1(p^*)X_1(q^*)) - \log_2 X_1(r^*) \tag{4.73}$$

For this problem, it turns out the the optimal location for a single sensor is at the free end of the cantilever beam. For the 2-sensor problem, it turns out that one sensor is always placed at the free end while the other is determined by the correlation

length parameter, λ. The information gain from the second sensor is a function of the correlation length parameter.

Figure 4.1 plots the optimal location, q^*, of the second sensor with respect to the prediction-error correlation length scale, λ. As λ increases, the second sensor is driven further away from the first, up to a limit beyond which it prefers place the sensor at the non-informative fixed end of the beam.

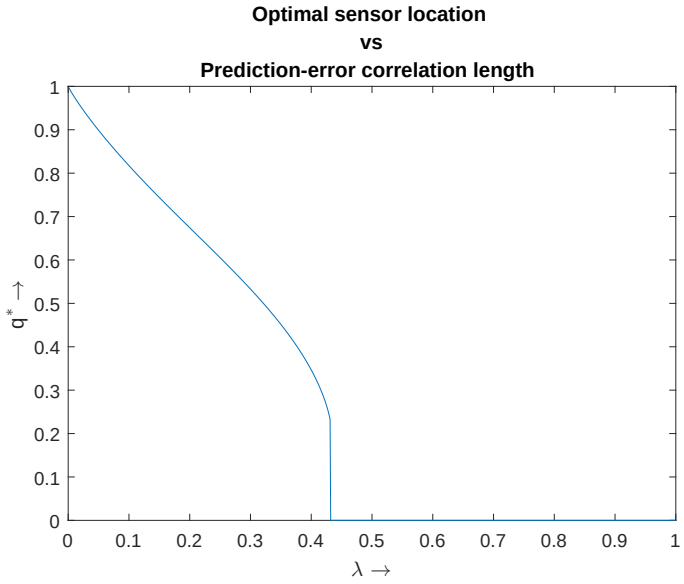

Figure 4.1: Optimal location of the second sensor as a function of the prediction-error correlation length scale (note that the first sensor is always at the free end of the cantilever beam)

Figure 4.2 is a plot of the additional information gain from the second sensor when it is placed optimally.

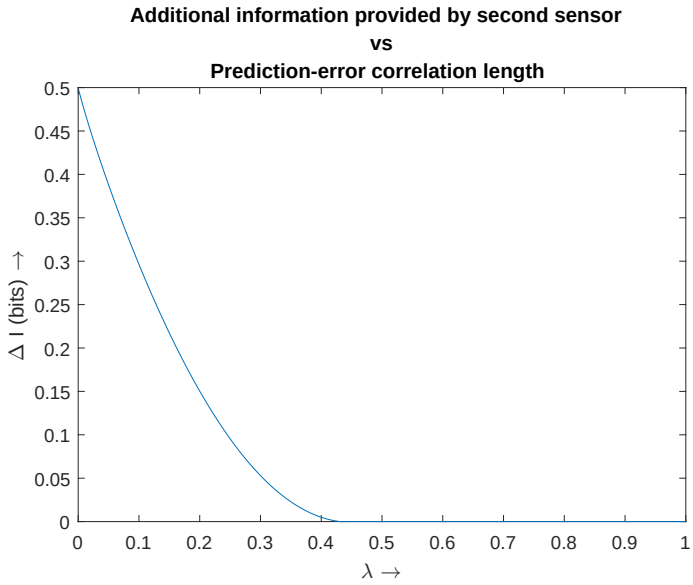

Figure 4.2: Additional bits of information from the optimally placed second sensor plotted as a function of the correlation length scale

At the critical correlation length scale of about $\lambda = 0.42$ as a fraction of the length of the beam, the second sensor becomes non-informative.

Through this example, we conclude that modeling prediction-error correlations is essential to prevent sensor clustering or clumping under mesh refinement. In the absence of a correlation provision, a predetermined set of possible sensor locations, spaced sufficiently apart, should be selected to prevent sensor clumping.

References

[1] C. Guestrin, A. Krause, A. P. Singh, "Near-optimal sensor placements in Gaussian processes," *Proceedings of the 22nd international conference on Machine learning - ICML '05*, New York, New York, USA: ACM Press, 2005, pp. 265–272, ISBN: 1595931805.

[2] D. Bayard, F. Hadaegh, D. Meldrum, "Optimal experiment design for identification of large space structures," *Automatica*, vol. 24, no. 3, pp. 357–364, May 1988.

[3] C. Papadimitriou, G. Lombaert, "The effect of prediction error correlation on optimal sensor placement in structural dynamics," *Mechanical Systems and Signal Processing*, vol. 28, pp. 105–127, 2012.

Chapter 5

BAYESIAN MODEL CLASS SELECTION USING THERMODYNAMIC INTEGRATION

5.1 Introduction

The goal of system identification is to find an optimal model that can make good predictions about the behavior of the system in the future. Given data from the system, making good predictions involves two aspects: data fit and Ockham's principle.

A sufficiently complex model can always fit the data at hand to any desired accuracy and therefore, as far as data-fit is concerned, there are many plausible models to choose from. Ockham's principle states that of these models, the simpler ones are more plausible for making reasonable predictions since they do not have a tendency to overfit the data.

Using Bayes' Theorem, one can obtain a posterior probability distribution over the uncertain parameters that quantifies the plausibility of each model in the class, by updating prior uncertainties in a manner consistent with the stochastic forward model. If several model classes are being considered, however, the challenging problem of model class selection arises. That is, one would like to have a way of comparing the predictive power of every chosen candidate model class.

It is known that the model evidence term, together with a prior over the model classes, allows one to evaluate the posterior probability of each model class, following which, model averaging or model class selection may be performed. The model evidence term respects Ockham's principle and penalizes more complex models that offer the same goodness-of-fit to the data ([1]).

In order to compare different model classes based on available data from a system using the Bayesian approach, computation of the evidence for each model class is required. Each evidence term can then be used to either choose the most plausible model or perform model averaging over quantities of interest [1].

5.2 Numerical Evaluation of Model Evidence

Consider a probability model class, $\mathcal{M}$, and available system output data, $\mathcal{D}$. Denote the uncertain parameters in $\mathcal{M}$ by θ. $\mathcal{M}$ specifies a stochastic forward-model, $p(Y|\theta, \mathcal{M})$, as well as a prior distribution, $p(\theta|\mathcal{M})$, over the uncertain parameters.

Of interest is the evidence for the model class under the data, $p(Y|\mathcal{M})$. The basic analytic expression for this quantity is given by Total Probability:

$$p(Y|\mathcal{M}) = \int p(Y|\theta, \mathcal{M})p(\theta|\mathcal{M})\, d\theta \tag{5.1}$$

Difficulties in computing model evidence

The likelihood function, which is the stochastic forward model with the actual data substituted, viewed as a function of the uncertain parameters, θ, can, in general, be a complex function of θ. Thus, except in the restricted case of priors algebraically conjugate to the likelihoods, it is very unlikely that an analytic expression for the evidence can be obtained.

Eqn. (5.1) can be viewed as the expected value of the likelihood function under the prior:

$$p(Y|\mathcal{M}) = \mathbb{E}_{\theta\sim p(\theta|\mathcal{M})}[p(Y|\theta, \mathcal{M})] \tag{5.2}$$

The question of how to go about evaluating this quantity needs to be addressed. One can imagine estimating the model evidence with a Monte-Carlo estimate using samples drawn from the prior density. The estimate would be given by:

$$p(Y|\mathcal{M}) \approx \frac{1}{N_s}\sum_{k=1}^{N_s} p(Y|\theta^{(k)}, \mathcal{M}) \tag{5.3}$$

For problems involving the stochastic embedding of dynamical systems with several uncertain parameters, estimates for the evidence term using prior samples alone can have unacceptably high variance. This is because the topography of likelihood function is typically different from that of the prior.

In the case of locally identifiable systems with a finite number of local *a posteriori* maxima, the Laplace approximation can be used to obtain a reasonable estimate for the log-evidence ([1]). The current endeavor is towards developing MCMC methods to handle more general probability model classes.

Thermodynamic Integration

In order to overcome the difficulties that arise when trying to compute the evidence directly, the concept of path sampling and thermodynamic integration for evaluation of model evidence are introduced along the lines of [2]. The path sampling estimate of evidence is obtained by integrating the log-likelihood function over distributions intermediate to the prior and the posterior. Such intermediate distributions arise naturally in annealing-based sampling schemes, such as AIMS. If these intermediate distributions are chosen carefully, then the aforementioned integration of the log-likelihood function can be performed numerically. Path sampling is also numerically stable since it deals with the logarithm of the likelihood function at every stage of the calculation, for which regular floating point arithmetic is sufficient [3].

Path sampling employs samples from distributions intermediate to the prior and posterior in order to reduce the variance of the evidence estimate. An exponent, $\beta \in [0, 1]$, controls the contribution of the likelihood function to the intermediate distribution. The following function is the first step in constructing the desired intermediate distribution:

$$q_\beta(Y, \theta) \stackrel{\Delta}{=} [p(Y|\theta, \mathcal{M}]^\beta p(\theta|\mathcal{M}) \tag{5.4}$$

Also of importance is the logarithm of this intermediate, un-normalized distribution:

$$\log q_\beta(Y, \theta) = \beta \log p(Y|\theta, \mathcal{M}) + \log p(\theta|\mathcal{M}) \tag{5.5}$$

The normalizing constant needed to turn this function into a distribution on θ is:

$$Z_\beta(Y) = \int q_\beta(Y, \theta) d\theta \tag{5.6}$$

We may thus construct the intermediate probability distribution as:

$$p_\beta(\theta|Y) = \frac{q_\beta(Y, \theta)}{Z_\beta(\theta)} \tag{5.7}$$

As is evident now, the role of β is to control the transition of the prior distribution into the posterior distribution as it ranges over $[0, 1]$. For $\beta = 0$, we can see that the intermediate distribution is equal to the prior: $p_{\beta=0}(Y, \theta) = p(\theta)$. At the other extreme, for $\beta = 1$, the intermediate distribution equals the posterior: $p_{\beta=1}(\theta|Y) = p(\theta|Y, \mathcal{M})$.

The quantity of interest to us is the normalizing constant $Z_{\beta=1}$, which corresponds to the evidence term for our model class:

$$Z_{\beta=1}(Y) = \int p(Y|\theta, \mathcal{M})p(\theta)d\theta \tag{5.8}$$

It turns out that the logarithm of this evidence constant can be determined by thermodynamic integration as follows [2]:

$$\begin{aligned}
\frac{\partial \log Z_\beta(Y)}{\partial \beta} &= \frac{1}{Z_\beta(Y)} \frac{\partial Z_\beta(Y)}{\partial \beta} && (5.9)\\
&= \frac{1}{Z_\beta(Y)} \frac{\partial}{\partial \beta} \int q_\beta(Y, \theta)\, d\theta && (5.10)\\
&= \frac{1}{Z_\beta(Y)} \int \frac{\partial q_\beta(Y, \theta)}{\partial \beta}\, d\theta && (5.11)\\
&= \int \frac{1}{Z_\beta(Y)} \frac{\partial q_\beta(Y, \theta)}{\partial \beta}\, d\theta && (5.12)\\
\Rightarrow \log Z_\beta(Y)\Big|_{\beta=0}^{1} &= \int_{\beta=0}^{1} \int \frac{1}{Z_\beta(Y)} \frac{\partial q_\beta(Y, \theta)}{\partial \beta}\, d\theta d\beta && (5.13)\\
\Rightarrow \log Z_{\beta=1}(Y) - \log Z_{\beta=0} &= \int_0^1 \int \frac{p_\beta(\theta|Y)}{q_\beta(D, \theta)} \frac{\partial q_\beta(Y, \theta)}{\partial \beta}\, d\theta d\beta && (5.14)\\
\Rightarrow \log Z_{\beta=1}(Y) - \log 1 &= \int_0^1 \mathbb{E}_{\theta \sim p_\beta(\theta|Y)} \left[\frac{1}{q_\beta(Y, \theta)} \frac{\partial q_\beta(Y, \theta)}{\partial \beta} \right] d\beta && (5.15)\\
\Rightarrow \log p(Y|\mathcal{M}) &= \int_{\beta=0}^{1} \mathbb{E}_{\theta \sim p_\beta(\theta|Y)} \left[\frac{\partial \log q_\beta(Y, \theta)}{\partial \beta} \right] d\beta && (5.16)
\end{aligned}$$

Using Equation (5.16) now have an expression for the log-evidence in terms of a path integral over the expected value of the log-likelihood function.

$$\log p(Y|\mathcal{M}) = \int_{\beta=0}^{1} \mathbb{E}_{\theta \sim p_\beta(\theta|D)} \left[\log p(Y|\theta, \mathcal{M}) \right] d\beta \tag{5.17}$$

Numerical estimates

Eqn. (5.17) can be estimated numerically. First, the integral over the β unit interval can be approximated by a numerical integration scheme such as the trapezoidal rule or the Gauss-Kronrod numerical integration scheme.

$$A(Y) \triangleq \log p(Y|\mathcal{M}) \approx \sum_{i=1}^{N_w} w_i \mathbb{E}_{\theta \sim p_{\beta_i}(\theta|Y)} \left[\log p(\mathcal{D}|\theta, \mathcal{M}) \right] \tag{5.18}$$

The question remains as to how to compute the integrand each discrete value of β. For each discrete value of β, a simple Monte Carlo estimate for the expected value using samples from an MCMC algorithm will suffice.

For a sampling algorithm such as AIMS that uses simulated annealing [4], the tempering schedule can be chosen to match that of the integration points for path sampling. The sampling algorithm would then automatically generate a set of samples from the intermediate distributions, which, in turn, could be used to determine the appropriate expected values of the log-likelihood in order to construct a numerical estimate of the path sampling integral. That is,

$$A(Y) \approx \sum_{i=1}^{N_w} \frac{w_i}{N_{s,i}} \sum_{k=1}^{N_{s,i}} \log p\left(Y|\theta^{(i,k)}, \mathcal{M}\right) \tag{5.19}$$

5.3 Duffing oscillator

For the system identification of structures using Bayesian updating, competing model classes can be compared in order to explain the observed response of the structure to excitation.

For instance, when the amplitude of excitation is low, a linear dynamical system with an appropriate stochastic embedding is typically used to explain the behavior of engineering structures, for example in ([5]). However, at larger amplitudes, non-linear geometric and material effects, such as plasticity, may come into play. Were it possible to efficiently evaluate the evidence term, then Bayesian model class selection would be applicable.

The Duffing oscillator is perhaps the most straightforward extension of the equation of motion for the linear oscillator to the non-linear case. This example will be used to demonstrate model class selection by evaluating the evidence term using AIMS-based path sampling.

The governing differential equation for a sinusoidally-driven Duffing oscillator may be taken as:

$$\ddot{x} + \delta\dot{x} + \beta x + \alpha x^3 = \gamma \cos \omega t \tag{5.20}$$

The restoring force therefore not only includes a spring-like displacement-proportional term but also a cubic non-linear term.

Non-linear behavior

The coefficient α in Eqn. (5.20) controls the type and extent of non-linear behavior in the Duffing oscillator. To gain insight into the specific nature of the non-linearity of this governing equation as in [6], we consider a simple case of the undamped oscillator with a potential given by:

$$V(x) = \frac{x^4}{4} - \frac{x^2}{2} \tag{5.21}$$

When subject to zero damping and input excitation, this is a conservative system with a double-well potential that can be described by the state-space differential equation:

$$\begin{Bmatrix} \dot{x} \\ \dot{v} \end{Bmatrix} = \begin{Bmatrix} v \\ x - x^3 \end{Bmatrix} \tag{5.22}$$

There are two stable equilibrium points at unit distance from the origin along the axis of displacement. The equilibrium point at the origin is unstable [7].

In state space, the trajectory of such a particle is closed due to the conservation of energy, as in Figure 5.1. Since the differential equation is non-linear, however, the solution is not sinusoidal. This system corresponds to the undamped free-vibration case with $\alpha = 1$ and $\beta = -1$ in Eqn. (5.20) with the damping and forcing coefficients zero. In this case, while the non-linear term is the restoring force, the linear term is a repelling force.

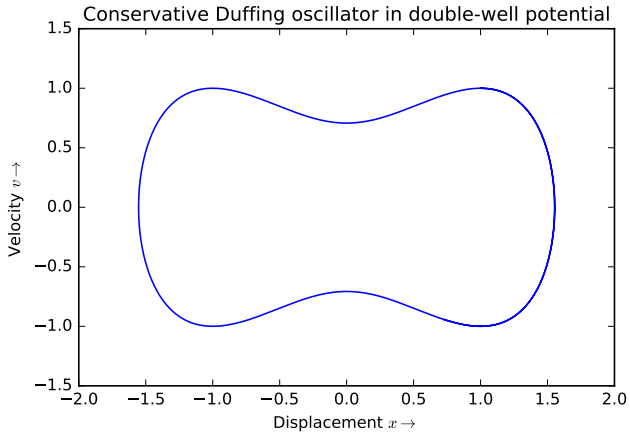

Figure 5.1: The closed state-space trajectory of a double-well Duffing oscillator

When this case is extended to include damping and forcing, the Duffing oscillator can behave in the rather complex fashion that non-linear systems commonly exhibit.

As the amplitude of the forcing input increases, the system can exhibit a period doubling bifurcation, as shown in Figure 5.2.

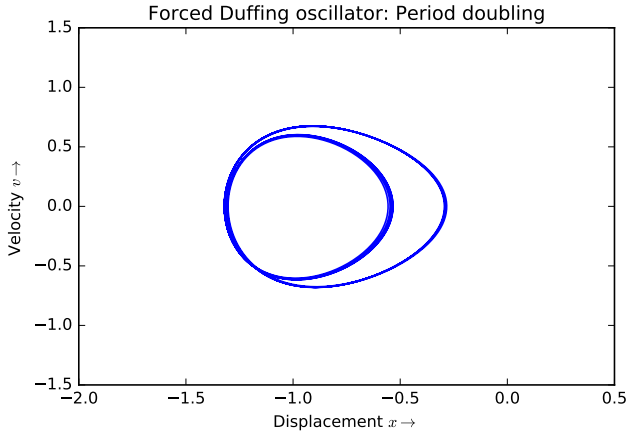

Figure 5.2: Period doubling bifurcation for forced, damped Duffing oscillator

With even further increases in the forcing amplitude, the system exhibits chaos and unpredictably spends its time near one or the other equilibrium point. An example of such a trajectory is shown in Figure 5.3. Note, however, that the figure only shows that a stable limit cycle is not achieved and does not reveal the intricacies of the underlying chaotic behavior, such as sensitivity to initial conditions.

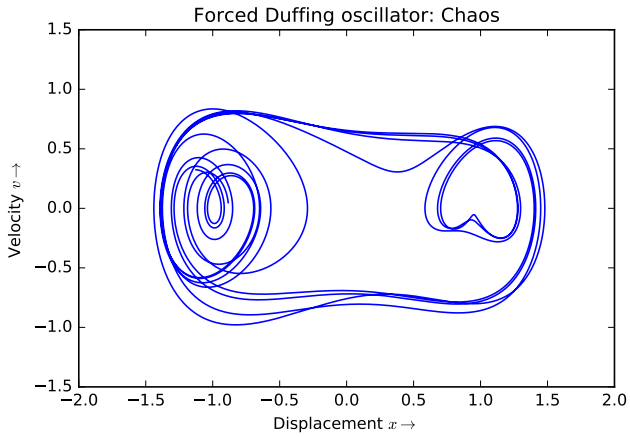

Figure 5.3: Chaotic trajectory for a forced, damped Duffing oscillator

In the case where $\alpha, \beta > 0$, there is no double-well potential. The only equilibrium point is at the origin and it is stable. This may be thought of as the case of a linear oscillator with an added cubic restoring force. When the forcing amplitude is high

relative to the damping, this system also exhibits period doubling bifurcations and eventually chaos.

In this section, a specific stochastic embedding of the Duffing oscillator is considered, according to how Bayesian inference might typically be done in practice. The focus of this work is on using Bayesian model class selection between the linear and Duffing stochastically embedded model classes. The bifurcation and chaotic regime of the Duffing oscillator will not be considered here since the likelihood function would vary quite sensitively with respect to the initial condition parameters. This warrants a separate investigation to verify that the posterior is indeed stable under the minute errors associated with finite-precision arithmetic.

5.4 Numerical example

Here, we consider a specific example. Let there exist two model classes, $\mathcal{M}^L$ and $\mathcal{M}^D$, that correspond to a linear and Duffing oscillator model classes respectively. We parameterize the model class with uncertain parameters $\zeta \in \mathbb{R}^5$ and $\theta \in \mathbb{R}^6$, respectively.

The uncertain parameters considered are: $\theta = [\delta, \beta, \gamma, \omega, \log\sigma, \alpha]$. ζ consists of the same uncertain parameters, except for the non-linear coefficient, α, which is omitted for the linear oscillator, which is equivalent to fixing it to zero.

Synthetic data generation

Denote the synthetic, noisy data generated from the linear and Duffing oscillators by $\mathcal{D}^L$ and $\mathcal{D}^D$ respectively.

A total of 500 data points, spaced equally apart in time at 40 ms intervals, are generated for either case. The values for the data-generating parameters are listed in Table 5.1.

The data at these time-points consists of the displacement prediction $x_n(\theta^*)$ or $x_n(\zeta^*)$ along with zero-mean independent Gaussian noise of standard deviation σ^* added to it, to mirror the imperfections that one might find in real-world data.

Stochastic embedding

Given data from the oscillator, on which Bayesian inference for the Duffing forward model is to be performed, the specification of a stochastic forward model along with an appropriate prior becomes necessary.

Table 5.1: Specification of the data-generating parameters for $\mathcal{M}^L$ and $\mathcal{M}^D$

Component	**Linear Model** (ζ^*)	**Duffing Model** (θ^*)
$\delta*$	0.1	0.1
$\beta*$	1.0	1.0
$\alpha*$	N/A	1.0
$\gamma*$	0.3	0.3
$\omega*$	1.4	1.4
$\log\sigma*$	-4.0	-4.0

For the stochastic forward model, we assume that displacement data are available at discrete points in time and that these measurements are equal to the predictions corrupted by stationary, Gaussian white noise. That is,

$$y_n = x_n(\theta) + \epsilon_n(\sigma), \text{ where } \epsilon_n \sim \mathcal{N}(\cdot|0, \sigma^2) \text{ i. i. d.} \tag{5.23}$$

The conditional probability of observing a certain data set of measurements, $\mathcal{D}_N = [y_n]_{n=1}^N$, is then:

$$p(\mathcal{D}_N|\theta, \sigma\mathcal{M}) = \prod_{n=1}^{N} \mathcal{N}(y_n|x_n(\theta), \sigma^2) \tag{5.24}$$

The prediction is obtained by numerically solving the trajectory of the Duffing oscillator:

$$\ddot{x} = -\theta_1\dot{x} - \theta_2 x - \theta_3 x^3 + \theta_4\cos(\theta_5 t), \text{ with } x(0) = 0 \text{ and } \dot{x}(0) = 0 \tag{5.25}$$

It is often desirable to look at the logarithm of the quantity in Eqn. (5.24) for several reasons: first, this is often the form in which computer programs can work with since the stochastic forward model, being a product of probability densities in this case, can be a very small number; and second, one can see how this stochastic forward model relates to non-linear least-squares estimation That is,

$$\mathcal{L} = \log p(\mathcal{D}_N|\theta, \mathcal{M}) = \frac{N}{2}\log(2\pi\sigma^2) + \frac{1}{2\sigma^2}\sum_{n=1}^{N}(\hat{y}_n - x_n(\theta))^2 \tag{5.26}$$

The additive noise model has been chosen here for its simplicity. More complex ones such as the additive-multiplicative noise model involve more free parameters and are harder to deal with both computationally and analytically. However, they

may be able to model the problem more realistically, since it works with relative values of prediction errors.

The other important component of the stochastic embedding that needs specification is the prior on the uncertain parameters. In engineering problems, the prior mean and variance of the system parameters are typically constrained to be in line with the engineer's judgment, which in turn would be consistent with the design specifications of the system. With these constraints, a prior can be chosen according to the principle of maximum entropy, which simply means that no unwarranted additional entropic certainty should be introduced in the prior other than the desired moments such as the mean and variance. This is to reflect the true state of knowledge about the parameters that does not incorporate unintended additional information that could possibly mislead one into obtaining an overly narrow posterior distribution over the parameters.

In the example considered in this section, synthetic data is generated from a known parameter vector. To reflect the uncertainty over the parameters in a real application of the problem, a prior distribution, $p(\theta|\mathcal{M})$ over the parameters, θ, is specified. The prior distribution over the parameters is chosen to be independent in each component. That is:

$$p(\theta|\mathcal{M}) = \prod_{i=1}^{N_\theta} p(\theta_i|\mathcal{M}) \tag{5.27}$$

The following describes the choice of bounds over the uncertain parameters for the Duffing oscillator. The distribution is the same for the linear oscillator, only the non-linear coefficient, α, is removed from the parameter vector. Table 5.2 details the prior for each component.

Table 5.2: Specification of priors for $\mathcal{M}^L$ and $\mathcal{M}^D$

Component	**Linear Model**	**Duffing Model**
δ	U[0.0, 0.3]	U[0.0, 0.3]
β	U[0.5, 1.5]	U[0.5, 1.5]
γ	U[0.0, 0.5]	U[0.0, 0.5]
ω	U[1.0, 1.5]	U[1.0, 1.5]
$\log \sigma$	U[-5.0, -3.0]	U[-5.0, -3.0]
α	N/A	ln $\mathcal{N}[\mu$=0.0, σ^2=0.1]

All parameters except for the non-linear coefficient, α, were chosen to take the same value range in both models.

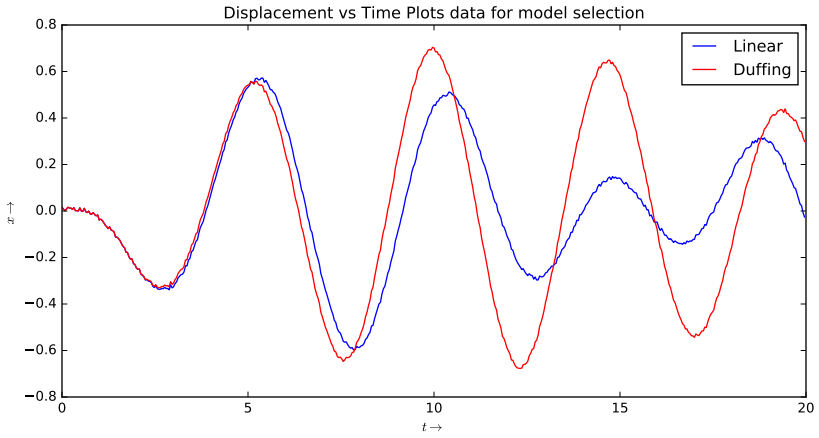

Figure 5.4: Displacement trajectory plots of the data data from a linear oscillator (blue) and a Duffing oscillator (red)

Figure 5.4 shows the plot of the displacement trajectories for the data-generating samples from either model class. Notice that the peaks and troughs for either trajectory begin to mis-align after one oscillation cycle about 7 seconds in. In addition, the amplitude of oscillation of the Duffing oscillator is higher beyond the second cycle.

Prior sensitivity analysis

We would like our estimates of evidence to not be sensitive to perturbations in the prior. In the example of the linear oscillator, the prior is a uniform distribution on a hyperrectangle in a d-dimensional parameter space:

$$\Omega = \prod_{i=1}^{d} [a_i, b_i] \tag{5.28}$$

The expression for the independent uniform prior is therefore:

$$p(\theta|\mathcal{M}) = \frac{1}{|\Omega|}\mathbf{1}_{\Omega} \tag{5.29}$$

Here, $|\Omega|$ denotes the volume of the region, Ω. If we denote the likelihood function by $f(\theta) \stackrel{\Delta}{=} p(\mathcal{D}|\theta, \mathcal{M})$, then we have an expression for the evidence term under the

original, unperturbed prior:

$$p(\mathcal{D}|\mathcal{M}) = \int f(\theta)p(\theta)d\theta \tag{5.30}$$

$$= \frac{1}{|\Omega|}\int_{\Omega} f(\theta)d\theta \tag{5.31}$$

Now, consider a perturbation of the prior, wherein the new prior is again uniform, but with the interval along each dimension extended in both directions by a perturbation factor. That is,

$$\Omega' = \prod_{i=1}^{d}[a_i - \epsilon_i, b_i + \epsilon_i] \tag{5.32}$$

Observe that we can approximate the new region volume as:

$$|\Omega'| \approx |\Omega| + 2d\sum_{i=1}^{d}\epsilon_i\prod_{j\neq i}(b_j - a_j) + O(\|\epsilon\|^2) \tag{5.33}$$

$$\approx |\Omega|\left(1 + 2d\sum_{i=1}^{d}\frac{\epsilon_i}{(b_i - a_i)}\right) \tag{5.34}$$

$$= |\Omega|(1 + c\cdot\epsilon) \tag{5.35}$$

Again, the definition of the density follows from the definition of the region:

$$p_\epsilon(\theta|\mathcal{M}) = \frac{1}{|\Omega'|}\mathbf{1}_{\Omega'} \approx \left(\frac{1 - c\cdot\epsilon}{|\Omega|}\right)\mathbf{1}_{\Omega'} \tag{5.36}$$

We therefore have an expression for the perturbed model evidence:

$$p_\epsilon(\mathcal{D}|\mathcal{M}) = \int f(\theta)p_\epsilon(\theta)d\theta \tag{5.37}$$

$$= \frac{1}{|\Omega'|}\int_{\Omega'} f(\theta)d\theta \tag{5.38}$$

$$\approx \frac{1 - c\cdot\epsilon}{|\Omega|}\int_{\Omega'} f(\theta)d\theta \tag{5.39}$$

We now analyze the integral. We can break it up as:

$$\int_{\Omega'} f(\theta)d\theta = \int_{\Omega} f(\theta)d\theta + \int_{\Omega'\setminus\Omega} f(\theta)d\theta \tag{5.40}$$

Now, let the likelihood function be bounded by M in the region $\mathbb{R}^d \setminus \Omega$. This would be a function of the data and so we may write $M(\mathcal{D})$ instead. Then it will also be

bounded by $M(\mathcal{D})$ in the region $\Omega' \setminus \Omega$. Also, the likelihood function is non-negative since it is a density when viewed as a function of the data. Hence, we have:

$$0 < f(\theta) < M(\mathcal{D}),\ \forall\, \theta \in \Omega' \setminus \Omega \tag{5.41}$$

Thus, we may bound the left hand side of Eqn. (5.40) as:

$$\int_{\Omega'} f(\theta) d\theta < \int_{\Omega} f(\theta) d\theta + M(\mathcal{D}) |\Omega' \setminus \Omega| \tag{5.42}$$

[The assumption that the likelihood function is bounded in a region outside that of the uniform prior is not an unreasonable one, since the stochastic forward model is usually bounded by choice. For instance, a common choice is:

$$p(\mathcal{D}|\theta, \mathcal{M}) = \prod_{n=1}^{N_t} \mathcal{N}(y_n | x_n(\theta), \sigma^2(\theta)) \tag{5.43}$$

So long as no choice of θ results in zero variance in the Gaussian distribution, we can always bound this quantity.]

Note that:

$$|\Omega' \setminus \Omega| \approx 2d \sum_{i=1}^{d} \epsilon_i \prod_{j \neq i} (b_j - a_j) + O(\|\epsilon\|^2) \tag{5.44}$$

$$= 2d|\Omega|(c \cdot \epsilon) + O(\|\epsilon\|^2) \tag{5.45}$$

Substituting this result in Eqn. (5.42), we get:

$$\int_{\Omega'} f(\theta) d\theta < \int_{\Omega} f(\theta) d\theta + 2dM(\mathcal{D})|\Omega|(c \cdot \epsilon) + O(\|\epsilon\|^2) \tag{5.46}$$

This may be substituted in the original expression for the perturbed evidence in Eqn. (5.39) to give:

$$p_\epsilon(\mathcal{D}|\mathcal{M}) < (1 - c \cdot \epsilon) \left[p(\mathcal{D}|\mathcal{M}) + 2dM(\mathcal{D})(c \cdot \epsilon) + O(\|\epsilon\|^2) \right] \tag{5.47}$$

$$= p(\mathcal{D}|\mathcal{M}) + O(\|\epsilon\|) \tag{5.48}$$

Thus, at least in this problem, perturbations to the prior within the same family of priors results in changes of the same order as the perturbation. This is a desirable property to have for a problem, since it implies that the value of the evidence is not brittle to minute changes in the prior ([8]).

5.5 Simulation Results

The posterior samples that AIMS produces are samples of dimension 7 or 8 depending on the model class. There is no good way of visualizing such sample sets. However, we can work with summary statistics as well as histograms of the marginal posterior distributions of the samples for each parameter. It is to be kept in mind, however, that such visualizations are limited and may not be able to convey the intricate correlations between the parameters that are introduced by the data upon performing the Bayesian update.

Figures 5.5 through 5.8 are histograms of the marginal distribution for each parameter dimension of the posterior samples generated by AIMS for the linear and Duffing oscillators. Superposed over each histogram is the data-generating parameter value in dashed blue as well as the maximum *a posteriori* (MAP) sample from AIMS in dotted red.

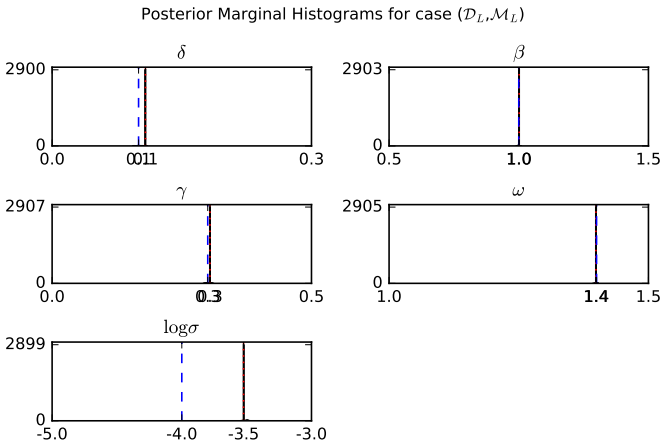

Figure 5.5: Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^L, \mathcal{M}^L)$

Figure 5.5 plots the marginal histograms for the linear oscillator data under the linear model class, that is $(\mathcal{D}^L, \mathcal{M}^L)$. We can see that the distribution is quite peaked and appears to be unimodal. The mode is not centered at the data-generating estimate. This is because in a problem with a moderate number of dimensions such as this, a sampling algorithm may get stuck in a local minimum if it does not generate a candidate near the data-generating sample.

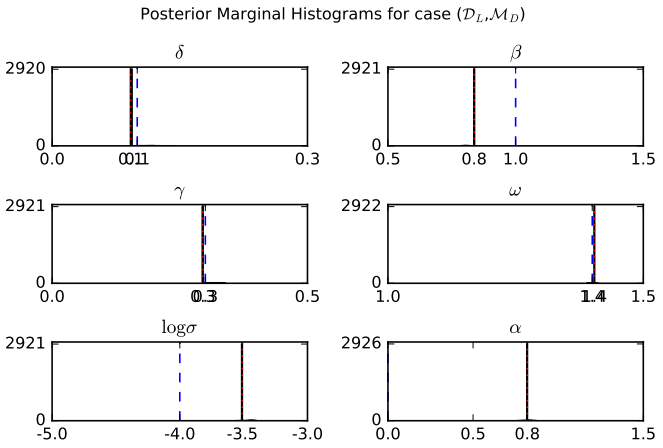

Figure 5.6: Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^L, \mathcal{M}^D)$

In contrast to Figure 5.5, we observe in Figure 5.6 that the marginal histograms for explaining the data from the linear oscillator ($\mathcal{D}^L$) using the Duffing model are not as peaked for some components. Note that in this case, since we presupposed non-linear behavior, the parameter α is not allowed to take the zero value. Because of this, no good data-fit is obtained and there is a correspondingly larger uncertainty associated with some of the parameters.

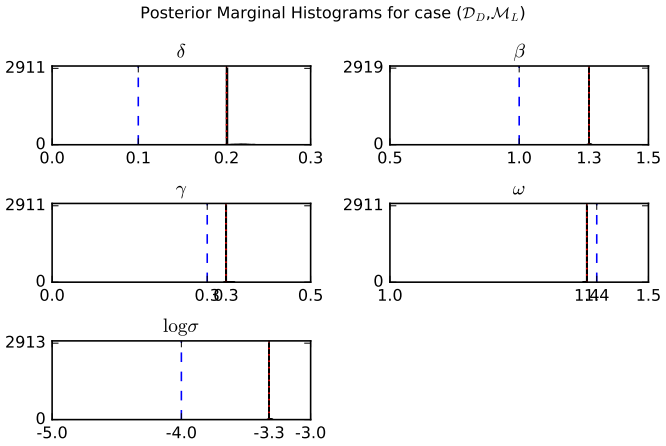

Figure 5.7: Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^D, \mathcal{M}^L)$

Figure 5.7 attempts to fit the data from a Duffing oscillator with a linear model. In this case, we do not anticipate a sample that might result in a very good fit at all data points, since we already know that the non-linear behavior of $\mathcal{D}^D$ cannot be captured by $\mathcal{M}^L$. For this specific 20 second window, however, the linear model is able to identify parameters that provide a good data fit. These do not correspond

to the data-generating parameters, however. The damping, δ and linear stiffness, β, have high marginal posterior probability since they provide a good data-fit.

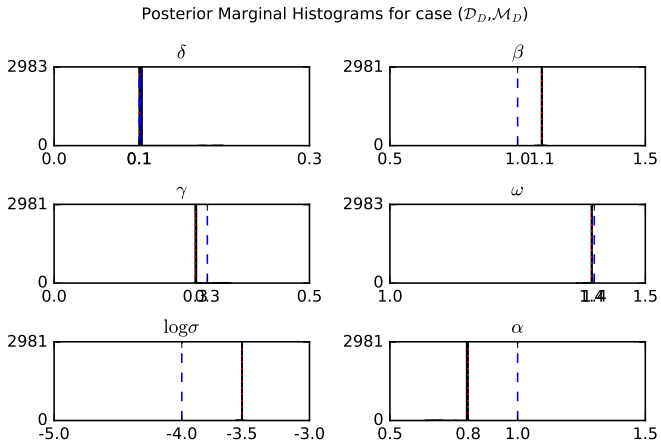

Figure 5.8: Posterior component-wise marginal histograms for $p(\zeta|\mathcal{D}^D, \mathcal{M}^D)$

Finally, Figure 5.8 is the result of trying to perform inference on $\mathcal{D}^D$ with the Duffing model class, $\mathcal{M}^D$. AIMS is able to find the posterior parameters that provide the best data-fit. Note that for this 20-second window, the most probable marginal damping and forcing parameters, δ, γ and ω, are nearly equal to their data-generating values. The stiffness and non-linear coefficients are noticeably different from the data-generating values and appear to have high marginal posterior probability.

The most probable samples from these histograms are now used to simulate the oscillator and the results are compared with the original data that was generated. Figures 5.9 through 5.13 compares the prediction of the maximum *a posteriori* sample with the observed data.

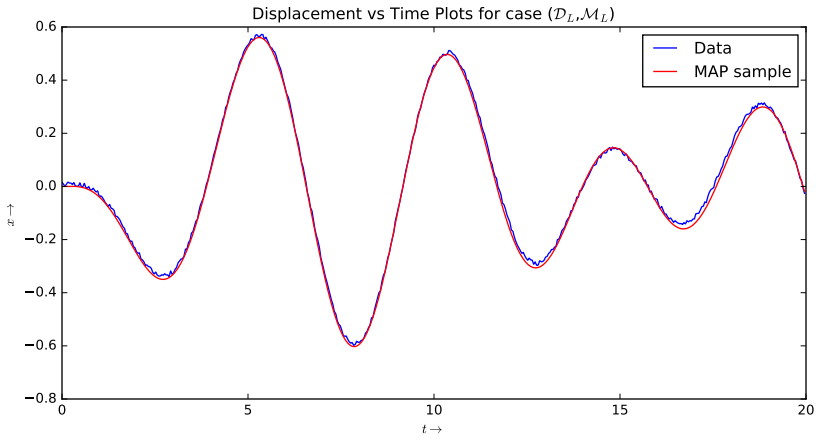

Figure 5.9: Displacement trajectory plots for data from a linear oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a linear oscillator model class

Figure 5.9 is a plot of the linear oscillator data against the displacement prediction for the most probable posterior sample. The prediction is a good fit to the data and values for the system parameters closely match the data-generating values.

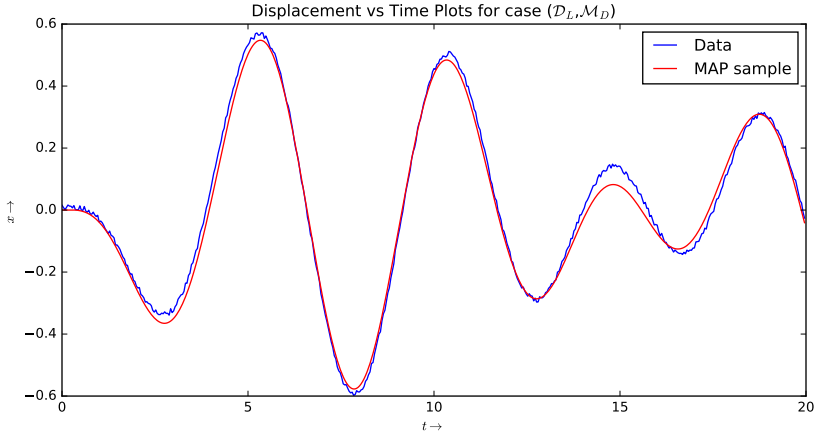

Figure 5.10: Displacement trajectory plots for data from a linear oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a Duffing oscillator model class

Figure 5.10 shows how the Duffing model class has trouble providing a good-data fit. Figure 5.11 is the plot of the prior marginal density for the non-linear parameter, α for the Duffing model class. The log-normal prior on the non-linear term penalizes values of α close to 0, and so the posterior contains low probability mass near the

data-generating parameter. Thus, the maximum a posterior sample contains a value of α that is consistent with the prior, that can provide a good data-fit.

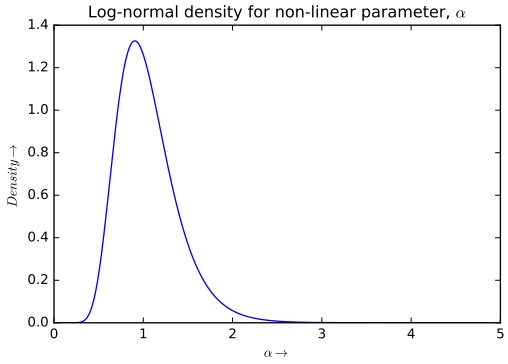

Figure 5.11: Probability density function for the log-normal prior marginal for the non-linear coefficient, α

Figure 5.12 is interesting because for this 20 second window, the posterior is able to find a sample that provides a good data-fit. However, the most probable posterior sample does not match the data-generating sample for the parameters very well.

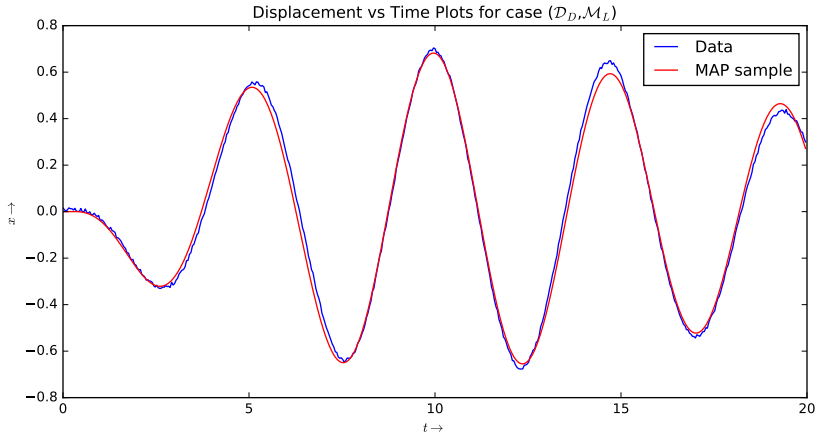

Figure 5.12: Displacement trajectory plots for data from a Duffing oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a linear oscillator model class

Finally, in Figure 5.13, we can see that the most probable sample is able to provide a good data-fit and match the oscillation phase. There appears to be some trade-off between the linear and nonlinear stiffness coefficients, since the most probably sample has noticeably different values for those parameters. This is probably because

within the tolerance of the noise added, these alternative posterior values also have high probability.

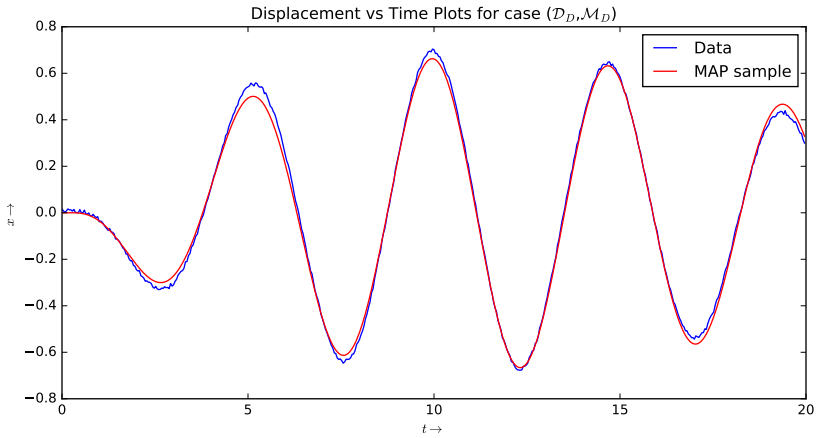

Figure 5.13: Displacement trajectory plots for data from a Duffing oscillator (blue) against the maximum *a posteriori* sample trajectory (red) from a Duffing oscillator model class

Table 5.3: Data generating samples and MAP estimates for each scenario

Quantity	Data-generating values	Linear model (Linear data)	Duffing model (Linear data)	Linear model (Duffing data)	Duffing model (Duffing data)
δ	0.1	0.1080	0.0924	0.2024	0.1000
β	1.0	1.0017	0.8375	1.2702	1.0951
γ	0.3	0.3043	0.2952	0.3362	0.2764
ω	1.4	1.3988	1.4043	1.3808	1.3950
$\log \sigma$	-4.0	-3.5239	-3.5106	-3.3213	-3.5298
α	1.0	N/A	0.8184	N/A	0.8061

Table 5.3 has the data-generating sample used for generating the synthetic data for the model selection example, along with the MAP estimates for the samples.

Now, in order to evaluate the log-evidence using path sampling, AIMS was used to draw samples from distributions intermediate to the prior and posterior along a thermodynamic integration curve parameterized by b (b is used here since β has been used for the stiffness of the oscillator in this section). The following AIMS parameters and schedules were selected:

Table 5.4: AIMS parameters

Parameter	Description	Value	Comment
γ_a	Relative effective sample size	0.9	Highest value greater than the recommended 0.5 and requiring reasonable computer time
c_j	Proposal kernel standard deviation	Varies	Adaptively chosen equal to mean sample nearest-neighbor distance
N_j	Number of samples per annealing stage	10,000	Highest constant permitted by computation time

The following values for the logarithm of the evidence were obtained for each case with different numerical integration schemes:

Table 5.5: Log-evidence values from thermodynamic integration

Measurements	Model	Trapezoid	Gauss	Kronrod
Linear ($\mathcal{D}^L$)	Linear ($\mathcal{M}^L$)	1204.87	1214.37	1206.41
Linear ($\mathcal{D}^L$)	Duffing ($\mathcal{M}^D$)	1001.2	996.356	1002.53
Duffing ($\mathcal{D}^D$)	Linear ($\mathcal{M}^L$)	983.415	994.515	988.207
Duffing ($\mathcal{D}^D$)	Duffing ($\mathcal{M}^D$)	1087.86	1103.93	991.423

Table 5.5 lists the values for the logarithm of the evidence obtained for each case. It should be noted that the values obtained using the Gauss/Kronrod quadrature are not as reliable as those obtained by Trapezoidal integration. This is because the quadrature methods do not give very accurate results for curves with sharp corners near the boundary: here, at $\beta = 0$. This, in turn, is because there is a steep rise in the value of expected log-likelihood in going from the prior stage to the first annealed stage. The thermodynamic integration curve for the linear case is considered in Figure 5.14.

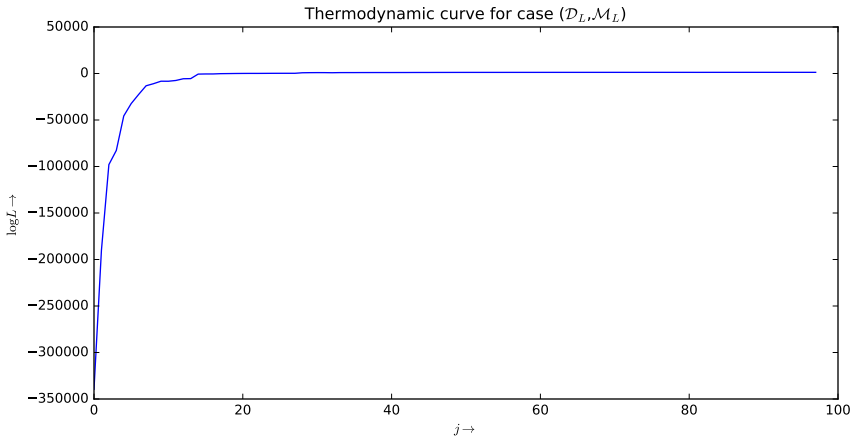

Figure 5.14: Thermodynamic integration curve for the ($\mathcal{D}^L$, $\mathcal{M}^L$) case

We also observe that when the the “correct" model class is employed, the model evidence is clearly much higher compared to the alternative.

5.6 Conclusion

These results show that path sampling with thermodynamic integration can be used to give reasonable estimates for model evidence in the case of a stochastically embedded dynamical model with a moderate number of dimensions. The AIMS algorithm is naturally suited to generating the samples required by the path sampling scheme.

In addition, it was found that when the model class containing the data-generating model is used along with a suitable prior, then the evidence for that model class turns out to be higher than that for competing classes.

References

[1] J. L. Beck, K. V. Yuen, “Model selection using response measurements: Bayesian probabilistic approach,” *Journal of Engineering Mechanics*, vol. 130, no. 2, pp. 192–203, 2004.

[2] B. Calderhead, M. Girolami, “Estimating Bayes factors via thermodynamic integration and population MCMC,” *Computational Statistics & Data Analysis*, vol. 53, no. 12, pp. 4028–4045, Oct. 2009.

[3] A. Gelman, X.-L. Meng, “Simulating normalizing constants: from importance sampling to bridge sampling to path sampling,” *Statistical Science*, vol. 13, no. 2, pp. 163–185, May 1998.

[4] J. L. Beck, K. M. Zuev, "ASYMPTOTICALLY INDEPENDENT MARKOV SAMPLING: A NEW MARKOV CHAIN MONTE CARLO SCHEME FOR BAYESIAN INFERENCE," *International Journal for Uncertainty Quantification*, vol. 3, no. 5, pp. 445–474, 2013.

[5] K.-V. Yuen, J. L. Beck, L. S. Katafygiotis, "Unified Probabilistic Approach for Model Updating and Damage Detection," *Journal of Applied Mechanics*, vol. 73, no. 4, p. 555, 2006.

[6] P. Holmes, D. Rand, "The bifurcations of duffing's equation: An application of catastrophe theory," *Journal of Sound and Vibration*, vol. 44, no. 2, pp. 237–253, Jan. 1976.

[7] C. Holmes, P. Holmes, "Second order averaging and bifurcations to subharmonics in duffing's equation," *Journal of Sound and Vibration*, vol. 78, no. 2, pp. 161–174, Sep. 1981.

[8] H. Owhadi, C. Scovel, T. Sullivan, "On the Brittleness of Bayesian Inference," *SIAM Review*, vol. 57, no. 4, pp. 566–582, Jan. 2015.

Chapter 6

CONCLUSIONS AND FUTURE WORK

The work in this thesis addresses the problem of optimal sensor placement for Bayesian parametric identification of structures. Several technical aspects of the problem were addressed, along with the development of a method of solution for a specific version of the problem.

An efficient method of solution to the log-determinant formulation of the sensor selection problem was presented by relaxing the original Boolean-vector independent optimization variable to one that can take any values on the unit interval. Since the resulting configuration is not Boolean, resolutions to interpretations of possible ambiguous results are presented. It should be noted that the solution is sometimes unambiguous despite the constraints and in this case this is the optimal solution to the original combinatorial problem.

When the list of possible sensor locations is over a continuum or a large number of points, then certain technical challenges in modeling the problem arise. A description of them is presented and possible resolutions to the problem are discussed, in the context set by the example of an elastic axial bar. In the absence of using a specified correlation model, it is best to choose points on a pre-defined grid sufficiently spaced apart for which prediction-error correlations are negligible. This prevents sensor clumping that occurs as the characteristic length scale of the sensors goes to zero.

When the simplifying assumptions that allow the use of the log-determinant problem formulation are violated, then the general entropy-based objective used as a criterion for sensor configuration optimality is the mutual information between the data variable and the uncertain parameters, given the model class. For practitioners, this quantity may be perceived as the expected gain in information entropy in going from the prior to the posterior distribution.

The mutual information in this formulation could not be expressed simply in terms of the sensitivity coefficients of the predictions. In the absence of analytical approaches to compute the objective, it became necessary to compute Monte Carlo estimates of the mutual information. Simple examples were considered and the mutual information estimate was obtained in these cases.

6.1 Future Work

The optimal sensor location problem in the context of Bayesian parametric identification continues to pose several interesting challenges.

The computational intractability of the general mutual information formulation is the most challenging aspect of the problem, in the absence of simplifying model assumptions about identifiability and shape of the posterior. This is because the computation of the entropy-of-evidence is very challenging if samples from the marginal distribution for the data predictions are used to compute a Monte Carlo estimate. Improvements in simulation techniques for approximating log-evidence would greatly ameliorate this problem.

In addition, a case was made to investigate the link between system identifiability and the mutual information of a particular sensor configuration.

When it comes to prediction-error correlations, there are two open problems. First, is properly characterizing prediction-error correlations and determining the length scales at which they occur, without making use of a preliminary sensor deployment. The second problem is to determine a method to compute relevant quantities efficiently, for a specified prediction-error correlation model. Even in the case of the formulation involving small prediction-errors or predictions slowly varying with respect to the parameters, if prediction-error correlations are included, then the objective does not neatly decompose into the log-determinant of the matrix sum of contributions from each point on the sensor location grid. The convex optimization method used, then does not apply directly to the optimal sensor location problem. It is necessary, therefore, to develop a new solution that does not involve a search over all possible configurations.

Appendix A

MATRIX CALCULUS IDENTITIES

A.1 Useful matrix calculus identities and properties

This appendix contains useful matrix and linear algebra identities.

Sum of rank-one matrices

For sample vectors $a_i \in \mathbb{R}^n, i = 1, \ldots, N$, the sum $A = \sum_{i=1}^{N} a_i a_i^T$:

- Is symmetric
- Is positive semi-definite
- Is positive definite if the samples a_i span $\mathbb{R}^n$

Gradient of log-determinant

The expression for the derivative of the log-determinant of a matrix parameterized by a vector is given by,

$$\frac{\partial \log\det Q(z)}{\partial z_i} = \operatorname{tr}\left(Q^{-1}\frac{\partial Q}{\partial z_i}\right) \tag{A.1}$$

Note that in the current problem, we have,

$$\frac{\partial Q(z)}{\partial z_i} = \frac{\partial}{\partial z_i}\left(\sum_{k=1}^{N_d} z_k Q(e^k)\right) \tag{A.2}$$

$$= Q(e^i) \tag{A.3}$$

Trace of matrix product

The trace of the product of two matrices, when they exist as expressed, is given by,

$$\operatorname{tr}(XY^T) = \operatorname{tr}(X^TY) = \operatorname{tr}(Y^TX) = \operatorname{tr}(YX^T) = \sum_{i=1}^{N}\sum_{j=1}^{N} X_{ij}Y_{ij} \tag{A.4}$$

In the current problem, this is applied in the following evaluation:

$$\mathrm{tr}\left(Q^{-1}(z)Q(e^i)\right) = \sum_{j=1}^{N}\sum_{k=1}^{N}\left(Q^{-1}\right)_{jk}\left(Q(e^i)\right)_{kj} \tag{A.5}$$

Gradient of matrix inverse

The derivative of the inverse of a matrix parameterized by a vector is given by,

$$\frac{\partial Q^{-1}(z)}{\partial z_i} = -Q^{-1}\frac{\partial Q}{\partial z_i}Q^{-1} \tag{A.6}$$

Appendix B

DERIVATIVES FOR CONVEX PROBLEM

B.1 Calculation of parametric-gradient for modal response

In order to evaluate the covariance shape matrix, Q, the gradient of the prediction $x_i(t)$ with respect to the parameters θ are required.

Since the modal matrix Φ is independent of the uncertain parameters, the gradient of the prediction may be written solely in terms of the gradient of each mode:

$$\nabla_{\theta_s} x_i(t_n) = \sum_{i=1}^{N_d} \Phi_{ij} \nabla_{\theta_s} q_j(t_n) \tag{B.1}$$

Now, $q_j(t_n)$ is given by Equation 3.49. Its partial derivatives with respect to each component of $\theta_s = [\omega_0, \alpha, \beta, a_0, \omega]$ are calculated analytically.

Breakdown of modal solution

Here, the modal solution is broken down into more elementary terms for ease of differentiation. The terms here are in shorthand and the explicit dependence on the

mode number, j, and time, t, is sometimes dropped. Define:

$$N_{1a} = \frac{\omega^3 + (2\zeta_j^2 - 1)\omega_j^2\omega}{\omega_{d_j}} \tag{B.2}$$

$$N_{1b} = \exp(-\zeta_j\omega_j t) \tag{B.3}$$

$$N_{1c} = \sin(\omega_{d_j} t) \tag{B.4}$$

$$N_{2a} = 2\zeta_j\omega_j\omega \tag{B.5}$$

$$N_{2b} = \exp(-\zeta_j\omega_j t) \tag{B.6}$$

$$N_{2c} = \cos(\omega_{d_j} t) \tag{B.7}$$

$$N_{3a} = (\omega_j^2 - \omega^2) \tag{B.8}$$

$$N_{3b} = \sin(\omega t) \tag{B.9}$$

$$N_{4a} = -2\zeta_j\omega_j\omega \tag{B.10}$$

$$N_{4b} = \cos(\omega t) \tag{B.11}$$

$$N_1 = N_{1a} \cdot N_{1b} \cdot N_{1c} \tag{B.12}$$

$$N_2 = N_{2a} \cdot N_{2b} \cdot N_{2c} \tag{B.13}$$

$$N_3 = N_{3a} \cdot N_{3b} \tag{B.14}$$

$$N_4 = N_{4a} \cdot N_{4b} \tag{B.15}$$

$$D_1 = (\omega_j^2 - \omega^2)^2 \tag{B.16}$$

$$D_2 = (2\zeta_j\omega_j\omega)^2 \tag{B.17}$$

$$N_j = N_1 + N_2 + N_3 + N_4 \tag{B.18}$$

$$D_j = D_1 + D_2 \tag{B.19}$$

$$q_j(t) = a_j \frac{N_j}{D_j} \tag{B.20}$$

The gradient with respect to the uncertain parameters θ_s may now be expressed as,

$$\nabla q_j(t) = q_j(t)\left(\frac{\nabla a_j}{a_j} + \frac{\nabla N_j}{N_j} - \frac{\nabla D_j}{D_j}\right) \tag{B.21}$$

$$\nabla N_j = \nabla N_1 + \nabla N_2 + \nabla N_3 + \nabla N_4 \tag{B.22}$$

$$\nabla D_j = \nabla D_1 + \nabla D_2 \tag{B.23}$$

$$\nabla N_1 = N_1\left(\frac{\nabla N_{1a}}{N_{1a}} + \frac{\nabla N_{1b}}{N_{1b}} + \frac{\nabla N_{1c}}{N_{1c}}\right) \tag{B.24}$$

$$\nabla N_2 = N_2\left(\frac{\nabla N_{2a}}{N_{2a}} + \frac{\nabla N_{2b}}{N_{2b}} + \frac{\nabla N_{2c}}{N_{2c}}\right) \tag{B.25}$$

$$\nabla N_3 = N_3\left(\frac{\nabla N_{3a}}{N_{3a}} + \frac{\nabla N_{3b}}{N_{3b}}\right) \tag{B.26}$$

$$\nabla N_4 = N_4\left(\frac{\nabla N_{4a}}{N_{4a}} + \frac{\nabla N_{4b}}{N_{4b}}\right) \tag{B.27}$$

Low-level derivatives

$$\frac{\partial \omega_j}{\partial \omega_0} = c_j \tag{B.28}$$

$$\frac{\partial \zeta_j}{\partial \omega_0} = \frac{c_j}{2}\left(\beta - \frac{\alpha}{\omega_j^2}\right) \tag{B.29}$$

$$\frac{\partial \zeta_j}{\partial \alpha} = \frac{1}{2\omega_j} \tag{B.30}$$

$$\frac{\partial \zeta_j}{\partial \beta} = \frac{\omega_j}{2} \tag{B.31}$$

$$\frac{\partial \omega_{d_j}}{\partial \omega_0} = c_j\sqrt{1-\zeta_j^2} - \frac{\zeta_j \omega_j}{\sqrt{1-\zeta_j^2}}\frac{\partial \zeta_j}{\partial \omega_0} \tag{B.32}$$

$$\frac{\partial \omega_{d_j}}{\partial \alpha \text{ or } \beta} = -\frac{\omega_j \zeta_j}{\sqrt{1-\zeta_j^2}}\frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \tag{B.33}$$

Derivatives with respect to forcing amplitude

The only relevant term in this case is,

$$\frac{\partial a_j}{\partial a_0} = -\frac{\Phi^T \mathbb{1}}{\mu_j} \tag{B.34}$$

Derivatives with respect to base natural frequency

The relevant terms in this case are those that depend on ω_j and ζ_j:

$$\frac{\partial N_{1a}}{\partial \omega_0} = \left(\frac{2\omega_j c_j \omega (2\zeta_j^2 - 1)}{\omega_{d_j}} + \frac{4\zeta_j \omega_j^2 \omega}{\omega_{d_j}} \frac{\partial \zeta_j}{\partial \omega_0} \right) - \left(\frac{\omega^3 + (2\zeta_j^2 - 1)\omega_j^2 \omega}{\omega_{d_j}^2} \frac{\partial \omega_{d_j}}{\omega_0} \right) \tag{B.35}$$

$$\frac{\partial N_{1b}}{\partial \omega_0} = -t\zeta_j \omega_j N_{1b} \left(\frac{1}{\zeta_j} \frac{\partial \zeta_j}{\partial \omega_0} + \frac{c_j}{\omega_j} \right) \tag{B.36}$$

$$\frac{\partial N_{1c}}{\partial \omega_0} = t \frac{\partial \omega_{d_j}}{\partial \omega_0} \cos(\omega_{d_j} t) \tag{B.37}$$

$$\frac{\partial N_{2a}}{\partial \omega_0} = 2\zeta_j \omega_j \omega \left(\frac{1}{\zeta_j} \frac{\partial \zeta_j}{\partial \omega_0} + \frac{c_j}{\omega_j} \right) \tag{B.38}$$

$$\frac{\partial N_{2b}}{\partial \omega_0} = \frac{\partial N_{1b}}{\partial \omega_0} \tag{B.39}$$

$$\frac{\partial N_{2c}}{\partial \omega_0} = -t \frac{\partial \omega_{d_j}}{\partial \omega_0} \sin(\omega_{d_j} t) \tag{B.40}$$

$$\frac{\partial N_{3a}}{\partial \omega_0} = 2c_j \omega_j \tag{B.41}$$

$$\frac{\partial N_{4a}}{\partial \omega_0} = -\frac{\partial N_{2a}}{\partial \omega_0} \tag{B.42}$$

$$\frac{\partial D_1}{\partial \omega_0} = 4\omega_j c_j (\omega_j^2 - \omega^2) \tag{B.43}$$

$$\frac{\partial D_2}{\partial \omega_0} = 2(2\zeta_j \omega_j \omega)^2 \left(\frac{1}{\zeta_j} \frac{\partial \zeta_j}{\partial \omega_0} + \frac{c_j}{\omega_j} \right) \tag{B.44}$$

Derivatives with respect to Rayleigh damping coefficients

The relevant terms here are those that depend on ζ_j.

$$\frac{\partial N_{1a}}{\partial \alpha \text{ or } \beta} = \left(\frac{4\zeta_j \omega_j^2 \omega}{\omega_{d_j}} \frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \right) - \frac{N_{1a}}{\omega_{d_j}} \frac{\partial \omega_{d_j}}{\partial \alpha \text{ or } \beta} \tag{B.45}$$

$$\frac{\partial N_{1b}}{\partial \alpha \text{ or } \beta} = -t\omega_j N_{1b} \frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \tag{B.46}$$

$$\frac{\partial N_{1c}}{\partial \alpha \text{ or } \beta} = t \cos(\omega_{d_j} t) \frac{\partial \omega_{d_j}}{\partial \alpha \text{ or } \beta} \tag{B.47}$$

$$\frac{\partial N_{2a}}{\partial \alpha \text{ or } \beta} = 2\omega_j \omega \frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \tag{B.48}$$

$$\frac{\partial N_{2b}}{\partial \alpha \text{ or } \beta} = -\omega_j t \exp(-\zeta_j \omega_j t) \frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \tag{B.49}$$

$$\frac{\partial N_{2c}}{\partial \alpha \text{ or } \beta} = -t \sin(\omega_{d_j} t) \frac{\partial \omega_{d_j}}{\partial \alpha \text{ or } \beta} \tag{B.50}$$

$$\frac{\partial N_{4a}}{\partial \alpha \text{ or } \beta} = -\frac{\partial N_{2a}}{\partial \alpha \text{ or } \beta} \tag{B.51}$$

$$\frac{\partial D_2}{\partial \alpha \text{ or } \beta} = 8\zeta_j \omega_j^2 \omega^2 \frac{\partial \zeta_j}{\partial \alpha \text{ or } \beta} \tag{B.52}$$

Derivatives with respect to forcing frequency

The following terms are relevant:

$$\frac{\partial N_{1a}}{\partial \omega} = \frac{3\omega^2 + (2\zeta_j^2 - 1)\omega_j^2}{\omega_{d_j}} \tag{B.53}$$

$$\frac{\partial N_{2a}}{\partial \omega} = 2\zeta_j \omega_j \tag{B.54}$$

$$\frac{\partial N_{3a}}{\partial \omega} = -2\omega \tag{B.55}$$

$$\frac{\partial N_{3b}}{\partial \omega} = t \cos(\omega t) \tag{B.56}$$

$$\frac{\partial N_{4a}}{\partial \omega} = -\frac{\partial N_{2a}}{\partial \omega} \tag{B.57}$$

$$\frac{\partial N_{4b}}{\partial \omega} = -t \sin(\omega t) \tag{B.58}$$

$$\frac{\partial D_1}{\partial \omega} = -4\omega(\omega_j^2 - \omega^2) \tag{B.59}$$

$$\frac{\partial D_2}{\partial \omega} = 8\zeta_j^2 \omega_j^2 \omega \tag{B.60}$$

Equations B.2 - B.60 can be used to compute the gradient of the displacements in Equation B.1.