On the Identifiability of Highly Parameterised Models of Physical Processes

Dhruva Venkita Raman

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2016

Abstract

This thesis is concerned with drawing out high-level insight from otherwise complex mathematical models of physical processes. This is achieved through detailed analysis of model behaviour as constituent parameters are varied. A particular focus is the well-posedness of parameter estimation from noisy data, and its relationship to the parametric sensitivity properties of the model. Other topics investigated include the verification of model performance properties over large ranges of parameters, and the simplification of models based upon their response to parameter perturbation. Several methodologies are proposed, which account for various model classes. However, shared features of the models considered include nonlinearity, parameters with considerable scope for variability, and experimental data corrupted by significant measurement uncertainty.

We begin by considering models described by systems of nonlinear ordinary differential equations with parameter dependence. Model output, in this case, can only be obtained by numerical integration of the relevant equations. Therefore, assessment of model behaviour over tracts of parameter space is usually carried out by repeated model simulation over a grid of parameter values. We instead reformulate this assessment as an algebraic problem, using polynomial programming techniques. The result is an algorithm that produces parameter-dependent algebraic functions that are guaranteed to bound user-defined aspects of model behaviour over parameter space.

We then consider more general classes of parameter-dependent model. A theoretical framework is constructed through which we can explore the duality between model sensitivity to non-local parameter perturbations, and the well-posedness of parameter estimation from significantly noisy data. This results in an algorithm that can uncover functional relations on parameter space over which model output is insensitive and parameters cannot be estimated. The methodology used derives from techniques of nonlinear optimal control. We use this algorithm to simplify benchmark models from the systems biology literature. Specifically, we uncover features such as fast-timescale subsystems and redundant model interactions, together with the sets of parameter values over which the features are valid.

We finally consider parameter estimation in models that are acknowledged to imperfectly describe the modelled process. We show that this invalidates standard statistical theory associated with uncertainty quantification of parameter estimates. Alternative theory that accounts for this situation is then developed, resulting in a

computationally tractable approximation of the covariance of a parameter estimate with respect to noise-induced fluctuation of experimental data.

Acknowledgements

The years I have spent at Oxford have been highly formative on a personal, as well as professional level. I owe this to many people. From those with whom I had a passing conversation that taught me an interesting lesson, to those who have been a constant source of guidance, wisdom and support, I thank you all. In the case of the latter category, I hope I have been able to personally convey my gratitude adequately. Above all I must thank my family (immediate and extended) for, well, everything. I'll try to come home/visit more often now my thesis is finished.

The avenues my doctoral research took ended up being quite nonlinear with some occasionally frustrating dead ends, as is probably natural. Throughout, I have benefited hugely from the support and expertise of my supervisors Antonis Papachristodoulou and James Anderson. Thanks for giving me both the freedom to find my own research topics, and the guidance without which I could not have effectively pursued them. Apologies for all the statistical theory, and for my initially arcane writing style, which I would not have improved without your painstaking corrections. I also must thank Giorgio Valmorbida, my academic uncle, for his many words of wisdom on both research and life. In particular, for instilling a more careful, considered approach to research than that to which I am naturally inclined, and for cooking the best barbeque I've ever had. Heather Harrington has also continually guided and encouraged me. Thanks, it was more important than you might realise.

I should probably thank everyone in the office, not least because my desk is otherwise in danger of sabotage. I've enjoyed the meandering, occasionally baffling conversations in the office, which have helped keep me relaxed and sociable during the write-up period, and before. I've also learnt a lot about control theory from you all, especially those aspects that can't be gleaned from a book. Thanks guys, its been great, and I hope that in future the gender ratio can become finite.

My first year in the Doctoral Training Centre introduced me to a variety of really interesting people. I hope and feel that the friendships I made there, as in the Control group, will be long-lasting. Sorry for not being so sociable recently! And thanks not only for the good times, but also for expanding my horizons from a research point of view. It was really valuable to talk research with people from such diverse backgrounds.

I would not be writing this thesis if not for the excellent undergraduate education I

was lucky enough to receive at the University of Warwick. The rigour of the course influenced my mindset and approach to problems not only in Mathematics, but in life (though thankfully not in all aspects, I hope!). The huge variety of modules available, and the communal atmosphere, were key factors in enthusing me to pursue a doctorate.

Finally, I thank my late violin teacher, Erik Houston. It was largely from you that I learned to be self-critical, which I believe has been a huge positive influence in my academic career and personal life. You'll be glad to know that I restarted violin playing in Oxford. I'm slowly, finally getting a hand on those beautiful Bach Partitas, which I was too immature to fully appreciate and play properly under your tutelage. I know that frustrated you, and I hope you would be pleased with my progress. I'm sorry I never got the chance to say goodbye properly.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 A Primer on Mathematical Modelling

Empiricism lies at the heart of all scientific research conducted today. Hypotheses are formulated, compared, strengthened, and discarded, based on their ability to account for the observed behaviour of the objects to which they pertain. The initial step of collecting relevant data on such observed behaviour is known as the experimental procedure. Some form of data analysis is then required to relate experimental results to existing hypotheses, or to uncover patterns that may lead to the formulation of new hypotheses. Science advances through an integrated cycle of experiment and analysis, with one informing the other.

Data analysis is often quantitative. By this we mean that the data has some form of numerical representation, and patterns in the data can therefore be conveniently exposed and analysed through the language of mathematics. At the simplest level, one might look for generic correlations (or lack thereof) between numerically representable quantities of interest. A famous example of the power of such an approach in informing the scientific method was published in the 19^{th} century by Gregor Mendel,

in his *Versuche über Pflanzen Hybriden* (translated as *Experiments in Plant Hybridisation*) [7]. Mendel considered seven binary characteristics (such as seed shape and colour) of the pea plant (*Pisum Sativum*). He obtained seven pairs of distinct pea populations, such that members of a single population displayed no variation in the considered characteristic, while each pair of populations differed in precisely a single characteristic. Members of the respective populations within a given pair were then cross-bred down two generations (labelled F_1 and F_2 respectively). A table of his results, taken from [118] is provided below:

Table 1.1: Characteristics of Mendel's pea plants, over two generations.

Parental characteristic	$\boldsymbol{F_1}$	$\boldsymbol{F_2}$	$\boldsymbol{F_2}$ **Ratio**
Round × wrinkled seeds	All round	5474 round; 1850 wrinkled	2.96:1
Yellow × green seeds	All yellow	6022 yellow; 2001 green	3.01:1
Purple × white petals	All purple	705 purple; 224 white	3.15:1
Inflated × pinched pods	All inflated	882 inflated; 299 pinched	2.95:1
Green × yellow pods	All green	428 green; 152 yellow	2.82:1
Axial × terminal flowers	All axial	651 axial; 207 terminal	3.14:1
Long × short stems	All long	787 long; 277 short	2.84:1

An observer to Mendel's experiment would have immediately recognised the lack of variation in the characteristics of F_1. Similarly obvious would have been the recovery of variation in F_2. However the clustering of the ratios of the displayed characteristics in F_2 around 3:1 is a quantitative pattern only retrievable by a mathematical analysis. It is this observation that was critical to the formulation of *Mendel's laws of inheritance*. Together, these provide a *model* of inheritance. In other words, they provide a simplified conceptual framework that largely explains the pertinent features of the data, and has some predictive power. Note that both the process of simplification and the choice of pertinent features depend largely on the aims of the modeller. For instance, a model predicting the survivability of a particular pea population would require completely different forms of data, and might not have to take account of the F_2 ratio given in the table above. The model of inheritance proposed by Mendel is

simple enough to be described verbally, with the aid of a diagram (see Figure 1.1).

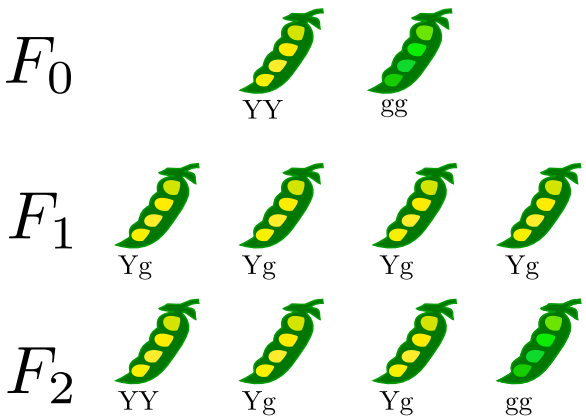

Figure 1.1: Mendel's Laws of Inheritance. *For each binary characteristic (in this case pea colour), there is an associated 'gene'. The gene comes in two varieties ('alleles'): dominant and recessive, denoted Y and g respectively in the figure. Each plant has two copies of the gene, together known as the 'genotype'. If at least one copy is of the dominant allele, then the form of the characteristic ('phenotype') associated with that allele is displayed. In this case it is yellow peas. A descendant of two parent plants inherits one randomly chosen copy of the gene from each parent. F_0 represents the founder population, while F_1 and F_2 represent successive generations. All plants in F_1 inherit a gene from the YY parent, so must be yellow. There is a 50% probability of a plant in F_2 inheriting the g allele from each parent, so one quarter of them have two copies of g and display green peas.*

Nowadays many models of physical processes are sufficiently complex that a verbal description is not convenient. Instead, objects involved in the process are represented numerically, and their interactions are specified through equations. A *mathematical model* of the process is created. This thesis focuses on the analysis of various classes of mathematical models, and consideration of the extent to which conclusions drawn from such a model are applicable to the physical process being modelled. Throughout a reading of this thesis, it is useful to keep in mind several core tenets of the modelling paradigm, which are detailed below.

All models are wrong, but some are useful [18]. Modelling entails a tradeoff between complexity and tractability. As the simplicity of a model increases, so too does the power of the analytical tools we can apply to extract information from it.

Conversely, the fidelity of the model in describing reality may decrease. Our current understanding of genetic inheritance is much more refined than that proposed by Mendel. Indeed, a complete description would require a book, rather than the single diagram of Figure 1.1. However, making predictions on the displayed characteristics of further generations of pea plants using the model of Figure 1.1, is both easy and largely correct. If we accounted for more recently appreciated issues such as the possibility of genetic mutation, or gene conversion during meiosis, this task becomes considerably more involved.

Sometimes, it is possible to decrease the complexity of a model without sacrificing any descriptive power. The simplified model can then be analysed more easily and/or comprehensively. This situation implies some level of redundancy in the original description of the model, which is excised in the simplified model. Chapters 3 and 5 introduce new methods of identifying redundancies of the form described, thus allowing for model simplification and improved analysis.

Models can be used to infer unobserved features of their associated physical process. This inference may not be exact. For Mendel, genes were wholly abstract entities, which existed inasmuch as they provided a convenient way of explaining his experimental data. Only much later was the existence of genes, and the verification of their function, validated experimentally.

Models often contain unobserved quantities and entities. If they are static (i.e. do not vary in time) we refer to them as parameters. The classification of time-varying unobserved quantities is more involved, and considered subsequently in the thesis. The genotypes of the parent, yellow-seeded plants, are static, and therefore parameters.

We can infer parameters of a model, from experimental data. These can directly correspond to unobserved properties of the process. This is known as *parameter estimation*. To take an example, let us temporarily accept the model of Figure 1.1.

In this case the data of Table 1.1 implies that the parent generation of yellow seeded plants are likely not to have possessed any recessive alleles coding for green seeds. Otherwise, some first generation (F_1 plants) would have been likely to inherit two copies of the recessive allele, and display the green-seeded phenotype. Note that our inference is probabilistic: it is possible that some of the parent generation possessed recessive alleles, but that none of these had passed on to the descendants. As soon as randomness enters a model, inference becomes probabilistic. In this case, randomness enters through the choice of allele passed down to the descendant. Another common source of randomness is that of measurement noise. It is safe to assume that Mendel's measurements were correct, as they were of binary characteristics. If he had measured the precise height of the plants, however, instead of classifying them as 'tall', or 'short', then measurement error would have become relevant. This thesis will be concerned with models that are deterministic, but where measurement noise corrupts the process of observation. Thus our methods are not applicable to models whose evolution over time is inherently stochastic, such as those describing Brownian motion.

We have used a model to infer the unobserved genotype of the parent generation of yellow seeded plants, with the caveat that our inference is probabilistic. However, one would reasonably conclude that the probabilistic inference is in fact highly likely. Why? The probability of flipping a (fair) coin 1000 times and never observing a tails is vanishingly small. In the same way, the large number of F_1 plants grown adds strength to the inference. This highlights the importance of experimental replicates in reducing the corrupting influence of randomness on model inference. Unfortunately in mathematical models with many parameters, it often happens that even a tiny amount of measurement noise can corrupt data in such a way that parameter inference becomes extremely imprecise. Chapters 3 and 4 introduce methods of quantifying the uncertainty associated with parameter inference in this situation. Meanwhile Chapter 6 considers the issues arising with quantifying such uncertainty when the model cannot

accurately represent the process for any values of the parameters. This is relevant when conducting analysis on highly simplified models.

Models should, but may not, have predictive power. Correspondingly, parameters should, but may not, be meaningful. We distinguish between two extremes of the modelling paradigm: mechanistic, and data-driven, while acknowledging that many models incorporate aspects of both. Our discussion will focus on mechanistic models, which curate hypotheses. They are built in such a way that each (hypothesised) interaction and entity in the process is represented in the model. In this case, model parameters represent unobserved (static) aspects of the process. Ultimately, the predictive power of such a model corresponds to the validity of the hypotheses. Let us demonstrate with an example of the trolley on a spring depicted in Figure 1.2.

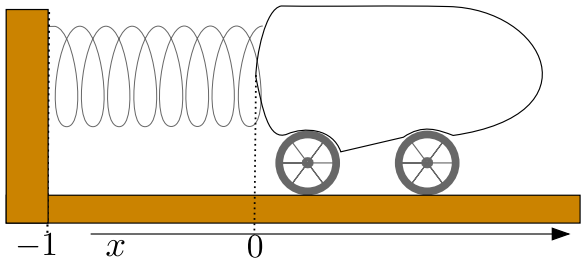

Figure 1.2: Trolley on a spring model. *The variable x denotes the difference between the length of the spring and its natural (unforced) length, at time t. We assume that $|x(0)| + |\dot{x}(0)| < 1$, to avoid collisions with the vertical beam.*

Suppose we wished to analyse the relationship between the starting position of the trolley in Figure 1.2, and its consequent trajectory. A simple, mechanistic approach to the problem could integrate two hypotheses: namely Newton's second law of motion, and Hooke's law of linear elasticity. Newton's second law of motion states that the acceleration of a particle, of constant mass, is directly proportional to the product of the force applied to it with its mass. Mathematically, this can be represented by the

equation

$$m\ddot{x}(t) = F(t),$$

where $x(t)$ is the position of the particle as a function of t, the time, and dotted variables denote time derivatives, with the number of dots equal to the order of the derivative. Meanwhile, m represents the particle's mass, and $F(t)$ the (time-dependent) force applied to it. Hooke's Law of Linear Elasticity proposes that the force exerted by a spring on its ends is linearly proportional to the relative extension of the spring. Mathematically,

$$-kx(t) = F(t),$$

where k is a constant of proportionality representing the stiffness of the spring, and known as the spring constant, and $x(t)$ is extension of the spring from its natural length. Note that Hooke's law is conventionally formulated in terms of the force required to extend a spring, which balances the force exerted by the spring on its ends. Thus the minus sign in the above equation disappears. Applying these laws to Figure 1.2, under the assumption that the only forces exerted on the trolley are due to the effects of the spring, we get the following equation for the motion of the trolley,

$$\ddot{x}(t) = -\frac{k}{m}x(t),$$

which can be integrated to yield

$$x(t) = x(0)\cos\left(\sqrt{\frac{k}{m}}t\right) + \dot{x}(0)\left(\sqrt{\frac{m}{k}}\right)\sin\left(\sqrt{\frac{k}{m}}t\right). \tag{1.1}$$

We can consider the initial position $x(0)$, initial velocity $\dot{x}(0)$, stiffness constant k,

and mass m, as parameters of the model. By evaluating the above equation for different values of the initial position and velocity, we can predict the behaviour of the true spring under similar circumstances. Note that this model is a fairly crude simplification of reality. However, we know that the regimes in which our model predictions fail to reflect reality are precisely those in which our hypotheses, or their underlying assumptions are not valid. This is an artefact of the mechanistic nature of our model building routine. For instance, we have ignored the effect of friction on the trolley's wheels in influencing the trajectory, so we can appreciate that our model will be more accurate in a low-friction setting. Nevertheless if our model predictions reflect experimental testing with reasonable fidelity, then it is likely that our hypotheses have some validity. This validity gives the parameters k and m, which encode information from the hypotheses, an intrinsic meaning. They encode some fundamental property of the objects to which they pertain, which predictably influence their behaviour under those conditions where their underlying hypotheses holds. So, for instance, the concept of mass ceases to be meaningful (through Newton's Laws, at least), when the relative velocity of different system components approaches the speed of light, at which point the model of Newtonian mechanics loses its predictive capabilities.

The sensitivity of model predictions to perturbations in the parameters is intimately related to the accuracy of parameter estimation from data. Let us revisit the model given in (1.1) of a trolley on a spring, shown in Figure 1.2. Note that the dynamics of the position depend only on the ratio of spring stiffness to the mass of the trolley, that is $\frac{k}{m}$. So any perturbations to k and m preserving this ratio have no effect on model output. We say that the parameters k and m are *structurally unidentifiable*, and that the model contains a *structural unidentifiability* over sets of parameters $\{k, m\}$ preserving the ratio $\frac{k}{m}$. Note that the expression $\frac{k}{m}$ is itself structurally identifiable: any change to $\frac{k}{m}$ will induce a change in model behaviour.

Structurally unidentifiable parameters cannot be estimated from data. Suppose that we obtained time-series data for the position of the trolley, and found a pair of values $\{k^*, m^*\}$ for the parameters, which yielded model predictions consistent with the data. We would have no way of discriminating between $\{k^*, m^*\}$ and any other pair of parameter values satisfying $\frac{k}{m} = \frac{k^*}{m^*}$. If we did not know of the structural unidentifiability, then we might have assumed that the values of k^* and m^* were meaningful through their individual values, rather than just their ratio. This may have informed erroneous hypotheses, such as an upper bound on the mass m. Thus it is important to delineate all structural unidentifiabilities in a model before conducting parameter estimation.

We have described how the sensitivity of model predictions to parameter perturbation is zero along structural unidentifiabilities. We have seen how this means that parameter estimation is impossible over them. What happens when a model has very low, but nonzero sensitivity to parameter perturbation along some direction in the space of parameters? In this case, a large set of parameters, extending along this direction, will induce very similar, but technically distinct, model predictions. As stated previously, parameter estimation becomes probabilistic as soon as experimental data is corrupted by measurement noise. Given such data, we cannot remove the measurement noise term. We can, however, see which model predictions would have been likely to generate the data, modulo measurement noise. If the model predictions induced by a particular parameter vector are consistent with noisy data, then it is likely that all 'similar' model predictions are also consistent. Thus it is hard to discriminate between potential parameters along directions of low sensitivity in parameter space. If we attempted to define a region in which it was almost certain the true parameters resided, that region would be very large. Where it is too large for the purposes of the modeller, the model is 'practically unidentifiable' from the data. We will provide a more technical overview of structural/practical unidentifiability and their implica-

tions for the parameter estimation problem in Chapter 4. In general, the delineation of (structural/practical) unidentifiabilities, and the verification of identifiability, are together referred to as the *identifiability problem*.

1.2 Thesis Motivation

The increasing availability of computing power has meant that it has become possible to computationally simulate highly detailed models of complex processes. Accordingly, such models have become increasingly common in fields as diverse as systems biology, climate science, and economics. For instance, a whole-cell model of the life cycle of the pathogen *Mycoplasma genitalium*, which accounts for the integrated function of each of its 525 genes, and contains over 1900 parameters, was recently published [62]. The analysis of such high-dimensional, often nonlinear, models raises many challenges. The contributions of this thesis form part of the extensive literature concerned with their solution.

The identification of (structurally and practically) unidentifiable parameters, and their associated unidentifiabilities, is a hard problem in general. The computational complexity of many available algorithms scales unfavourably with increasing model size, to the extent that they are intractable in models with many parameters. In this case not only is estimation of every parameter from data impossible, but we cannot say which parameters, or functions thereof, can and cannot be estimated. Increasingly scalable algorithms for specifying structurally and practically unidentifiable parameters have been entering the literature in recent years [138, 52, 4, 112]. Less progress has been made in finding the form of their associated unidentifiabilities, which provides additional information. For instance in our previous trolley-spring example, it was useful to know that both k and m were structurally unidentifiable, but it was

even more useful to see that model predictions depended solely on their ratio: $\frac{k}{m}$. Knowledge of such unidentifiabilities paves the way for model simplification, and provides qualitative insight into the characteristics of the process. Furthermore, many parameter estimation routines, such as those based on Markov-chain Monte Carlo, become numerically ill-conditioned in the presence of unidentifiabilities, and fail to provide output [46, 24, 124]. *A priori* specification of the unidentifiabilities avoids these issues. Chapter 5 contributes a new, scalable algorithm for identifying both structurally and practically unidentifiable parameters and their associated unidentifiabilities. We demonstrate through example how this can pave the way for model simplification. For instance, we can identify approximable model features such as timescale-separated subsystems and redundant mechanisms, as well as the regimes in which such potential simplifications are valid.

Once identifiability is verified, interest shifts to precise quantification of the accuracy of a parameter estimate. That is, we want to provide bounds within which each 'true' parameter value is likely to reside, up to some desired degree of confidence. The set of parameter values satisfying such bounds is often referred to as the *uncertainty region* of the parameter estimate, and depends both on the form of measurement noise by which the data is corrupted, and the parametric sensitivity characteristics of the model. The motivation for quantification of the uncertainty region is that the set of conclusions one can derive from a parameter estimate are predicated on the accuracy of the estimate. For example, we could estimate the acceleration due to gravity on various planets by measuring the trajectory of a falling object, and fitting the data to a model derived from Newton's Law of Universal Gravitation. If we performed such an experiment on Venus, and concluded that the best-guess of acceleration was $8.8ms^{-2}$, but that any estimate within two units of this guess was plausible (i.e. we estimated acceleration as $8.8 \pm 2ms^{-2}$), then we could not reasonably verify the hypothesis that gravitional acceleration on Venus was smaller than that on Earth

(where it is $9.81ms^{-2}$, to two decimal places).

Precise specification of the uncertainty region associated with a parameter estimate is a hard task. One approach is to simulate the model over a large number of discrete parameter values, and test whether the model prediction arising from each simulation is within the uncertainty region (see e.g. [111] for a review). This suffers from the *curse of dimensionality*: a linear increase in the number of parameters results in an exponential increase in the number of sample points required to cover the space of parameters at a given density. So massive numbers of simulations are required to analyse models with many parameters. Another approach is to formulate quadratic approximations of the uncertainty region [21]. This allows for computationally efficient characterisation, but becomes less realistic as the signal to noise ratio of the data decreases. In Chapter 4 we provide a new method of quantifying the uncertainty region, which relies neither on simulation nor approximation. In Chapter 6, we investigate the consequences of removing the conventional statistical assumption that the model, for some parameter values, can perfectly recreate the modelled physical process. This assumption is central to much of the literature on uncertainty quantification, but is not always justified. We then provide a new framework for uncertainty quantification that is independent of the assumption.

We have so far focused on challenges related to the estimation of fixed parameter values from experimental data, so as to infer unobserved characteristics of the modelled process. Model parameters can also represent quantities that vary with the environmental conditions experienced by the modelled process. Alternatively, possible discrepancies between model predictions and the true behaviour of the process under consideration are accounted for by inserting variable parameters into the model. In either case, quantitative analysis must account for a range of potential parameter values. We now present a motivating example, illustrated in Figure 1.3.

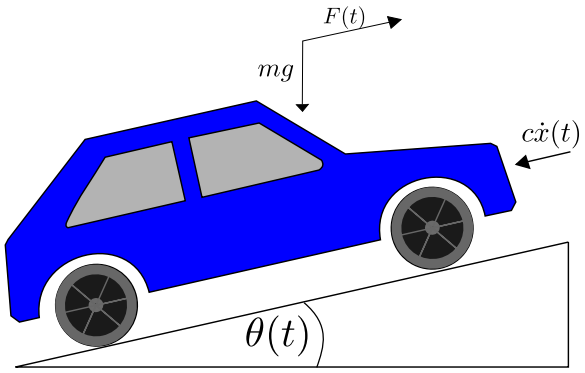

Figure 1.3: Simplified model of cruise control system. *Sensors measure the slope angle in real time, and the cruise control uses this information, along with the current speed, to calculate a force (acceleration or braking) to apply. The calculation is known as a control law, and must obey certain performance criteria over a range of possible values of the drag coefficient c, as this varies dependent on factors such as the number of passengers.*

Consider a cruise controller governing the acceleration and braking of a car. The cruise controller tells the engine to exert a force $F(t)$ in order to maintain a velocity close to some setpoint v^*. Other forces on the car are drag (proportional to velocity, via a coefficient c), and gravity (proportional to the slope angle θ of the road). The equation of motion for the car is therefore $m\ddot{x}(t)+c\dot{x}(t) = F(t)-mg\sin\theta(t)$, where $x(t)$ denotes displacement, m denotes car mass, and g is acceleration due to gravity. The formula for calculating $F(t)$ (known as the *control law*) must obey certain performance criteria. For instance we may want to place an upper bound on the absolute difference between the true velocity $\dot{x}(t)$, and the desired velocity v^*, in order that the driver does not inadvertently break speed limits on e.g. a downhill slope. The control law for $F(t)$ is calculated based on real-time estimation of $\theta(t)$, using sensors. The drag-coefficient c is harder to estimate in this way, despite being variable, due to unmodelled effects such as the number of passengers. An alternative is to provide upper and lower bounds on c, and formulate the control law for calculating $F(t)$ such that all relevant performance criteria are satisfied over the relevant range of possible c values.

For many classes of model, analysis over continuous ranges of parameters is not pos-

sible. Instead, the model is simulated over a discrete set of parameter values covering the relevant range, and analysed individually over each element of the set. So if we wanted to analyse a model incorporating a parameter that potentially varied between 0 and 1, we could fix the parameter at a range of values between 0 and 1, and analyse the model in each instance. Such an analysis possesses potential shortcomings. Important qualitative behaviours of the model may be exhibited only at non-simulated values of the parameters. An example is provided in Figure 1.4, which demonstrates that a simulation based analysis of the structural identifiability of the unidentifiable trolley-spring system shown in Figure 1.2 will always yield an erroneous conclusion. Another problem is that one cannot verify a model property holds *for all* parameters in a set, by demonstrating that it holds *for some* parameters. We therefore see the need for mathematical tools that can analyse wider classes of models over continuous ranges of parameter space, providing guaranteed verification of model-related properties such as performance and identifiability. Chapter 3 introduces new tools of this type. Specifically, methods are presented for quantifying and bounding the discrepancy in behaviour between pairs of nonlinear systems of Ordinary Differential Equations (ODEs), over continuous ranges of parameters.

1.3 Thesis Outline and Contributions

Chapter 2 provides a concise summary of mathematical concepts, notation, and theory relevant to the thesis. We first summarise the statistical reasoning behind parameter estimation from noisy data, and the quantification of the uncertainty surrounding a parameter estimate. We compare Bayesian and Frequentist approaches to parameter estimation, which provide different philosophical meanings to uncertainty quantification, but are nevertheless both amenable to the methods developed in this thesis. The identifiability problem is presented. We then provide a background on convex

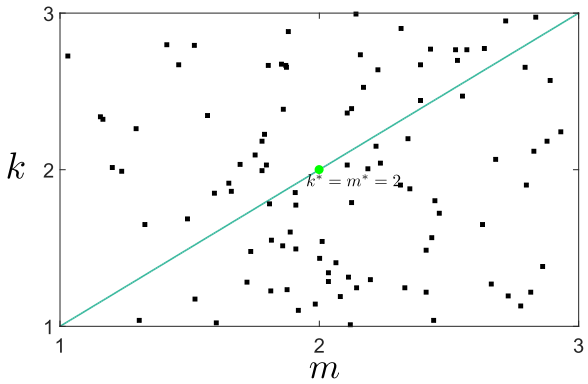

Figure 1.4: Simulation-based analysis of the unidentifiability of the trolley-spring system *We revisit the model depicted in Figure 1.2, modelled by equation* (1.1)*. We fix* $k^* = m^* = 2$ *as nominal parameters. To verify that* k^* *and* m^* *are structurally identifiable within the set of parameters* $k \in [1, 3]$*,* $m \in [1, 3]$*, we simulate the model at a large number of randomly chosen parameter values, drawn from a uniform distribution on the set. At each pair of randomly chosen* $\{k, m\}$*, we check whether the model predictions are identical to those induced by* $\{k^*, m^*\}$*, which would imply unidentifiability. Given a random value of* k*, we require that* $m = \frac{km^*}{k^*}$ *for* $\{k, m\}$ *to lie on the same structural unidentifiability as* $\{k^*, m^*\}$*. But the probability of a uniformly distributed random variable taking an exact value is zero. Graphically, this corresponds to one of the randomly generated points on the right-half figure lying exactly on the line* $\frac{k}{m} = 1$*. So this will never be the case.*

optimisation and the sum-of-squares decomposition. This is necessary to understand the theory and results underlying Chapter 3.

In Chapter 3 we consider polynomially nonlinear systems of ODEs, with polynomial parameter dependence. These are often used in the modelling of biochemical and engineered processes. We then consider the problem of constructing two-sided, algebraic bounds on user-defined measures of the discrepancy between the parameter-dependent model, and some reference trajectory. These bounds are rational functions of the initial conditions and parameters of the model. When constructed, they can provide guarantees of various aspects of model behaviour over continuous ranges of parameter space. For instance, they can verify structural identifiability of the model over a range of parameters. We formulate construction of the bounds as a computationally tractable convex optimisation problem. Theoretical advances on the choice of

objective function used in this optimisation ensure that the bounds are the 'best possible' over the space of possible bounds. Examples of application are then provided. The results of this chapter were published in [102].

Chapter 4 formulates a new quantification of the sensitivity of model predictions to parameter perturbations. We refer to this as *multiscale sloppiness*. This explicitly takes into account the magnitude of parameter perturbation considered, and incorporates existing experimental data, if the latter exists. We show how multiscale sloppiness relates to the geometry of uncertainty regions associated with parameter estimates made from data corrupted by non-infinitesimal measurement noise. In the infinitesimal case, multiscale sloppiness relates to the Fisher Information Matrix, and the existing definition of *model sloppiness*. Using the theoretical framework developed, we show how intuition on the well-posedness and accuracy of parameter estimation depends strongly on the signal-to-noise ratio of data, and existing approximations taken in the limit of increasing ratio can induce erroneous conclusions. The results of this chapter were published in [100].

Chapter 5 combines the formalism of multiscale sloppiness with methods of optimal control to construct a new algorithm for determining the structural and practical unidentifiabilities of a model. This algorithm is applicable to a large number of model classes, but we specifically apply it to benchmark examples from the systems biology literature, which take the form of nonlinear ODEs. Since unidentifiabilities correspond to expressions in parameter space over which model output is invariant, their functional form can identify redundant mechanisms within the model. Indeed, we obtain concrete information related to the biology of the studied processes through deduction of such functional forms. For instance, we isolate hidden, timescale-separated subsystems within a model, identify redundant reaction mechanisms, and ascertain chains of reactions that can be lumped without sacrificing model fidelity. The results

of this chapter were published in [100].

Chapter 6 is concerned with the situation in which there exists a possible discrepancy between model output and the true nature of the modelled process, regardless of parameter values. Much of the statistical theory associated with the uncertainty quantification of parameter estimates relies on this situation not occurring, an assumption which we call the *perfect model assumption*. An important example is the standard method of estimating the covariance of a parameter estimate with respect to the distribution of measurement noise on data, which consists of computing the Fisher Information Matrix, and invoking the famous Cramer-Rao inequality. We construct an alternative estimate of the covariance of a parameter estimate, amenable to a common class of models, that remains valid when the perfect model assumption fails to hold. Nevertheless when it does hold, our estimate agrees with the standard one based on the Fisher Information Matrix, in the limit of increasing data quantity. We show that the sources of error in the two methods of covariance estimation are complementary in some sense, and it is thus desirable to utilise both when flagging an ill-conditioned parameter estimation problem. The results of this chapter were published in [101].

Chapter 7 summarises the contributions of the thesis. Potential future avenues of research are then detailed.

Appendix B solves a specific class of optimal control problem necessary to the results of Chapter 5. Optimal control is concerned with choosing time-varying inputs to a dynamical system that enable it to optimise particular performance criteria subject to constraints. An example would be finding the pattern of acceleration and braking that enables a car to end up stationary at a red light, while minimising the time taken to achieve this.

Appendices C and D provide details of benchmark models of metabolic reaction net-

works from the literature. These are analysed in Chapter 5.

Chapter 2

Preliminaries

In this chapter we provide an account of the mathematical background on which the results of this thesis are predicated. We further introduce a consistent language and framework by which we can both analyse mathematical models and relate them to the physical systems they approximate.

2.1 Notation and Nomenclature

2.1.1 Basic Notation

We begin by outlining notation associated with vector quantities. $\mathbb{R}^n$ denotes the set of real-valued vectors with n entries. The n-superscript is omitted in the case $n = 1$, and analogous conventions are adopted for vectors of integers $\mathbb{Z}^n$, and natural numbers $\mathbb{N}^n$ (which we take to include zero). $\mathbb{R}^+$ is the restriction of $\mathbb{R}$ to the non-negative numbers.

The i^{th} component of an object x with i components will be denoted x_i, as long as it is scalar. If it is a vector quantity, then the subscript is bracketed (i.e. $x_{(i)}$). So if

we have an m-length sequence, denoted $\underline{\mathrm{x}}$, of n-vectors, then the i^{th} sequence element is denoted $\underline{\mathrm{x}}_{(i)} \in \mathbb{R}^n$, while the j^{th} component of $\underline{\mathrm{x}}_{(i)}$ is $\underline{\mathrm{x}}_{(i),j} \in \mathbb{R}$. The l_p-norm of a vector $x \in \mathbb{R}^n$ is defined to be $(\sum_{i=1}^n x_i^p)^{\frac{1}{p}}$, and is denoted by $\|x\|_p$.

We frequently encounter matrices $A \in R^{n \times m}$. The operator $vec : \mathbb{R}^{n \times m} \to \mathbb{R}^{nm \times 1}$ converts such a matrix into a vector by stacking the matrix columns on top of one another, such that $vec(A)_{(j-1)n+i} = A_{ij}$. If A is square, that is $n = m$, we denote the trace of A as $Tr(A)$. If A is also symmetric, then $A \in \mathbb{S}^n$, the space of symmetric matrices. The operator $\otimes$ denotes the standard Kronecker product on matrices, while $I_n \in \mathbb{R}^{n \times n}$ denotes the identity matrix of dimension n.

The Cartesian product of two sets X and Y is given as

$$X \times Y = \{(x, y) : x \in X; y \in Y\}.$$

The boundary of $X \in \mathbb{R}^n$ is denoted ∂X. This represents the set of $x \in \mathbb{R}^n$ such that any open $U \ni x$ has a nonempty intersection with both X and $\mathbb{R}^n / X$.

The space of continuous, scalar-valued functions on a domain X, with k continuous derivatives, is denoted $\mathcal{C}^k(X)$. These derivatives are taken to be partial. We will notationally distinguish between the partial and total derivatives of a function. The partial derivative of a function $f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^o \to \mathbb{R}^p$, at a point (x, y, z), and with respect to a variable $x \in \mathbb{R}^n$, is denoted $\frac{\partial f}{\partial x}(x, y, z)$. Alternatively, we may denote it $\nabla_x f(x, y, z)$. This alternative notation is useful when we wish to describe vectors of partial derivatives. For instance, $\nabla_{x,y} f(x, y, z)$ is taken as the concatenation $\{\nabla_x f(x, y, z), \nabla_y f(x, y, z)\}$. The total derivative of f is denoted $\frac{\mathrm{d}f}{\mathrm{d}x}(x, y, z)$, and accounts for possible dependencies of x on y and z.

The p-norm of a function f, is given as

$$\|f\|_{\mathcal{L}^p} = \left(\int_{\mathbb{R}^n} \sum_i |f_i(x)|^p \, dx \right)^{\frac{1}{p}}.$$

For the case in which $p = \infty$, we take

$$\|f\|_{\mathcal{L}^\infty} = \sup_{x \in \mathbb{R}^n} f(x).$$

We will denote by $\mathcal{L}^p$ the space of functions whose p-norm is defined, and finite.

We now define some standard operators. The inner product of two elements a, b in a Hilbert space $\mathcal{H}$ is denoted $\langle a, b \rangle$. If $\mathcal{H} = \mathbb{R}^n$, then $\langle a, b \rangle = \sum_{i=1}^n a_i b_i$. If $\mathcal{H} = \mathbb{R}^{n \times m}$, then $\langle A, B \rangle = Tr(A^T B)$. If $\mathcal{H} = \mathcal{L}^2$, then the inner product is given as

$$\langle f, g \rangle_{\mathcal{L}^2} = \int_{\mathbb{R}^n} \sum_i^n f_i(x) g_i(x) \, dx,$$

and we note that $\langle f, f \rangle = \|f\|_{\mathcal{L}^2}^2$

Landau notation is used to bound the behaviour of functions as their arguments reach asymptotic limits. Specifically,

$$f(x) = \mathcal{O}\big(g(x)\big) \text{ as } x \to a$$

means that there exist $M > 0$ and $\delta > 0$ such that

$$\|f(x)\|_2 \leq M \|g(x)\|_2, \quad \forall \ \|x - a\|_2 \leq \delta.$$

Note that the choice of l_2-norm in the above equation is arbitrary: all norms are equivalent in $\mathbb{R}^n$ and, more generally, finite-dimensional vector spaces, [117].

2.1.2 Models, Parameters, Data, and Measurement Noise

This thesis is concerned with mathematical models containing several unknown parameters. They model a real-world phenomenon, which we refer to as a **process**. The collection of data on a particular process is known as an **experiment**. We only deal with deterministic models of processes, but we frequently assume that experiments are corrupted by some form of **measurement noise**, which takes the form of a random variable. Thus we distinguish between the deterministic **process model** (which we refer to as simply the model), and the stochastic **experimental model**, which models the data, rather than just the process, by accounting for measurement noise.

Let us take q as the number of parameters contained within a model. The full set of q parameters is listed in a **parameter vector**, which is denoted θ. Typically, each parameter is represented as a real number, so we have that $\theta \in \mathbb{R}^q$. However, there may be additional, hard constraints on the allowable values that parameters can take. For instance if a parameter represents the mass of an object, it cannot be negative. We therefore define $\Theta \subseteq \mathbb{R}^q$, as the **parameter space**. Θ represents the set of allowable parameter vectors, given the relevant hard constraints.

Model output is assumed to be real-valued, and will be denoted $\mathcal{Y}(\theta)$. Output can either be *finite* or *infinite* dimensional. In this thesis, the latter case generally occurs only when model output is a time-varying function. For instance, we frequently encounter models described by systems of parameterised ordinary differential equations (ODEs):

Example 2.1. ***Standard ordinary differential equation model***

We consider a system of equations of the form:

$$\dot{x}(t) = f(x(t), \theta), \qquad y(t) = h(x(t), \theta) \tag{2.1}$$

$$x(0) = x_0 \qquad \theta \in \Theta. \tag{2.2}$$

*Here $y(t) \in \mathbb{R}^m$ is known as the output vector of the ODE system. This is a continuous time representation of that aspect of the process measured by the experimental procedure. $x(t)$, meanwhile, is known as the **state** of the system at time t. Its meaning lies in the fact that it completely determines the future behaviour of the system, for a fixed parameter θ. Given a state $x(t)$, we can solve for a future state $x(t + \delta t)$ by integrating the left-hand side (LHS) of (2.1), over time. Mathematically,*

$$x(t + \delta t) = \int_t^{t+\delta t} \dot{x}(t)\ dt = \int_t^{t+\delta t} f(x(t), \theta)\ dt.$$

*Note that we assume existence and uniqueness of this integral, for any times $t, \delta t \in \mathbb{R}^+$ and parameters $\theta \in \Theta$. Sufficient conditions on such an assumption, through e.g. the Picard-Lindelof Theorem, are given in [20]. Future output $y(t + \delta t)$ can then be obtained by evaluating the output map $h(x(t + \delta t), \theta)$. This implies that we can completely determine the behaviour of an ODE of the form (2.1) by fixing a parameter θ, and an initial state x_0. Consequently we define the **state trajectory** of an ODE, denoted $x(t; \theta, x_0)$, as the state vector attained at time t, given a parameter vector θ, and initial state x_0. Analogously, we have the **output trajectory** $y(t; \theta, x_0)$. Note that x_0 is often referred to as the **initial condition** of the ODE. As a static quantity that affects model predictions, it can be considered as part of the parameter vector. The notational delineation between θ and x_0 stems from the fact that initial conditions satisfy particular properties not shared with more general parameters, and these properties are exploited later in the thesis.*

ODE systems of the form (2.1) are often used to model physical processes. We can take the model output $\mathcal{Y}(\theta, x_0)$ to be the output trajectory $y(t; \theta, x_0)$. In this case, model output is infinite dimensional, with $\mathcal{Y}(\theta, x_0)$ defined as a function from $\mathbb{R}^+$ (indexing time), to $\mathbb{R}^m$ (representing the output trajectory $y(t; \theta, x_0)$ at the inputted time). Alternatively we may define model output to be the evaluation of the output trajectory at a finite vector of timepoints $T = \{t_1, \ldots, t_F\}$. This would be justified if the experimental procedure generated time-series data on the output trajectory $y(t; \theta, x_0)$. In this case, we would have finite dimensional model output $\mathcal{Y}(\theta; x_0) \in \mathbb{R}^m \times \mathbb{R}^F$, associating with each timepoint t_i an output vector $y(t_i; \theta, x_0)$.

Frequently, data on a process is corrupted by some form of measurement noise. In such a case, data would not exactly correspond to the output of the process model. This motivates the experimental model, which is then a random variable predicated on (parameter-dependent) process model output, and denoted $\mathcal{D}(\theta)$. The probability density function associated with this random variable is denoted $\psi_{\mathcal{Y}(\theta)}$. So if we have finite-dimentional data, then for any region $\Omega \subseteq \mathbb{R}^m$:

$$\mathbb{P}[\mathcal{D}(\theta) \in \Omega] = \int_\Omega \psi_{\mathcal{Y}(\theta)}(x)\ dx.$$

If actual experimental data exists, then we denote it by $\mathcal{D}$. So if the experimental model perfectly described both the process and the form of experimental measurement noise, then $\mathcal{D}$ would be a realisation of the random variable $\mathcal{D}(\theta)$.

2.2 Parameter Estimation and Statistics

This section provides theory necessary to understand the statistical basis of parameter estimation from data. We place particular emphasis on quantifying the uncertainty surrounding a parameter estimate.

2.2.1 Probability Basics

Here we provide a basic overview of the mathematical quantification of stochastic, real-world events. For an in-depth treatment, see e.g.[136].

Definition 2.1. ***Probability Spaces***

A probability space is a triple $\{\Omega, \mathbb{F}, \mathbb{P}\}$*, consisting of a sample space, event space, and probability measure, respectively. The sample space* Ω *consists of the set of possible outcomes given one performance of the random process. Elements of the event space* $\mathbb{F}$*, known as events, are formed from the outcomes. In particular,* $\mathbb{F}$ *must have a* σ*-algebra structure on* Ω *(see [136] for more details). So given an event* $f_1 \in \mathbb{F}$*,* Ω / f_1 *is also in* $\mathbb{F}$*. Given an additional event* $f_2 \in \mathbb{F}$*, the union* $f_1 \cup f_2$ *is in* $\mathbb{F}$*. By assumption, every single outcome* $\omega \in \Omega$ *is an event, as is the set of all possible outcomes:* Ω*. Finally the probability measure* $\mathbb{P}$ *is a function whose domain is* $\mathbb{F}$*, and whose range is the unit interval* $[0, 1]$*. So it assigns a probablity to every event.* $\mathbb{P}$ *must satisfy certain regularity assumptions listed in [136]. Importantly we have that* $\mathbb{P}(\Omega) = 1$*, i.e. the probability that some outcome happens is one.*

It is important to note that there may exist events that are strict subsets of Ω, but nevertheless have probability measure one. If we take such an event $f \in \mathbb{F}$, with $\mathbb{P}(f) = 1$, then we refer to f as an **almost-sure** (a.s.) event. An implication is that $\mathbb{P}(\Omega / f) = 0$.

We next define random variables, which assign values, or vectors of values, to each possible event. We shall restrict our attention to the case in which such values are real, although they can more generally be part of any measurable space [136].

Definition 2.2. ***Random Variables***

A real-valued ***random variable*** X *on a probability space* $\{\Omega, \mathbb{F}, \mathbb{P}\}$ *is a function*

$X : \mathbb{F} \to \mathbb{R}^n$, *such that the pre-image under* X *of any measurable set on* $\mathbb{R}^n$ *is an event.*

Although a random variable X can only be defined with respect to an underlying probability space $\{\Omega, \mathbb{F}, \mathbb{P}\}$, we will often omit any explicit specification of the latter, and merely refer to a random variable $X \in \mathbb{R}^n$.

Example 2.2. ***Rolling a die***

Suppose we are interested in the outcome of rolling a six-sided die. The sample space Ω *is then the set* $\{1, 2, 3, 4, 5, 6\}$*, since these are the possible outcomes of the roll. An example of an event* $f \in \mathbb{F}$ *is* $f = \{1, 2\}$*. So* f *defines the event in which the die rolls a* 1 *or a* 2*: it is the union of the outcomes* 1 *and* 2*. If the die is fair, then the probability associated with any individual outcome is* $\frac{1}{6}$*. In this case,* $\mathbb{P}(f) = \frac{1}{3}$*. We can define a random variable* $X : \mathbb{F} \to \mathbb{R}$ *so that* $X(f) = 1$*, and* $X(g) = 0$ *for all* $g \in \mathbb{F}$ *such that* $g \neq f$*. So* X *is nonzero if and only if the die rolls a* 1 *or a* 2*.* X *is non-negative for any event, so the pre-image of* $(-\infty, 0)$ *under* X *is the empty event.*

Definition 2.3. ***Cumulative distribution functions***

Given a random variable $X \in \mathbb{R}^n$*, we define the* ***cumulative distribution function*** *as*

$$F_X(x) := \mathbb{P}[\{\omega \in \Omega : X(\omega) \leq x\}].$$

A random variable is defined as **continuous** if the cumulative distribution function is continuous in $\mathbb{R}^n$. Note a key implication of continuity: $\mathbb{P}[X = a] = 0$, for any $a \in \mathbb{R}^n$.

Definition 2.4. ***Probability density functions and independence of random variables***

The probability density of a continuous random variable $X \in \mathbb{R}^n$ *is a function* $f_X(x)$:

$\mathbb{R}^n \to \mathbb{R}$ *satisfying, for any* $A \subseteq \mathbb{R}^n$*:*

$$\mathbb{P}[X \in A] := \int_A f_X(x)\ dx.$$

Further conditions on the existence of a probability density are given in [136]. We only deal with random variables that possess probability densities in this thesis. For such random variables, the ***support*** *is defined as the closure of the set of vectors* $x \in \mathbb{R}^n$ *for which* $f_X(x) > 0$*. So it is the smallest closed set in* $\mathbb{R}^n$ *containing all the values that* X *can take with positive probability. This is denoted* $\mathrm{supp}(x)$*.*

Given an additional continuous random variable $Y \in \mathbb{R}^m$*, we can also define the* ***joint probability density*** *of* X *and* Y *as a function* $f_{XY}(x, y) : \mathbb{R}^{n+m} \to \mathbb{R}$ *satisfying, for* $A \in \mathbb{R}^n$ *and* $B \in \mathbb{R}^m$*,*

$$\mathbb{P}[X \in A; Y \in B] := \int_{A \times B} f_{XY}(x, y)\ dx\ dy.$$

If, for all A *and* B *such that* $\mathbb{P}(A) > 0$ *and* $\mathbb{P}(B) > 0$*, we have that*

$$\int_{A \times B} f_{XY}(x, y)\ dx\ dy = \left(\int_A f_X(x)\ dx \right) \left(\int_B f_Y(y)\ dy \right),$$

then X *and* Y *are known as* ***independent*** *random variables. The outcome of the experiment determining* X *does not have any bearing on that determining* Y*, and vice versa. If, furthermore,* $f_X(x) = f_Y(y)$ *almost surely (i.e. on a set of probability one), then* X *and* Y *are known as* ***independent, identically distributed*** *(i.i.d) random variables.*

Example 2.3. ***Uniformly distributed random variables***

Let us consider some compact interval $[a, b] \in \mathbb{R}$*. A random variable* $X \in \mathbb{R}^n$ *is uniformly distributed over* $[a, b]$ *if it has the probability density* $f_X(x) = \frac{1}{b-a}$ *for* $x \in$

$[a, b]$, *and* $f_X(x) = 0$ *otherwise. The cumulative distribution is then*

$$F_X(x) = \frac{x-a}{b-a}, \quad x > a$$
$$F_X(a) = 0.$$

We see this as

$$\frac{x-a}{b-a} = \int_a^x \frac{1}{b-a}\, dx,$$

for $x \in [a, b]$.

Definition 2.5. ***Convergence of random variables***

We consider an infinite sequence $\{X^k\}_{k=1}^{\infty}$ *of random variables, with associated cumulative distributions* F_{X^k}. *The sequence* ***converges in distribution*** *to some other random variable* X *if*

$$\lim_{k\to\infty} F_{X^k}(x) = F_X(x)$$

at all points where $F_X(x)$ *is continuous. We use the notation*

$$X^k \to_D X.$$

A stronger property is ***convergence in probability****. This holds if,* $\forall \epsilon > 0$,

$$\lim_{k\to\infty} \mathbb{P}[\|X^k - X\|_2 > \epsilon] = 0.$$

Convergence in probability implies convergence in distribution, but the converse is not true.

Definition 2.6. ***Expected value of random variables***

We consider a random variable $X \in \mathbb{R}^n$ on a probability space $\{\Omega, \mathbb{F}, \mathbb{P}\}$. Given any (measurable) function $h : \mathbb{R}^n \to \mathbb{R}^n$, we define the expected value of $h(X)$ as

$$\mathbb{E}[h(X)] := \int_{\mathbb{R}^n} h(x) f_X(x)\ dx.$$

Note that the expected value of X itself is obtained by taking h as the identity function. The meaning behind the expected value comes from the following argument: If we run repeated trials of the experiment, and denote the value of $h(X)$ on the i^{th} run by $h(X^i)$, then, with probability one, we have that:

$$\lim_{k\to\infty} \frac{1}{k} \sum_{i=1}^{k} h(X^i) = \mathbb{E}[h(X)].$$

In other words, the arithmetic mean of repeated trials of a function of a random variable eventually converges towards the expected value. This is known as the law of large numbers, and explains why the expected value is often alternatively known as the **mean**. Another important property of $\mathbb{E}$ is that it is a linear functional. So $\mathbb{E}[aX+bY] = a\mathbb{E}[X]+b\mathbb{E}[Y]$, for constants $a, b \in \mathbb{R}$ and an additional random variable $Y \in \mathbb{R}^n$.

Definition 2.7. ***Covariance of random variables***

The covariance of two scalar random variables $X, Y \in \mathbb{R}$ is given by

$$\mathrm{Cov}[X, Y] := \int_{\mathbb{R}^n} xy f_{XY}(x, y)\ dx\ dy - \left(\int_{\mathbb{R}^n} x f_X(x)\ dx\right)\left(\int_{\mathbb{R}^n} y f_Y(y)\ dy\right).$$

*The **covariance matrix** Σ of two vector-valued random variables $X, Y \in \mathbb{R}^n$ is an $n \times n$ matrix, such that $\Sigma_{ij} = \mathrm{Cov}[X_i, Y_j]$.*

*The covariance of a random variable with itself, i.e. $\mathrm{Cov}[X, X]$, is termed the **variance** of X.*

Note that covariance can be reformulated as

$$\mathrm{Cov}[X, Y] = \mathbb{E}\big[X - \mathbb{E}[X]\big]\mathbb{E}\big[Y - \mathbb{E}[Y]\big].$$

Thus it is a description of the tendency of random variables to deviate from their expected value in a correlated manner. When $\mathrm{Cov}[X, Y]$ is strictly positive and we know that X is larger than its expected value, then we would expect the same to be true of Y.

Example 2.4. ***Multivariate Gaussian random variables***

A random variable $X \in \mathbb{R}^n$ is (multivariate) Gaussian, with expected value $\mu \in \mathbb{R}^n$ and covariance $\Sigma \in \mathbb{R}^{n\times n}$, if its associated probability density is expressed as

$$f_X(x) = (2\pi)^{\frac{-n}{2}} |\Sigma|^{\frac{-1}{2}} e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)},$$

where $|\Sigma|$ denotes the determinant of Σ. We write $X \sim \mathcal{N}(\mu, \Sigma)$.

Gaussian random variables occur frequently throughout statistics due to the ***Central Limit Theorem*** *[136]. This concerns the properties of any sequence X^i of i.i.d. random variables, with expected value μ, and variance Σ. Specifically, if we take $S_k = \frac{1}{k}\sum_{i=1}^{k} X^i$, then*

$$S_k \to_D X,$$

where $X \sim \mathcal{N}(\mu, \Sigma)$ as before.

2.2.2 Parameter Estimation

We consider a situation in which there exists both experimental data $\mathcal{D}$, and a parameterised process model with output $\mathcal{Y}(\theta)$. Parameter estimation consists of finding

the set of parameter vectors that *optimally* describe the process, given data. We now provide the statistical framework necessary to formulate this problem. All definitions and theory provided are taken from [25], although the notation differs. Example 2.6 is unique to the thesis.

We commence by considering the well-posedness of parameter estimation. If two separate parameter vectors induce the same model output, then they clearly cannot be discriminated on the basis of experimental data. This motivates the concept of structural identifiability:

Definition 2.8. ***Structural Identifiability*** *Consider a model with output $\mathcal{Y}(\theta)$, together with a choice θ^* of parameter vector. If, for all $\theta \in \Theta$, we have that*

$$\mathcal{Y}(\theta^*) \neq \mathcal{Y}(\theta),$$

then the parameter vector is ***globally structurally identifiable****. If a parameter vector is not globally structurally identifiable, it may still be* ***locally structurally identifiable****. We define this to be the situation in which*

$$\exists \epsilon > 0 \text{ such that } \|\theta - \theta^*\|_2^2 < \epsilon \Rightarrow \mathcal{Y}(\hat{\theta}) \neq \mathcal{Y}(\theta^*).$$

If all possible parameter vectors are (locally/globally) structurally identifiable, then we say that the model itself is (locally/globally) structurally identifiable. If a parameter vector does not satisfy local structural unidentifiability, then we say it is ***unidentifiable****. In particular, suppose that we can construct a continuous function $u_{\theta^*}(\theta) : \Theta \to \mathbb{R}$ such that*

$$u_{\theta^*}(\theta) = 0 \Rightarrow \mathcal{Y}(\theta) = \mathcal{Y}(\theta^*). \tag{2.3}$$

In this case, we say that the model has a ***structural unidentifiability*** *along $u_{\theta^*}(\theta)$.*

A more refined classification of specific types of structural unidentifiability is provided in [38], but is unnecessary for the purposes of this thesis. To gain insight into the concept of structural identifiability, let us recall the trolley-on-a-spring model illustrated in Figure 1.2, whose equations of motion were given by

$$x(t) = x(0)\cos\left(\sqrt{\frac{k}{m}}t\right) + \sqrt{\frac{m}{k}}\dot{x}(0)\sin\left(\sqrt{\frac{k}{m}}t\right),$$

with parameters $\theta = \{k, m\}$ denoting the spring stiffness and trolley mass respectively. Suppose that model output $\mathcal{Y}(\theta)$ corresponds to observations of $x(t)$ at an arbitrary number and choice of timepoints. We see that the model is unidentifiable, as any measure of model variation over parameter space will be null over cross-sections of parameter space preserving the ratio $\frac{k}{m}$. A structural unidentifiability of the form $u_{\theta^*}(\theta) = \frac{k}{m} - \frac{k^*}{m^*}$ therefore exists. While this means that the individual values of any estimates of k and m made from experimental data on the trolley motion are meaningless, their ratio is a globally structurally identifiable parameter combination, and thus meaningful.

We also briefly consider an example of a model that is locally, but not globally, structurally identifiable. Suppose that model output is given as

$$\mathcal{Y}(\theta) = [\exp(\theta_1 + \theta_2), \exp\big(2(\theta_1 + \theta_2)\big), \exp\big(3(\theta_1 + \theta_2)\big)],$$

with $\theta_1 \neq \theta_2$. This corresponds to observation of a mixed exponential decay at three separate timepoints. Note that if we swapped the values of θ_1 and θ_2, output would not be changed, ruling out global structural identifiability. However, any other local perturbation to the parameters is guaranteed to disrupt model output, implying local structural identifiability.

Structural identifiability verification is a hard problem for many classes of model,

such as nonlinear systems of ordinary differential equations. So too is the detection of structural unidentifiabilities. A review of the literature in which these problems are considered is found in Chapter 5.

Accurate parameter estimation may not even be possible for structurally identifiable models. We previously explained how we would not necessarily expect experimental data to perfectly correspond to the output of a deterministic process model, due to the confounding effects of measurement noise. In this thesis, hypothetical data is therefore taken as a random variable, whose probability density relies both on the output of the parameterised process model, and measurement noise. We recall that this random variable is denoted $\mathcal{D}(\theta)$ and is called the experimental model. The probability density associated with $\mathcal{D}(\theta)$ is denoted $\psi_{\mathcal{Y}(\theta)}$. Data space $\mathfrak{D}$ is taken as the set of possible data realisations. Mathematically, this can be described as $\mathfrak{D} = \bigcup_{\theta \in \Theta} \operatorname{supp} \mathcal{D}(\theta)$.

Experimental replicates are commonly taken to mitigate the effects of measurement noise and identify anomalies. **E-replicated data** is denoted $\mathcal{D}^{[E]}$. This represents the concatenation of E datasets on the same experiment. Formally,

$$\mathcal{D}^{[E]} = \{\mathcal{D}^1, \mathcal{D}^2, \ldots, \mathcal{D}^E\},$$

where $\mathcal{D}^i$ denotes the i^{th} replicate of experimental data. Correspondingly, the **E-replicated experimental model** is denoted $\mathcal{D}^{[E]}(\theta)$. Since each replicate must be independent and identically distributed, $\mathcal{D}^{[E]}(\theta)$ is the Cartesian product of E copies of the random variable $\mathcal{D}(\theta)$. We often remove the term 'E-replicated' from our description of $\mathcal{D}^{[E]}$ and $\mathcal{D}^{[E]}(\theta)$ for conciseness, as the distinction between these terms

and $\mathcal{D}$ and $\mathcal{D}(\theta)$ will be clear from the context. The **mean data** is taken as

$$\bar{\mathcal{D}}^{[E]}(\theta) := \frac{1}{E}\sum_{i=1}^{E}\mathcal{D}^i(\theta).$$

We assume throughout that experimental replicates are independent, in that the results of one have no bearing on the other. This allows us to represent the probability density of the E-replicated data as

$$\psi^{[E]}_{\mathcal{Y}(\theta)} := \prod_{i=1}^{E}\psi_{\mathcal{Y}(\theta)}(\mathcal{D}^i),$$

using the definition of independence provided in Definition 2.4.

Definition 2.9. ***Likelihood functions***

We define the likelihood as

$$\mathcal{L}(\theta;\mathcal{D}^{[E]}) := .\psi^{[E]}_{\mathcal{Y}(\theta)}(\mathcal{D}^{[E]}).$$

A ***maximum likelihood estimate*** *(MLE) with respect to $\mathcal{D}$ is defined as*

$$\hat{\theta}(\mathcal{D}^{[E]}) = \max_{\theta\in\Theta}\mathcal{L}(\theta|\mathcal{D}^{[E]}).$$

Note that, due to the monotonicity of the logarithm function, an equivalent expression for the MLE is given by

$$\hat{\theta}(\mathcal{D}^{[E]}) = \min_{\theta\in\Theta} -\log\mathcal{L}(\theta|\mathcal{D}^{[E]}).$$

We see that the likelihood of a parameter vector, given data, is the same as the probability density of the data, given an experimental model $\mathcal{D}^{[E]}(\theta)$. The intuition is that, as the likelihood of a parameter vector increases, the form of the data $\mathcal{D}^{[E]}$

becomes more plausible as a realisation of $\mathcal{D}^{[E]}(\theta)$.

Much of the statistical theory associated with maximum likelihood estimation rests on two common assumptions, which are provided below.

1. **Consistency:** For some **optimal parameter** $\theta^* \in \Theta$, we have that

$$\mathbb{P}\left[\lim_{E\to\infty} \hat{\theta}(\mathcal{D}^{[E]}) = \theta^*\right] = 1 \tag{2.4a}$$

2. **Asymptotic normality:** For a consistent estimator with optimal parameter θ^*, the following holds:

$$\sqrt{E}(\hat{\theta}(\mathcal{D}^{[E]}) - \theta^*) \to_D \mathcal{N}(0, \Sigma), \tag{2.4b}$$

   where Σ is some covariance matrix.

Sufficient conditions for consistency to hold are that

1. $\psi_{\mathcal{Y}(\theta')} = \psi_{\mathcal{Y}(\theta)} \Leftrightarrow \theta = \theta'$.
2. $\psi_{\mathcal{Y}(\theta)}$ is continuous on a set of probability one.
3. Θ is a compact set.
4. There exists some bounding function $\mathcal{U}(x)$ such that $|\log\left(\psi_{\mathcal{Y}(\theta)}\right)| < \mathcal{U}(x)$ for all $\theta \in \Theta$.

Given a consistent estimator, additional sufficient conditions for asymptotic normality to hold are that

1. The support of $\mathcal{D}(\theta)$ is independent of θ. This means that data space $\mathfrak{D}$ is equal to $\operatorname{supp}\mathcal{D}(\theta)$, for all $\theta \in \Theta$.
2. The optimal parameter θ^* must not be on the boundary of Θ.

Proofs of the sufficient conditions listed are given in [25]. We assume that they hold for models considered in the thesis unless otherwise stated. An exception is the condition $\psi_{\mathcal{Y}(\theta')} = \psi_{\mathcal{Y}(\theta)} \Leftrightarrow \theta = \theta'$, which will often be violated in the thesis, and is related to the subsequently discussed topic of model identifiability.

Example 2.5. ***Additive Gaussian measurement noise***

Suppose we assume that observations of a process are corrupted by additive Gaussian measurement noise, with mean μ and covariance Σ. Then

$$\mathcal{D}(\theta) \sim \mathcal{N}(\mathcal{Y}(\theta), \Sigma).$$

Note that the support of any Gaussian random variable occupies $\mathbb{R}^n$, which therefore corresponds to data space. The E-replicated experimental model is, by definition, given as the Cartesian product of E experimental models.Meanwhile the mean data has distribution

$$\bar{\mathcal{D}}^{[E]}(\theta) \sim \mathcal{N}\left(\mathcal{Y}(\theta), \frac{\Sigma}{E}\right).$$

The covariance of the mean data decreases as more experimental replicates are added. Thus adding experimental replicates can be used as a proxy for reducing the effect of measurement noise.

Now let us consider the likelihood function, given data $\mathcal{D}^E$. This can be calculated as

$$\mathcal{L}(\theta; \mathcal{D}^{[E]}) = E\langle \bar{\mathcal{D}}^{[E]} - \mathcal{Y}(\theta), \Sigma^{-1}[\bar{\mathcal{D}}^{[E]} - \mathcal{Y}(\theta)]\rangle.$$

For this example, we see that the likelihood is a function solely of the mean data. In fact it takes the form of a weighted sum of the squared residuals between the output of the process model, and the data.

Interpretation of the likelihood function follows two distinct philosophical avenues. The **Frequentist** approach assumes the existence of an **optimal parameter** vector θ^*, such that $\mathcal{Y}(\theta^*)$ corresponds exactly to the process. An MLE $\hat{\theta}$ is an estimate of the optimal parameter vector, which may not be correct due to the effect of measurement noise. It is a *hypothesis* that assumes both the existence of an optimal parameter vector, and its values. The goal of the Frequentist is then to obtain a **confidence region** for θ^*, along with the MLE. If we ran the experiment (including replicates) a high number of times, each time calculating an n-percent confidence region, then the confidence region would contain the true (unchanging) optimal parameter n-percent of the time. Mathematically, the region must be constructed such that

$$\mathbb{P}\left(\theta \in \mathrm{Conf}_n\big(\mathcal{D}^{[E]}(\theta)\big)\right) = \frac{n}{100}, \quad \forall \theta \in \Theta.$$

The **Bayesian** approach does not assume the existence of a fixed optimal parameter vector θ^*, which induces randomly distributed data $\mathcal{D}^{[E]}$, each with its own confidence region. Instead, data is taken as the fixed quantity, and the parameter vector used to generate the data is regarded as a random variable, with a probability density known as the **posterior density**. One must assume the existence of, and provide, a **prior density** on parameter space, which is independent of the data and describes the inherent likelihood of different parameter vectors. So areas of parameter space with a higher prior density are taken as inherently more plausible, regardless of the data. The posterior density is given by **Bayes' rule**, namely

$$\mathbb{P}_{post}(\theta; \mathcal{D}^{[E]}) \propto \mathbb{P}_{prior}(\theta) \prod_{i=1}^{E} \psi_{\mathcal{Y}(\theta)}(\mathcal{D}^i). \tag{2.5}$$

The goal of the Bayesian is then to obtain a **credible region** for θ^*. Given the fixed data $\mathcal{D}^{[E]}$, if we drew a large number of samples from the posterior density, then

n-percent of them would be within the n-percent credible region. Mathematically,

$$\frac{\int_{\mathrm{Cred}_n(\mathcal{D}^{[E]})} \mathbb{P}_{post}(\theta; \mathcal{D}^{[E]})\ d\theta}{\int_{\Theta} \mathbb{P}_{post}(\theta; \mathcal{D}^{[E]})\ d\theta} = \frac{n}{100}. \tag{2.6}$$

Note that the denominator of (2.6) should be one, as it is the integral of a probability density over its support. However, it is often hard to find the appropriate constant of proportionality for (2.5) that thus normalises $\mathbb{P}_{post}$. If $\mathbb{P}_{post}$ is not normalised (and thus technically not a probability density), (2.6) still holds.

Example 2.6. ***Coin tossing***

There are two types of coin: A and B. They are both biased, as shown in the table below

Table 2.1: Coin-tossing example

	Coin A	Coin B
Percentage probability of heads	91	11
Percentage probability of tails	9	89

Suppose that we are given an (unlabelled coin), and toss it. Our parameter inference problem consists of estimating whether the coin is of type A or type B, given the result of the toss. Note that the likelihood function is given as

$$\mathcal{L}(heads|A) = 0.91 \qquad \mathcal{L}(tails|A) = 0.09$$

$$\mathcal{L}(heads|B) = 0.11 \qquad \mathcal{L}(tails|B) = 0.89.$$

The Frequentist strategy would take the confidence region as $\{A\}$ if we draw heads, and $\{A, B\}$ if we draw tails: If the true coin-type (i.e. optimal parameter) was A, then over 90% of our tosses would draw heads, and our confidence region would contain the true coin-type over 90% of the time. If the true coin type was B, under 90% of our tosses would draw tails, so our confidence region would have to include more than

B to include the correct coin type over 90% of the time.

The Bayesian strategy to creating 90% credible regions would differ. We first have to assume a prior probability to the chances of receiving coin type A or B. Let us assume that the probability of receiving A is 50%, although this may not be justified. Suppose we draw heads. We then must ask what the probability of the coin being type A is, assuming that we draw heads, and that the prior probability of having type A is 50%. This probability is $\frac{91}{91+11}$, which is less than 90%. Therefore our 90% credible region must encompass both A and B. Now suppose that we draw tails. Coin type B would have been the type drawing tails in $\frac{89}{9+89}$ percent of the cases. This is over 90%, so our credible region contains only B. In summary, our strategy is to take the credible region as $\{A, B\}$ is we draw heads, and $\{B\}$ if we draw tails, in direct opposition to the Frequentist case.

Neither credible nor confidence regions are unique in a given inference problem. Intuitively speaking however, a smaller region is better, as one would want to rule out as large a portion of parameter space as possible. In Chapter 4 we will highlight well-established notions of optimality, for such regions. We show that, for the problems considered in the chapter, optimal confidence and credible regions in fact have equivalent expressions.

2.3 Sum-of-squares Optimisation and Semidefinite Programming

2.3.1 Optimisation Basics

A (continuous) **optimisation problem** is a mathematical problem of the form

$$\underset{x\in\mathbb{R}^n}{\text{Minimise}}\ f(x) \tag{2.7a}$$

$$\text{subject to } g_i(x) \leq 0; \quad \forall i \in \{1,\dots,m\}, \tag{2.7b}$$

$$h_i(x) = 0; \quad \forall i \in \{1,\dots,l\}, \tag{2.7c}$$

where $f : \mathbb{R}^n \to R$ is known as the **objective function**, g_i are the **inequality constraints**, and h_i are the **equality constraints**. Any vector $x \in \mathbb{R}^n$ satisfying the constraints is a **feasible point**, and the feasible points together make up the **feasible set**. This thesis is solely concerned with the case in which the functions f, g_i, and h_i are continous.

A solution x^* of the problem (2.7) is known as an **optimal point** or **global minimum**, and denoted x^*. The set of optimal points is known as the **optimal set**. All global minima are also **local minima**, although the converse may not be true.

Definition 2.10. ***Local minimum***

A point $\hat{x} \in \mathbb{R}^n$ is a local minimum of the optimisation problem (2.7) *if there exists some $\epsilon > 0$ such that for all $x \in \mathbb{R}^n$ satisfying the constraints* (2.7b) *and* (2.7c)*, as well as the condition $\|x - \hat{x}\|_2^2 < \epsilon$, we have that $f(x) \geq f(\hat{x})$. It is furthermore a **strict** local minimum if instead we have that $f(x) > f(\hat{x})$.*

Intuitively, a local minimum minimises the objective function among close-by feasible points. The **Karush-Kuhn-Tucker** (KKT) conditions provide necessary conditions

for local minimality of a feasible point, for an optimisation problem of the form (2.7), satisfying regularity conditions given in [19]. We omit regularity conditions for the general form of (2.7), but subsequently provide them for specific classes of optimisation problems.

Lemma 2.1. ***Karush-Kuhn-Tucker conditions [19]***

If some feasible $x^* \in \mathbb{R}^n$ is a local minimum of a problem (2.7)*, satisfying regularity conditions provided in [19], then the following conditions must hold:*

1. ***Stationarity****. There exist* ***Lagrange multipliers*** *$\lambda_i \in \mathbb{R}$, and $\mu_i \in \mathbb{R}^+$, such that*

$$-\frac{\mathrm{d}}{\mathrm{d}x} f(x) = \sum_i^m \mu_i \frac{\mathrm{d}}{\mathrm{d}x} g_i(x) + \sum_i^l \lambda_i \frac{\mathrm{d}}{\mathrm{d}x} h_i(x) = 0.$$

2. ***Complementary slackness***

$$\mu_i g_i(x) = 0 \ \forall i \in \{1, \ldots, m\}.$$

Alternatively, we can define a function L as

$$L(x, \mu, \lambda) = f(x) + \sum_i^m \mu_i g_i(x) + \sum_i^l \lambda_i h_i(x).$$

The stationarity requirement is then equivalent to demanding that

$$\frac{\mathrm{d}}{\mathrm{d}x} L(x, \mu, \lambda) = 0.$$

2.3.2 Convexity, and Convex Optimisation

A large class of real-world problems can be cast as optimisation problems. Various algorithmic approaches then exist for computational solution of the problems. However, the most general optimisation problems are computationally intensive or intractable even when involving a small number of variables. One of the major research themes in optimisation is the discovery and exploitation of structure in the objective and constraint functions, in order to allow for the implementation of more efficient numerical solution schemes. A commonly exploited such structure is convexity. A **Convex optimisation problem** (or convex program) is a problem of the form (2.7), where the objective and constraint functions are all convex, and the equality constraints are additionally affine (these properties are defined subsequently). Every local minimum of a convex optimisation is a global minimum. This is an important property: the verification of global minimality for a local minimum is an extremely hard problem in the most general cases of (2.7). Furthermore, highly efficient computational schemes for the solution of convex programs exist [19]. The numerical solutions provided in this thesis all rely on interior point methods. This class of algorithm was originally developed to solve linear programs [61], a special class of convex program in which the objective and constrants are affine functions. However, advances presented in [82] extended the domain of applicability of these methods to general convex programs. We will not delve further into the details of currently-used algorithms in optimisation, but the interested reader could refer to [11].

Definition 2.11. ***Convex sets and functions***

A ***convex set*** $C \in \mathbb{R}^n$ *is defined as any set satisfying the following: for any* $v_1, \ldots v_k \in C$*, we have*

$$\sum_{i=1}^{k} \lambda_i v_i \in C, \text{ where } \lambda_1, \ldots, \lambda_k \in \mathbb{R}^+, \quad \sum_{i=1}^{k} \lambda_i = 1.$$

*If the above property additionally holds for $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$ (i.e. where λ_i are not required to be non-negative), then we refer to C as an **affine set**.*

*A **convex function** $f : C \to \mathbb{R}$, where C is a convex set, satisfies the following:*

$$f\big(tx + (1-t)y\big) \leq tf(x) + (1-t)f(y),\ \forall x, y \in C; t \in [0,1].$$

*If we can furthermore replace the above inequality with equality, the function is additionally referred to as **affine**.*

Program 2.1. ***Convex optimisation problem***

A convex optimisation problem (or convex program) takes the form

$$\underset{x \in \mathbb{R}^n}{\text{Minimise}}\ f(x) \tag{2.8a}$$

$$\text{subject to } g_i(x) \leq 0; \quad \forall i \in \{1, \ldots, m\}, \tag{2.8b}$$

$$h_i(x) = 0; \quad \forall i \in \{1, \ldots, l\}, \tag{2.8c}$$

*where $f(x)$ and $\{g_i(x)\}$ are convex functions, and $\{h_i(x)\}$ are affine functions. This is known as the **primal problem**. We can further define the **dual problem** as*

$$\underset{\mu, \lambda}{\text{Maximise}}\ \inf_{x \in \mathbb{R}^n} \left(f(x) + \sum_{i=1}^{m} \mu_i g_i(x) + \sum_{i=1}^{l} \lambda_i h_i(x) \right) \tag{2.9a}$$

$$\text{subject to } \mu_i \in \mathbb{R}^+; \qquad \forall i \in \{1, \ldots, m\}, \tag{2.9b}$$

$$\lambda_i \in \mathbb{R}\ ; \qquad \forall i \in \{1, \ldots, l\}. \tag{2.9c}$$

Note an important property of the dual problem: if we take any point in the feasible set of the primal problem (2.8), then the summation terms in (2.9a) cannot be strictly positive, regardless of the value of μ and λ. Therefore the value of the objective at an optimal point x^*, in the primal problem, will be bounded from below by the value

of the objective at an optimal point $\{\mu^*, \lambda^*\}$, in the dual problem. This result, which assumes feasibility of both the primal and dual, is known as **weak duality**, and does not exploit convexity. A stronger result, which applies specifically to convex programs, is given below:

Lemma 2.2. ***Slater condition***

Consider a convex program of the form (2.1)*. For the primal problem, suppose there exists some feasible point x satisfying*

$$g_i(x) < 0; \quad \forall i \in \{1, \ldots, m\}.$$

Then the dual is also feasible. Moreover, if we take p^* as the value of the primal objective at an optimal point, and d^* as the value of the dual objective at an optimal point, we have

1. *$p^* = d^*$ (**strong duality**).*
2. *The KKT conditions of Lemma 2.1 hold, without any additional regularity conditions.*

We now present a specific convex set: the space of positive-semidefinite matrices.

Definition 2.12. ***Positive-semidefinite matrices***

Let us some $M \in \mathbb{S}^n$, the space of symmetric matrices. Then the following conditions are equivalent

1. *$z^T M z \geq 0 \quad \forall\ z \in \mathbb{R}^n$.*
2. *All eigenvalues of M (which must be real by symmetricity of M) are non-negative.*

*Any matrix M satisfying the above is referred to as **positive-semidefinite**. If the matrix is furthermore invertible, then it is **positive-definite**. We represent these*

properties mathematically as $M \succeq 0$ *and* $M \succ 0$ *respectively.*

The space of PSD matrices is convex. If $A \succeq 0$ and $B \succeq 0$, then any interpolation of the form $tA + (1-t)B$ is also PSD, for $t \in [0, 1]$. This motivates the formulation of a particular type of convex program, known as a **semidefinite program**.

Program 2.2. ***General semidefinite program (SDP)***

A semidefinite program is a specific instance of Program 2.1, which takes the form

$$\underset{X \in \mathbb{S}^n}{\textit{Minimise}} \ \langle C, X \rangle \tag{2.10a}$$

$$\textit{subject to} \ \langle A_i, X \rangle \leq b_i; \ \forall i \in \{1, \ldots, m\} \tag{2.10b}$$

$$X \succeq 0, \tag{2.10c}$$

where $C \in \mathbb{S}^n$, $A_i \in \mathbb{R}^n$, *and* $b_i \in \mathbb{R}$. *Note that the affine equality constraints of Program 2.1 are replaced by affine inequality constraints of* (2.10b). *These are equivalent, as* (2.10b) *is akin to demanding that* $\langle A_i, X \rangle + \lambda_i = b_i$, *for some undetermined* $\lambda_i > 0$ *known as a slack variable.*

The SDP of (2.10) *is in primal form. The dual SDP is given as*

$$\begin{aligned} \underset{y \in \mathbb{R}^m}{\textit{maximise}} \quad & \langle b, y \rangle \\ \textit{subject to} \quad & C - \sum_{i=1}^{m} y_i A_i \succeq 0. \end{aligned} \tag{2.11}$$

Many mathematical problems, including several in this thesis, can be formulated as or approximated by SDPs. Numerical solution, in this thesis, is always carried out using the freely available SeDuMi [115] solver, although several others exist.

2.3.3 Sum-of-squares Polynomials and Semidefinite Programming

We now illustrate how questions concerning the existence of non-negative polynomials can be framed through SDPs. In this section, given vectors $x \in \mathbb{R}^n$ and $v \in \mathbb{N}^n$, we take $x^v = \prod_i^n x_i^{v_i}$. We restrict our attention to polynomials with coefficients and indeterminates that are real.

Definition 2.13. ***The algebraic structure of polynomials***

We define a ***monomial*** *in the indeterminate $x \in \mathbb{R}^n$, as an expression of the form x^v, where $v \in \mathbb{N}^n$. We take the* ***degree*** *of a monomial as $\sum_i^n v_i$.*

A ***polynomial*** *is the sum of a finite number of monomials. Formally, the set of polynomials in an indeterminate $x \in \mathbb{R}^n$ is taken as*

$$\mathbb{R}[x] = \left\{ \sum_i c_i x^{v^{(i)}};\ c_i \in \mathbb{R};\ v^{(i)} \in \mathbb{N}^n \right\}. \tag{2.12}$$

The nomenclature $\mathbb{R}[x]$ reflects the fact that this set is an algebraic ring over the field $\mathbb{R}$, with the additive and multiplicative identities respectively given as the constant polynomials 0 and 1. The ***degree*** *of a polynomial is the maximal degree of its component monomials. By convention, the summation given in* (2.12) *is sorted in descending order with respect to the degree of the monomials $v^{(i)}$. Monomials with the same degree are sorted according to the degree of the first component x_1, then the second component x_2, and so on, until there is no ambiguity in the ordering of $\{v^{(i)}\}$.*

We often deal with constrained subsets of $\mathbb{R}[x]$. One type of constraint is on the monomial structure. Given a sequence of vectors $\underline{v}$, with each sequence component

in $\mathbb{N}^n$, we take

$$\mathbb{R}_{\underline{\mathrm{V}}}[x] = \left\{ \sum_{i=1}^{l} c_i x_{(i)}^{v^{(i)}};\ \{v^{(1)}, \ldots, v^{(l)}\} = \underline{\mathrm{v}}; c_i \in \mathbb{R} \right\}.$$

We say that the vector sequence $\underline{\mathrm{v}}$ represents the monomial structure. Note that $\mathbb{R}_{\underline{\mathrm{V}}}[x]$ is a vector space under addition.

Another type of constraint is on the set of values that polynomials can take. Given a domain $\Omega \in \mathbb{R}^n$, we define $\mathcal{P}_{\underline{\mathrm{V}},\Omega}[x]$ as

$$\mathcal{P}_{\underline{\mathrm{V}},\Omega}[x] = \{p \in \mathbb{R}_{\underline{\mathrm{V}}}[x] : p(x) \geq 0 \ \ \forall x \in \Omega\}.$$

We now introduce the set of **sum-of-squares** (SOS) polynomials.

Definition 2.14. ***Sum-of-squares polynomials***

A polynomial $p \in \mathbb{R}[x]$, of degree $2d$ is sum-of-squares if it satisfies the two following (equivalent) conditions [14]:

1. $p(x) = \sum_{i}^{l} \big(g_i(x)\big)^2$, *for some* $g_i(x) \in \mathbb{R}[x]$ *and* $i \in \{1, \ldots, l\}$.
2. $p(x) = Z^T(x) Q Z(x)$, *where* $Z(x)$ *is the vector containing all monomials of degree at most* d, *and* $Q \succeq 0$ *(i.e.* Q *is PSD).*

We denote the set of SOS polynomials as $\Sigma[x] \subset \mathbb{R}[x]$. The set of SOS polynomials sharing a fixed monomial structure $\underline{v}$ is denoted $\Sigma_{\underline{v}}[x]$. Note that this set is empty for many choices of v: the highest and lowest degree monomial vectors represented in $\underline{v}$ must both be even for the SOS property to hold.

Note that SOS polynomials are non-negative for all $x \in \mathbb{R}^n$. If not we would have $Z^T(x) Q Z(x) < 0$, for some $x \in \mathbb{R}^n$, contradicting positive semi-definiteness of Q. However, there are non-negative polynomials that are not SOS, the Motzkin poly-

nomial being a famous example [14]. There are three special cases in which non-negativity and the SOS property are known to be equivalent [106]. These are

- Polynomials in a single variable (i.e. $\mathbb{R}[x]$ where $x \in \mathbb{R}$).
- Polynomials of degree two.
- Polynomials of degree four in three variables.

Nevertheless SOS polynomials are dense in the space of non-negative polynomials, regardless of degree or dimensionality. So given a non-negative polynomial $p(x)$, any compact set $\Omega \subseteq \mathbb{R}^n$, and any $\epsilon > 0$, there exists a SOS polynomial $s(x)$ such that $|s(x) - p(x)| \leq \epsilon$ over Ω [65].

Ascertaining whether a given polynomial $p(x)$ has a SOS decomposition of the form $Z^T(x)QZ(x)$, for $Q \succeq 0$, is a semidefinite program. We demonstrate by example:

Example 2.7. ***Verifying the SOS property (taken from [90])***

Let us take $p(x) = 2x^4 + 2x^3y - x^2y^2 + 5y^4$. Is $p(x)$ a sum-of-squares (and therefore globally non-negative)?

This is equivalent to asking if there is a decomposition of the form $p(x) = Z^T(x)QZ(x)$, where $Z(x)$ is a vector containing all monomials of up to degree 2 in x and y, and Q is a PSD matrix. Mathematically, we demand

$$p(x) = \begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix}^T \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix},$$

where $Q \succeq 0$. Clearly $q_{11} = 2$ is required, in order to match the coefficient of x^4 in $p(x)$. Similarly we can define further affine equality constraints on the elements q_{ij}

of Q:

$$2q_{13} = 2, \qquad q_{33} + 2q_{12} = -1, \qquad q_{22} = 5.$$

Finding such a Q is now a semidefinite program.

It is important to note that a $Q \succeq 0$ satisfying the affine constraints is not necessarily unique. This is due to the fact that the terms of the monomial $Z(x)$ are not algebraically independent. For instance, the x^2y^2 coefficient of $p(x)$ is equal to $q_{33} + 2q_{12}$. In the previously described language of optimisation, the feasible set is not necessarily a single point in $\mathbb{S}^n$, the space of symmetric matrices. It is, however, an affine space, as every coefficient-matching constraint is linear in the entries q_{ij}. We can minimise linear objective functions of the coefficients q_{ij} among the space of feasible matrices, while remaining within the SDP framework.

A particular $Q \succeq 0$ within the feasibility region of this example is given by

$$Q = \begin{bmatrix} 2 & -3 & 1 \\ -3 & 5 & 0 \\ 1 & 0 & 5 \end{bmatrix}.$$

Taking a Cholesky decomposition of Q allows us to obtain a lower-triangular matrix L satsifying $Q = LL^T$. This gives us $p(x) = \big(L^T Z(x)\big)^T \big(LZ(x)\big)$, resulting in the explicit representation of $p(x)$ in terms of the sums of squared polynomials. In this example, L is given by

$$L = \frac{1}{\sqrt{2}} \begin{bmatrix} 2 & -3 & 1 \\ 0 & 1 & 3 \end{bmatrix}.$$

Therefore, we obtain

$$p(x) = \frac{1}{2}(2x^2 - 3y^2 + xy)^2 + \frac{1}{2}(y^2 + 3xy)^2.$$

The SDP constructed in Example 2.7 does not have an objective function. This can be incorporated in the framework of (2.10) by taking $C = 0$. If the SDP is feasible, the minimum value attained by the objective function is then zero. This implies that we can use the dual SDP (see (2.11)) to *invalidate* the possibility of a polynomial being SOS. If we run the dual SDP, and obtain a strictly positive value of the objective, then the primal cannot be feasible without violating weak duality (recall the discussion accompanying the presentation of Program 2.2).

We now provide the most general form of optimisation problem concerning sum-of-squares polynomials that can be converted to an SDP.

Program 2.3. ***Convex sum-of-squares (SOS) program (taken from [88])***

Minimise the linear objective function

$$w^T c,$$

where c is a vector formed from the (unknown) coefficients of

$$\begin{aligned} &\textit{polynomials } p_i(x) && \forall i \in \{1, 2, \ldots, \hat{N}\}, \\ &\textit{sum-of-squares polynomials } p_i(x) && \forall i \in \{\hat{N}+1, \hat{N}+2, \ldots N\} \end{aligned}$$

such that

$$a_{0,j}(x) + \sum_{i=1}^{N} a_{i,j}(x)p_i(x) = 0; \qquad \forall j \in \{1, \ldots, \hat{J}\}$$

$$a_{0,j}(x) + \sum_{i=1}^{N} a_{i,j}(x)p_i(x) \in \Sigma[x]; \qquad \forall j \in \{\hat{J}+1, \ldots, J\}.$$

Note that the monomial structure of each polymomials $p_i(x)$ in Program 2.3 must be fixed *a priori*, although different polynomials $p_i(x)$ may have different such structures. In other words, for each $i \in \{1, 2, \ldots, \hat{N}\}$, we must specify a monomial structure $\underline{\mathrm{v}}$ such that

$$p_i(x) \in \mathbb{R}_{\underline{\mathrm{V}}}[x].$$

Similarly, for any $i \in \{\hat{N}+1, \hat{N}+2, \ldots N\}$, we must specify a monomial structure $\underline{\mathrm{v}}$ such that

$$p_i(x) \in \Sigma_{\underline{\mathrm{V}}}[x].$$

Note that some literature deals with the approximate, numerical solution of more general, non-convex SOS programs than can be described by Program 2.3 (see e.g. [53]). These are not considered in the thesis. A schematic for the computational solution of SOS programs, as carried out in this thesis, is provided in Figure 2.1.

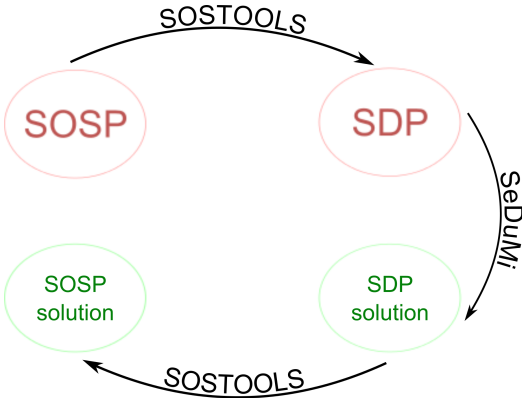

Figure 2.1: Conversion and solution of sum-of-squares programs: *A SOS program in the form of Program 2.3 is written using MATLAB code. The SOSTOOLS package [88] then converts it to an SDP of the form given in Program 2.2. An SDP solver (in our case, SeDuMi [115]) solves the SDP. SOSTOOLS then converts the SDP output back into the form of a SOS program.*

Chapter 3

On the Performance of Nonlinear Dynamical Systems under Parameter Perturbation

This chapter is concerned with the analysis of the behaviour of certain types of nonlinear ODEs over regions of parameter space. In particular, we consider user-defined measures of the discrepancy (over time) between a parameterised model and some reference trajectory, as parameters vary. We formulate algorithms resulting in two-sided bounds on such measures. These bounds are algebraic functions of the initial conditions and parameters of the model. As such, they are valid over continuous regions of parameter space. This distinguishes the work of this chapter from more standard, simulation-based approaches for the parametric analysis of nonlinear systems such as those based on Monte Carlo sampling [47], elementary effects [111], or random parametric variation [40], which can only provide guarantees of system behaviour over discrete sets of parameter vectors. The continuity of the constructed bounds allows the modeller to formulate guarantees such as *'performance property x is satisfied by*

the system for ALL parameter values in the region y'.

The theoretical underpinnings of this chapter build on a rich vein of literature, principally centred in the field of control theory, and relating to the stability and dissipativity analysis of nonlinear ODE systems. We present several theoretical results of direct relevance to research in these areas. The development of **storage functions** as a tool for systems analysis dates back to [135]. These can be used to upper bound system-related performance criteria such as the state-to-output gain. Moreover, the construction of storage functions can often be cast as a computationally tractable convex program. We provide a novel methodology by which lower bounds to such performance criteria can be constructed, also as the solution of a convex program. We foresee that minor variations on our methodology could result in algorithms for lower-bounding a wide variety of performance criteria for which only upper bounds are currently computable, using storage-function based analysis.

A natural measure of the quality of bounds to an unknown quantity is the gap between the upper and lower bound. The algorithms with which we bound the trajectory deviation of nonlinear systems in this chapter, therefore require some objective function that minimises this gap, over parameter space. The objective, moreover, must be linear, in order to preserve convexity of the algorithms, despite the fact that the gap is likely to be a complicated, nonlinear function of the parameters and initial conditions. We present a new type of linear objective function that comes with theoretical guarantees ensuring that it minimises the integral of the gap over the space of parameters and initial conditions. We foresee that existing convex programs used for synthesising storage function-based bounds to performance criteria could be adapted to incorporate this objective, thereby improving the accuracy of the bounds. This work is published in [102].

3.1 Motivation

A key property of any controlled system is asymptotic convergence to a desired equilibrium, for all system initialisations likely to be encountered. Accordingly, much work has been devoted to assessing the regions of stability and attraction for ODE models of such systems. In the case that dynamics are nonlinear, and parametric uncertainty is incorporated into the model, this is not trivial. Assessment methods for the stability of linear systems with parametric uncertainty have a long history in the literature, see e.g. [57]. More recently, approaches based on the sum-of-squares (SOS) relaxation have proved particularly fruitful in the case of polynomial nonlinearities [3, 29, 119, 122, 129].

Once a controller exhibiting satisfactory stability properties is synthesised, interest concentrates on questions of performance. Typical performance measures considered in the literature include $\mathcal{L}_2 \to \mathcal{L}_2$, $\mathcal{L}_2 \to \mathcal{L}_\infty$-induced norms, input-to-state and state-to-output properties. Model nonlinearity and uncertainty, the latter of which we will take to be time-invariant, again complicate matters. To deal with uncertainty, it is customary to construct a performance guarantee valid over the entire parameter range considered, as in [3, 30, 120, 137, 51].

While qualitative guarantees on system dynamics over an entire uncertainty set are useful, they do not consider the manner in which system dynamics vary quantitatively within the uncertainty set. For instance, it may be desirable to find those regions of parameter space eliciting poor transient responses, rather than just bound the quality of transient response over the whole set. Refinement of the uncertainty model, through experimental invalidation or controller redesign, could then concentrate on these regions. Note that initial conditions also play an important role in the transient response, and these regions may be dependent on them. Previous literature has

Another circumstance in which quantitative information on system dynamics within an uncertainty set is of particular interest is in gain scheduled control [86]. Here, a family of linear, time-invariant (LTI) controllers is constructed from the linearisations of the open-loop system on a grid of operating points. A scheduling algorithm then switches between the LTI controllers, dependent on the value of the scheduling variables that determine the operating point [13]. Poor transient response for a region in the space of scheduling variables implies that an additional operating point should be added in this region, or that the region is tied to a suboptimal operating point.

Related questions arise in the model discrimination problem. Here, one typically considers two parameterised systems modelling the same process. Mechanistic models often incorporate parameters of physical significance. This means that some of the parameters of two differing models are often functionally related, if not identical. For instance, any model of a turbine in a wind tunnel will incorporate windspeed as a parameter. In this case, it is natural to consider which choices of shared parameters result in maximally differing trajectories for the two systems. These parameters can then be incorporated in the physical system, increasing the discriminatory capacity of consequent data. Indeed, there is a long history of experimental design for model discrimination in the statistics literature based on this principle, [23, 58]. More recently, these design algorithms have been extended to deal with nonlinear dynamic processes [28, 64]. However, existing algorithms suffer from the necessity of optimisation over spaces of dynamic processes: each cost function evaluation requires numerical simulation of a system model. Note that algorithms of the form outlined are undertaken prior to the generation of experimental data. (In)validation of competing model structures, given existing experimental data, has also been approached in the literature [2, 83, 109, 110, 114].

Model discrimination is only feasible if competing models are structurally distin-

guishable, a concept proposed in [48]. Two models are structurally indistinguishable if any parameterisation of one model is related to a parameterisation of the second model in that their trajectories are identical. In such a case, there is no way to discriminate between the competing model structures. However, testing structural indistinguishability is hard, although progress has been made using techniques based on the observability rank condition [43, 42].

The methods proposed in this chapter relate to the structural indistinguishability problem posed in [48], for the case of models with an input term, and [42], for uncontrolled models. Two models are structurally indistinguishable if any parameterisation of one model is related to a parameterisation of the second model in that their trajectories are identical. In such a case, there is no way to discriminate between the competing model structures. However, testing structural indistinguishability is hard,although progress has been made using techniques based on the observability rank condition [43, 42] Our methods result in a polynomial lower bound, which is a function of both the parameters and initial conditions, between the trajectory deviation of two models. Therefore structural distinguishability is verified if we can prove strict positivity of the lower bound, for fixed parameters in one model, over the space of parameters and initial conditions of the other model. This is an algebraic problem, that could be formulated as an SOS problem.

A closely related problem to that considered in this chapter is surveyed in [131]. Here, algorithms are presented that provide inner and outer approximations to the set S of parameters whose model output deviates from a nominal trajectory by an amount within some user-specified upper and lower bounds. The purpose is to find the set of model parameters that 'reasonably fit' the nominal trajectory, which may be experimental data. This problem can be approached using the methods of the chapter. First, a (potentially black-box), fixed-parameter model that fits the nominal

trajectory of the form described, and adheres to the assumptions of the chapter, would have to be found. One could then construct algebraic bounds on the discrepancy between this fixed model and the model proposed for fitting, over the parameters of the latter model. The set of parameters over which the constructed algebraic upper bound was smaller than a desired critical value and the constructed lower bound exceeded some critical value, would form an inner approximation to S. The set of parameters over which the constructed algebraic lower bound was smaller than a desired critical value and the constructed upper bound exceeded some critical value, would form an outer approximation. The key limitation of applying our methods would be the necessity of synthesising a fixed model of the form described. The set approximations would be sublevel sets of a polynomial in the parameters and initial conditions of the fitting model. This could be less restrictive than the approximations proposed in [131], which are necessarily polygonal or ellipsoidal. Furthermore, our methods can be directly applied to polynomially nonlinear models, whereas the methods of [131] have limited applicability to nonlinear models, relying on repeated linearisation of dynamics.

The aim of this chapter, as previously stated, is to construct two-sided bounds on a quantity that reflects the extent to which the trajectories of two parameterised system models differ. These bounds are rational functions of both the initial conditions and parameters, and thus meaningful with regards to the problems previously outlined. Computation of trajectory deviation, given initial conditions and parameters, thereby requires only an algebraic evaluation rather than a computationally intensive simulation of a system model. Finding the parameters and initial conditions which result in maximally diverging trajectories, as required in the model discrimination problem described, is furthermore a tractable algebraic optimisation problem. In the context of gain scheduling, regions in the space of scheduling variables and initial conditions where the transient response deviates highly from the nominal trajectory

can be expressed through algebraic inequalities.

We consider pairs of dynamical systems taking the form

$$\dot{x}(t) = f(x(t), \theta), \qquad y(t) = h_1(x(t), \theta) \tag{3.1a}$$

$$\dot{z}(t) = g(z(t), \theta), \qquad r(t) = h_2(z(t), \theta) \tag{3.1b}$$

where $x(t) \in \mathbb{R}^{n_x}$ and $z(t) \in \mathbb{R}^{n_z}$ are the state vectors of (3.1a) and (3.1b) respectively. The trajectory of (3.1a) at time t, with initial condition x_0 and parameterisation θ is denoted $x(t; \theta, x_0)$, and analogous notation is used for (3.1b). Note that we do not require $n_x = n_z$, but the dimensions of the outputs must be identical, hence we take $y(t), r(t) \in \mathbb{R}^m$. The two systems share the same compact, predefined parameter space $\Theta \subseteq \mathbb{R}^q$, from which we take θ. Note that this still allows for the systems to have different parameters, as we can concatenate the parameter vectors of both systems into a single parameter vector shared between them. Both systems have parameter-dependent state equilibria: $x^*(\theta)$ and $z^*(\theta)$, such that $f(x^*(\theta), \theta)$ and $g(z^*(\theta), \theta)$ are 0. The methods of this chapter are amenable to systems with multiple equilibria, but we only consider a single equilibrium, as described subsequently. However we allow this equilibrium to be parameter-dependent. We take

$$\nu(x, z, \theta) = \|h_1(x, \theta) - h_2(z, \theta)\|_2^2$$

as the instantaneous difference between output trajectories. Our aim, in this chapter, is to produce functions V_u and V_l satisfying:

$$V_u(x_0, z_0, \theta) \geq \int_0^\infty \nu(x(t; \theta, x_0), z(t; \theta, z_0), \theta)\ dt \tag{3.2}$$

$$V_l(x_0, z_0, \theta) \leq \int_0^\infty \nu(x(t; \theta, x_0), z(t; \theta, z_0), \theta)\ dt, \tag{3.3}$$

Of course, a necessary condition for the above integrals to be finite is that both output maps converge to the same equilibrium point. As such, validity of (3.2) and (3.3) is restricted to the subset of Θ on which this is true, which we denote $\tilde{\Theta}$. In fact, sufficiency also holds for the class of systems we consider. For linear systems, this is simple to prove, and the Hartmann theorem [94] can be used to extend sufficiency to the case in which f, g, h_1, h_2 are $\mathcal{C}^2$ continuous, with nonsingular linearised dynamics around the equilibrium. We provide a proof of this claim in Appendix A.

It is important to note that V_u in (3.2) is a storage function with respect to the supply rate $-\|y(t;\theta,x_0) - r(t;\theta,z_0)\|_2^2$ [135]. Construction of storage function/supply rate pairs for nonlinear systems analysis has been approached previously in the literature [3, 89, 97, 98, 137, 120]. Recent consideration has been given to the additional complication of parameter-varying equilibria [55]. In some cases, one optimises over a restricted space of possible supply rates in order to elucidate a qualitative aspect of the system, such as a bound on the input-to-state gain. In this chapter, the supply rate is given, and optimisation occurs over the space of storage functions in order to minimise the conservatism in the dissipation inequality. Thus tight approximations are provided to the RHS of (3.2) and (3.3). Previous examples of this type of optimisation exist, c.f. [97, 98]. In these, the storage function is necessarily a SOS polynomial. By definition, the storage function is then expressible in the form $n^T(x)Qn(x)$, where $n(x)$ is a predefined vector of monomials, and $Q \succeq 0$ is a matrix (see Section 2.3.3). The trace of Q is then taken as the linear objective function to minimise.

We provide an optimisation routine that operates on the space of general (rather than SOS) polynomial storage functions. Having a larger space of feasible storage functions is likely to yield a less conservative approximation to (3.2). In fact our optimisation routine has theoretical optimality guarantees highlighted in Theorem

3.2. It is also empirically tight, as shown in the Examples section (see Table 3.2), while the SOS routine required to synthesise the storage function requires only a single SOS constraint (see Program 3.3). Finally, this method of optimisation allows for construction of a lower bound of the type (3.3). Two sided bounds to dissipation inequalities based on SOS programming are extremely useful from a quantitative perspective.

Thus far we have considered nonlinear systems, but it would be useful to understand the problems posed in the linear case first. In particular, we will deal with finding regions of initial conditions and parameters that result in large transient response deviations, and express these algebraically, for a class of linear systems. We provide an analytic result, where (3.1a) is a Linear Time Invariant Parameter Dependent (LTIPD) system. In this case, we can find a symbolic function V_e satisfying:

$$V_e(x_0, z_0, \theta) = \int_0^\infty \nu(x(t;\theta,x_0), z(t;\theta,z_0), \theta)\ dt, \tag{3.4}$$

i.e. V_e is exact and not a bound as in the polynomial case. This result is derived in Section 3.3. The nonlinear problem is subsequently dealt with in Section 3.4. Finally, numerical examples, provided in Section 3.5, conclude the chapter.

3.2 Preliminaries

We use some notation specific to this chapter, which is now outlined. For $k \in \mathbb{R}^n$, we define the (symmetric) hyper-rectangle specified by k as

$$\{x \in \mathbb{R}^n : x_i \in (-k_i, k_i)\ \forall i \in \{1, \ldots, n\}\}.$$

We use the symbol $\bullet$ to denote a choice of element from a discrete set without nu-

merical indices, e.g. $\{+,-\}$. So the expression $a \bullet b$, with $\bullet \in \{+,-\}$ is equivalent to either $a+b$ or $a-b$.

We will subsequently create mathematical objects that are defined with respect to a particular dynamical system. Since, in general, we will be dealing with two systems, (3.1a) and (3.1b), we will append a superscripted x, or z, to specify which of these a given object is defined with respect to. So if $\mathcal{R}$ was defined as the region of attraction (ROA) for a dynamical system equilibrium point, we would write $\mathcal{R}^x$ and $\mathcal{R}^z$ to denote the ROA of (3.1a) and (3.1b) respectively.

We take $\zeta(x,k)$, for $x \in \mathbb{R}^n$ and $k \in \mathbb{N}^p$, to denote the set of multivariate monomials in x including all terms with degree k_i. This is analogous to the '*monomials*' command in the SOSTOOLS package [88].

3.3 Trajectory Deviation in LTIPD Systems

We commence by considering the trajectory deviation of LTIPD systems, with polynomial parameter dependence, over the space of initial conditions and parameters. We provide a computable quadratic form in the initial conditions, with rational dependence on the parameters, that corresponds exactly to equation (3.4). Thus, the output of this quadratic form exactly determines the deviation in performance between the two systems, for a given parameter and initial condition. As mentioned previously, it is customary to construct robust performance guarantees that hold for all parameters within a compact, pre-specified set. This has a rich literature, particularly with regards to Linear Matrix Inequality (LMI) based methods [9, 87, 91, 30]. Our methods provide parameter-dependent performance quantification. This makes it possible to find the areas of parameter space whose inclusion in the model most degrade overall performance properties. Improved model/component design could then

concentrate on eliminating inclusion of such areas.

Note that a pre-requisite to our methods is the guaranteed stability of both systems over the parameter domain considered. This can be achieved using the methods of e.g. [15, 45, 141]. Our quadratic form is constructed symbolically by solving multiple Sylvester equations. The symbolic solution of the related algebraic Riccati Equations for controller analysis and synthesis, has been considered previously [60, 113]. Here we solve a parametrically dependent linear matrix equation, yielding a solution rational in the parameters. In [16], it is shown that continuously parameter-dependent LMIs always admit a branch of polynomially parameter-dependent solutions. In our computation, we make use of the following, well-known Lemma, found, for example, in [143].

Lemma 3.1.

The solution to the equation $AX + XB = C$*, for matrices* $A, B, C, X \in \mathbb{R}^{n\times n}$*, where* X *is unknown, is unique if and only if* A *and* $-B$ *share no common eigenvalues. In this case, the matrix* $(I_n \otimes A + B^T \otimes I_n)$ *is non-singular, and* X *is given by* $vec(X) = (I_n \otimes A + B^T \otimes I_n)^{-1} vec(C)$*.*

Equation (3.4) can be calculated symbolically using Lemma 3.1, and takes the form of a ratio of polynomials. We deal with a special case of the set of systems described in (3.1a) and (3.1b), which is provided below:

$$\begin{aligned} \dot{x}(t) &= A_1(\theta)x(t) & \dot{z}(t) &= A_2(\theta)z(t) \\ h_1(x(t),\theta) &= C_1(\theta)x(t) & h_2(z(t),\theta) &= C_2(\theta)z(t). \end{aligned} \tag{3.5}$$

We demand that the entries of all the matrices in (3.5) are polynomial in the param-

eters, so

$$A_{1,ij}(\theta),\ A_{2,ij}(\theta),\ C_{1,ij}(\theta),\ C_{2,ij}(\theta) \in \mathbb{R}[\theta],$$

for all relevant $i, j \in \mathbb{N}$. We see that the systems under consideration are linear, for any fixed parameter vector θ. However the form of the matrices dictating state evolution and the state-to-output mapping depend polynomially on the parameter vector.

In the notation of (3.1a) and (3.1b), we have that

$$\begin{aligned} f(x(t),\theta) &= A_1(\theta)x(t) & g(z(t),\theta) &= A_2(\theta)z(t) \\ h_1(x(t),\theta) &= C_1(\theta)x(t) & h_2(x(t),\theta) &= C_2(\theta)z(t). \end{aligned}$$

We restrict attention to the case that $A_1(\theta)$ and $A_2(\theta)$ are Hurwitz for all $\theta \in \Theta$.

Theorem 3.1.

For a system of the form (3.5)*, and where $A_1(\theta)$ and $A_2(\theta)$ are Hurwitz for all $\theta \in \Theta$, we can express* (3.4) *as:*

$$V_e(x_0, z_0, \theta) = x_0^T W_1(\theta) x_0 + 2 z_0^T W_2(\theta) x_0 + z_0^T W_3(\theta) z_0 \tag{3.6}$$

where $W_i(\theta)$ are the unique solutions to the following Sylvester equations:

$$\begin{aligned} A_1^T(\theta)W_1(\theta) + W_1(\theta)A_1(\theta) &= -C_1^T(\theta)C_1(\theta) \\ A_2^T(\theta)W_2(\theta) + W_2(\theta)A_1(\theta) &= C_2^T(\theta)C_1(\theta) \\ A_2^T(\theta)W_3(\theta) + W_3(\theta)A_2^T(\theta) &= -C_2^T(\theta)C_2(\theta). \end{aligned} \tag{3.7}$$

Moreover, V_e takes the form of a ratio of polynomials in its arguments (i.e. is a

rational function). These polynomials can be explicitly constructed.

Proof.

First observe that, due to linearity of an LTIPD system for any fixed parameterisation, (3.4) can be written

$$V_e(x_0, z_0, \theta) = \int_0^\infty \|C_1(\theta)e^{A_1(\theta)t}x_0 - C_2(\theta)e^{A_2(\theta)t}z_0\|_2^2 \, dt.$$

If we define

$$\begin{aligned} W_1(\theta) &= \int_0^\infty e^{A_1(\theta)^T t}C_1^T(\theta)C_1(\theta)e^{A_1(\theta)t} \, dt \\ W_2(\theta) &= -\int_0^\infty e^{A_2^T(\theta)t}C_2^T(\theta)C_1(\theta)e^{A(\theta)t} \, dt \\ W_3(\theta) &= \int_0^\infty e^{A_2^T(\theta)t}C_2^T(\theta)C_2(\theta)e^{A_2(\theta)t} \, dt, \end{aligned} \tag{3.8}$$

then Equation (3.6) follows by straightforward expansion. Differentiation, with respect to time, of the equalities of (3.8) results in the Sylvester Equations (3.7). Each of these equations can be solved symbolically in turn. Uniqueness of the solutions follows from Lemma 3.1, using the fact that, by assumption, all eigenvalues of both $A_1(\theta)$ and $A_2(\theta)$ have negative real part for all $\theta \in \Theta$, since both matrices are Hurwitz. In particular, each solution takes the form of a ratio of polynomials in the elements of θ. To see this for the function $W_1(\theta)$, note that its explicit solution is $vec(W_1(\theta)) = \left(I_n \otimes A_1(\theta)^T + A_1(\theta)^T \otimes I_n\right)^{-1} vec(-C_1^T(\theta)C_1(\theta))$. The linearity of the matrix $I_n \otimes A_1(\theta)^T + A_1(\theta)^T \otimes I_n$ in θ assures us that its adjugate matrix is composed of polynomial entries in θ, while its determinant is a polynomial in θ. Expressibility of elements of the matrix $W_1(\theta)$ as ratios of polynomials then follows from Cramer's rule. This rule also provides bounds on the maximal degree of these polynomials. Let us take $|A_1(\theta)|$ as the maximal degree of the polynomials comprising the entries of $A_1(\theta)$, and so on for the remaining system matrices. Then the maximal degree of the

polynomial entries comprising the numerator and denominator are given by

$$\begin{aligned} &2|C_1(\theta)| + (n_x - 1)|A_1(\theta)| && \text{(numerator)} \\ &2|C_1(\theta)| + n_x|A_1(\theta)| && \text{(denominator).} \end{aligned}$$

This follows as the degree of a product of polynomials is the sum of their respective maximal degrees. We then upper bound the degree of every entry in $A(\theta)$ and $C(\theta)$ by $|A(\theta)|$ and $|C(\theta)|$ respectively, and multiply through. Analogous results can be gained for $W_2(\theta)$ and $W_3(\theta)$, which, together with Equation (3.6), implies that V_e too is a ratio of polynomials in its arguments. □

In Section 3.5.1 a wind turbine model of LTIPD form from the literature is presented, for which the function V_e described in this section was constructed.

3.4 Trajectory Deviation Between Pairs of Nonlinear Systems

The aim of this section is to construct algebraic bounds V_u and V_l satisfying (3.2) and (3.3). Sufficient conditions for the existence of these bounds are provided in Lemma 3.2, and SOS programs whose feasibility results in their construction are provided (Programs 3.3 and 3.4). Several intermediate steps are, however, required. In order to make the bounds as tight as possible, an objective function is added to the SOS program. The formulation of this objective requires some preliminary theory, with which we commence the section. Next, we face the issue that the bounds are valid on an invariant subset of the state space with particular properties. Sufficient conditions for these properties to hold are given in Theorem 3.3, and semialgebraic inner approximations are constructed using the SOS program provided in Program

3.1.

First we define a functional on the space of polynomials, which takes a polynomial, discards all monomials with odd degree terms (and their associated coefficients), and sums a weighted form of the coefficients of the remaining monomial terms. Weights are proportional to the integral of the relevant monomial over a hyper-rectangle. This functional is linear in the coefficients of its polynomial argument, and it therefore can be used as an objective function in a semidefinite program. Moreoever, Theorem 3.2 shows that minimising this functional, subject to arbitrary constraints, is equivalent to minimising the integral of the polynomial over a symmetric hyper-rectangle, subject to the same constraints. We use it subsequently to minimise the integral of the error between the functions V_u and V_l, and the quantity they bound from above and below, in a theoretically justified manner.

Consider a polynomial $p \in \mathbb{R}_{\underline{v}}[x]$, with $p = \sum_{i=1}^{l} c_i x^{v^{(i)}}$. Let $\hat{\underline{v}}$ be the subsequence of $\underline{v}$ consisting of vectors containing only even terms. Given an arbitrary $k \in \mathbb{R}^n$, let us define $\hat{c} \in \mathbb{R}^l$ such that:

$$\hat{c}_i = \begin{cases} 2c_i \prod_j (v_j^{(i)} + 1)^{-1} k^{v^{(i)}+1} & \text{if } v^{(i)} \in \hat{\underline{v}} \\ 0 & \text{otherwise.} \end{cases} \tag{3.9}$$

Then we define $SqC(p,k) = \sum_{i=1}^{l} \hat{c}_i$. Notice that $\hat{c}_i = c_i \int_\Omega x^{v^{(i)}} dx$, where Ω is the hyper-rectangle specified by k.

As an example, consider $p(x) = 3x_1^4 x_2^2 + 5x_1 x_2^3$. Let $k = [1,1]$. Then $SqC(p,k) = 3\int_{-1}^{1}\int_{-1}^{1} x_1^4 x_2^2 \, dx_1 dx_2$.

Theorem 3.2.

For some $k \in \mathbb{R}^n$, let $\Omega \subset \mathbb{R}^n$ be the hyper-rectangle specified by k. Take an arbitrary set of polynomials $\mathcal{P}^{con}_{\underline{v},\Omega}[x]$, such that $\mathcal{P}^{con}_{\underline{v},\Omega}[x] \subseteq \mathcal{P}_{\underline{v},\Omega}[x]$. As before, $\underline{v}$ is an arbitrary

l-length sequence of vectors in $\mathbb{N}^n$*. Then*

$$\arg\min_{p\in\mathcal{P}^{con}_{\underline{v},\Omega}} \int_\Omega p(x)dx = \arg\min_{p\in\mathcal{P}^{con}_{\underline{v},\Omega}} SqC(p,k) \geq 0.$$

Proof. The symmetry in Ω implies that any monomial with odd degree terms integrates to 0 over the domain. Thus, if we take arbitrary $p \in \mathcal{P}^{con}_{\underline{v},\Omega}[x]$ and decompose it to get $p(x) = \sum_{i=1}^{l} c_i x^{v^{(i)}}$, we have:

$$\int_\Omega \sum_{i=1}^{l} c_i x^{v^{(i)}} dx = \sum_{i=1}^{l} c_i \int_\Omega x^{v^{(i)}} dx = SqC(p,k).$$

Equality of the integral with $SqC(p,k)$ ensures that they have the same minimiser. Positivity of all $p \in \mathcal{P}_{\underline{v},\Omega}[x]$ over the domain Ω ensures non-negativity of the integral. □

Remark 3.1. *Note that* $SqC(p,k)$ *is linear in the coefficients of* p*. Therefore it can be added as an objective within an SOS program without affecting convexity.*

We now provide the assumptions placed on systems (3.1a) and (3.1b) for the remainder of the chapter. All assumptions except A3 are formulated with respect to system (3.1a) for notational clarity, but also apply to system (3.1b).

Assumptions:

A1 The initial conditions and parameters considered both belong to compact hyperrectangles. These are denoted R^x (for initial conditions) and Θ (for parameters).

A2 $f(x,\theta)$ and $h_1(x,\theta)$ are polynomial in their arguments. For each $\theta \in \Theta$, there is only a single $x \in R^x$ such that $f(x,\theta) = 0$. This is denoted $x^*(\theta)$.

A3 We can construct a non-negative polynomial $\mathcal{C} : \Theta \to \mathbb{R}^d$, for some $d \in \mathbb{N}$, such that $\mathcal{C}(\theta) = 0 \Rightarrow h_1(x^*(\theta),\theta) = h_2(z^*(\theta),\theta)$.

We translate the states and parameters of the system such that R^x and Θ are both symmetric around 0, without loss of generality. Note that A2 does not forbid additional equilibria $\tilde{x}$ such that $f(\tilde{x}, \theta) = 0$, if $\tilde{x} \notin R^x$. Note also that in the case of equilibrium outputs that do not vary with the parameters (i.e. $h_1(x^*(\theta), \theta)$ and $h_2(z^*(\theta), \theta)$ are equal and constant over Θ), we can take $\mathcal{C}(\theta) := 0$.

Definitions:

D1 $\tilde{\Theta} = \{\theta \in \Theta : \mathcal{C}(\theta) = 0\}$. $\tilde{\Theta}$ specifies parameters for which (3.2) and (3.3) will be valid, as discussed in the introduction.

D2 Let $\mathcal{R}^x(\theta) \subseteq R^x$ to be the maximal positively invariant set in R^x, for $\theta \in \tilde{\Theta}$. That is, given $x_0 \in \mathcal{R}^x(\theta)$, we have that $x(t; \theta, x_0) \in \mathcal{R}^x(\theta)$, for all $t \in \mathbb{R}^+$.

D3 Let $j^x \in \mathbb{R}^{n_x}$ and $j^\theta \in \mathbb{R}^q$ specify the hyper-rectangles R^x and Θ respectively. Then we define polynomials $J_i^x(x) = (x_i - j_i^x)(x_i + j_i^x)$, for $i \in \{1, \ldots, n_x\}$. Similarly, $J_i^\theta(\theta) = (\theta_i - j_i^\theta)(\theta_i + j_i^\theta)$ for $i \in \{1, \ldots, q\}$.

Note that for $\theta \in \Theta$ and $x \in R^x$, we necessarily have that $J_i^\theta(\theta) \leq 0$ and $J_i^x(x) \leq 0$.

We add a further condition on the system:

*A4 For all $x_0 \in \mathcal{R}^x(\theta)$ and $\theta \in \Theta$ such that $\mathcal{C}(\theta) = 0$: $\lim_{t\to\infty} x(t; \theta, x_0) = x^*(\theta)$.

This is starred as it is not necessary to verify a priori. Feasible solution of Program 3.1 verifies this assumption.

3.4.1 Approximation of Constrained Regions of Attraction

Definition D2 defines $\mathcal{R}^x(\theta)$ as the maximal positively invariant subset of the ROA of $x^*(\theta)$ within R^x. This set is important as the bounds (3.2) and (3.3) we construct will only hold for $x_0 \in \mathcal{R}^x(\theta)$. The same is true of $\mathcal{R}^z(\theta)$, with respect to initial conditions z_0 of (3.1b). Feasible solution of Programs 3.1 or 3.2 result in an

explicit semi-algebraic inner approximation to $\mathcal{R}^x(\theta)$. This semi-algebraic approximation is parameter-dependent, thus minimising conservatism. It takes the form of a parameterised sublevel set of a parameterised Lyapunov function. Similar approaches have been taken in [119, 121, 122], for constant sublevel sets. Computation, in these cases, requires use of the SOS decomposition along with non-convex, bilinear, matrix inequalities. These formulations do not carry the constraint that the positively invariant subsets must be restricted to a pre-specified portion of the state space. In our case, this constraint holds, as $\mathcal{R}^x(\theta)$ is confined to R^x. This allows for convex computation, taking the form of an SOS program that can be expressed as an SDP (Program 3.1). In Program 3.2, we further provide a linear objective that heuristically minimises the conservatism of the sublevel sets as inner approximations to $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$, by pushing their shape, for all $\theta \in \tilde{\Theta}$, towards that of R^x and R^z.

Theorem 3.3.

Consider a system of the form (3.1a), satisfying Assumptions A1 and A2. Take Γ *as an arbitrary subset of* Θ*. If there exist functions* $\gamma : \Theta \to \mathbb{R}^+$ *and continuously differentiable* $E : \mathbb{R}^{n_x} \times \mathbb{R}^q \to \mathbb{R}$ *that satisfy:*

1. $\dot{E}(x, \theta) \leq 0 \quad \forall (x, \theta) \in R^x \times \Gamma$
2. $E(x, \theta) > \gamma(\theta) \quad \forall x \in \partial R^x, \theta \in \tilde{\Theta}$
3. $E(x^*(0), 0) \leq 0.$

Then, for any $\theta \in \Gamma$*, the following inclusion is satisfied:*

$$\{x \in R^x : E(x, \theta) < \gamma(\theta)\} \subseteq \mathcal{R}^x(\theta). \tag{3.10}$$

Assumption A3.4 is additionally guaranteed to hold if:

4. $\dot{E}(x, \theta) = 0 \Rightarrow x = x^*(\theta) \quad \forall (x, \theta) \in R^x \times \Gamma.$

Proof.

Define $\mathcal{L}_\theta = \{x \in R^x : E(x,\theta) \leq \gamma(\theta)\}$. Note that Condition *(3)* ensures that $\mathcal{L}_0$ is nonempty, but this is not the case for arbitrary $\theta \in \Gamma$. Fix $\theta \in \Gamma$ and assume $\mathcal{L}_\theta$ is nonempty. $\mathcal{L}_\theta$ is contained within the interior of R^x due to Condition *(2)*. It is also positively invariant: if a trajectory $x(t;\theta,x_0)$ leaves the set, then there exist t_1, t_2 such that $t_1 < t_2$ and $E(x(t_1;\theta,x_0)) < E(x(t_2;\theta,x_0))$. However, this contradicts Condition *(1)*. Finally, asymptotic stability to $x^*(\theta)$ holds on $\mathcal{L}_\theta$ due to strict negativity of $\dot{E}(x,\theta)$ on $x \neq x^*(\theta)$ (Condition *(4)*). This verifies Assumption A3.4. □

We now provide chapter-specific notation used in a subsequent computational routine that produces functions satisfying the assumptions of Theorem 3.3.

Suppose that we have a function: $E : \mathbb{R}^n \times \mathbb{R}^q \to \mathbb{R}$, where $R^x \subset \mathbb{R}^n$ is the hyper-rectangle specified by $j^x \in \mathbb{R}^n$. ∂R^x consists of $2n$ faces. We enumerate the faces by taking $\partial R^x_{i,+}$ and $\partial R^x_{i,-}$ to be the subsets of vectors in R^x for which the i^{th} component is restricted to be j^x_i or $-j^x_i$ respectively. Thus:

$$\bigcup_{i=1}^{n} \bigcup_{\bullet \in \{+,-\}} \partial R^x_{i,\bullet} = \partial R^x.$$

We define functions $\{E^\bullet_i\}$ to be the restrictions of the first argument of E to each of the faces of ∂R^x. Notice that we obtain an explicit form for a function $E^\bullet_i$ by substituting $\pm j^x_i$ into the i^{th} component of the unknown in the function E.

To guarantee strict negativity of $\dot{E}$ outside equilibrium (Condition *(4)* of Theorem 3.3) we require a polynomial $D(x,\theta)$ satisfying:

$$D(x,\theta) \geq 0; D(x,\theta) = 0 \Leftrightarrow x = x^*(\theta) \qquad \forall x \in R^x, \theta \in \tilde{\Theta}.$$

We may trivially take $D(x,\theta) = \|f(x,\theta)\|_2^2$. However lower degree choices of D, if

they exist, reduce the computational complexity of subsequent algorithms.

The following program establishes a function pair $E \in \mathbb{R}[x, \theta]$ and $\gamma(\theta) \in \Sigma[\theta]$ satisfying the assumptions of Theorem 3.3, for a system of the form (3.1a). In particular we take the set Γ of the theorem as the zero set of a positive polynomial C. C can be chosen arbitrarily, but we take it to be as defined in Assumption A3. This means that $\Gamma = \tilde{\Theta}$.

Program 3.1.

Consider a system of the form (3.1a) *obeying Assumptions A1 to A3. If A3.4 does not hold, pick a small constant* $\epsilon > 0$*, else set* $\epsilon = 0$*. Search for a polynomial:*

$$E \in \mathbb{R}[x, \theta]$$
$$E(x^*(0), 0) = -1$$

and SOS polynomials:

$$\{S_i\}_{i=1}^{n_x} \in \Sigma[x] \qquad \{S_i\}_{i=n_x+1}^{n_x+q} \in \Sigma[\theta]$$
$$\{Q_i\}_{i=1}^{d} \in \Sigma[\theta] \qquad \{U_{ij}\}_{i,j=1}^{i=d,j=n_x} \in \Sigma[\theta]$$
$$\{T_{ij\bullet}\}_{i,j,\bullet\in\{+,-\}}^{i=n_x-1,j=n_x} \in \Sigma[x] \qquad \{R_{ij\bullet}\}_{i,j,\bullet\in\{+,-\}}^{i=q,j=n_x} \in \Sigma[\theta]$$
$$\gamma \in \Sigma[\theta]$$

such that the following are SOS polynomials:

$$\text{(a)} \quad -\dot{E}(x,\theta) + \sum_{i=1}^{n_x} J_i^x(x) S_i(x) + \sum_{i=1}^{q} J_i^\theta(\theta) S_{n_x+i}(\theta) + \sum_{i=1}^{d} C_i(\theta) Q_i(\theta) + \epsilon D(x,\theta)$$

(b) $E_j^{\bullet}(x,\theta) + \sum_{i=1, i\neq j}^{n} J_i^x(x) T_{ij\bullet}(x) + \sum_{i=1}^{q} J_i^{\theta}(\theta) R_{ij\bullet}(\theta) +$

$\sum_{i=1}^{d} C_i(\theta) U_{ij}(\theta) - \gamma(\theta) \quad \forall j \in \{1,\ldots,n_x\}, \bullet \in \{+,-\}.$

Condition *(a)* of Program 3.1 ensures that Condition *(1)* of Theorem 3.3 is satisfied by E. To see this, note that each function $C_i(\theta)$ takes a value of 0 within $\tilde{\Theta}$, each product $J_i^x(x)S_i(x)$ is non-positive within R^x, and each product $J_i^{\theta}(\theta)S_i(\theta)$ is non-positive in Θ. Condition *(4)* is fulfilled where $\epsilon > 0$ as the term $\epsilon D(x,\theta)$ is then zero for $x \in R^x$ if and only if $x = x^*(\theta)$, due to uniqueness of the equilibrium (Assumption A2). Condition *(2)* in Theorem 3.3 is implied by Condition *(b)* of the program, since each summation term in Condition *(b)* is negative over $R^x \times \tilde{\Theta}$. Thus all assumptions of Theorem 3.3 are satisfied and so (3.10) holds.

Ideally the set inclusion (3.10) should approach equality, in order that a minimally conservative semialgebraic approximation of $\mathcal{R}^x(\theta)$ is synthesised. Program 3.1 does not account for this. Therefore we provide Program 3.2, which adds an objective function for this purpose to Program 3.1. Justification of the objective is given subsequently.

Program 3.2.

Consider a system of the form (3.1a) *obeying Assumptions A1 to A3.4. Define*

$$O(x,\theta) = \left(\sum_{j=1}^{n_x} \sum_{\bullet\in\{+,-\}} \left(E_j^{\bullet}(x,\theta) - \gamma(\theta) \right) \right) + \sum_{j=1}^{n_x} \sum_{i=1}^{d} C_i(\theta) U_{ij}(\theta). \tag{3.11}$$

Add the following to Program 3.1, while setting $\epsilon = 0$*:*

$$\text{Minimise } SqC(O(x,\theta), [j^x, j^{\theta}]). \tag{3.12}$$

Recall that $[j^x, j^\theta]$ *specifies the hyper-rectangle formed by the Cartesian product of* R^x *with* Θ.

The following inequalities motivate the objective function:

$$O(x,\theta) \geq E_j^\bullet(x,\theta) - E_k^\bullet(x,\theta) \tag{3.13}$$

$$O(x,\theta) \geq E_j^\bullet(x,\theta) - \gamma(\theta) \tag{3.14}$$

$$\forall (x,\theta) \in R^x \times \tilde{\Theta}; j,k \in \{1,\ldots,n_x\}; \bullet \in \{+,-\} \tag{3.15}$$

Theorem 3.2 ensures that Program 3.2 minimises the volume integral $\int_{R^x \times \Theta} O(x,\theta)$. Positivity of $O(x,\theta)$ on $R^x \times \Theta$ is guaranteed due to Condition *(b)* of Program 3.1. Thus Program 3.2 pushes the values of $O(x,\theta)$ on $R^x \times \Theta$ towards zero, which is desirable as it would ensure, through (3.13), that ∂R^x is a level set of E. It would additionally ensure, through (3.14), that this sublevel set was defined by $\gamma(\theta)$. These in turn would imply equality of the set inclusion (3.10).

To see that (3.13) and (3.14) hold true, note first that each term of the form $E_j^\bullet(x,\theta) - \gamma(\theta)$ in the summation on the RHS of (3.11) is positive over the domain (3.15), thanks to Condition *(b)* of Program 3.1. This implies (3.14). We recover (3.13) as, over the domain (3.15), we have

$$\begin{aligned}
E_j^\bullet(x,\theta) - E_k^\bullet(x,\theta) &= \Big(E_j^\bullet(x,\theta) - \gamma(\theta)\Big) - \Big(E_k^\bullet(x,\theta) - \gamma(\theta)\Big) \\
&\leq \Big(E_j^\bullet(x,\theta) - \gamma(\theta)\Big) + \Big(E_k^\bullet(x,\theta) - \gamma(\theta)\Big) \\
&\leq \left(\sum_{j=1}^{n_x} \sum_{\bullet \in \{+,-\}} \Big(E_j^\bullet(x,\theta) - \gamma(\theta)\Big)\right) \leq O(x,\theta).
\end{aligned}$$

Running Programs 3.1 and 3.2 twice, with respect to both dynamical systems (3.1a) and (3.1b), provides us with function pairs $(E^x(\theta), \gamma^x(\theta))$ and $(E^z(\theta), \gamma^z(\theta))$, such that inner approximations to $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$ are algebraically specified through the

inequality (3.10).

3.4.2 Construction of Bounds on the Transient Deviation between Systems

We now provide theory related to the construction of the bounds (3.2) and (3.3) themselves. A computational routine whose successful completion results in these bounds is then provided.

Lemma 3.2.

Given systems (3.1a) *and* (3.1b) *obeying Assumptions A1 to A3.4, suppose there exist continuously differentiable functions W_u, W_l such that:*

$$W_\bullet : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^q \to \mathbb{R} \quad \bullet \in \{u, l\}$$

$$-\dot{W}_u(x, z, \theta) \geq \nu(x, z, \theta) \tag{3.16}$$

$$-\dot{W}_l(x, z, \theta) \leq \nu(x, z, \theta)$$

$$\forall (x, z, \theta) \in \mathcal{R}^x(\theta) \times \mathcal{R}^z(\theta) \times \tilde{\Theta}. \tag{3.17}$$

Then take

$$V_\bullet(x, z, \theta) = W_\bullet(x, z, \theta) - W_\bullet(x^*(\theta), z^*(\theta), \theta).$$

We then have that

$$V_u(x_0, z_0, \theta) \geq \int_{t=0}^{\infty} \nu\big(x(t; \theta, x_0), z(t; \theta, z_0), \theta\big)\ dt \tag{3.18}$$

$$V_l(x_0, z_0, \theta) \leq \int_{t=0}^{\infty} \nu\big(x(t; \theta, x_0), z(t; \theta, z_0), \theta\big)\ dt$$

where $(x_0, z_0, \theta) \in \mathcal{R}^x(\theta) \times \mathcal{R}^z(\theta) \times \tilde{\Theta}$.

Proof. The result follow straightforwardly for W_u by integration of (3.16) over time. Note that invariance of (3.1a) and (3.1b) within $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$ respectively, ensures that the trajectory remains in the domain (3.17), and the inequality (3.16) therefore holds over all time points. An analogous argument can be used with respect to W_l. □

The next two programs constructively establish a V_u and V_l satisfying the specifications of Lemma 3.2.

Program 3.3. *Search for a polynomial*

$$W_u \in \mathbb{R}[x, z, \theta]$$

and SOS polynomials:

$$\{S_i\}_{i=1}^{n_x} \in \Sigma[x] \qquad \{S_i\}_{i=n_x+1}^{n_x+n_z} \in \Sigma[z]$$

$$\{S_i\}_{i=n_x+n_z+1}^{n_x+n_z+c} \in \Sigma[\theta] \qquad \{Q_i\}_{i=1}^{d} \in \Sigma[\theta]$$

that minimise

$$SqC(B^a(x, z, \theta), [j^x, j^z, j^\theta]) \tag{3.19}$$

subject to

$$B^b(x, z, \theta) \in \Sigma[x, z, \theta].$$

Here, $[j^x, j^z, j^\theta]$ *specifies the hyper-rectangle* $R^x \times R^z \times \Theta$*, and we define*

(a) $B^a(x,z,\theta) := -\dot{W}_u(x,z,\theta) - \nu(x,z,\theta) + \sum_{i=1}^{d} C_i(\theta)Q_i(\theta)$

(b) $B^b(x,z,\theta) := B^a(x,z,\theta) + \left[\sum_{i=1}^{n_x} J_i^x(x)S_i(x) + \sum_{i=1}^{n_z} J_i^z(z)S_{n_x+i}(z) + \sum_{i=1}^{q} J_i^\theta(\theta)S_{n_x+n_z+i}(\theta) \right].$

As in Lemma 3.2, we take

$$V_u(x,z,\theta) = W_u(x,z,\theta) - W_u(x^*(\theta), z^*(\theta), \theta),$$

and note that V_u satisfies the inequality (3.18)*.*

Since *(b)* is SOS, the inequality $B^a(x,z,\theta) \geq 0$ holds on $R^x \times R^z \times \tilde{\Theta}$, as all summation terms in the condition are zero or negative within this space. Thus W_u fulfils the assumptions of Lemma 3.2, implying that V_u satisfies (3.18). Note that the conservatism of (3.18) at a point (x_0, z_0, θ) is expressed as the line integral $\int_0^\infty B^a(x(t;\theta,x_0), z(t;\theta,z_0), \theta)\ dt$, by straightforward integration. The path this line integral will take is dependent on initial conditions and parameters. Therefore, in order to minimise all line integrals of the form described, we minimise the volume integral $\int_{R^x \times R^z \times \Theta} B^a(x,z,\theta)\ dx\ dz\ d\theta$. Theorem 3.2 provides us with the objective function necessary to carry this out, motivating the choice of minimiser.

Note that only a single SOS constraint is required in the entire Program. We are guaranteed through Theorem 3.2 that our objective is necessarily positive. Nullity of the objective would indicate equality of (3.2).

Program 3.4. *This program follows Program 3.3. However, replace W_u with W_l, and the term $-\dot{W}_u(x,z,\theta) - \nu(x,z,\theta)$, which appears in* (a)*, with: $\dot{W}_l(x,z,\theta) + \nu(x,z,\theta)$.*

This has the effect of producing a lower bound of the form (3.3).

Remark 3.2. *Any polynomial p in an SOS program that is constrained to be SOS, requires a decomposition of the form $p(x) = Z^T(x)QZ(x)$, as given in Definition 2.14. While Q is returned as the output of the program, the monomial vector $Z(x)$ must be pre-specified. Providing high-dimensional vectors $Z(x)$ with monomials of higher maximal degree results in a program that is less conservative. However, the resulting program is more computationally intensive and prone to numerical error. Determining optimal monomial choices is an ongoing research topic [92, 95]. Our heuristic for choosing monomial sets for Lyapunov/Storage functions (i.e. those denoted E or W in the programs of the chapter) was to fix parameters at nominal values, and find a parameter independent Lyapunov function for the consequent system, using the SOS techniques of e.g. [3, 89]. Denoting by $\mathcal{M}$ the set of monomials from which such a Lyapunov function is comprised, we typically took $\mathcal{M} \cup (\mathcal{M} \otimes \zeta(\theta, [1, \ldots, w])$ as the set of monomials for our desired parameterised Lyapunov/Storage function, where w was set to be 2 or 3. Those SOS polynomials directly attached to the indicator functions defined in D3, which were denoted R, Q, S or T, were given a monomial vector $Z(x)$ with terms of up to degree 2 in the relevant variables.*

3.5 Examples

3.5.1 Wind Energy Conversion System Dynamics

We now provide numerical examples. The first example describes dynamics of a fixed-pitch wind energy conversion system (WECS) derived in *Bianchi et al* [13]. We ignore short time-scale, noisy fluctuations of windspeed from its steady state. The closed

loop nonlinear model is of the form:

$$\begin{aligned} \dot{x}(t) &= a(x(t),\theta) + Bu(t) \\ x(t) &= [\hat{\psi}_s, \hat{\Omega}_r, \hat{\Omega}_g]^T, \quad \theta = [\bar{V}, \bar{\Omega}]^T, \quad u(t) = \Omega_z. \end{aligned} \tag{3.20}$$

Here x, θ, and u denote the state, uncertain parameter vector and control input respectively, while the drift dynamics are

$$\underbrace{\begin{pmatrix} 0 & 1 & -1 \\ \dfrac{-K_s}{J_r} & -\dfrac{B_s}{J_r} & \dfrac{B_s}{J_r} \\ \dfrac{K_s}{J_g} & \dfrac{B_s}{J_g} & -\dfrac{B_s+B_g}{J_g} \end{pmatrix} \begin{pmatrix} \hat{\psi}_s \\ \hat{\Omega}_r \\ \hat{\Omega}_g \end{pmatrix} + \begin{pmatrix} 0 \\ \dfrac{T_r(\lambda) - T_r(\bar{\lambda})}{J_r} \\ 0 \end{pmatrix}}_{a(x,\theta)},$$

where $\lambda = \frac{(\bar{\Omega}+\hat{\Omega}_r)R}{\bar{V}}$ denotes the tipspeed-windspeed-ratio for the turbine and $C_Q(\lambda) = C_{Qmax} - k_Q(\lambda - \lambda_{Qmax})^2$ is a quadratic approximation of the torque coefficient. The aerodynamic torque is defined as $Tr(\lambda) = \frac{1}{2}\rho\pi R^3 C_Q(\lambda)\bar{V}^2$ and finally the control matrix is $B = \left(0, 0, \frac{B_g}{J_g}\right)^T$. In the above model, an overbar on a variable indicates steady state for that variable, while a hat indicates deviation from steady state. From now on, we assume that the first state, $\hat{\psi}_s$ has been rescaled by a factor of 100, so that $x_1 = 0.01\hat{\psi}_s$. Additional model nomenclature is summarised in Table 3.1. Nominal values for the constants are collected from [12, 13]. Note that a nominal value is not provided for the states and parameters of the model, as our subsequent model analysis will be valid for a (subsequently specified) range of values of these model quantities. Note that the parameters of the model, $\bar{V}$ and $\bar{\Omega}$, denote steady state windspeed and rotor speed, respectively.

In [13], (3.20) is converted to an LTIPD form through the first-order Taylor approxi-

Table 3.1: Variables of wind energy conversion system model

Variable	Units	Nominal value	Meaning
K_s	Tm^2/s^2	7500	Drivetrain Stiffness
J_g	$kg\ m^2$	70	Generator Inertia
J_r	$kg\ m^2$	160	Rotor Inertia
B_s	$kg\ m^2$	50	Drivetrain Damping
B_g	$kg\ m^2$	10	Generator Damping
B_r	kgm/s	38.44	Intrinsic Rotor Damping
ρ	kg/m^3	1.225	Density of Air
R	m	17.5	Rotor Radius
C_{Qmax}	1	0.066	Maximum Torque Coefficient
λ_{Qmax}	1	6.95	Tip-windspeed ratio at C_{Qmax}
$k_{r,V}$	1	15.1	Rotor Torque/Windspeed Gain
$\hat{\Omega}_r$	rad/s	N.A.	Rotor Speed Variation
$\hat{\Omega}_g$	rad/s	N.A.	Generator Speed Variation
ψ_s	rad	N.A.	Torsion Angle: $\Omega_r - \Omega_g$
$\bar{V}$	m/s	N.A.	Steady State Windspeed
$\bar{\Omega}$	rad/s	N.A.	Steady State Rotor Speed
Ω_z	rad/s	N.A.	Zero Torque Generator Speed

mation:

$$T_r(\lambda) - T_r(\bar{\lambda}) \approx \hat{\Omega}_r \frac{\partial T_r}{\partial \hat{\Omega}_r}(\bar{\lambda}),$$

which induces linearity of the model with respect to both the parameters $(\bar{V}, \bar{\Omega})$, and the state $\hat{\Omega}_r$. Thus the LTIPD system can be described as:

$$\dot{z}(t) = A(\theta)z(t) + Bu(t). \tag{3.21}$$

The LTIPD conversion described is conducted in [13] so as to facilitate gain-scheduled control of the wind turbine. In gain-scheduled control, a family of linear, time-invariant (LTI) controllers is constructed from the linearisations of the open-loop system on a grid of operating points in parameter space. A scheduling algorithm then switches between the LTI controllers, dependent on the value of the scheduling

variables that determine the operating point. Poor transient response for a region in the space of scheduling variables imply that an additional operating point should be added in this region, or that the region is tied to a suboptimal operating point. The methods of the chapter allow us to construct algebraic bounds on the discrepancy in dynamics between the LTIPD and fully nonlinear system, as a function of the initial conditions and parameters. Extra operating points can then be placed in regions of the parameter space where this discrepancy is large.

We now compare the two models given in (3.20) and (3.21) respectively, using the methods of the chapter. First, in order to obtain the closed loop dynamics, we added a simple, nominal, state feedback controller of the form $u = -Kx$. K was chosen using Ackermann's formula so as to ensure that the eigenvalues of the linearised system were well inside the left half plane.

The steady states of both models describing WECS dynamics have a parametrically invariant steady state of 0. We therefore took $R^x = [-1, 1]^3$ and $R^z = [-1, 1]^3$ as reasonable intervals of initial conditions for which to investigate the two models. We then chose $\Theta = [-0.5, 0.5]^2$ as a reasonable set of steady-state windspeeds and rotor speeds to investigate. We used Program 3.1 to verify that the described sets Θ, R^x, and R^z satisfy Assumption A3.4 imposed in Section 3.4. Program 3.2 constructed function pairs $\{E^x, \gamma^x\}$ and $\{E^z, \gamma^z\}$, each satisfying (3.10). The sizes of the SDPs described, as represented through the dimensions of the constraint matrix, were no more than 14841×2437 [1]

The relative volume of our approximations to $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$ within R^x and R^z was then assessed. 10000 random triples (x, z, θ) were drawn uniformly from $R^x \times R^z \times \Theta$. Of these, 43.21% of doubles (x, θ) satisfied (3.10), and 43.56% of doubles (z, θ) satisfied the corresponding inequality for system (3.1b). Without incorporation of the objective

[1] An SDP is of the form: $\min c^T x$ subject to $Ax = b$ and x is in the cone of positive semidefinite matrices. The size given is that of the matrix A.

function provided in Program 2, only 3.74% of doubles (x, θ) and 4.06% of doubles (z, θ) satisfied (3.10). This demonstrates the utility of the objective function provided in Program 2. Note that the true volume of $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$ within $\mathcal{R}^x$ and $\mathcal{R}^z$ respectively, is a property of the system rather than our algorithm.

Programs 3.3 and 3.4 were then run with respect to the systems (3.20) and (3.21), resulting in upper and lower bounds of the form (3.2), (3.3), where the functions h_1 and h_2, defined in (3.1a) and (3.1b), were both the identity function. Thus our quantifier of model discrepancy at a fixed time was given as

$$\nu(x(t;\theta,x_0), z(t;\theta,z_0), \theta) = \|x(t;\theta,x_0) - z(t;\theta,z_0)\|_2^2.$$

Recall that this is a user-defined function, and could have been replaced by any polynomial function of the states $x(t;\theta,x_0)$ and $z(t;\theta,z_0)$ by appropriate modification of h_1 and h_2. In particular we could have transformed these states by a polynomial output map, or omitted particular state components.

The sizes of the SDPs solved by these programs, as represented through the dimensions of the constraint matrix, were both equal to 29457×3003. To test the conservatism of V_u and V_l with respect to the integral they approximate, 1000 triples (x_0, z_0, θ) were drawn uniformly from the set $R^x(\theta) \times R^z(\theta) \times \Theta$. For each triple, the true value of the gain (ie the RHS of (3.2)) was calculated numerically, and compared to the corresponding values of the functions V_u and V_l. Results are summarised in Table 3.2. Numerical results are displayed with accuracy to 4 decimal places. We see that both V_u and V_l are highly accurate, being on average within one decimal place of the numerical solution. In relative terms, the lower bound V_l attains on average 95.66% of the upper bound V_u.

Finding a parameterisation that maximises the $\mathcal{L}_2$ deviation between (3.20) and

Table 3.2: Evaluation of bound quality for wind energy conversion system model

	Upper Bound	Numerical	Lower Bound
Average magnitude	6.6372	6.3932	6.3494
Least accurate point: upper bound (relative)	0.1107	0.0190	0.0039
Least accurate point: lower bound (relative)	0.1107	0.0190	0.0039
Least accurate point: upper bound (absolute)	53.5159	51.2330	51.2082
Least accurate point: lower bound (absolute)	55.4051	54.2410	53.0904

(3.21), for given initial conditions $\hat{x}_0$ and $\hat{z}_0$, can now be relaxed to finding:

$$\max_{\theta\in\Theta} V_\bullet(\hat{x}_0, \hat{z}_0, \theta), \bullet \in \{u, l\}.$$

Areas of $\mathcal{R}^x(\theta) \times \mathcal{R}^z(\theta) \times \Theta$ in which the two models are guaranteed to correspond poorly can be expressed as $\{(x_0, z_0, \theta) : V_l(x_0, z_0, \theta) > \lambda\}$, for some user-defined λ. This could useful in gain scheduling. Here, one has a single nonlinear model, and several competing linearised models at different operating points. Storage functions of the form (3.2) can be repeatedly constructed with respect to the nonlinear model, and each of the linearised models in turn, with Θ denoting the space of scheduling variables. At a given value of the system state and scheduling variables (x, z, θ), the storage function with the lowest value would then determine the ideal operating point.

Theorem 3.1 was applied to the LTIPD system (3.21), to construct a function V_e satisfying

$$V_e(z_0, \hat{z}_0, \theta) = \int_0^\infty \|(z(t; \theta, z_0) - z(t; 0, \hat{z}_0), \theta)\|_2^2 \, dt.$$

In other words, V_e charts, as a function of the model parameters and initial conditions,

the deviation in trajectory from the nominally parameterised dynamics.

V_e, as stated in Theorem 3.1, took the form of a sum of quadratic forms in the initial conditions, with rational parameter dependence. The maximal degree of the numerator polynomial over all the rational matrix elements of these quadratic forms was 7. The corresponding value for the denominator was 6. Completion of the symbolic computations necessary to construct V_e took 4.64 seconds on a PC with Intel Core i7 3.40GHz processor and 16Gb RAM. Therefore, finding 'worst case' parameterisations $\theta \in \Theta$ maximising the $\mathcal{L}_2$ norm of the output transient is converted to a nonlinear algebraic optimisation, which consists of maximising V_e. As stated previously, the symbolically constructed V_e is valid exactly where the parameterised model is asymptotically stable. Inner approximations to the set of parameters for which this is valid can be easily provided using the methods of e.g. [45]. At a single parameter vector, one can check asymptotic stability merely by testing whether the state transition matrix (i.e. $A(\theta)$) is Hurwitz.

3.5.2 Guaranteed Boundedness of the SqC Operator

The objective function of optimisation routines throughout this chapter (Programs 3.2, 3.3 and 3.4) has been the output of the SqC operator applied to a polynomial. The SqC functional takes a weighted sum of the coefficients of the polynomial. In this section we demonstrate the importance of this weighting.

Consider the objective function (3.19), in Program 3.3. Here, the argument of SqC is a polynomial that is constrained to be positive in a hyper-rectangle. Therefore its integral is also positive, and Theorem 3.2 then guarantees that the objective to be minimised is positive. The theorem also provides optimality guarantees on the integral. Suppose we now modified the SqC functional to take an unweighted sum of monomial coefficients. Using the notation introduced immediately prior to Theorem

(3.2), we would change (3.9) to:

$$\hat{c}_i = \begin{cases} c_i & \text{if } v^{(i)} \in \underline{\hat{v}} \\ 0 & \text{otherwise.} \end{cases}$$

Guaranteed positivity of the objective is now lost, as are the optimality guarantees of Theorem 3.2. Note that the existence of a feasible solution with negative objective results in an unbounded optimisation routine: we could iteratively scale the expression:

$$-\dot{W}_u(x, z, \theta) + \sum_{i=1}^{d} C_i(\theta) Q_i(\theta),$$

as found in definition *(a)* of Program 3.3, by two. This would not affect feasibility, but would reduce the value of the objective (3.19). Indeed, application of Programs 3.2, 3.3 and 3.4, with this modified SqC functional, to Example 3.5.1 results in optimisation routines that all have unbounded objective, resulting in infeasibility.

3.5.3 Parameter-dependent Equilibria

We now provide an academic example that highlights the applicability of the methods of the chapter to models that have parameter-dependent equilibria, and nonlinear parameter-dependence. We compare two models of forced mass-spring systems with nonlinear springs. Take the following system:

$$\begin{aligned} &\dot{x}_1 = x_2 && \dot{z}_1 = z_2 \\ &m\dot{x}_2 = \theta_1 - cx_2 - \theta_2^2 x_1^5 && m\dot{z}_2 = \theta_1 - cz_2 - \theta_2 z_1^3 \\ &h_1(x) = x_1 && h_2(z) = z_1. \end{aligned} \tag{3.22}$$

Here, m and c denote mass and damping coefficients, θ_1 represents a constant forcing term, and θ_2 alters the behaviour of the springs. Note that we could easily make the model linear in the parameters by taking $\theta_3 = \theta_2^2$, but we wish to emphasise the applicability of the approach to models that are nonlinear in both parameters and states. Each model has a single, parameter-dependent, equilibrium, respectively given by $x_1^5 = \frac{\theta_1}{\theta_2^2}$, $x_2 = 0$; and $z_1^3 = \frac{\theta_1}{\theta_2}$, $z_2 = 0$. Moreover, the polynomial $\mathcal{C}(\theta) = (\theta_2\theta_1^2 - 1)^2$ satisfies Assumption A3. In other words, the models share a (parameter-dependent) equilibrium when $\mathcal{C}(\theta) = 0$. We took $R^x = [0.5, 1.5] \times [-0.5, 0.5]$, $R^z = [0.2, 1.2] \times [-0.5, 0.5]$, and $\Theta = [0.6, 1] \times [1, 2]$. These represent the sets of parameters and initial conditions over which our eventually constructed algebraic bounds will be valid. This choice was arbitrary. To ensure validity, we first had to run Program 3.1 twice with respect to the x and z-systems, Both programs yielded a feasible solution to verify Assumption A3.4. Program 3.2 was then performed, resulting in function pairs $\{E^x, \gamma^x\}$ and $\{E^z, \gamma^z\}$, with each pair satisfying (3.10). This resulted in inner approximations to $\mathcal{R}^x(\theta)$ and $\mathcal{R}^z(\theta)$. 1000 triples (x, z, θ) were randomly sampled from $R^x \times R^z \times \tilde{\Theta}$. Of these, the inner approximations guaranteed that 41.92% of pairs (x, θ) were within $\mathcal{R}^x(\theta)$, and 73.59% of pairs (z, θ) were within $\mathcal{R}^z(\theta)$.

Programs 3.3 and 3.4 were further run, resulting in bounds of the form (3.2) and (3.3). Accuracy of the bounds was checked by numerical calculation of the RHS of (3.2) and (3.3) at 1000 sample points. These values were checked against those of the bounds V_u and V_l. V_u yielded a value on average 11.69% larger than that of the numerical value, while V_l was on average 15.16% smaller.

Monomials for the programs were chosen as in Remark 3.2. The polynomials W_u and W_l of Programs 3.3 and 3.4 were of degree 8. The sizes of the consequent SDPs, as represented by the dimensions of the constraint matrix were no greater than 30894×2991 for Program 3.2, and were equal to 71835×6314 for Programs 3.3 and 3.4.

3.6 Summary

In this chapter we have first considered LTIPD systems, and provided an exact method for quantifying the effect of perturbation in parameters and initial conditions on transient dynamics. This converts regional sensitivity analysis of such a system into an algebraic problem. A convex SOS program was then provided that constructed a parametrically dependent, invariant subset of the ROA for a polynomial system. In particular, this subset was guaranteed to lie within a pre-specified set, and was optimised to take up as large a proportion of the pre-specified set as possible. Existence of such a subset was necessary for the final part of our chapter, which generalised our exact LTIPD related results to the comparison of separate, polynomial, parameterised systems. A convex optimisation routine was constructed resulting in algebraic upper and lower bounds to trajectory deviation between the systems, over time. These bounds were dependent on parameters and initial conditions of the two systems. The optimisation routine came with theoretical optimality guarantees on the conservatism of the bounds. Examples were provided demonstrating the methods of the chapter and quantifying their efficacy.

Chapter 4

On Parametric Sensitivity and the Well-Posedness of the Parameter Estimation Problem

This chapter is concerned with the parameter estimation problem, in which the parameters of a model are fitted so that model output best corresponds to noisy experimental data. Generally, many parameter vectors are 'reasonably consistent' with the data, and these together form the ϵ-*uncertainty region*, where ϵ is a numeric quantifier (derived subsequently) of what constitutes 'reasonable'. We consider geometrical aspects of the uncertainty region, with a view to both determining the accuracy with which it is possible to make parameter estimates, and to gaining fundamental mechanistic insight into the modelled process that cannot be gained by parameter estimation alone. Recent literature has focused on the concept of *model sloppiness*, a model property in which the uncertainty region has a high aspect ratio. In a sloppy model, some combinations of parameters can be estimated with high accuracy, and some not at all. However the quantification of sloppiness assumes that measurement noise is

infinitesimally small, and approximates the uncertainty region based on this assumption. We formulate the novel concept of **multiscale sloppiness**, which quantifies the aspect ratio of the uncertainty region given a nonzero, user-defined magnitude of measurement noise. We show, through example, that multiscale sloppiness and standard sloppiness often correspond poorly. Moreover, the multiscale sloppiness of a model can be highly dependent on the magnitude of measurement noise corrupting experimental data. Thus we demonstrate that labelling combinations of parameters easy or hard to estimate accurately from data on the basis of model sloppiness is not justified, and offer an alternative. We then highlight theoretical properties of multiscale sloppiness, including its relationship to likelihood-ratio testing and the Fisher Information Matrix. We show that the existing concept of model identifiability (both structural and practical) can be incorporated into our framework. The results of this chapter are published in [100], and parts of [38].

4.1 Motivation

The previous chapter introduced methods of algebraically quantifying the disruptive effect of changing parameters in ODE models. This gave us a highly detailed, simulation-free understanding of how model behaviour varies over parameter space. The measure of variation in model behaviour was user-defined to some extent, but constrained to be an integral over time of the instantaneous difference between the model output and some nominal trajectory. This type of measure is common when undertaking a model-based analysis in the absence of experimental data.

In this chapter, we analyse how the geometrical characteristics of model variation in parameter space influence the amount of information associated with a parameter vector that optimally fits (noisy) experimental data. As noted in Definition 2.8 of

Chapter 2, parameter vectors along a structural unidentifiability can never be discriminated on the basis of experimental data, for any model, since they induce identical outputs. If the data is noisy, then even globally structurally *identifiable* models can have various parameter vectors that cannot be discriminated, as we now explain.

When measurement noise corrupts observations of a process, it is impossible to disentangle which portion of the data was attributable to noise, and which was due to the process. We can, however, say that the noise-free data is more likely to have attained some values than others. Often, error bars are placed around experimental data in order to signify the regions of data space in which the noise-free data is assumed to reside, with some desired degree of confidence. In this case, any parameter vector inducing model output that remains within the error bars is 'plausible' in some sense. These plausible parameter vectors constitute the **uncertainty region** associated with a parameter estimate, to some level of confidence. A more formal definition will subsequently be given, and an illustrative diagram is provided in Figure 4.1. We see that any structural unidentifiability associated with a parameter estimate will always be within its uncertainty region, regardless of the signal-to-noise ratio of the data, and the desired confidence level. If the uncertainty region of a locally structurally identifiable model is also too large for the purposes of the modeller, then the estimate is termed **practically unidentifiable**. A functional relation in parameter space, analogous to (2.3), over which model output changes so little as to be indistinguishable through noise-corrupted observation is denoted a **practical unidentifiability**. A more rigorous definition of practical unidentifiabilities is provided in Chapter 5.

The geometry of the uncertainty region can be investigated by assessing the sensitivity of model output to parameter perturbation, as there is an underlying duality between this sensitivity and the uncertainty associated with parameter estimation in the presence of measurement uncertainty. For instance, it is hard to accurately

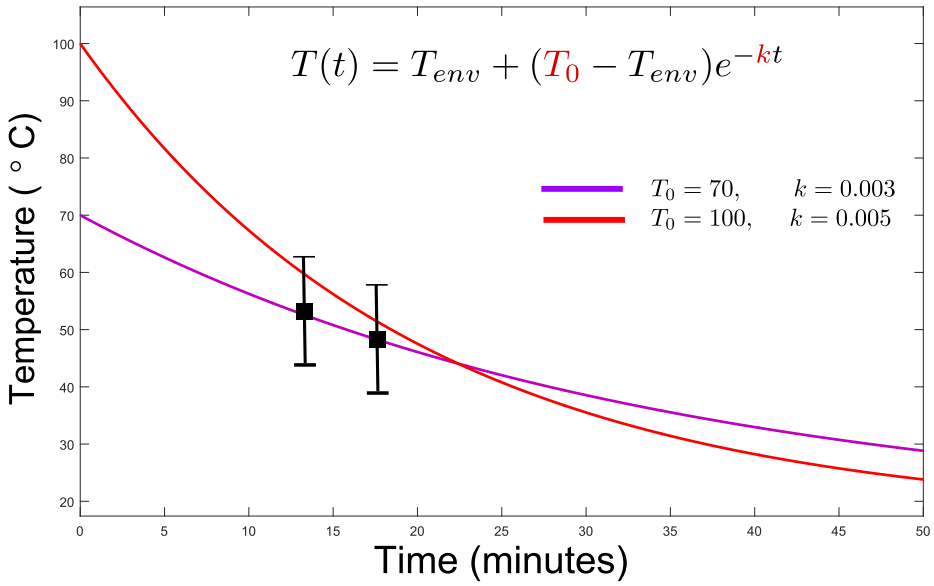

Figure 4.1: Model predictions and error bars. *Model of the temperature of a cooling body. T_{env} is the environmental temperature. T_0 is the initial temperature of the body. k is proportional to the heat transfer coefficient. Two measurements are taken at distinct timepoints (black squares) and the model is globally structurally identifiable. However the measurements are noisy, and we consequently associate error bars containing the range of plausible noise-free temperature values. The temperature trajectories of the model at two different parameter vectors are drawn. The purple trajectory ($T_0 = 70$, $k = 0.003$) best fits the data, in that it is a perfect interpolation. However, the red trajectory ($T_0 = 100$, $k = 0.005$) is also plausible, as its predictions stay within the error bars associated with the noisy measurements.*

estimate the model parameters from noisy data if a large change in some (combinations of) parameters induces only a small change in model output. Recently, this duality has been explored through the concept of model *sloppiness* [21, 50, 132, 123]. Here, the sensitivity of model output to infinitesimal parameter perturbation is highly anisotropic, in that the effect of perturbation in sensitive ('stiff') directions exceeds that of insensitive ('sloppy') directions by many orders of magnitude. It has been shown that sloppiness is a common feature in Systems Biology models [50]. A sloppy direction suggests the existence of a practical unidentifiability, although the precise correspondence has been debated [6, 32].

Existing literature [59, 128, 105, 56], has suggested that the local sensitivity char-

acteristics of nonlinear models often poorly approximate their sensitivity to larger perturbations. This suggests that model sloppiness may not be informative in determining the uncertainty associated with parameter estimates when measurement uncertainty is non-infinitesimal. Dually, it implies, when simplifying models, that the most appropriate reduced model is highly dependent on the range of parameters over which validity is intended. This case has been made in e.g. [2, 93]. We introduce a new notion, *multiscale sloppiness*, that quantifies sensitivity anisotropy as a function of the length scale of perturbation considered, relative to a fixed dataset/model prediction. We show how it furthermore relates to the geometry of the set of parameters satisfying a particular likelihood-ratio hypothesis test. Multiscale sloppiness asymptotically corresponds to the standard formulation of model sloppiness in the limit of decreasing length-scale. We find that the sensitivity characteristics of models can alter drastically as perturbations of increasing magnitude are considered. In addition, both sloppiness and multiscale sloppiness can be highly dependent on the particular parameter vector considered, as we demonstrate subsequently by example. This suggests that caution must be exercised in labelling an entire model structure sloppy, based on analysis at a single parameter vector. Multiscale sloppiness allows for analysis of how the uncertainty region associated with a parameter estimate changes as a function of the signal-to-noise ratio of the data.

The previous discussion demonstrates that determination of (structurally/practically) unidentifiable parameters, and their associated unidentifiabilities, is important when evaluating the accuracy of a parameter estimate. We can also gain mechanistic insight into the studied process from such an analysis. Indeed, if model sensitivity to parameter perturbation is dominated by a few 'stiff' directions in parameter space, it is possible that an underlying, macroscopic model simplification exists [125]. To take a simple example, we can again revisit the trolley-on-a-spring model (Figure 1.2), in which we concluded that dynamics of the process were dependent on the ratio

of spring stiffness to trolley mass, rather than their individual values. This insight allows us to tune the spring stiffness to maintain dynamics when changing the mass of the trolley. It also suggests that the model can be simplified, as we could replace the individual parameters with their ratio without sacrificing model informativeness. For structural unidentifiabilities, this process of replacement is known as a locally identifiable reparameterisation, and relevant algorithms have been provided in e.g. [27, 41, 78]. This further motivates the requirement for algorithms, and accompanying theoretical frameworks, that can accurately determine structural and practical unidentifiabilities arising in the parameter space of a model.

4.2 A General Framework for Quantifying Variability in Model Output over Parameter Space

We first require a cost function in parameter space that quantifies the discrepancy between parameterised model outputs $\mathcal{Y}(\theta)$, and either a nominal output $\mathcal{Y}(\theta^*)$, or (E-replicated) experimental data $\mathcal{D}^{[E]}$ if it exists. We will denote this cost function $C_{\mathcal{D}^{[E]}}(\theta)$ or $C_{\mathcal{Y}(\theta^*)}(\theta)$, depending on the choice of reference. When data is not replicated (i.e. $E = 1$), then we use the shorthand $\mathcal{D}$ to denote the data.

For the remainder of the thesis, we will use the subscript $\bullet$ to denote quantities that may either be formulated with respect to data $\mathcal{D}$, or model predictions $\mathcal{Y}(\theta^*)$. So $C_\bullet(\theta)$ denotes *either* $C_{\mathcal{Y}(\theta^*)}(\theta)$ *or* $C_{\mathcal{D}}(\theta)$. A consistent choice of $\mathcal{D}$ or $\mathcal{Y}(\theta^*)$ must be made over a particular equation. For instance, consider the equation $U_\bullet(\epsilon) = \{\theta : C_\bullet(\theta) \leq \epsilon\}$. We demand that both bullets in this equation refer to the same quantity.

The lower the value of the cost function, the better that model output at θ represents the nominal output. The quantity $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)$, which is used extensively in the sequel, will be referred to as the *Cost Hessian*. We make the following assumptions

for the remainder of the chapter:

A1 The cost functions $C_{\mathcal{D}^{[E]}}(\theta)$ and $C_{\mathcal{Y}(\theta^*)}(\theta)$, as well as model output $\mathcal{Y}(\theta)$, are twice differentiable with respect to θ, for all $\theta \in \Theta$.

A2 $\mathcal{Y}(\theta) = \mathcal{Y}(\theta^*) \Leftrightarrow \theta = \min_{\theta \in \Theta} C_{\mathcal{Y}(\theta^*)}(\theta)$. This minimiser does not have to be unique in Θ.

Recall from Section 2.2 that we distinguish between the (deterministic) *process model* $\mathcal{Y}(\theta)$, and the *experimental model* $\mathcal{D}(\theta)$, which also accounts for the data-collection method and therefore includes a stochastic term representing measurement noise. Specifically $\mathcal{D}(\theta)$ is a random variable predicated on $\mathcal{Y}(\theta)$, with probability density $\psi_{\mathcal{Y}(\theta)}$. Our final assumption is that:

A3 The support of $\psi_{\mathcal{Y}(\theta)}$ is independent of θ.

The methods of this chapter are amenable to any model and cost function pair satisfying the above. However we now describe a common, statistically motivated choice of cost function used in the parameter estimation problem. The methods of this chapter are not restricted to this choice, but it lends insight. First note that $\psi_{\mathcal{Y}(\theta)}(\mathcal{D})$ is numerically equal to the likelihood of θ, given data $\mathcal{D}$, by Definition 2.9. As demonstrated in Section 2.2, we can modify $\psi_{\mathcal{Y}(\theta)}$ to account for E-replicated data, by taking

$$\psi^{[E]}_{\mathcal{Y}(\theta)} := \prod_{i=1}^{E} \psi_{\mathcal{Y}(\theta)}(\mathcal{D}^i),$$

but do not do so in the following discussion, for notational clarity. If we take

$$C_{\mathcal{D}}(\theta) = -\log \psi_{\mathcal{Y}(\theta)}(\mathcal{D}) + R(\mathcal{D}), \tag{4.1a}$$

where $R(\mathcal{D})$ is some offset term, then the cost of parameterised model output is proportional to the *negative log-likelihood* of θ, given the data. The minimiser of θ is

then the *maximum likelihood estimate*.

In the case that data does not exist, but we wish to explore model sensitivity around a nominal parameter vector θ^*, we can instead take

$$C_{\mathcal{Y}(\theta^*)}(\theta) = \mathbb{E}_{\theta^*}[-\log \psi_{\mathcal{Y}(\theta)}] - \mathbb{E}_{\theta^*}[-\log \psi_{\mathcal{Y}(\theta^*)}]. \tag{4.1b}$$

This is the expected value of (4.1a), assuming that data $\mathcal{D}$ was drawn from the experimental model $\mathcal{D}(\theta^*)$, and that $R(\mathcal{D}) = -\mathbb{E}_{\theta^*}[-\log \psi_{\mathcal{Y}(\theta^*)}]$ (so that $C_{\mathcal{Y}(\theta^*)}(\theta^*) = 0$). This choice of $C_{\mathcal{Y}(\theta^*)}(\theta)$ is known as the **Kullback-Leibler (KL) Divergence**, or relative entropy, between the random variables $\mathcal{D}(\theta^*)$ and $\mathcal{D}(\theta)$. Moreover the Cost Hessian $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)$ then corresponds to the **Fisher Information Matrix** at θ^*. Both cost functions of (4.1) can be trivially modified to account for E-replicated data.

Remark 4.1. ***Cost functions and Bayesian probability***

We can modify the choices of cost function given in (4.1) *to incorporate a prior probability distribution* $\mathbb{P}_{prior}(\theta)$ *on parameter space, if one exists. We can do this by replacing the quantity* $\psi_{\mathcal{Y}(\theta)}$ *(the likelihood function), where it enters* (4.1)*, with* $\psi_{\mathcal{Y}(\theta)}\mathbb{P}_{prior}(\theta)$ *(the posterior density). In this case, minimisers of* $C_{\mathcal{D}}(\theta)$ *are no longer maximum likelihood estimates, but instead are* ***maximum a postieriori*** *estimates. However, the links with the Kullback-Leibler divergence and the Fisher information matrix are lost.*

We now take a concrete example to provide context to the previous discussion.

Example 4.1.

Suppose that we have finite-dimensional model output $\mathcal{Y}(\theta)$*. Furthermore, assume that measurement noise is additive, and distributed as a Gaussian random variable.*

Then

$$\mathcal{D}(\theta) \sim \mathcal{N}(\mathcal{Y}(\theta), \Sigma),$$

for some covariance matrix Σ*. This implies that the probability density of* $\mathcal{D}(\theta)$ *is*

$$\psi_{\mathcal{Y}(\theta)}(X) = (2\pi)^{\frac{-n}{2}} |\Sigma|^{\frac{-1}{2}} e^{-\frac{1}{2}(X-\mathcal{Y}(\theta))^T \Sigma^{-1}(X-\mathcal{Y}(\theta))}.$$

In this case, the choices of cost function given in (4.1) *(with respect to* E*-replicated data, for extra clarity) are respectively given by*

$$C_{\mathcal{D}^{[E]}}(\theta) = \sum_{i=1}^{E} \langle \big(\mathcal{D}^i - \mathcal{Y}(\theta)\big), \Sigma^{-1} \big(\mathcal{D}^i - \mathcal{Y}(\theta)\big) \rangle \tag{4.2a}$$

$$C_{\mathcal{Y}^{[E]}(\theta^*)}(\theta) = E\langle \mathcal{Y}(\theta^*) - \mathcal{Y}(\theta), \Sigma^{-1}(\mathcal{Y}(\theta^*) - \mathcal{Y}(\theta)) \rangle. \tag{4.2b}$$

The costs take the form of a weighted sum of the squared residuals between the nominal output $\mathcal{D}^{[E]}$ *or* $\mathcal{Y}(\theta^*)$*, and the model output* $\mathcal{Y}(\theta)$*. Maximum likelihood estimation from data* $\mathcal{D}^{[E]}$*, which equates to minimising the cost function* $C_{\mathcal{D}^{[E]}}(\theta)$*, is then a least squares regression.*

We have assumed that θ^* is a local minimum of the function $C_{\mathcal{Y}(\theta^*)}$. In the case of the specific choice (4.2b) of cost function, this can be seen explicitly. More generally, this holds for any reasonable choice of cost function, since $\mathcal{Y}(\theta^*)$ approximates the output of $\mathcal{Y}(\theta^*)$ (itself) at least as well as any other parameter vector θ. This implies

- $\nabla_\theta C_{\theta^*}(\theta^*) = 0$, as it is the gradient of a function evaluated at a local minimum.
- $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*) \succeq 0$ (it is positive-semidefinite), as it is the Hessian of a function evaluated at a local minimum.

These points motivate the definition of model sloppiness.

Definition 4.1. ***Model sloppiness*** *[21]*

The sloppiness of a parameter θ^* is taken as the condition number of the Cost Hessian $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)$. Since the Hessian matrix is symmetric (by the equality of mixed partial derivatives), this equates to

$$Slop(\theta^*) = \frac{\lambda_{max}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)}{\lambda_{min}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)}, \tag{4.3}$$

where λ_{max} and λ_{min} denote the maximal and minimal magnitude eigenvectors, respectively. If the denominator of (4.3) *is zero, then we say that the model is* ***infinitely sloppy*** *at θ^*. A sufficient condition for infinite sloppiness is unidentifiability. However, locally identifiable models can still be infinitely sloppy, as will be demonstrated in Example 4.4.*

The ***stiffest*** *direction of parameter perturbation is the eigenvector associated with $\lambda_{max}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)$, while the* ***sloppiest*** *direction is that associated with $\lambda_{min}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)$.*

Sloppiness quantifies the anisotropy in output disruption when infinitesimal parameter perturbations are applied to a nominal parameter vector θ^*. Specifically, the degree of sloppiness is the ratio between the degree of disruption in the most and least sensitive directions. Geometrically, it is the aspect ratio of the hyper-ellipse of perturbations $\delta\theta$ satisfying

$$\langle \delta\theta, [\nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)]\delta\theta \rangle = 1.$$

The major and minor axes of this hyper-ellipse respectively represent the sloppiest and stiffest directions in parameter space (see Figure 4.2).

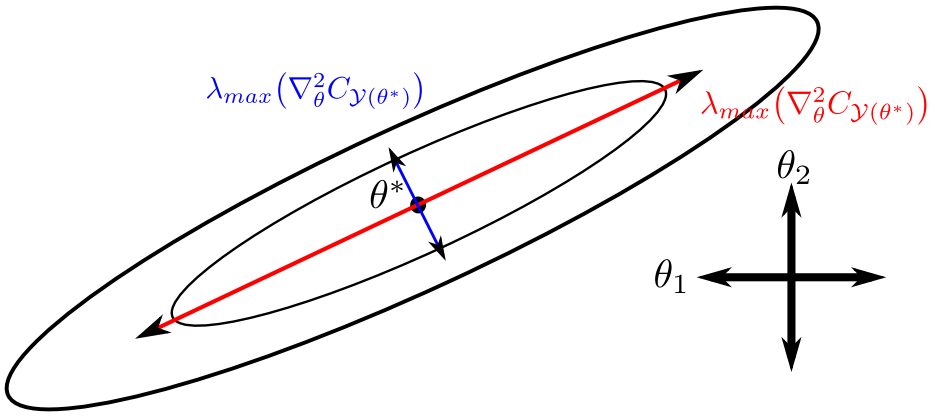

Figure 4.2: [Effect of infinitesimal perturbations on a parameter vector θ^*. *The ellipses represent ϵ-level sets of the cost function for small $\epsilon > 0$. A large perturbation in the direction of the red arrow (the 'sloppiest' eigenvector) is required to induce the same disruptive effect as a small perturbation in the direction of the blue arrow (the 'stiffest' eigenvector). These eigenvectors are not necessarily aligned with the co-ordinate axes, as we see in the diagram.*

Remark 4.2. *In the literature (see e.g. [21]), it is implicitly assumed that measurement noise is additive and Gaussian (with zero mean), as in Example 4.1. In this case, the Cost Hessian is given by*

$$\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*) = \langle \nabla_\theta \mathcal{Y}(\theta^*), \Sigma^{-1}(\nabla_\theta \mathcal{Y}(\theta^*)) \rangle \tag{4.4}$$

and, for this specific choice of Cost Hessian, sloppiness is calculated as in (4.3). *Definition 4.1 is therefore a minor generalisation of the literature.*

Sloppiness is often considered as a fundamental feature of model structure [132, 50], rather than a property of a particular parameter vector. Caution must be exercised in intuiting the former from the latter. We illustrate with a simple, novel example.

Example 4.2. ***Non-dimensionalised, damped, harmonic oscillator***

The differential equation for a damped, harmonic oscillator is given by

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = 0,$$

where $y(t)$ and m respectively denote the oscillator's position and mass, c is the damping coefficient, and k is the spring constant. This constitutes an extension of the trolley-spring model given in Figure 1.2, and equation (1.1) *to include a damping term. If we take the parameter vector as $\theta = [m, c, k, y(0), \dot{y}(0)]$, then the model contains a structural unidentifiability over areas of parameter space preserving the ratios $\frac{c}{m}$ and $\frac{k}{m}$. A common, identifiable reparameterisation of the model is given as*

$$\ddot{y}(t) + 2\zeta\omega_0\dot{y}(t) + \omega_0^2 y(t) = 0. \tag{4.5}$$

Here ω_0 and ζ, known as the natural frequency and damping ratio respectively, are given by $\omega_0 = \sqrt{\frac{k}{m}}$, $\zeta = \frac{c}{2\sqrt{mk}}$. The reparameterised model is now globally structurally identifiable, with parameter vector $\theta = [\zeta, \omega_0, y(0), \dot{y}(0)]$.

Note that the oscillator position, which will be our measured variable, is a function of time. We denote the parameter-dependent trajectory as $y(t;\theta)$. As it is not finite-dimensional, the cost functions given in Example 4.1 are not valid. One solution is to assume that position is measured at a finite number of timepoints. However, sloppiness is then dependent on the time-points chosen, which may be undesirable for some purposes. A more common approach used in the literature (see e.g. [50]) is to take a cost function of the form

$$C_{\mathcal{Y}(\theta^*)}(\theta) = \int_0^\infty \left\langle y(t;\theta) - y(t;\theta^*), \Sigma^{-1}\big(y(t;\theta) - y(t;\theta^*)\big)\right\rangle \, dt, \tag{4.6}$$

which is a continuous-time analogue of the cost given in Example 4.1.

Using the cost function (4.6)*, the degrees of sloppiness at the two nominal parameter*

vectors $\theta^{1,*} = [0.5, 1, 1, 0]$*;* $\theta^{2,*} = [5, 1, 1, 0]$*, are* 94.0 *and* 3.41×10^7 *respectively, a separation of six orders of magnitude. So the model is very sloppy at some parameter vectors, and not at all at others. Note that the oscillator is* underdamped *at* $\theta^{1,*}$*, and* overdamped *at* $\theta^{2,*}$*.*

We briefly mentioned the duality between the parametric sensitivity of model output at a parameter θ^* and the uncertainty associated with parameter estimation when data is taken as a noisy realisation of $\mathcal{D}(\theta^*)$. Our reasoning was that if sensitivity is low along some direction in parameter space, then, by definition, model outputs along this direction wiill be very similar, and thus hard to distinguish between on the basis of noisy data. We can formalise this statement by appealing to the famous **Cramer-Rao inequality** [103].

Lemma 4.1. ***Cramer-Rao Inequality*** *[103, 25]*

Suppose that E-replicated data $\mathcal{D}^{[E]}$ *is generated as a noisy realisation of* $\mathcal{D}(\theta^*)$*, for some* $\theta^* \in \Theta$*. Take the choices of cost function* $C_{\mathcal{Y}(\theta^*)}$ *presented in* (4.1)*. If the Cost Hessian is invertible, then, for any* ***estimator*** $\hat{\theta}(\mathcal{D}^{[E]})$*, which estimates the true value of* θ^* *given the data, the following holds:*

$$\lim_{E\to\infty} cov\left[\hat{\theta}(D^{[E]})\right] \succeq \left[\nabla^2_\theta C_{\mathcal{Y}^{[E]}(\theta^*)}(\theta^*)\right]^{-1}. \tag{4.7}$$

Moreover, equality is attained if we use the maximum likelihood estimator, given as,

$$\hat{\theta}(D^{[E]}) = \min_{\theta\in\Theta} C_{\mathcal{D}^{[E]}}(\theta),$$

or, more generally, any estimator satisfying asymptotic normality*, as defined in equation* (2.4) *of Section 2.2.*

The assumptions placed at the beginning of this section constitute a sufficient condi-

tion for the Cramer-Rao inequality to hold. However, more specific regularity conditions can be found in [25]. Note that the Cost Hessian, whose inverse constitutes the RHS of (4.7) is commonly known as the **Fisher Information Matrix**.

The Cramer-Rao inequality warrants further discussion. First note that the vector θ^* in Lemma 4.1 is unknown. Therefore, so too is the true value of the RHS of (4.7), which is a function of θ^*. It is commonly approximated by replacing θ^* with the estimator $\hat{\theta}(D^{[E]})$.

We also assumed, in Lemma 4.1, that the Cost Hessian was invertible at θ^*. Indeed, the inverse is used in equation (4.7). If there is a structural unidentifiability emanating from θ^*, then, by definition, model output is invariant over the unidentifiability. Assumption A2 then implies that $C_{\mathcal{Y}(\theta^*)}(\theta) = 0$ over the unidentifiability, which means that any tangent vector $\delta\theta$, evaluated at θ^* and directed along the unidentifiability, must satisfy

$$\langle \delta\theta, \nabla_\theta^2 C_{\mathcal{Y}^{[E]}(\theta^*)}(\theta^*)\delta\theta \rangle = 0.$$

This implies rank deficiency of the Cost Hessian, and failure of the Cramer-Rao inequality to hold. It also implies infinite sloppiness.

Finally, the Cramer-Rao inequality is an asymptotic result, valid in the limit of increasing experimental replicates. This corresponds to the limit of decreasing measurement uncertainty, as drawing limitless numbers of samples from $\mathcal{D}(\theta^*)$ allows us to perfectly recreate its distribution, and thereby infer the form of the noise-free data $\mathcal{Y}(\theta^*)$. Therefore model sloppiness, and the Cost Hessian, only meaningfully determine the uncertainty associated with parameter estimation when that uncertainty is already infinitesimal.

We now present a novel framework for parameter estimation in the context of non-

infinitesimal measurement uncertainty, which is valid regardless of parameter identifiability. Let us define the ϵ-uncertainty regions associated with a parameter estimate θ^*, or data $\mathcal{D}$ as

$$U_{\mathcal{D}}(\epsilon) = \{\theta : C_{\mathcal{D}}(\theta) \leq \epsilon\} \tag{4.8a}$$

$$U_{\mathcal{Y}(\theta^*)}(\epsilon) = \{\theta : C_{\mathcal{Y}(\theta^*)}(\theta) \leq \epsilon\}.. \tag{4.8b}$$

A definition of (4.8b), which is specialised to the case of additive Gaussian measurement noise presented in Example 4.1, is given in [127]. More generally if we take the choices of cost function given in (4.1), then ϵ-uncertainty regions correspond to **likelihood-based confidence regions** [25].

Intuitively, an ϵ-uncertainty region captures the set of parameters that approximate the reference output 'reasonably' well, where ϵ quantifies what is reasonable. As with model sloppiness, it is based on the sensitivity properties of model output in parameter space. Some extra work is therefore required to directly link the properties of the ϵ-uncertainty region with the accuracy of parameter estimation from noisy data. Our first step is to outline the correspondence between the ϵ-uncertainty region and the Cost Hessian, in the limit of decreasing ϵ.

Lemma 4.2. ***Limiting behaviour of*** $\mathbf{U}_{\mathcal{Y}(\theta^*)}(\epsilon)$

Given a globally structurally identifiable parameter θ^*, and a compact parameter space Θ, consider the volume-normalised version of $U_{\mathcal{Y}(\theta^*)}(\epsilon)$ given by

$$\tilde{U}_{\mathcal{Y}(\theta^*)}(\epsilon) = \left\{\theta^* + \frac{\delta\theta}{\sqrt{2\epsilon}} : \theta^* + \delta\theta \in U_{\mathcal{Y}(\theta^*)}(\epsilon)\right\}. \tag{4.9}$$

Then

$$\lim_{\epsilon \to 0} \tilde{U}_{\mathcal{Y}(\theta^*)}(\epsilon) = \theta^* + \left\{ \delta\theta : \langle \delta\theta, \nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*) \delta\theta \rangle \leq 1 \right\}.$$

Proof. Since θ^* is structurally identifiable, we know that $U_{\theta^*}(0) = \{\theta^*\}$. Our first step is to show, for sufficiently small ϵ, that $U_{\theta^*}(\epsilon)$ is a connected set. If not, then there exists a sequence $\{\theta^{(i)}\}_{i=1}^{\infty}$ such that, for all $i \in \mathbb{N}$:

$$\exists \kappa > 0 : \|\theta^{(i)} - \theta^*\|_2^2 > \kappa \text{ and } C_{\theta^*}(\theta^{(i)}) < \frac{1}{i}.$$

The Extreme Value Theorem [117] then guarantees the existence of a limiting parameter vector $\theta^{(\infty)} \in \Theta$ satisfying

$$\exists \kappa > 0 : \|\theta^{(\infty)} - \theta^*\|_2^2 > \kappa \text{ and } C_{\theta^*}(\theta^{(\infty)}) = 0,$$

contradicting structural identifiability.

Since $U_{\theta^*}(\epsilon)$ is connected for sufficiently small ϵ, and $C_{\mathcal{Y}(\theta^*)}$ is continuous, we know that

$$\lim_{\epsilon \to 0} \left[\max_{\delta\theta \in \Theta} \|\delta\theta\|_2^2 : C_{\theta^*}(\theta^* + \delta\theta) < \epsilon \right] = 0.$$

Next, taking a Taylor expansion of (4.8b), and using the definition of $\tilde{U}_{\mathcal{Y}(\theta^*)}$ from (4.9), we see that

$$\tilde{U}_{\mathcal{Y}(\theta^*)}(\epsilon) = \theta^* + \left\{ \delta\theta : \left(\frac{(\sqrt{2\epsilon})^2}{2} \langle \delta\theta, [\nabla_\theta^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)] \delta\theta \rangle \right) + \mathcal{O}\left(\|2\sqrt{\epsilon}\delta\theta\|_2^3 \right) \leq \epsilon \right\}.$$

We can divide the above equation through by ϵ, and note, by definition, that

$$\lim_{\epsilon \to 0} \frac{\mathcal{O}\left(\|2\sqrt{\epsilon}\delta\theta\|_2^3\right)}{\epsilon} = 0,$$

to get

$$\lim_{\epsilon \to 0} \tilde{U}_{\mathcal{Y}(\theta^*)}(\epsilon) = \theta^* + \left\{\delta\theta : \left(\langle \delta\theta, [\nabla_\theta^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)]\delta\theta\rangle\right) \leq 1\right\}$$

as required. □

Note that while we assumed global structural identifiability of θ^*, the theorem can easily be tailored to account for local structural identifiability. We need only to restrict parameter space Θ to a small enough region around θ^* that global structural unidentifiability is attained. Since the theorem quantifies the local geometry of the uncertainty region, and there is necessarily a minimum distance between a locally identifiable θ^* and its mutually unidentifiable parameter vectors, the proof is not affected. A graphical illustration of Lemma 4.2 is provided in Figure 4.3.

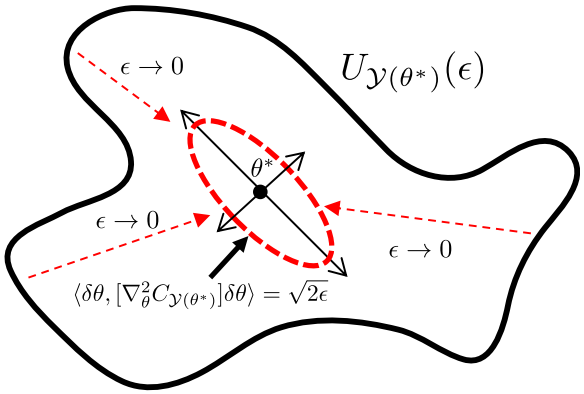

Figure 4.3: Limiting behaviour of $U_{\mathcal{Y}(\theta^*)}(\epsilon)$. *As ϵ tends to zero, the shape of the ϵ-uncertainty region approximates level sets of the Cost Hessian.*

In summary, the Fisher Information Matrix (which is the Cost Hessian for a particular choice of cost function) predicts the uncertainty associated with a parameter estimate

made from noisy data, in the limit of increasing experimental data replicates. This was shown in Lemma 4.1. The ϵ-uncertainty region gives the same prediction, in the limit of decreasing ϵ, through Lemma 4.2. However the discussion has assumed both local structural identifiability and infinitesimal measurement uncertainty (i.e. unbounded experimental replicates). We next demonstrate that the formalism of ϵ-uncertainty regions allows for the quantification of uncertainty in a parameter estimate even when measurement noise is significant. To do so, we first summarise some well-known results on confidence testing in both the Bayesian and Frequentist frameworks. Recall from Section 2.2 that an n-percent (Bayesian) credible region is an area of parameter space satisfying

$$\frac{\int_{\mathrm{Cred}_n(\mathcal{D}^{[E]})} \mathbb{P}_{post}(\theta; \mathcal{D}^{[E]})\ d\theta}{\int_{\Theta} \mathbb{P}_{post}(\theta; \mathcal{D}^{[E]})\ d\theta} = \frac{n}{100},$$

where $\mathbb{P}_{post}(\theta; \mathcal{D}^{[E]})$ is the (possibly un-normalised) posterior distribution on parameter space, given $\mathcal{D}^{[E]}$. So if one drew parameter vectors randomly from the posterior distribution, the samples would lie within the credible region n times out of a hundred. Of course credible regions are not unique, even for fixed n. However, the smaller such a region is, the better: typically one wants to rule out as large a volume of parameter space as possible. This motivates the concept of a **highest posterior density** credible region. This is a minimiser of the volume

$$\int_{\mathrm{Cred}_n(\mathcal{D}^{[E]})} 1\ d\theta,$$

among the space of n-percent credible regions. As long as the likelihood function is continuous in both data and parameters, a highest posterior density can therefore be formulated as

$$\mathrm{Cred}_n(\mathcal{D}^{[E]}) = \left\{\theta \in \Theta : \mathbb{P}_{post}(\theta; D^{[E]}) > \gamma\right\},$$

for some $\gamma \in \mathbb{R}^+$. If we take $C_{D^{[E]}}(\theta) \propto -\log\left[\mathbb{P}_{post}(\theta; D^{[E]})\right]$, then we see that a highest posterior density region corresponds to an ϵ-uncertainty region. Indeed, if the prior distribution is uniform on parameter space, then

$$\mathbb{P}_{post}(\theta; D^{[E]}) = \psi^{[E]}_{\mathcal{Y}(\theta)}(D^{[E]}),$$

and so the cost $C_{D^{[E]}}(\theta)$ given in (4.1a) is just such a choice. Meanwhile the choice $C_{\mathcal{Y}(\theta)}$ given in (4.1b) is proportional to the expected value of the choice (4.1a) of $C_{D^{[E]}}(\theta)$, assuming that data $\mathcal{D}$ is generated according to the experimental model $\mathcal{D}(\theta)$. In this case, $U_{\mathcal{Y}(\theta^*)}(\epsilon)$ contains the parameter vectors *expected* to be in the highest posterior density region.

We now provide a Frequentist interpretation of ϵ-uncertainty regions. To do so we first introduce some basic instruments of hypothesis testing.

Definition 4.2. ***Likelihood ratio test***

Suppose we have

1. *A null hypothesis $\mathcal{H}_0$, specifying that an unknown parameter vector θ^* belongs to some set Θ_0.*
2. *An alternative hypothesis $\mathcal{H}_1$, specifying that θ^* belongs to a set Θ_0^c, the complement of Θ_0 in Θ.*
3. *A likelihood function $\mathcal{L}(\theta; \mathcal{D}^{[E]})$ that assigns to each parameter vector a likelihood based on data $\mathcal{D}^{[E]}$.*

Then any statistical test that rejects $\mathcal{H}_0$ (and, equivalently, accepts $\mathcal{H}_1$) based on a

criterion of the form

$$\Lambda(\mathcal{D}^{[E]}) = \frac{\sup_{\theta\in\Theta_0} \mathcal{L}(\theta; \mathcal{D}^{[E]})}{\sup_{\theta\in\Theta} \mathcal{L}(\theta; \mathcal{D}^{[E]})} \leq k^*,$$

for some critical value $k^* \in [0, 1]$*, is known as a likelihood ratio test. The function* $\Lambda(\mathcal{D}^{[E]})$ *is known as the* ***likelihood ratio test statistic****. The significance level of a likelihood ratio test is given by* α*, where*

$$\mathbb{P}[\Lambda(\mathcal{D}^{[E]}) \leq k^* | \mathcal{H}_0] = \alpha. \tag{4.10}$$

Note that the significance level α corresponds to the probability of rejecting $\mathcal{H}_0$ when it is in fact true. The likelihood ratio test statistic can be formulated in terms of our cost function $C_{\mathcal{D}}(\theta)$ given in (4.1a), since, by definition

$$\mathcal{L}(\theta; \mathcal{D}^{[E]}) = \psi^{[E]}_{\mathcal{Y}(\theta)}(\mathcal{D}^{[E]}) = \prod_{i=1}^{E} \psi_{\mathcal{Y}(\theta)}(\mathcal{D}^i).$$

We get

$$\log\left(\Lambda(\mathcal{D}^{[E]})\right) = \inf_{\theta\in\Theta} C_{\mathcal{D}^{[E]}}(\theta) - \inf_{\theta\in\Theta_0} C_{\mathcal{D}^{[E]}}(\theta).$$

Frequentist confidence regions (see Section 2.2) can be constructed using the likelihood ratio test. Suppose the null hypothesis $\mathcal{H}_0$ is that data is generated by the experimental model $\mathcal{D}(\hat{\theta})$, for some fixed $\hat{\theta} \in \Theta$. In the terminology of Definition 4.2, $\Theta_0 = \hat{\theta}$. The alternative hypothesis $\mathcal{H}_1$ is then that data is generated by some $\theta \in \Theta$ such that $\theta \neq \hat{\theta}$. The appropriate likelihood ratio test statistic is given by

$$\Lambda(\mathcal{D}^{[E]}) = \frac{\mathcal{L}(\hat{\theta}; \mathcal{D}^{[E]})}{\sup_{\theta\in\Theta} \mathcal{L}(\theta; \mathcal{D}^{[E]})}, \tag{4.11}$$

and we reject the null hypothesis when this falls below a critical value k^*. Given some data $\mathcal{D}^{[E]}$, the set of parameters $\theta \in \Theta$ that would have been accepted as a null hypothesis is then given by

$$\big\{\hat{\theta} \in \Theta : C_{\mathcal{D}^{[E]}}(\hat{\theta}) \leq \inf_{\theta \in \Theta} C_{\mathcal{D}^{[E]}}(\theta) - \log(k^*)\big\}. \tag{4.12a}$$

$$= \; U_{\mathcal{D}}(\epsilon), \text{ for } \epsilon = \inf_{\theta \in \Theta} C_{\mathcal{D}^{[E]}}(\theta) - \log(k^*). \tag{4.12b}$$

Note that the probability of the true parameter not being in this region is given by (4.10) as α. So (4.12b) is an 100α-percent confidence region, obtained through *inversion* of the likelihood ratio test.

We can evaluate model sensitivity relative to some parameter θ^*, rather than to data $D^{[E]}$, by slight modification of the previous arguments. We can suppose that hypothetical data $\mathcal{D}$ was distributed as a realisation of the random variable $\mathcal{D}(\theta^*)$, and consider the expected behaviour of the likelihood ratio test statistic under this assumption. Given a null hypothesis $\hat{\theta} = \theta^*$, the expected value of (4.11) is given by

$$\mathbb{E}_{\theta^*}\big[\Lambda(\mathcal{D}^{[E]})\big] = \mathbb{E}_{\theta^*}\left[\frac{\mathcal{L}(\hat{\theta}; \mathcal{D}^{[E]})}{\sup_{\theta \in \Theta} \mathcal{L}(\theta; \mathcal{D}^{[E]})}\right].$$

Using the choice (4.1b) of $C_{\mathcal{Y}(\theta^*)}$ and linearity of the expectation operator, this can equivalently be stated as

$$\mathbb{E}_{\theta^*}\big[\log\big(\Lambda(\mathcal{D}^{[E]})\big)\big] = \mathbb{E}_{\theta^*}\big[\inf_{\theta \in \Theta} C_{\mathcal{D}^{[E]}}(\theta)\big] - C_{\mathcal{Y}^{[E]}(\theta^*)}(\hat{\theta}).$$

The set of parameters that would be in confidence region (4.12b) in the case that

$\Lambda(\mathcal{D}^{[E]})$ obtained its expected value, is then given by

$$\begin{aligned}&\Big\{\hat{\theta} \in \Theta : C_{\mathcal{Y}^{[E]}(\theta^*)}(\hat{\theta}) \leq \mathbb{E}_{\theta^*}\big[\inf_{\theta\in\Theta} C_{\mathcal{D}^{[E]}}(\theta)\big] - \log(k^*)\Big\}\\ =\ & U_{\mathcal{Y}^{[E]}(\theta^*)}(\epsilon),\ \text{for } \epsilon = \mathbb{E}_{\theta^*}\big[\inf_{\theta\in\Theta} C_{\mathcal{D}^{[E]}}(\theta)\big] - \log(k^*).\end{aligned}$$

We have now established the relationship between the ϵ-uncertainty regions formulated in this chapter, and standard, statistical confidence regions. We have shown that the value ϵ is related to the significance level of the confidence region through a constant k^*. It remains to describe how a k^* satisfying (4.10) can be found.

First note that k^* is monotonically increasing with respect to the significance level α. This is a clear consequence of (4.10): increasing k^* means that the set of likelihood ratio statistic values upper-bounded by k^* expands, and consequently has greater (or equal) probability measure. So as we decrease ϵ (consequently increasing k^* through the relation (4.12b)) , the quantity $U_{\mathcal{D}}(\epsilon)$ represents an increasingly significant confidence region. When k^* takes the values zero and one respectively, so too must α. Unfortunately, numerically relating k^* to intermediate values of α is a hard problem in the general case. One asymptotic approximation is given by the following theorem.

Theorem 4.1. ***Wilks' Theorem*** *[134]*

Consider a model that is locally structurally identifiable over a set Θ_0 of parameters encoding the null hypothesis. Let us denote by v the difference in dimensionality between Θ_0 and Θ_0^c. Suppose that noisy data is generated according to $\mathcal{D}(\theta^*)$, in such a way that the maximum likelihood estimator is consistent and asymptotically normal (definitions provided in (2.4)*). Then, as the number of experimental replicates E tend*

to infinity, the following convergence in distribution holds:

$$-2\log\left(\Lambda(D^{[E]})\right) \to_D \chi^2_v,$$

where χ^2_v denotes the chi-squared distribution with v degrees of freedom.

Wilks' theorem shows that the probability distribution informing (4.10) is approximable by a chi-squared distribution, which has known cumulative distribution. This allows us to form approximate values of k^*, for a desired significance level. As with the Cramer-Rao inequality (Lemma 4.1), the approximation is only valid in the limit of increasing experimental replicates, and for locally structurally identifiable models. Note, however, that the ϵ-uncertainty regions are exact characterisations of likelihood ratio test-based confidence regions, and we are only approximating the conversion from ϵ to the significance level of the confidence region. Reasoned lists of the advantages possessed by the approximation based on Wilks' Theorem, as opposed to other available methods for setting the significance level, are provided in [76, 81].

4.3 Multiscale Sloppiness

One of the topics covered in the previous section was model sloppiness, which quantifies the anisotropy of model sensitivity to infinitesimal parameter perturbations at a fixed parameter vector. We now provide a new notion of multiscale sloppiness. While sloppiness possesses a duality with the uncertainty of parameter estimation when data is corrupted with infinitesimal measurement uncertainty, multiscale sloppiness possesses this duality where measurement uncertainty is significant. In the limit of decreasing measurement uncertainty, multiscale sloppiness and sloppiness are theoretically guaranteed to correspond. However, we see through example that the multiscale sloppiness characteristics of even simple linear models are highly dependent

on the degree of measurement uncertainty considered.

Definition 4.3. ***Maximally and minimally disruptive parameters***

We respectively define the maximally and minimally disruptive parameters, relative to some θ^* and at length-scale δ, as

$$B_\bullet^{max}(\delta) = \arg\max_{\theta\in\Theta} C_\bullet(\theta) : \|\theta - \theta^*\|_2^2 \leq \delta \tag{4.13}$$

$$B_\bullet^{min}(\delta) = \arg\min_{\theta\in\Theta} C_\bullet(\theta) : \|\theta - \theta^*\|_2^2 = \delta. \tag{4.14}$$

Note that (4.13) and (4.14) are set-valued functions: the optimisers may not be unique. We see from the definition that the extremally disruptive parameters can be formulated with respect to experimental data $\mathcal{D}$. Nevertheless, they still require some fixed θ^*, which could be the maximum likelihood estimate of the data.

The maximally and minimally disruptive parameters represent the most and least disruptive perturbations that can be applied to θ^*, among the set of perturbations of fixed magnitude δ. The dual, estimation-based interpretation follows naturally. Let us take $\bar{U}_\bullet(\epsilon)$ as the connected component of $U_\bullet(\epsilon)$ around θ^*. Choose δ_1 and δ_2 such that $C_\bullet(B_{\theta^*}^{min}(\delta_1)) = \epsilon$, and $C_\bullet(B_{\theta^*}^{max}(\delta_2)) = \epsilon$. Then

$$B_\bullet^{min}(\delta_1) = \arg\max_{\theta} \|\theta - \theta^*\|_2^2 : \theta \in \bar{U}_\bullet(\epsilon)$$

$$B_\bullet^{max}(\delta^2) = \arg\min_{\theta} \|\theta - \theta^*\|_2^2 : \theta \notin \bar{U}_\bullet(\epsilon).$$

Thus the maximally and minimally disruptive parameters respectively signify the furthest parameter vectors inside, and the closest parameter vectors outside, the connected component of an ϵ-uncertainty region. This is illustrated in Figure 4.4.

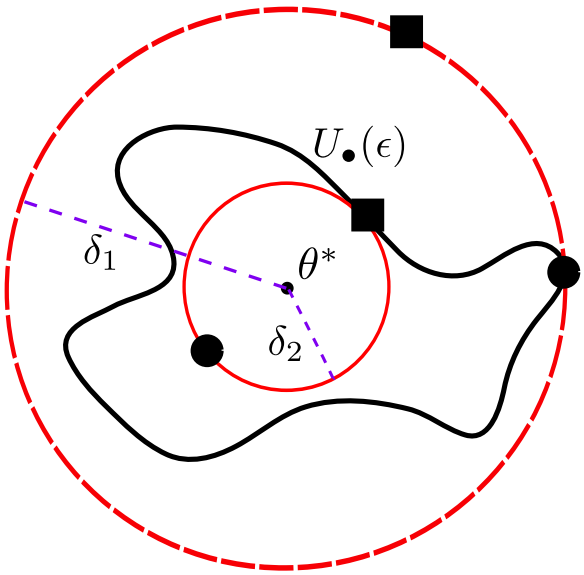

Figure 4.4: Extremally disruptive parameters and the ϵ-uncertainty region. *We pick balls of radius δ_1 and δ_2 (dashed) such that $\mathcal{S}_\bullet^{min}(\delta_1) = \mathcal{S}_\bullet^{max}(\delta_2)) = \epsilon$. Square and circular dots respectively indicate locations of $D_\bullet^{max}(\delta_i)$ and $D_\bullet^{min}(\delta_i)$. The balls intersect the uncertainty region $U_\bullet(\epsilon)$ at $D_\bullet^{max}(\delta_1)$ and $D_\bullet^{min}(\delta_2)$.*

Definition 4.4. ***Multiscale sloppiness***

We define the multiscale sloppiness of a parameter, at length scale δ, relative to θ^*, as

$$\mathcal{S}_\bullet(\delta) := \frac{\mathcal{S}_\bullet^{max}(\delta)}{\mathcal{S}_\bullet^{min}(\delta)} := \frac{\max_{\theta\in\Theta} C_\bullet(\theta) : \|\theta - \theta^*\|_2^2 \leq \delta}{\min_{\theta\in\Theta} C_\bullet(\theta) : \|\theta - \theta^*\|_2^2 = \delta}. \tag{4.15a}$$

$$\mathcal{S}_\bullet(0) := \lim_{\delta\to 0} \mathcal{S}_\bullet(\delta). \tag{4.15b}$$

We see that the multiscale sloppiness of a parameter represents the ratio of the costs of the maximally and minimally disruptive parameters. This corresponds (in the absence of data) to traditional sloppiness as the length scale δ of parameter perturbation tends to zero, as shown in the following Lemma.

Lemma 4.3. ***The local properties of multiscale sloppiness***

We claim that

$$\lim_{\delta \to 0} \mathcal{S}_{\mathcal{Y}(\theta^*)}(\delta) = \frac{\lambda_{max}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)}{\lambda_{min}\big(\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*)\big)}.$$

Proof. Since $C_{\mathcal{Y}(\theta^*)}(\theta^*)$ is both zero and locally minimal, we have both that $\nabla_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*) = 0$, and that $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta^*) \succeq 0$. If we consider a perturbation $\delta\theta$, the remainder theorem ensures

$$C_{\mathcal{Y}(\theta^*)}(\theta^* + \delta\theta) = \langle \delta\theta, \nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)\delta\theta\rangle + \mathcal{O}(\|\delta\theta\|_2^3). \tag{4.16}$$

Therefore, for $\|\delta\theta\|_2^2 \leq \delta$, we have

$$\lim_{\delta \to 0} \frac{\mathcal{S}^{max}_{\theta^*}(\delta)}{\mathcal{S}^{min}_{\theta^*}(\delta)} = \lim_{\delta \to 0} \frac{\max_{\delta\theta}\langle \delta\theta, \nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)\delta\theta\rangle}{\min_{\delta\theta}\langle \delta\theta, \nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)\delta\theta\rangle} : \tag{4.17}$$

Clearly the numerator and denominator of (4.17) must respectively lie in the eigenspace of the maximal and minimal eigenvectors of $\nabla^2 C_{\mathcal{Y}(\theta^*)}(\theta^*)$, as the highest order Taylor term dominates for sufficiently small δ. □

There is no analogous relationship between $\mathcal{S}_{\mathcal{D}}(\delta)$, the data-dependent variant, and the traditional degree of sloppiness, as the latter cannot account for experimental data.

Multiscale sloppiness can be highly dependent on δ, as we demonstrate by example.

Example 4.3. ***The damped, harmonic oscillator revisited***

Recall the damped, harmonic oscillator of Example 4.2. In the example, we took a parameter $\theta^{*,1} = [0.5, 1, 1, 0]$*, and found its degree of sloppiness to be* 94.0*, according to the cost function of equation* (4.6)*. Figure 4.5 depicts the multiscale sloppiness of*

$\theta^{*,1}$, as a function of the length-scale δ, according to the same cost function. This was calculable using the results of Section 3.3 of Chapter 3, which allowed us to formulate the cost function (4.6) *as a rational function in the parameters. We see that multiscale sloppiness increases exponentially with the length-scale of parameter perturbation δ. In fact, at $\delta = 0.25$, an asymptote is reached: the maximally disruptive parameter $\theta = [0, 1, 1, 0]$ corresponds to an undamped system, as the first parameter component represents the damping ratio. In this case, the cost function of* (4.6) *becomes infinite, as the undamped system oscillates forever, while the underdamped system induced by $\theta^{*,1}$ reaches equilibrium in finite time.*

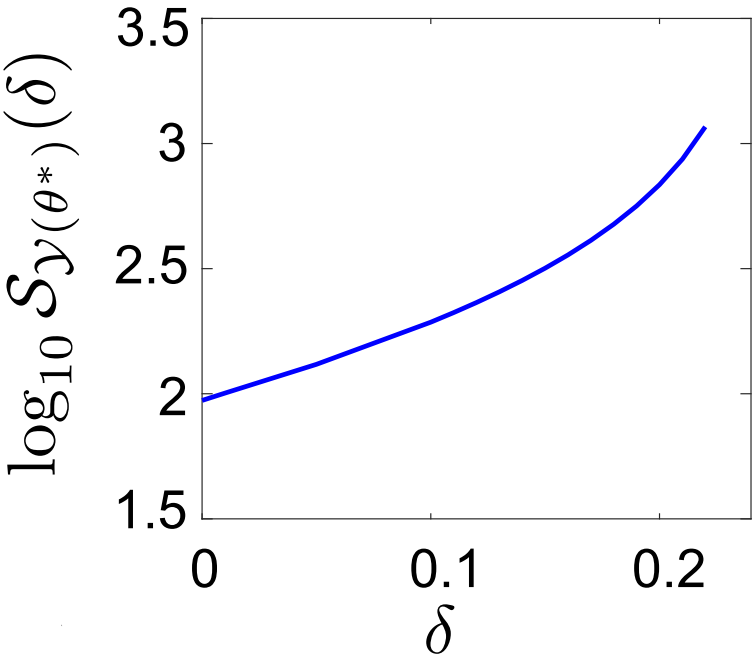

Figure 4.5: Multiscale sloppiness of the damped, harmonic oscillator. *As the length scale of parameter perturbation increases linearly, the multiscale sloppiness increases exponentially. An asymptote necessarily exists (although it is not depicted) at $\delta =$* 0.25.

Suppose that the parameter vector $\theta^{*,1}$ given in Example 4.3 was synthesised as a parameter estimate from noisy data (with non-infinitesimal measurement noise). Our analysis of the multiscale sloppiness reveals that the (traditional) degree of sloppiness at $\theta^{*,1}$ would have highly underestimated the aspect ratio of the uncertainty region.

Our next example is also extremely simple, in that it is also a linear model with linear parameter dependence. However, we find parameter vectors at which the multiscale sloppiness characteristics display several important traits.

Example 4.4. ***Multiscale sloppiness of a linear system with linear parameter dependence***

Consider the model

$$\begin{bmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \end{bmatrix} = \begin{bmatrix} -1 & \phi_1 \\ \phi_1 + \phi_2 & \phi_2 - 4 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} \tag{4.18a}$$

$$y(t) = x_1(t) + x_2(t) \tag{4.18b}$$

$$\theta = [\phi_1, \phi_2, x_1(0), x_2(0)]. \tag{4.18c}$$

We take three choice of nominal parameter:

$$\theta^{*,1} = [0, 0, 1, 2], \qquad \theta^{*,2} = [0.001, 0, 1, 2], \qquad \theta^{*,3} = [-0.2, -0.2, 1, 2].$$

This model is globally structurally identifiable, for all parameter vectors with non-zero initial conditions. This can be verified using the techniques of [10]. Specifically, one can take a Laplace transform of the differential equations, and then algebraically invalidate non-uniqueness of the resultant algebraic equations in the frequency domain. We use the cost function (4.6) *provided in Example 4.2, with the same justification. As before, we can use the results of Section 3.3 of Chapter 3 to express this cost as a rational function of the parameters.*

A graph of multiscale sloppiness against length scale for this model is given in Figure 4.6, for the three parameter vectors $\theta^{*,i}$. The parameter vector $\theta^{*,1}$ is globally struc-

turally identifiable (through the method of [10]), yet 'infinitely' sloppy, as the Cost Hessian $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta)$ is rank deficient. However as we increase the length-scale of parameter perturbation considered, the relative sloppiness decreases at a faster than exponential rate. The parameter vector $\theta^{*,2}$ is a small perturbation of the first, and would thus be expected to have similar sensitivity characteristics. The multiscale sloppiness of $\theta^{*,2}$ is indeed very similar to $\theta^{*,1}$, for nonzero length-scale δ, as we see in the figure. However at $\delta = 0$, where multiscale sloppiness corresponds to the traditional definition of sloppiness, there is a qualitative difference, in that $\theta^{*,2}$ has a non-singular Cost Hessian with 'finite' sloppiness.

The parameter vector $\theta^{*,3}$ is much less (traditionally) sloppy, in that the Cost Hessian is well conditioned compared to $\theta^{*,1}$ and $\theta^{*,2}$. However, as the length-scale δ grows, the multiscale sloppiness increases and eventually overtakes that of the other two parameter vectors considered. So if data were generated from $\theta^{*,3}$ with enough measurement noise, the confidence region of the consequent parameter estimate would be expected to be more anisotropic than for $\theta^{*,1}$ and $\theta^{*,2}$, despite its lower level of sloppiness.

The structure of the model in Example 4.4 was very simple, in that it was linear in both states and parameters. Nevertheless, the multiscale sloppiness was highly variable between different parameter vectors and over different length scales, as seen in Figure 4.6. This demonstrates the degree to which parametric sensitivity characteristics are affected by the length-scale of parameter perturbation. Of course, analysing multiscale sloppiness by hand, as was done in the previous example, is not possible with models of realistic complexity. This motivates the next chapter, which constructs an algorithm for probing the non-local sensitivity of large-scale models.

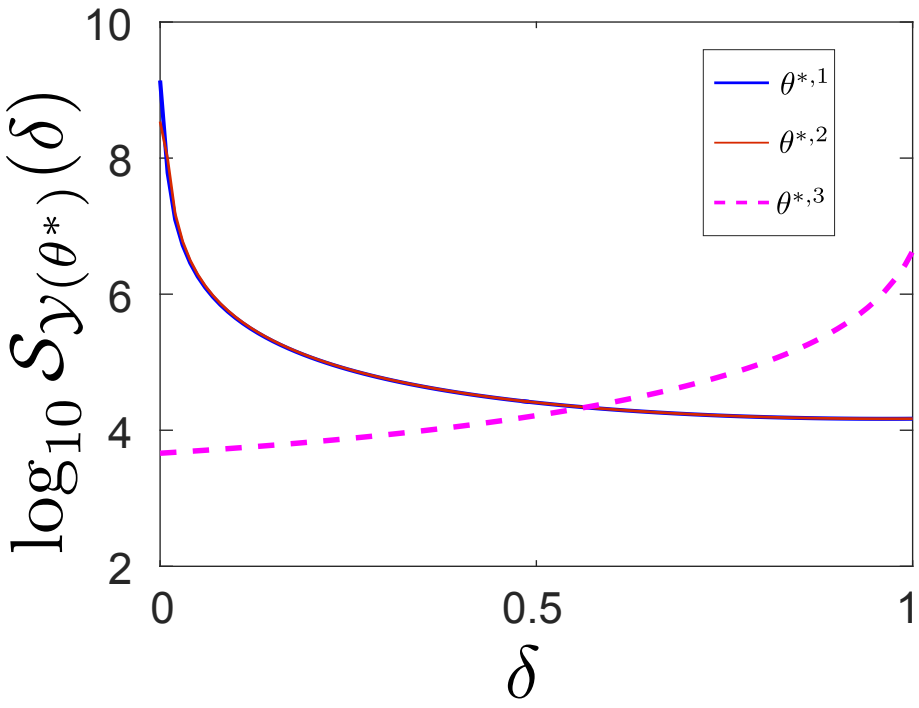

Figure 4.6: Multiscale sloppiness of a linear system. *Graphs of* $\mathcal{S}_{\theta^*}(\delta)$ *against* δ *for the model of* (4.18b). *The right-hand figure has three choices of nominal parameter vector:* $\theta^{*,1} = [0, 0, 1, 2]$, $\theta^{*,2} = [0.001, 0, 1, 2]$, *and* $\theta^{*,3} = [-0.2, -0.2, 1, 2]$. *Not shown on the right-hand graph is that for* $\theta^{*,1}$, *we have* $\mathcal{S}_{\theta^*}(0) = \infty$.

4.4 Summary

In this chapter, we have developed a framework for the quantification of the parametric sensitivity of models to nonlocal parameter perturbation. Theoretical links were provided between this sensitivity and the uncertainty associated with parameter estimation from noisy data. We have demonstrated that parametric sensitivity can be highly affected by the degree of nonlocality in the perturbation. Thus local quantifications of model sensitivity, such as model sloppiness, may be insufficient to answer questions on the fidelity of parameter estimation from noisy data when the noise is not infinitesimally small. We have provided a nonlocal alternative.

Chapter 5

Delineating Parameter Unidentifiabilities through Optimal Control

This chapter describes a scalable method for detecting (structural and practical) unidentifiabilities in parameter space, for generic models. This allows for model simplification, and appreciation of which parameters (or functions thereof) cannot be estimated from data. We use methods derived from the theory of optimal control to evolve curves in parameter space over which model output varies little. Theoretical guarantees on optimality of the generated curves is provided. Our algorithm can identify features such as redundant mechanisms and fast timescale subsystems, as well as the regimes in parameter space over which such approximations are valid. We base our algorithm on the theoretical framework of multiscale sloppiness introduced in the previous chapter. The approach is automatic, in that no prior analysis of model unidentifiabilities is required. We finally apply our methods to several models, including a large-scale, benchmark Systems Biology model of NF-κB, in which we uncover

previously unknown unidentifiabilities. The results of this chapter are published in [100], and parts of [38].

5.1 Motivation

The previous chapter established a duality between the non-local parametric sensitivity characteristics of models, and the uncertainty associated with estimating model parameters from noisy data. This chapter is chiefly concerned with computational methods for calculating non-local parametric sensitivity. In particular, we provide methods for uncovering functional relations between parameters over which model output varies little or not at all. Recall from the the preliminary chapter (Definition 2.8) that a functional relation over which model output is (almost) completely invariant is denoted a (practical) structural unidentifiability. As previously discussed, parameter estimation is theoretically impossible along structural unidentifiabilities, and practically impossible over practical unidentifiabilities. Therefore, it is useful to know their functional form as, given a parameter estimate, we can then precisely specify alternative parameter vectors fitting the data equally well. Knowledge of the structural unidentifiabilities in a model additionally provides mechanistic insight into the modelled process, and paves the way for model simplification, as demonstrated on the trolley-on-a-spring model (Figure 1.2). However, even finding the structurally unidentifiable parameters (i.e. those involved in the functional relation defining the structural unidentifiability) is a hard problem.

The determination of structurally unidentifiable parameters in various classes of models has been an ongoing research topic for several decades. Parameter redundancy techniques, as developed in [26, 33] exactly determine the number of structural unidentifiabilities in a model, as well as potentially finding identifiable reparameterisations.

However, they require symbolic calculation of the derivative matrix of model output with respect to parameters, and are thus unsuitable for differential equations models than cannot be solved analytically. For the latter class of model, multiple algebraic approaches that directly analyse the system equations, have been suggested, e.g. [54, 96, 126, 72, 138, 78, 8]. Several such approaches rely on the fact that the identifiability problem can be viewed as a special case of the observability problem [54], where parameters are considered as time-invariant states to be estimated. Research on observability in nonlinear systems often makes use of the Observability Rank Condition [54], which analyses the properties (specifically the rank) of a matrix of Lie derivatives of model output over time to determine unobservable model states [66, 69, 133]. Correspondingly, identifiability has been analysed using Differential Geometry-based approaches that exploit parametric symmetries inherent in the aforementioned matrix of Lie derivatives [126, 74, 43, 36, 37, 139]. These make use of the local state isomorphism theorem [116] to verify global structural identifiability through the solution of a partial differential equation, which can be simplified in various cases. This is known as the similarity transformation approach. Many of the approaches mentioned can potentially determine the exact functional form of the unidentifiabilities [27, 41, 78, 33, 77], but in general all suffer from a lack of general applicability and scalability, see e.g. [31]. The aforementioned issue has motivated the implementation of numerical approaches based on model simulation at multiple points in parameter space, e.g. [4, 112]. Such approaches can be applied to much larger models, but cannot deal with general nonlinearities, are not guaranteed to find all unidentifiabilities, and are not suitable for global analyses.

More recently, data-based approaches to identifiability analysis have been developed. The profile likelihood method [105] is a scalable numerical method that detects both unidentifiable parameters and associated unidentifiabilities, and has found widespread popularity in the Systems Biology community. It is also able to detect practical

unidentifiabilities [127], in which large parameter perturbations induce small but nonzero changes in model output. However, it relies on iteratively moving individual parameters, while each time re-optimising model output over the other parameters. As such, a separate analysis is required for each parameter, and only one-dimensional unidentifiabilities are detected. The approach of [52] can flag the probable existence of more general unidentifiabilities through extensive model simulation over parameter space. However unidentifiabilities are provided only as probable functionally related groups of parameters; the functional relations themselves are not determined.

The algorithm presented in this chapter uncovers both structural and practical unidentifiabilities (along with their associated functional forms). Moreover, it is highly scalable, as shown subsequently by example. The scalability stems from the fact that the algorithm does not require optimisation of model output over parameter space, which can be computationally expensive and non-convex. The idea of analysing model structure through evolving particles in parameter space has been previously considered in [125]. Here, particles evolve along geodesics of the Fisher Information Metric. These trajectories have different properties to ours: movement along such a curve changes the nominal parameter vector, while always considering infinitesimal magnitude parameter perturbations. Movement along our curves, by contrast, changes the length scale of parameter perturbation considered, relative to a fixed dataset or nominal parameter vector θ^*. This builds on the theory relating nonlocal parameter sensitivity and (structural or practical) identifiability introduced in the previous chapter. More insight on this distinction is provided in Example 5.4.1. The method of [125] is, moreover, inapplicable when the model under analysis is structurally unidentifiable, and ill-conditioned when it has sloppy parameter vectors: calculation of the geodesic acceleration requires inversion of the Fisher Information Matrix.

5.2 Minimally Disruptive Curves

In the previous chapter we highlighted the relationship between minimally disruptive parameter vectors, practically unidentifiable subspaces (as quantified through the ϵ-uncertainty region), and multiscale sloppiness. However, examples were restricted to models for which the cost functions had a convenient form. In the general case, finding the minimally/maximally disruptive parameter vectors for a non-zero length-scale equates to a nonlinear, non-convex optimisation which may be very difficult to solve. However, suppose a branch of minimally disruptive parameters vary continuously in parameter space as a function of the length-scale δ (see Figure 5.1). In this case, we can evolve particles in parameter space, emanating from θ^*, as the solution of a set of Hamiltonian equations constructed so as to ensure that the particles trace over these branches.

We first define the notion of a 'continuous branch' of minimally disruptive parameter vectors more concretely. Let us consider any continuous functions Ξ^{Δ}_{min} satisfying:

$$\Xi^{\Delta}_{min}\big(\delta\big) \in B^{min}_{\bullet}(\delta) \;\forall \delta \in [0, \Delta], \tag{5.1}$$

where Δ is the largest considered length-scale. So such a Ξ^{Δ}_{min} exists when there is a continuous curve of minimally disruptive parameter vectors with strictly monotonically increasing distance δ from θ^*, up to distance Δ. We can pick an analogous function Ξ^{Δ}_{max} with the same properties, but formulated with respect to $B^{max}_{\theta^*}$ (as defined in Definition 4.3) and the maximally disruptive parameters. Note that the quantities

$$\frac{\mathrm{d}}{\mathrm{d}\delta}\left[\Xi^{\Delta}_{min}(\delta)\right]_{\delta=0}, \text{ and } \frac{\mathrm{d}}{\mathrm{d}\delta}\left[\Xi^{\Delta}_{max}(\delta)\right]_{\delta=0}$$

lie in the eigenspace of the sloppiest and stiffest eigenvalues of the Cost Hessian, re-

spectively.

Definition 5.1. ***Minimally disruptive curves***

A minimally disruptive curve of length F, relative to θ^*, is defined as a solution of the following optimisation problem:

$$\gamma^* = \min_{\gamma \in \Gamma^F_{\theta^*}} : \int_\gamma C_\bullet : \tag{5.2a}$$

$$\Gamma^F_{\theta^*} = \left\{ \gamma : [0, F] \to \Theta : \begin{array}{ll} \gamma(0) = \theta^*; & \frac{\mathrm{d}}{\mathrm{d}s}\Big[\|\gamma(s) - \theta^*\|_2\Big]_{s>0} > 0; \\ \big\|\frac{\mathrm{d}}{\mathrm{d}s}\gamma(s)\big\|_2 = 1; & C_\bullet\big(\gamma(F)\big) = B^{min}_\bullet\big(\|\gamma(F) - \theta^*\|_2\big) \end{array} \right\}. \tag{5.2b}$$

*A **weakly** minimally disruptive curve of length F, relative to θ^*, is defined as a solution of the following optimisation problem:*

$$\gamma^* = \min_{\gamma \in \tilde{\Gamma}^F_{\theta^*}} : \int_\gamma C_\bullet : \tag{5.3a}$$

$$\tilde{\Gamma}^F_{\theta^*} = \left\{ \gamma : [0, F] \to \Theta : \begin{array}{ll} \gamma(0) = \theta^*; & \frac{\mathrm{d}}{\mathrm{d}s}\Big[\|\gamma(s) - \theta^*\|_2\Big]_{s>0} \geq 0; \\ \big\|\frac{\mathrm{d}}{\mathrm{d}s}\gamma(s)\big\|_2 = 1; & C_\bullet\big(\gamma(F)\big) = B^{min}_\bullet\big(\|\gamma(F) - \theta^*\|_2\big) \end{array} \right\}. \tag{5.3b}$$

Notice that the only difference between a minimally disruptive curve and a weakly minimally disruptive curve is strictness of the inequality $\|\gamma(s) - \theta^*\|_2 > 0$, for $s > 0$. It turns out that minimally disruptive curves have stronger theoretical properties, while weakly minimally disruptive curves are more easily derived computationally. In practice, however, almost any weakly minimally disruptive found computationally maintains strictness of the aforementioned inequality, and so is minimally disruptive.

In Definition 5.1, minimally disruptive curves are parameterised in such a way that

their velocity is constant (and equal to one), whereas in equation (5.1), Ξ^{Δ}_{min} is parameterised so that the rate of change of Euclidean distance from θ^* is constant (and equal to one). As a consequence of this constant velocity, every curve in $\Gamma^F_{\theta^*}$ must have length F.

Lemma 5.1. *If a function satisfying* (5.1) *exists, then it is a minimally disruptive curve.*

Proof. First note that for any $\gamma \in \Gamma^F_{\theta^*}$, we can take the reparameterisation

$$\delta(s) = \|\gamma(s) - \theta^*\|_2. \tag{5.4}$$

We have, for $s > 0$, that $\frac{\mathrm{d}}{\mathrm{d}s}\|\gamma(s) - \theta^*\|_2 > 0$ (specified in (5.2b)). This ensures that $\delta(s)$ is a bijective function, and so it has an inverse. We can therefore parameterise the line integral of (5.2a) using either δ or s. Note however that the reparameterisation $\delta(s)$ will depend on the specific curve γ.

Suppose that the lemma is false. Then a function Ξ^{Δ}_{min} satisfying (5.1) exists, but is not a minimiser of (5.2a) (when reparameterised so that its velocity is constant). We take the actual optimiser of (5.2a) as γ^* . Reparameterisations of the form (5.4) will be different for Ξ^{Δ}_{min} and γ^*, so we denote them $\delta^1(s)$ and $\delta^2(s)$ respectively. Now by assumption, we have that

$$\int_0^{\delta^1(F)} C_\bullet\big(\gamma^*(\delta)\big) \left\| \frac{\mathrm{d}}{\mathrm{d}\delta}\big[\gamma^*(\delta)\big] \right\|_2 d\delta < \int_0^{\delta^2(F)} C_\bullet\big(\Xi^{\Delta}_{min}(\delta)\big) \left\| \frac{\mathrm{d}}{\mathrm{d}\delta}\left[\Xi^{\Delta}_{min}(\delta)\right] \right\|_2 d\delta. \tag{5.5}$$

Since Ξ^{Δ}_{min}, by definition, only traces over minimally disruptive parameters, we know,

for any $\delta \in \big(0, \delta^2(F)\big]$, that

$$C_\bullet\big(\gamma^*(\delta)\big) \geq C_\bullet\big(\Xi^\Delta_{min}(\delta)\big).$$

Furthermore the inequality is strict for at least one such δ, in order that γ^* itself does not only trace over minimally disruptive parameters. For (5.5) to hold, we therefore require

$$\int_0^{\delta^1(F)} \left\| \frac{\mathrm{d}}{\mathrm{d}\delta}\big[\gamma^*(\delta)\big] \right\|_2 d\delta < \int_0^{\delta^2(F)} \left\| \frac{\mathrm{d}}{\mathrm{d}\delta}\left[\Xi^\Delta_{min}(\delta)\right] \right\|_2 d\delta.$$

The LHS and RHS of the above inequality are equal to the lengths of the curves γ^* and Ξ^Δ_{min}, respectively. So γ^* must be strictly shorter than Ξ^Δ_{min}. However, this contradicts (5.2b), which demands that both γ^* and Ξ^Δ_{min} have length F. So γ^* cannot exist. □

Note that Lemma 5.1 does not specify the length F of the minimally disruptive curve. When we subsequently formulate an algorithm uncovering (weakly) minimally disruptive curves, it turnes out that F does not need to be specified *a priori*, so this does not become an issue.

Remark 5.1. *If a structural unidentifiability intersecting θ^* exists, and furthermore has monotonically increasing length-scale from θ^*, then it must be a weakly minimally disruptive curve. This follows from the global minimality, with respect to the cost function, of each parameter vector on the structural unidentifiability. Hence by uncovering weakly minimally disruptive curves, we can find the structural unidentifiabilities of a model. Of course, minimally disruptive curves exist independently of the structural identifiability properties of θ^*.*

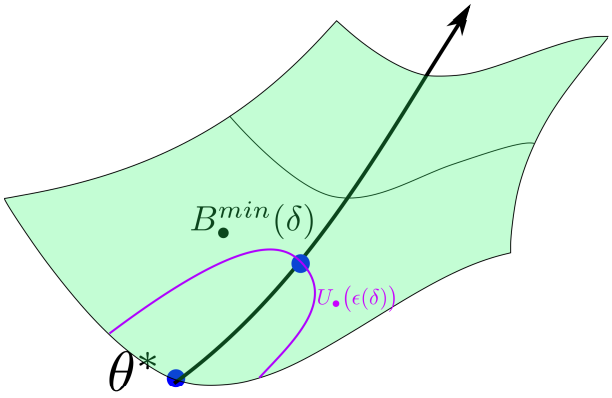

Figure 5.1: Minimally disruptive curve travelling along the cost landscape. *When a branch of minimally disruptive parameters $B_\bullet^{min}(\delta)$ vary continuously in δ, their union is a minimally disruptive curve. Each minimally disruptive parameter on the curve is the (non-unique) furthest point from θ^* within some (connected) ϵ-uncertainty region.*

5.3 Evolution of Minimally Disruptive Curves using Optimal Control

This section is concerned with reformulating the optimisation problem defining (weakly) minimally disruptive curves (Definition 5.1) in a computationally tractable way.

Remark 5.2. *A naive approach to solution of* (5.2a) *would be a gradient descent style algorithm, where the derivative of $\gamma(s)$ points along the direction of steepest descent of $C_\bullet\big(\gamma(s)\big)$, subject to appropriate constraints. However in a structurally unidentifiable model, $C_{\mathcal{Y}(\theta^*)}(\gamma(s))$ is always a local minimum of $C_{\mathcal{Y}(\theta^*)}$ with null Jacobian $\nabla_\theta C_{\mathcal{Y}(\theta^*)}(\gamma(s))$. Since exact nullity is never realised in a numerical context, a steepest descent direction will always exist and be highly sensitive to perturbation around the local minima, making numerical evolution ill-conditioned (see Figure 5.2). The same issue arises in structurally identifiable but practically unidentifiable models, where the Jacobian is effectively null. Instead we can consider* (5.2a) *as a constrained variational problem, and obtain necessary conditions on $\gamma(s)$ to satisfy* (5.2a) *by ap-*

plication of Pontryagin's Minimum Principle.

A maximally disruptive curve can be defined by turning the minimisation of (5.2a) *into a maximisation. An analogue of Lemma 5.1 relating such a curve to* Ξ^{Δ}_{max}*, if it exists, can be formulated. However it is of less interest from the point of view of unidentifiability characterisation. Note that in this case, a gradient-ascent style approach is sufficient from a computational viewpoint.*

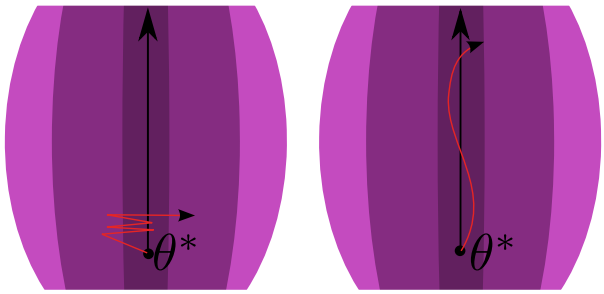

Figure 5.2: Topographic map of minimally disruptive curves evolving along parameter space *A bird's eye view of level sets of the cost function is shown. Darker areas have lower cost, while the black arrow traces the root of the cost (i.e. a structural unidentifiability). Red lines denote the trajectory of a particle attempting to evolve numerically along the structural unidentifiability. On the left, a steepest descent style algorithm fails due to the stiffness of the direction of descent around a local minimum. On the right: the inherent momentum associated with trajectory evolution via solution of Hamiltonian equations ensures that numerical error does not completely disrupt evolution.*

First, let us consider the optimal control problem solved in Problem B.1 of Appendix B. For a dynamical system of the form

$$\dot{x}(t) = f(x(t), u(t)) \qquad x(0) = x_0 \tag{5.6a}$$

$$x(t) \in \mathbb{R}^n \qquad u(t) \in \mathbb{R}^p, \tag{5.6b}$$

the problem was given as

$$\begin{aligned} &\text{Minimise} \quad J(u) = \int_0^{t_F} L\big(x(t), u(t)\big)\ dt \\ &\text{subject to} \quad K_1\big(x(t), u(t)\big) \leq 0; \qquad\qquad K_2\big(x(t), u(t)\big) = 0. \end{aligned} \tag{5.7}$$

Now suppose we (temporarily) alter the optimisation problem (5.3) defining a weakly minimally disruptive curve to *omit* the following constraint in $\tilde{\Gamma}^F_{\theta^*}$:

$$C_\bullet\big(\gamma(F)\big) = B^{min}_\bullet\big(\|\gamma(F) - \theta^*\|_2\big). \tag{5.8}$$

Then the optimisation problem (5.3) defining a weakly minimally disruptive curve is a special case of (5.7). Specifically, the dynamical system given in (5.6) becomes

$$\dot{\gamma}(s) = u(s) \qquad\qquad \gamma(0) = \theta^* \tag{5.9a}$$

$$\gamma(s) \in \mathbb{R}^q \qquad\qquad u(s) \in \mathbb{R}^q, \tag{5.9b}$$

while the optimisation problem (5.7) becomes

$$\begin{aligned} &\text{Minimise} \quad J(u) = \int_0^{F} C_\bullet\big(\gamma(s)\big)\ ds \\ &\text{subject to} \quad -\langle u(s), \gamma(s) - \theta^* \rangle \leq 0; \qquad \|u(s)\|_2 - 1 = 0. \end{aligned} \tag{5.10}$$

In Appendix B, necessary conditions on a local minimum of (5.7) in terms of a Hamiltonian function were provided. We now list the relevant quantities and necessary conditions given in Appendix B, having adapted them to the specific minimisation problem of (5.10). Firstly, the (augmented) Hamiltonian becomes

$$\begin{aligned} \bar{H}(\gamma(s), u(s), \lambda(s)) &= C_\bullet\big(\gamma(s)\big) + \langle \lambda(s), u(s) \rangle \\ &- \mu_1(s) \langle < u(s), \gamma(s) - \theta^* \rangle > + \mu_2(s)\big(\|u\|_2 - 1\big). \end{aligned}$$

Recall that $\lambda(s)$ denotes the costate of the system, while μ_1 and μ_2 denote s-varying Lagrange multipliers. We require

$$\begin{aligned} &\dot{\lambda}(s) = \mu_1(s)u(s) - \frac{\partial C_\bullet}{\partial \gamma}\big(\gamma(s)\big); && \lambda(F) = 0 \\ &\mu_1(s) \geq 0. \end{aligned} \tag{5.11}$$

Next, the necessary conditions on a local minimum are given as

$$\frac{\partial \bar{H}}{\partial u} = 2\mu_2(s)u(s) - \mu_1(s)\big(\gamma(s) - \theta^*\big) + \lambda(s) = 0 \tag{5.12a}$$

$$\mu_1(s)\langle \dot{\gamma}(s), \gamma(s) - \theta^* \rangle = 0 \tag{5.12b}$$

$$\bar{H}\big(\gamma(s), u(s), \lambda(s), \mu(s)\big) = \bar{H}\big(\gamma(0), u(0), \lambda(0), \mu(0)\big), \text{ for all } s \in [0, F]. \tag{5.12c}$$

These can be simplified. First let us take the inner product of (5.12a) with $u(s)$, while heeding the constraints of (5.10). This gives

$$2\mu_2(s) = C_\bullet(\gamma(s)) - \tilde{H}\big(\gamma(s), u(s), \lambda(s)\big).$$

Next, the time invariance of the Hamiltonian combined with the condition that $\lambda(F) = 0$, imply that

$$\bar{H}\big(\gamma(s), u(s), \lambda(s), \mu(s)\big) = C_\bullet\big(\gamma(F)\big).$$

Substituting these simplifications into (5.12), we can re-express the necessary conditions as

$$C_\bullet\big(\gamma(s)\big) + \langle \lambda(s), u(s) \rangle - C_\bullet\big(\gamma(F)\big) = 0 \tag{5.13a}$$

$$\mu_1(s)(\gamma(s) - \theta^*) + \Big(C_\bullet(\gamma(F)) - C_\bullet(\gamma(s))\Big)u(s) = \lambda(s). \tag{5.13b}$$

Given known values for $\gamma(s)$, $\lambda(s)$, $C_\bullet\big(\gamma(s)\big)$, and $C_\bullet\big(\gamma(F)\big)$, the necessary conditions in (5.13) become linear in $u(s)$ and $\mu_1(s)$, and we can find these quantities by solving the consequent system of linear equations. Meanwhile, we can find $\dot{\lambda}(s)$ using (5.11). Therefore, $\gamma(s)$ and $\lambda(s)$ evolve as a system of coupled differential equations. It remains to find the initial values $\gamma(0)$, $\lambda(0)$, and the final cost $C_\bullet\big(\gamma(F)\big)$.

Note that we do not need to provide an *a priori* value for F, as it does not explicitly enter any of the necessary conditions provided. This means that we are in fact free to set $C_\bullet\big(\gamma(F)\big)$ manually, and calculate F only after the curve $\gamma(s)$ has been generated. This explains the omission of the constraint given in (5.8).

We can then find $\lambda(0)$ in terms of $u(0)$ by evaluating (5.12a) at $s = 0$ to get:

$$\lambda(0) = \big[C_\bullet(\gamma(F)) - C_\bullet\big(\theta^*\big)\big]\, u(0).$$

Our final task is then to set $u(0)$. As s approaches zero, $\gamma(s)$ approaches θ^*. Consequently the behaviour of $C_\bullet\big(\gamma(s)\big)$ becomes dominated by the Cost Hessian $\nabla^2_\theta C_\bullet(\theta^*)$, and the initial direction $\dot{\gamma}(0)$ of the curve must lie within the eigenspace of the sloppiest eigenvector of the Cost Hessian. If the dimension of this eigenspace is e, then we have $2e$ choices of direction for $u(0)$, as the curve can travel in the positive or negative direction of a given eigenvector. Note that the magnitude of $u(0)$ is normalised by the constraints of the problem. These $2e$ choices can be iteratively tested.

We have now arrived at a method of evolving curves in parameter space that satisfy necessary conditions for being weakly minimally disruptive. In fact there are only a finite number of possible curves that can satisfy these necessary conditions (the number being twice the dimension of the eigenspace of the minimal eigenvector of the Cost Hessian, as explained in the previous paragraph). Therefore, if we evolve each of these curves, one must in fact be a weakly minimally disruptive curve. Furthermore,

if such a curve satisfies the strict inequality

$$\|\gamma(s) - \theta^*\|_2 > 0, \text{ for } s > 0,$$

as opposed to the weak inequality demanded by the necessary conditions, then it is (strongly) minimally disruptive. In practice, we have observed that the strict inequality is generally satisfied.

Note that a strong correspondence exists between the methods of optimal control used in this section, and those of Hamiltonian mechanics. The Principle of Least Action [22] states that the path a mechanical particle takes between two points minimises some functional known as the action. Thus, finding the appropriate path is an optimisation problem of a very similar form to (5.7). In Hamiltonian mechanics, the costate vector $\lambda(s)$ that has featured prominently in this section has a physical meaning: it is the momentum of the particle. Recall that in our problem, we are free to set the initial magnitude of the costate, which corresponds to the initial momentum magnitude, and also equals the final cost $C_\bullet\big(\gamma(F)\big)$. It is this momentum that allows our solution method to avoid the pitfalls associated with a gradient descent approach, which were described in Remark 5.2 and illustrated in Figure 5.2.

The freedom to set $C_\bullet\big(\gamma(F)\big)$ raises another issue. It may happen that the cost of a minimally disruptive curve never reaches the value chosen. For instance if the model possesses a structural unidentifiability intersecting θ^*, then the cost may be identically zero along the curve. Note however that the differential equations dictating curve evolution at some s^* are 'blind' to the behaviour of the cost function outside the ball of radius $\|\gamma(s^*) - \theta^*\|_2$ around θ^*, except through the assumption of some terminal point $\gamma(F)$ with predetermined cost. Thus the cost function outside of this ball could be arbitrarily modified without influencing the behaviour of the curve $\gamma(s)$, for $s \leq s^*$. One such modification could provide a point $\gamma(F)$ with the predetermined

cost required. Practically speaking, we can therefore terminate the curve prematurely at s^*. We provide an example of the situation just described in Section 5.4.2.

5.4 Examples

We now provide several examples of application for the methods developed in this chapter. The first two illustrate concepts introduced in the chapter, while highlighting the importance of non-locality in sensitivity quantification. We then verify our methods on a benchmark model used in identifiability analysis. The final example provides a more detailed analysis of a different benchmark model from the identifiability literature, specifically a metabolic reaction network. Novel structural and practical unidentifiabilities are identified. It is shown how our method can be used to gain additional mechanistic insight. Specifically, we uncover timescale-separated subsystems (and the regimes in which the timescale separation is valid), unnecessary mechanisms, and model approximations (together with their regimes of validity).

Each example involves the evolution of minimally disruptive curve(s). We choose our terminal cost as $C_{\theta^*}\big(\gamma(F)\big) = 0.1$ throughout, although curve evolution is often terminated prematurely. We observed that the numerical results generated by our algorithm appeared to be robust to tuning of this quantity. However, increasing the terminal cost by several orders of magnitude appeared to slightly reduce the computational time associated with curve evolution, while sacrificing some numerical accuracy. Since the terminal cost essentially corresponds to the momentum of the particle in parameter space, it could be that the added momentum hinders conformation to the minimally disruptive curves.

All the examples in this section have finite-dimensional output (typically the evaluation of a system of ordinary differential equations at discrete timepoints). All cost

functions used are of the form

$$C_{\mathcal{Y}(\theta^*)}(\theta) = \|\mathcal{Y}(\theta) - \mathcal{Y}(\theta^*)\|_2^2.$$

From Example 4.1 of Chapter 4 we see that this choice is equal to the Kullback-Leibler Divergence between the experimental models $\mathcal{D}(\theta)$ and $\mathcal{D}(\theta^*)$, in the case that

$$\mathcal{D}(\theta) \sim \mathcal{N}\big(\mathcal{Y}(\theta), I\big),$$

and where I represents the identity matrix of appropriate dimension.

5.4.1 Mixed Exponential Model

We first use a simple example to illustrate the concepts introduced in this chapter, and compare them against the existing sloppiness framework. The model, previously considered in [124, 125], is given as:

$$\mathcal{Y}(\theta) = \left[e^{\frac{-\theta_1}{3}} + e^{\frac{-\theta_2}{3}}, e^{-\theta_1} + e^{-\theta_2}, e^{-3\theta_1} + e^{-3\theta_2}\right]. \tag{5.14}$$

We take parameter space as $\Theta = [0, 4] \times [0, 4]$. We consider a nominal parameter vector $\theta^* = [4, 0.5]$. We deliberately place θ^* on the boundary of Θ so that only one direction of the sloppiest eigenvector of $\nabla^2_{\mathcal{Y}(\theta)} C_{\theta^*}(\theta)$ need be considered. This halves the analysis required without sacrificing illustrative power.

Note that model output (5.14) corresponds to observing a mixed exponential decay of the form $e^{-\theta_1 t} + e^{-\theta_2 t}$, at the timepoints $t = \{\frac{1}{3}, 1, 3\}$, and so model output is invariant with respect to permutation of the parameters. The model at θ^* is therefore locally, but not globally, structurally identifiable (for nonzero parameter vectors). In such a case, a minimally disruptive curve is not guaranteed to uncover the structurally

unidentifiable parameter vector $[0.5, 4]$: the absence of a structural unidentifiability as described in Remark 5.1 does not imply global structural identifiability. In this example, by chance, the minimally disruptive curve does in fact trace over $[0.5, 4]$.

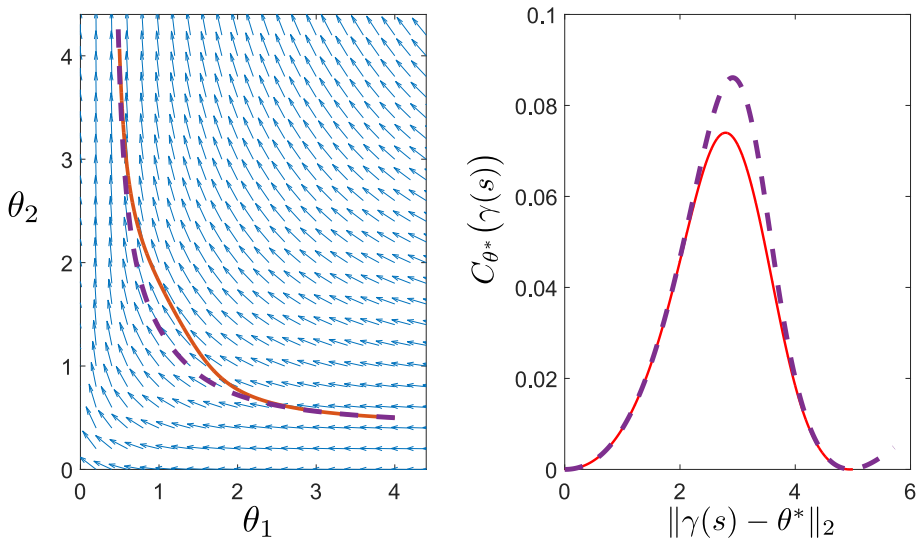

Figure 5.3: Discrepancy between minimally disruptive curves and flow of the minimal eigenvalues of the Cost Hessian. *Left: Vector field formed by sloppiest eigenvectors of the Cost Hessian (one direction only), for the model* (5.14)*. Overlaid are (solid line, red) a curve of minimally disruptive parameters, and (broken line, purple) the flow of the vector field. Right: The costs associated with these two curves. The x-axis charts Euclidean distance from* θ^**, while the y-axis charts the associated cost.*

We demonstrate in Figure 5.3 the differences inherent to evolving a particle in parameter space along its sloppiest eigenvector (which retains no memory of the initial vector θ^*), as compared to evolving a minimally disruptive curve (whose evolution is explicitly dependent on θ^*). To these ends, we first construct a vector field that, when evaluated at some $\theta \in \Theta$, points in the direction of the sloppiest eigenvector of $\nabla^2_\theta C_{\mathcal{Y}(\theta)}(\theta)$. The flow of this vector field, initialised from some θ^* (we take $\theta^* = [4, 0.5]$), is compared to a minimally disruptive curve emanating from the same θ^*. We see that the former, although by definition always pointing in the sloppiest direction (to infinitesimal magnitude), diverges from the latter. Indeed, the flow of the vector field takes a higher cost route, as shown on the figure, tracing over parameter vectors with more highly diverging output (relative to $\mathcal{Y}(\theta^*)$) than the minimally

disruptive curve (as required by definition). This discrepancy arises from the fact that the flow of the vector field, at some $\theta \in \Theta$ is determined independently of θ^*, and by analysis of the matrix $\nabla^2_\theta C_{\mathcal{Y}(\theta^*)}(\theta)$. The direction minimising local change in $\nabla^2_\theta C_{\mathcal{Y}(\theta)}(\theta)$, is not necessarily the direction minimising non-local change in $C_{\mathcal{Y}(\theta^*)}(\theta)$. Furthermore the minimal eigenvector of the Cost Hessian at θ^* (the sloppiest eigenvector) is a poor marker of the direction of minimal sensitivity to non-infinitesimal parameter perturbation.

5.4.2 Simple Algebraic Model

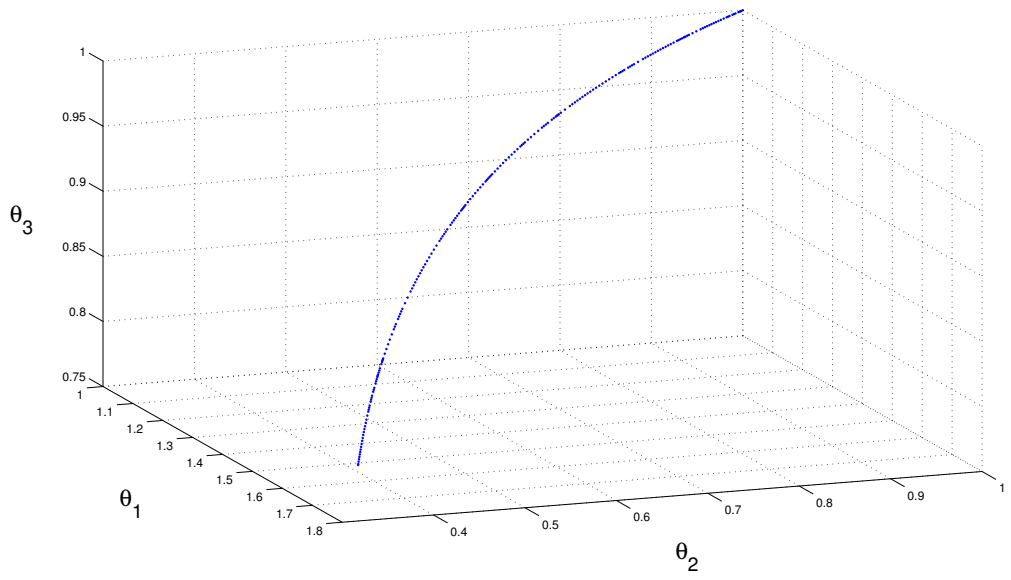

Figure 5.4: Minimally disruptive curve of a simple algebraic model. *We consider the simple algebraic model given in* (5.15)*, with* $\theta^* = [1, 1, 1]$*. The minimally disruptive curve travels along the root of the polynomial* $\theta_1^5\theta_2^3 - \theta_3^2 = 0$*.*

We next consider a toy model of the form

$$\mathcal{Y}(\theta) = (\theta_1^5\theta_2^3 - \theta_3^2)(\theta_1^3 + 4\theta_2^2 + 6\theta_3\theta_2), \tag{5.15}$$

and take an arbitrary choice of nominal parameter vector: $\theta^* = [1, 1, 1]$. This does not

correspond to a model of any real physical process, but allows us to gain insight into the concept of minimally disruptive curves. We see that θ^* is a root of the first factor of $\mathcal{Y}(\theta)$, i.e. $(\theta_1^5\theta_2^3 - \theta_3^2)$. The model therefore contains a structurally unidentifiable subspace emanating from θ^*, defined by the set of roots of this factor. We evolve a weakly minimally disruptive curve γ, while taking our *a priori* choice of terminal cost as $C_{\mathcal{Y}(\theta^*)}\big(\gamma(F)\big) = 0.1$. Note that this choice cannot be realised: structural unidentifiability ensures that $C_{\mathcal{Y}(\theta^*)}(\gamma(s))$ should be zero along the curve. However, structural unidentifiability is generally not known a priori. We terminate curve evolution when $\|\gamma(s) - \theta^*\|_2 = 1$. As expected, our curve stays within the structurally unidentifiable subspace to within numerical error, maintaining the property: $C_{\mathcal{Y}(\theta^*)}(\gamma(s)) < 10^{-5}$. The curve is plotted in Figure 5.4.

5.4.3 IL13-induced JAK-STAT Pathway

We next verify our methods on a differential equation model of the IL13-induced JAK-STAT pathway [99], which was used as a benchmark for identifiability analysis in [104]. The model has a 23-dimensional parameter space together with published estimates of each parameter, constituting θ^*. Details on the model equations, and observed timepoints, are provided in Appendix C. In [104], two structural unidentifiabilities were found: $C_{\mathcal{Y}(\theta^*)}(\theta)$ was zero where

$$\theta_{17}^*\theta_{22}^* = \theta_{17}\theta_{22}; \tag{5.16a}$$

$$\theta_{15}^*\theta_{21}^* = \theta_{15}\theta_{21} \text{ and } \theta_{15}^*\theta_{11} = \theta_{11}^*\theta_{15}. \tag{5.16b}$$

In the comparison of [104], only the method of [105] succeeded in detecting both the parameters involved in (5.16a), (5.16b), and their functional relations. However recent, unpublished work suggests that the Taylor series-based method of [96] may also

be effective. The profile likelihood methos of [105] requires evolution of a separate curve for each of the 23 parameters. To generate the i^{th} curve, θ_i^* is varied incrementally, and the other parameters iteratively re-optimised so as to minimise the effect on the cost function. This has the advantage of not requiring computation of $\nabla_\theta C_\bullet$ (unlike our method), and also generates componentwise confidence intervals. However the re-optimisation steps can be expensive. Our method recovered both structural unidentifiabilities given in (5.16), requiring only two curve generations (evolvable in parallel). Their initial velocities $\dot{\gamma}(0)$ were respectively aligned with the two minimal eigenvalues of the Cost Hessian. These curves are shown in Figure 5.5.

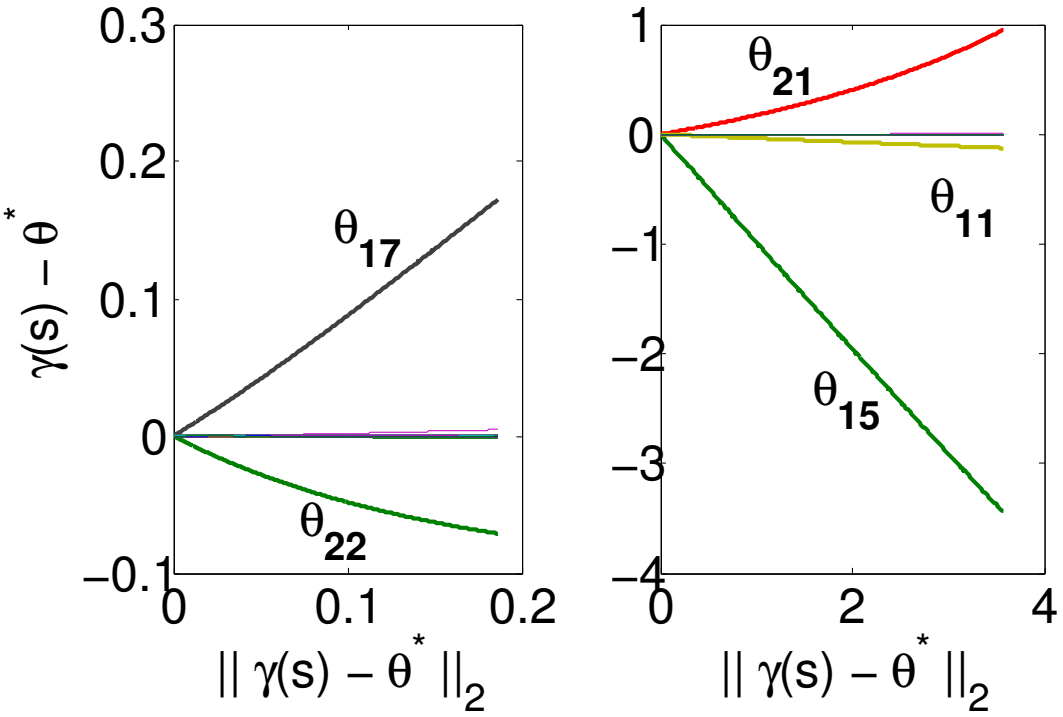

Figure 5.5: Minimally disruptive curves of the IL13-induced JAK-STAT pathway model *Each curve, on both diagrams, represents the change in one element of the minimally disruptive curve $\gamma(s)$, as s increases. So in the left-hand figure, we see, as $\gamma(s)$ evolves, that the parameter θ_{17} increases, while the parameter θ_{22} decreases. Furthermore, the changes in these components are co-ordinated such that the product $\theta_{17}\theta_{22}$ is preserved over the trajectory. Similarly, on the right-hand diagram, the products $\theta_{11}\theta_{15}$ and $\theta_{15}\theta_{21}$ are preserved. The cost $C_{\mathcal{Y}(\theta^*)}(\gamma(s))$ remains at numerical zero as the curve evolves, implying structural unidentifiabilities of the form given in* (5.16).

5.4.4 NF-κB Regulatory Module

We next find novel structural and practical unidentifiabilities in a benchmark differential equation model used for identifiablity analysis in the literature: the 29-parameter model of the NF-κB regulatory module [68]. The model describes an intracellular mechanism mediating dynamics of the NF-κB protein complex. A schematic diagram, based on one provided in [68], is given in Figure 5.6, and details on the system equations and parameters are provided in Appendix D. The system is excited by extracellular TNF. Previous analyses [31, 5] declared the model locally structurally identifiable, under relaxed assumptions necessary to maintain viability of the respective algorithms. Specifically, identifiability held so long as all parameters could be discriminated on the basis of some (unknown) time-course of TNF excitation, rather than the binary TNF signal assumed in [68]. In [31], sixteen parameters were furthermore fixed to maintain computational tractability. We analyse the full model, taking θ^* as the (complete) set of published parameter estimates/assumptions provided in [68].

We evolved two curves curves satisfying (5.13) in parallel, setting $\dot{\gamma}(0)$ in both the positive and negative directions of the eigenvector corresponding to the minimal eigenvalue of the Cost Hessian. The result is shown in Figure 5.7a, which connects both curves at θ^*. The top half of the figure shows the log-space evolution of each component of $\gamma(s)$, relative to θ^*. The components corresponding to the rate constants k_2 and c_4 are marked, and we see that the summation $\ln k_2 + \ln c_4$ is preserved. The lower half of the figure shows, meanwhile, that the cost $C_{\mathcal{Y}(\theta^*)}(\gamma(s))$ is zero, to a numerical precision of 10^{-8}. This implies a structural unidentifiability over areas of parameter space preserving the product $k_2 c_4$. This makes sense, as k_2 and c_4 respectively represent the translation (i.e. production) rate of the protein $A20$, and the inactivation rate of the kinase IKKa caused by $A20$. Since $A20$ and its transcript are not directly

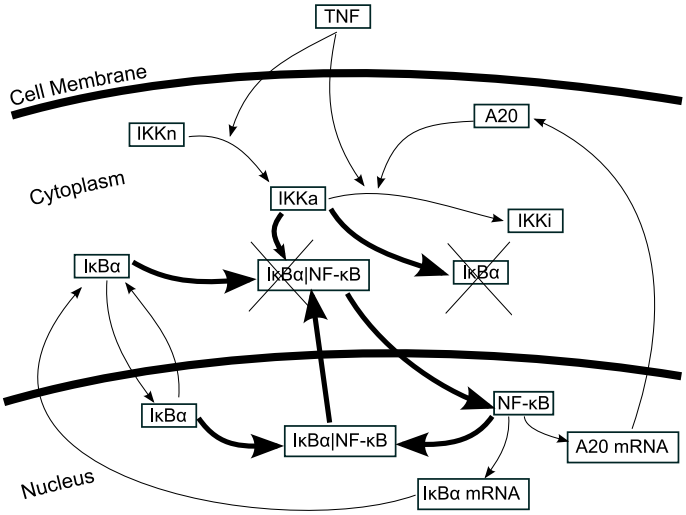

Figure 5.6: Schematic of the NF-κB regulatory module from [68], based upon Figure 2 in the same paper. *Bold arrows denote fast-timescale reactions. Crosses denote degradations. Extracellular Tumour Necrosis Factor (TNF) transforms neutral IκB kinase (IKKn) into its active form IKKa. The IKKa catalyses degradation of IκBa, both in its free state, and when bound in the IκBa|NF-κB complex. Catalysis of the previous complex frees cytoplasmic NFκB, which is then shuttled to the nucleus. Nuclear NF-κB promotes both IκBa and A20 transcripts. The translated A20 inactivates IKKa, thus indirectly promoting cytoplasmic IκBa. The translated IκBa binds to NF-κB, making the latter inert and shuttling it to the cytoplasm. So nuclear NF-κb is involved in a negative feedback loop with IκBa, and a positive feedback loop with A20.*

observed, it is clear that $A20$ only enters the model through its effect on IKKa. This refines the conclusion of [68], where the effect of $A20$ on dynamics was proposed to be a combination of its *concentration* (rather than translation rate), and the inactivation rate.

We next rewrote k_2 as $k_2 = \frac{k_2^* c_4^*}{c_4}$, thus reducing the dimensionality of parameter space by one, and eliminating the previous structural unidentifiability. Two minimally disruptive curves were again evolved in parallel, with initial directions in the positive and negative directions of the sloppiest eigenvector. Results are shown in Figure 5.7b. Three parameters show significant movement, and are involved in a practical unidentifiability. They are listed below:

- $e2a$ represents the rate of transport of the IκBα|NF-κB complex from the cyto-

plasm to the nucleus. This was appreciated as a fast-timescale reaction in [68], as shown in Figure 5.6.

- c_{5a} is the constitutive degradation rate of free cytoplasmic IκBα.
- t_1 is the rate constant for catalysis of the IκBα-IKKa complex into IKKa, also described as a fast timescale reaction in [68].

By varying these three parameters in turn, we see that each individual parameter has a close to zero effect on model output when varied, and the unidentifiability does not depend on their correlated change. Since e_{2a} and t_1 are fast-timescale reaction rates, the involved reactants and products are equilibrated near-instantaneously at θ^*, and increasing the rates makes no difference to observed model dynamics. Similarly, Figure 5.6 implies that IκBα-IKKa degradation is dominated by terms other than constitutive degradation, and decreasing c_{5a} will therefore barely affect observed dynamics. Thus, elimination of the reaction corresponding to the rate c_{5a}, and instant equilibration of the reactions corresponding to t_1 and c_{5a}, would negligibly affect dynamics at θ^*.

Note from Figure 5.6 that as the minimally disruptive curve travels in the direction of decreasing e_{2a}, decreasing t_1, and increasing c_{5a}, the associated cost starts to increase, until a step change in behaviour occurs. This can be explained mechanistically. Decrease in t_1 and e_{2a} slows the timescale of their associated reactions, until instant equilibration is no longer a valid approximation and observed dynamics are affected. Meanwhile, as c_{5a} increases it starts to make a meaningful contribution to IκBα degradation, and thereby also affects observed dynamics. Crucially, our algorithm finds at what parameter values these step changes occur, thereby defining the areas of parameter space over which the instant equilibration/zero degradation approximations are valid.

We fixed the parameters c_{5a}, e_{2a}, and t_1 and again evolved minimally disruptive curves

as previously. Results are shown in Figure 5.7c. We see that in a region around θ^*, dynamics are insensitive to parameter changes preserving the product $c_{1a}c_{4a}$. The former is the rate constant for promotion of the mRNA transcript for IκBα by free nuclear NF-κB. The latter is the translation rate for IκBα. Since the production rate of a protein is the product of its transcription and translation rates, then, as long as mRNA transcription is dominated by c_{1a}, it makes sense that preservation of the product $c_{1a}c_{4a}$ should not affect observed dynamics. However, the figure shows that this unidentifiability breaks down when c_{1a} decreases sufficiently. At this point, the cost starts to increase, and a step change in the practical unidentifiability is observed. The reason is that the model incorporates constitutive transcription of IκBα, which becomes the dominant contributor to overall transcription in this regime.

Again, we removed the practical unidentifiability in $c_{1a}c_{4a}$ by substituting $c_{1a} = \frac{c^*_{1a}c^*_{4a}}{c_{4a}}$, and evolved minimally disruptive curves (Figure 5.7d). The resulting curve involved a_2, which represents the association rate of IKKa and IκBα in the nucleus. In [68], it is depicted as a fast timescale reaction, with a nominal value of 0.2 obtained by assumption. In fact, as shown in Figure 5.7d, we can decrease a_2 by over twenty orders of magnitude without affecting nominal dynamics, so long as other parameters are changed in a compensatory manner. Thus we could switch off this mechanism with only a negligible change in nominal dynamics.

For our final curve evolution, we fixed a_2, with results shown in Figure 5.7e. This results in a complicated practical unidentifiability involving many parameters. We see that several parameters can be changed by five or more orders of magnitude without affecting dynamics, demonstrating that even with all the previous unidentifiabilities factored out, parameter estimation without prior assumptions on particular parameter values would be a challenging task in this model.

5.5 Discussion

We have provided an automatic, numerical approach to elucidation of both unidentifiable parameters (whether structural or practical) and their associated unidentifiabilities. Note that, as a numerical method, our algorithm cannot theoretically guarantee structural unidentifiabilities, and is only suitable for local identifiability analysis. However in this regard it can complement symbolic approaches to finding unidentifiable parameters, which additionally provide theoretical guarantees, but can suffer from scalability issues. In particular, algebraic approaches based on the Observability Rank Condition [54] are in general computationally intensive even for modestly-sized models, but work well when there are few parameters and many outputs under consideration. Our method can be used as a preconditioning tool for such algorithms, as it can highlight groups of parameters likely to be structurally unidentifiable. All other parameters can then be fixed, and unidentifiability of the simplified model checked algebraically. For instance, in the example of the NF-κB network from [68], we fixed all parameters except k_2 and c_4 (which we discovered were structurally unidentifiable through our algorithm, as discussed previously and shown in Figure 5.7a). The reduced, two-parameter model could then be easily analysed algebraically, and verified as structurally unidentifiable by direct application of the Observability Rank Condition. The same, integrated procedure was also performed successfully to verify the structurally unidentifiable parameters of the JAK-STAT model presented previously.

Our method additionally inferred the functional form of the structural unidentifiabilities, through analysis of the trajectories generated by the algorithm. In general, this inference may not be trivial. It is considerably simplified by considering the trajectories in log-space, as depicted in Figure 5.7, so long as the unidentifiability is rational. The reason is that if there is an unidentifiability over a relation of the

form $p(\theta) = 0$, where $p(\theta)$ is rational, then a linear function of the logarithms of the parameters is preserved, and inference of the functional relation becomes a linear regression problem.

A key theme of this chapter is that the parametric sensitivity characteristics of a mathematical model can depend highly on both the parameter vector considered and the length-scale of parameter perturbation. We saw in the NF-κB network example that there were several practical unidentifiabilities that held in a certain region of parameter space, but had a distinct boundary at which a step change in the unidentifiability occurred (see e.g. Figures 5.7b and 5.7c). These corresponded to the regions in which particular model approximations were valid. For instance, some such unidentifiabilities involved the rate constants of fast-timescale reactions, which were unidentifiable as long as a timescale separation existed between the reaction dynamics and observed model dynamics, but regained identifiability once they fell below a certain level. Our algorithm identified both the practical unidentifiabilities, and the boundaries in parameter space at which the form of the unidentifiabilities experienced a qualitative change. It is therefore of use in the model reduction problem, both for identifying redundant mechanisms, and quantifying the regions of parameter space at which they remain redundant. Note that poor choices of unit for parameters result in practical unidentifiabilities. Given a parameter θ_1 in a model, if we transformed parameter space by taking $\hat{\theta}_1 = 1000\theta_1$, then the sensitivity of the model to perturbations in its first parameter would be reduced by a factor of a thousand. The converse implication is not true, as demonstrated through our discussion of the practical unidentifiabilities shown in Figures 5.7b and 5.7c. In both cases, model behaviour was sensitive to perturbation in a particular direction in some areas of parameter space, and insensitive to perturbation in the same direction in other areas of parameter space (i.e. where the practical unidentifiability held). Scaling parameters would not be able to change the discrepancy between the sensitivities in the regimes where the practical unidenti-

fiability did and did not hold. Thus the practical unidentifiabilities were not merely an artefact of poorly scaled parameters.

When the algorithm of this chapter is performed, the majority of computational time is taken up simulating the model at different points in parameter space, in order to calculate the associated cost, and cost gradient. The density of simulations on a minimally disruptive curve varies with the 'momentum' (i.e. the choice of magnitude of the initial costate $\lambda(0)$), the choice of ODE solver used to evolve the Hamiltonian flow (we used the 'ode113' solver from the MATLAB toolbox [75]), and the model in question. We took $C_{\mathcal{Y}(\theta^*)}(\gamma(F)) = 0.1$ to set our momentum, for all curves generated in the NF-κB example. 440 internal evaluations of the cost and its gradient were required by the ode113 solver to produce Figure 5.7a, while Figure 5.7e required 1363 evaluations. Each combined evaluation of the cost and its gradient took about five seconds on a single-core, 3.40 GHz processor. The gradient was calculated by solving the sensitivity equations associated with the model (see [63]), rather than by finite-difference methods.

5.6 Summary

In this chapter, we have devised and applied an algorithm for delineating curves in parameter space over which model output varies minimally, relative to data or some nominal output. The algorithm has several purposes, which were demonstrated upon benchmark models of metabolic reaction networks from the literature. Specifically, it was used to identify structural and practical unidentifiabilities in a model, which cannot be estimated from data. Furthermore, hidden model features were extracted, including redundant mechanisms and timescale separated subsystems, together with the regimes in which such features existed.

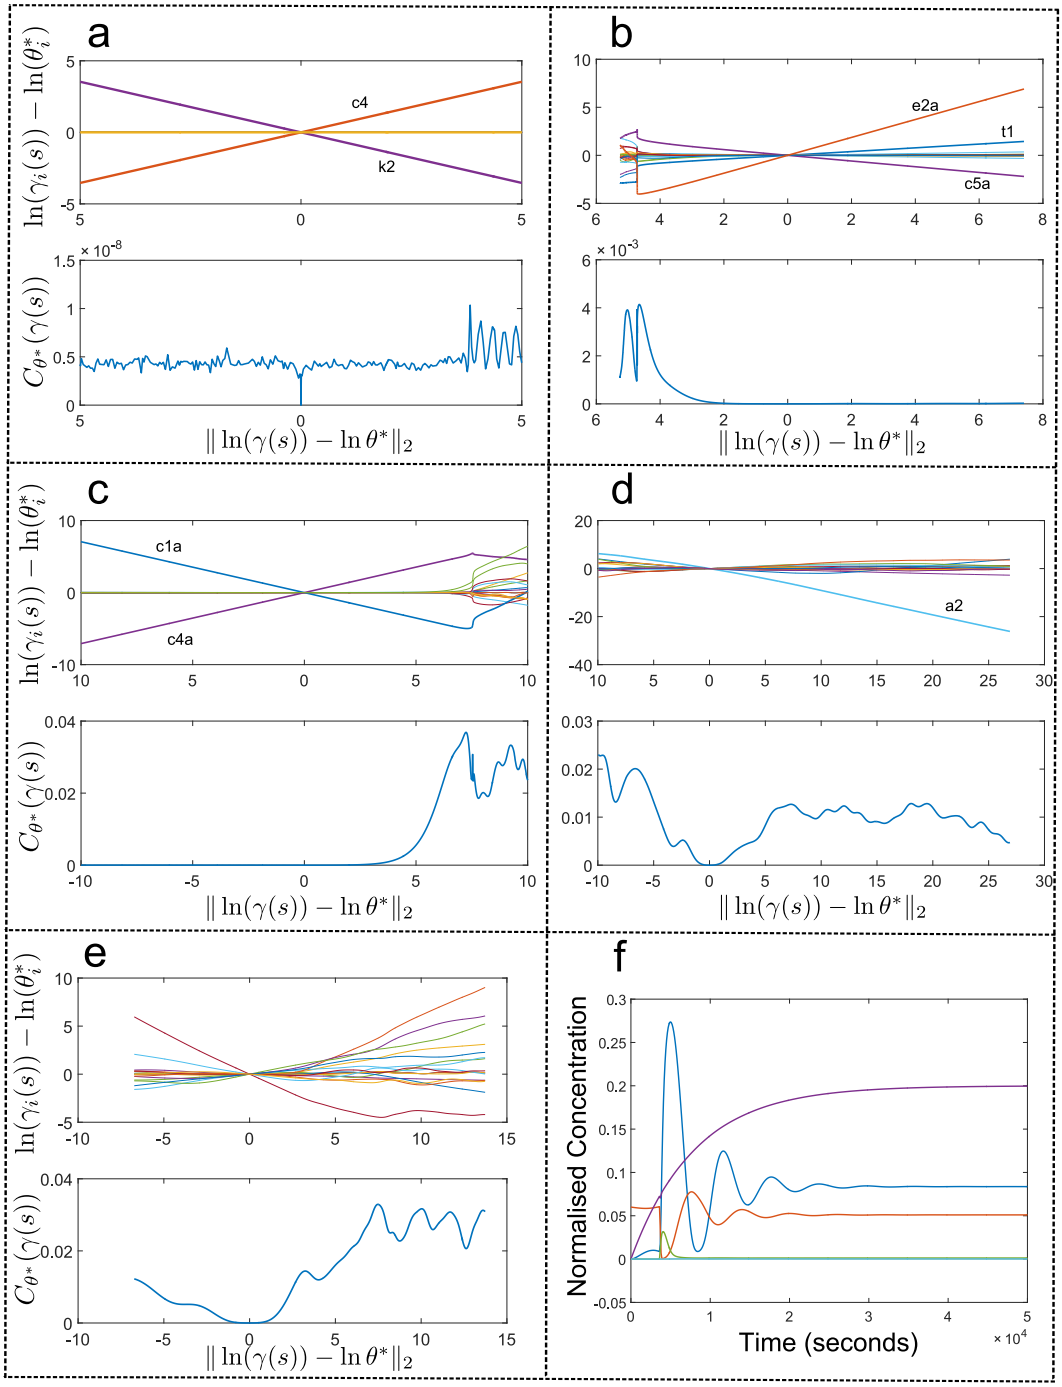

Figure 5.7: Analysis of an NF-κB regulatory module model through evolution of minimally disruptive curves. *Panels a-e depict structural and practical unidentifiabilities in the NF-κB model from [68]. The x-axes of all plots in these panels depict the distance, in log-space, between the nominal parameter vector θ^*, and that of $\gamma(s)$, as s varies. Each line on the top-half figures of these panels represents the relative change of a particular parameter component θ_i, from θ_i^*, as s varies. The bottom half figures represent the cost associated with the parameter vector $\gamma(s)$, relative to θ^*, as s varies. Panel f provides the normalised time-course of the observed quantities in the model.*

Chapter 6

Uncertainty Quantification of Parameter Estimation in the Presence of Model Error

This chapter is concerned with the *perfect model assumption*, which underlies much of the theory concerning parameter estimation. This assumption is that the model output induced by some (unknown) optimal parameter vector(s) perfectly recreates process dynamics. However it is often the case that a model can never perfectly recreate the modelled process, due to the existence of unmodelled or badly modelled process features. In this case, the 'perfect model' assumption does not hold. However, we can still find a parameter vector that best approximates process dynamics, and it may indeed be meaningful despite the deficiencies of the model. In this chapter we argue, however, that quantification of the uncertainty associated with estimating this parameter vector changes in the absence of the perfect model assumption.

The models studied in this chapter are ODE models, whose parameters are estimated from time-series data corrupted by Gaussian measurement noise. In this context, we

propose a new method for quantifying the covariance of a parameter estimate. It is in many respects similar to the standard covariance estimate obtained by inverting the Fisher Information Matrix. Indeed it agrees with this estimate, in the limit of increasing experimental replicates, and when the perfect model assumption holds. Critically, our covariance quantifier does not rely on the perfect model assumption to maintain theoretical validity. In fact, no quantification of the discrepancy between the dynamics of the process, and those of the (optimally parameterised) model, is required. So our method is valid even when dealing with very crude models of complex processes. The results of the chapter are published in [101].

6.0.1 Notation

Throughout this chapter we will use the notation $\lambda(A)$ to denote the Lebesgue measure of a set A. See e.g. [107] for a rigorous presentation of the Lebesgue measure. Meanwhile, $\|A-B\|_F$ will denote the Frobenius distance between matrices A and B. Recall that the Frobenius distance is given by $Tr[(A-B)^T(A-B)]$.

6.1 Motivation

The previous chapters have dealt with the problem of parameter estimation. Here, the model structure is fixed, and so the parameters enter the model in a precisely specified way. The challenge lies in providing each parameter with a numeric value, or region of values. In this chapter, we are additionally concerned with the appropriateness of the model structure itself. Therefore this chapter falls within the field of **System Identification**, which deals with the synthesis of both a model structure and model parameters that accurately represent the modelled process. Note that system identification has three subfields. White-box system identification attempts to completely

derive a model from mechanistic principles, ending up with a parameter-free model. Grey-box system identification partially derives a model from mechanistic principles, but also incorporates free parameter values that must be fitted to data, due to lack of sufficient *a priori* knowledge of the process. Black-box system identification attempts to create a function that approximates the input-output characteristics of the process, with no regard given to the mechanistic principles underlying the process. The entirety of this thesis is related to grey-box system identification.

Mathematical models are only approximations of the physical processes they attempt to recreate, and, as such, there is a mismatch between dynamics of the model and the true process. The discrepancy between the dynamics of the process model $\mathcal{Y}(\theta)$, and the observed data is typically known as the **model error** [70], and stems from two sources. One is the data-corruption caused by measurement noise, which we have discussed extensively at previous points in the thesis. The other is **undermodelling**, which is caused by the fundamental inability of the model structure to replicate all the intricacies of the process dynamics, regardless of the considered parameter values.

When a process is modelled, there is an inherent tradeoff between the complexity of the model structure and the fidelity of the model. Traditional system identification has been mainly concerned with simultaneous minimisation of the complexity of the model structure being searched over, and the model error. In other words, the objective is to find a simple model structure that faithfully recreates process dynamics. The numeric values of any parameters within the model structure, given these objectives, are not the focus. Of more interest is the (lack of) model error induced by the best-fit parameters, which has a diverse literature [85, 71, 73, 84, 49]. In this context, issues such as the identifiability of parameters, and the uncertainty associated with a parameter estimate from noisy data, are not important.

In this thesis, system identification has a different objective. We are not only con-

cerned with having a model (and associated parameter values) that accurately describe the process, but also with the estimation of certain parameters with physical significance, e.g. the thermal conductivity of a material or the rate constant of a chemical reaction. Accordingly, the quantification of uncertainty in a parameter estimate is important. We saw in Chapter 4 that fitting parameters to noisy data is a statistical problem. A best-fitting parameter vector is known as a **nominal** parameter vector, and it is an estimator of the (unknown) **optimal** parameter vector, which would best fit the data in the absence of measurement noise. In this chapter we only consider the maximum likelihood estimator (MLE), as defined in Definition 2.9 of the Preliminaries. This is an extremely common choice of estimator, due to its various attractive theoretical and practical properties (see equations (2.4) of the Preliminaries).

This chapter deals with finding the covariance of the MLE. A standard approximation of this covariance was previously provided through the well-known Cramer-Rao Inequality (see Lemma 4.1). This is widely used as a heuristic approximation of the MLE covariance where parameter estimation is carried out on ODE models (see e.g. [39, 79, 130, 140]). Indeed, inferences on the inherent well-posedness of the parameter estimation problem are made on the basis of this matrix. Some of the well known limitations of the Cramer-Rao inequality were considered in the discussion accompanying Lemma 4.1. An additional concern, which is our focus in this chapter, is that the Cramer-Rao inequality is not valid when the process is undermodelled. This is a particular concern in fields such as systems biology, where a high dimensional parameter space coupled with inaccurate measurements means that the validity of published parameter estimates is often questioned [35, 44, 50]. A need therefore exists to quantify MLE covariance when undermodelling is an issue.

In this chapter, we derive a new approximation of the covariance of the maximum

likelihood estimator with several attractive properties. Firstly, it is mathematically valid even when the process is assumed to be undermodelled. In fact, no quantification of the magnitude or form of undermodelling is required. Secondly, it converges asymptotically to the true covariance under regularity conditions that are less onerous than those required by the inverse FIM. Finally, the approximation error in our method is a function of the distance between the sample means of the datapoints, and their (unknown) true values. By contrast, the approximation error in the FIM-based method is a function of the distance between the nominal and (unknown) optimal parameter vectors. Both methods also possess linearisation-derived errors. Thus the two methods are complementary in that for problems in which one approximation is inaccurate, the other may be accurate. Our method is valid when considering discrete-time observations of an unknown dynamic process corrupted by Gaussian measurement noise. Furthermore, the model can be nonlinear.

6.2 Parameter Estimation in the Absence of a Perfect Model

Much of the mathematical theory and notation used in this chapter has been introduced at previous points in the thesis. However, the incorporation of undermodelling into our mathematical framework induces some subtle, but important differences in the way we treat models and experimental data both mathematically and philosophically.

We will consider parameterised process models with finite-dimensional output $\mathcal{Y}(\theta)$. Note that the results of previous chapters have allowed for infinite dimensional models. However, the relationship between the process model, the data, and the assumed form of measurement noise change in this chapter. Previously, we modelled the data $\mathcal{D}(\theta)$

as a noise-corrupted realisation of the process model $\mathcal{Y}(\theta)$. This was valid in the absence of undermodelling. For instance, if noise was assumed Gaussian, then we had

$$\mathcal{D}(\theta) \sim \mathcal{N}\big(\mathcal{Y}(\theta), \Sigma\big). \tag{6.1}$$

We will continue to use the random variable $\mathcal{D}(\theta)$, but we will no longer assume that $\mathcal{D}(\theta)$ accurately represents the noise-corrupted process. Instead, we will assume that data is a noise-corrupted realisation of an **unmodelled** process, whose noise-free output is not given by $\mathcal{Y}(\theta)$, regardless of the choice of θ. We denote the noise-free output of this process as $\mathcal{G}$. So we have that

$$\mathcal{D}(\mathcal{G}) \sim \mathcal{N}\big(\mathcal{G}, \Sigma\big). \tag{6.2}$$

This raises a philosophical issue as to the meaning of maximum-likelihood estimation. As explained in the preliminaries (Definition 2.9), maximum likelihood estimation consists of finding a parameter vector $\hat{\theta}$ that maximises the probability density

$$\mathbb{P}[\mathcal{D}^{[E]}; \mathcal{Y}(\theta)].$$

However, the probability density of hypothetical data is now independent of the process model $\mathcal{Y}(\theta)$. Numerically we can of course still evaluate $\mathbb{P}[\mathcal{D}^{[E]}(\hat{\mathcal{G}}); \mathcal{Y}(\theta)]$, and minimise it over the set of possible $\theta \in \Theta$ to get a best-fit parameter. Philosophically, a more natural information-theoretic interpretation of the maximum likelihood estimator exists in this context: Given data $\mathcal{D}^{[E]}$, we can minimise the **Kullback-Leibler divergence** between the empirical distribution of the data, and $\mathcal{D}(\theta)$, over θ. The Kullback-Leibler divergence between random variables F and G, with associated

probability density functions $f(x)$ and $g(x)$, is given by

$$\begin{aligned} KL\big[F||G\big] &= \int_X f(x) \log\left(\frac{f(x)}{g(x)}\right)\, dx \\ &= \mathbb{E}_F[\log f(x)] - \mathbb{E}_F[\log g(x)]. \end{aligned}$$

It is a non-symmetric measure of the information loss associated with approximating the random variable F by G [34]. Recall that a standard choice of cost function given in equation (4.1b) of Chapter 4 was

$$\begin{aligned} C_{\mathcal{Y}(\theta^*)}(\theta) &= \mathbb{E}_{\theta^*}[-\log \psi_{\mathcal{Y}(\theta)}] - \mathbb{E}_{\theta^*}[-\log \psi_{\mathcal{Y}(\theta^*)}] \\ &= KL\big[\mathcal{D}(\theta^*)||\mathcal{D}(\theta)\big], \end{aligned} \tag{6.3}$$

where $\psi_{\mathcal{Y}(\theta)}$ denoted the probability density function of the random variable $\mathcal{D}(\theta)$.

Given E-replicated data $\mathcal{D}^{[E]}$, we can define the **empirical distribution** of the data. This random variable, which we will denote $\mathcal{E}(\mathcal{D}^{[E]})$, ascribes a probability $\frac{1}{E}$ to each of the E data replicates. We can denote the probability density of this distribution as

$$f_{\mathcal{D}^{[E]}}(x) = \frac{1}{E}\sum_{i=1}^{E} \delta(x - \mathcal{D}^i),$$

where δ denotes the Dirac delta-function. Then we get that

$$\min_{\theta\in\Theta} KL\big[\mathcal{E}(\mathcal{D}^{[E]})]||\mathcal{D}(\theta)\big] = \min_{\theta\in\Theta}\left[\mathbb{E}_{\mathcal{E}(\mathcal{D}^{[E]})}\big[-\log \psi_{\mathcal{Y}(\theta)}\big] - \mathbb{E}_{\mathcal{E}(\mathcal{D}^{[E]})}\big[-\log f_{\hat{\mathcal{G}}}(x)\big]\right].$$

Eliminating θ-independent terms in the above equation gives

$$\begin{aligned}\min_{\theta\in\Theta} KL\big[\mathcal{E}(\mathcal{D}^{[E]})]||\mathcal{D}(\theta)\big] &= \min_{\theta\in\Theta} \mathbb{E}_{\mathcal{E}(\mathcal{D}^{[E]})}\big[-\log\psi_{\mathcal{Y}(\theta)}\big] \\ &= \min_{\theta\in\Theta} \sum_{i=1}^{E} -\log\psi_{\mathcal{Y}(\theta)}\big(\mathcal{D}^i\big) \\ &= \min_{\theta\in\Theta} -\log\mathcal{L}(\theta;\mathcal{D}).\end{aligned}$$

We see that this equates to minimisation of the negative log-likelihood of the data over parameter space. So maximum likelihood estimation and minimisation of the Kullback-Leibler divergence coincide [34]. The latter, however, places no causal restrictions on the relationship between the experimental model $\mathcal{D}(\theta)$, and the actual data. We will use the notation

$$\hat{\theta}\big(\mathcal{D}^{[E]}\big) = \min_{\theta\in\Theta} KL\big[\mathcal{E}(\mathcal{D}^{[E]})]||\mathcal{D}(\theta)\big]$$

for the remainder of this chapter. So $\hat{\theta}$ corresponds to the maximum likelihood estimator.

Recall that the data $\mathcal{D}$, in this chapter, is assumed to be a realisation of some unknown process with noise-free output $\mathcal{G}$ (see (6.2)). As well as estimating the maximum likelihood parameter $\hat{\theta}$, we will on occasion estimate the true value of $\mathcal{G}$, given data. Our estimate of $\mathcal{G}$ will be denoted $\mathcal{G}$. Note that the maximum likelihood estimate, assuming Gaussian measurement noise of the form (6.2), is given by

$$\hat{\mathcal{G}} = \frac{1}{E}\sum_{i=1}^{E}\mathcal{D}^i.$$

The distribution of $\hat{\mathcal{G}}$ is then given as

$$\hat{G} \sim \mathcal{N}\left(\mathcal{G}, \frac{\Sigma}{E}\right), \tag{6.4}$$

as a result of (6.2). Furthermore the sample mean $\hat{\mathcal{G}}$ is a *sufficient statistic* [25]. So the maximum likelihood estimate can be calculated solely from $\hat{\mathcal{G}}$, without full knowledge of the the data. We will therefore interchangeably use $\hat{G}$ and $\mathcal{D}^{[E]}$ as the argument of the maximum likelihood estimator $\hat{\theta}$, depending on the context.

6.3 Asymptotic Distribution of Estimated Model Output

In this section we formulate theory relating to the asymptotic properties of $\hat{G}$, in the limit of increasing experimental replicates. We commence by providing some assumptions necessary for the results of the chapter.

A1 Consistency:

$$\mathbb{P}\left[\lim_{E\to\infty} \hat{\theta}(\mathcal{D}^{[E]}) = \theta^*\right] = 1,$$

A2 $\mathcal{C}^2$ Differentiability: $\nabla_\theta \mathcal{Y}(\theta)$ and $\nabla^2_\theta \mathcal{Y}(\theta)$ are defined for all $\theta \in \mathbb{R}^q$.

A3 Unique maximum likelihood parameters: Let

$$W = \left\{ x \in \mathbb{R}^n : \arg\min_{\theta\in\Theta} \mathcal{L}(\theta; x) \text{ is not unique} \right\}$$

Then $\lambda(W) = 0$.

A4 Let

$$U = \{x \in \mathbb{R}^{n\times r} : \hat{\theta}(x) \notin \mathcal{C}^2\}.$$

Then $\lambda(U) = 0$ and $\theta^* \notin U$.

Assumption A1, along with sufficient conditions for it to hold true, was introduced in the Preliminaries (see equation (2.4a)). Note that the discussion in the Preliminaries assumed no undermodelling of the process. One of the sufficient conditions was global structural identifiability ($\mathcal{Y}(\theta) = \mathcal{Y}(\theta^*) \Rightarrow \theta = \theta^*$). We will see subsequently how this is not a sufficiently strong condition for consistency when undermodelling is allowed.

Remark 6.1. *Note that assumptions A3 and A4 are not restrictive conditions. Any probability measure that has a probability density function is absolutely continuous with respect to the Lebesgue measure, and vice versa [136]. Therefore a necessary and sufficient condition for the probability of a sample mean to not possess a property with probability one, is that this property occurs on a set of Lebesgue measure zero. If $\mathcal{C}^2$ differentiability (and/or uniqueness) of the mapping held on a set of positive Lebesgue measure, then we would be assured of the existence of a ball in $\mathbb{R}^{n\times r}$, on which θ_{est} was nowhere differentiable (and/or unique). This would be a very strange occurence.*

We will continue to refer to θ^*, defined as the limiting maximum likelihood estimator in Assumption A1, as the optimal parameter vector. This is despite the absence of the perfect model assumption , which means that $\mathcal{Y}(\theta^*)$ no longer necessarily corresponds to the true process output $\mathcal{G}$. Note that $\theta^* = \hat{\theta}(\mathcal{G})$, the parameter estimate given noise-free data. This follows from continuity of the maximum-likelihood estimator at θ^*, combined with Assumption A1.

Lemma 6.1.

Suppose assumptions A1 to A3 hold. Then

$$\sqrt{E}\big[\hat{\theta}(\mathcal{D}^{[E]}) - \hat{\theta}(\mathcal{G})\big] \to_D \mathcal{N}\Big(0, \nabla\hat{\theta}(\mathcal{G})\Sigma\nabla\hat{\theta}(\mathcal{G})^T\Big). \tag{6.5}$$

Proof. By definition, we have $\hat{\theta}(\hat{\mathcal{G}}) = \hat{\theta}(\mathcal{D}^{[E]})$. Let us take a Taylor Expansion of $\hat{\theta}$ around $\mathcal{G}$. We get

$$\hat{\theta}(\hat{\mathcal{G}}) - \theta^* = [\nabla\hat{\theta}(\mathcal{G})]^T[\hat{\mathcal{G}} - \mathcal{G}] + \mathcal{O}([\hat{\mathcal{G}} - \mathcal{G}]^T[\hat{\mathcal{G}} - \mathcal{G}])$$

The distribution of the first order term follows from (6.4):

$$\nabla\hat{\theta}(\mathcal{G})]^T[\hat{\mathcal{G}} - \mathcal{G}] \sim \mathcal{N}\Big(0, \nabla\hat{\theta}(\mathcal{G})\frac{\Sigma}{E}\nabla\hat{\theta}(\mathcal{G})^T\Big).$$

It remains to show that the high-order terms in the Taylor Expansion are unimportant in the limit of increasing data. The Law of Large Numbers, recalling (6.4), gives

$$\hat{\mathcal{G}} - \mathcal{G} \xrightarrow[E\to\infty]{P} 0.$$

This allows for application of the multivariate Delta Method to each component $\hat{\theta}_i(\hat{\mathcal{G}})$ of the maximum likelihood estimate, providing the result. The Delta Method (see e.g. [25], p242) is a generalisation of the Central Limit Theorem, which ensures that high order terms in $\hat{\mathcal{G}} - \mathcal{G}$ do not affect the limiting distribution as this quantity tends to zero in probability. □

Theorem 6.1.

Suppose we have data $\mathcal{D}$ such that $\hat{\theta}(\mathcal{D})$ is $\mathcal{C}^1$ differentiable in an open neighbourhoud

$N \ni \mathcal{D}$. Take $\hat{\theta} = \hat{\theta}(\mathcal{D})$. Then, if

$$|\nabla^2_\theta \mathcal{L}(\hat{\theta}, \mathcal{D})| \neq 0, \tag{6.6}$$

we have

$$\nabla\hat{\theta}(\mathcal{D}) = \Big(\nabla^2_\theta \mathcal{L}(\hat{\theta}, \mathcal{D})\Big)^{-1} \nabla_{\mathcal{D}} \nabla_\theta \mathcal{L}(\hat{\theta}, \mathcal{D}). \tag{6.7}$$

Moreover, if $\nabla_\theta \mathcal{L}(\hat{\theta}, \mathcal{D})$ at $\mathcal{D}$ is $\mathcal{C}^k$ differentiable for $k > 1$, then so too is $\hat{\theta}(\mathcal{D})$.

Proof. First order optimality conditions imply that $\nabla_\theta \mathcal{L}(\hat{\theta}, \mathcal{D}) = 0$. The implicit function theorem [94] guarantees the existence of a unique, continuously differentiable function $\psi : \mathbb{R}^n \to \mathbb{R}^q$, and an open set $U \ni \mathcal{D}$, such that:

$$\nabla_\theta \mathcal{L}(\psi(z), z) = 0 \quad \forall z \in U.$$

Without loss of generality we can assume $U \subseteq N$. By first order optimality conditions, we have

$$\nabla_\theta \mathcal{L}\big(\hat{\theta}(z), z\big) = 0 \ \ \forall z \ \in U.$$

Uniqueness of ψ then implies that $\hat{\theta}(z) = \psi(z)$ on U. Their gradients at $\mathcal{D}$ must therefore correspond, giving (6.7).

Another consequence of the implicit function theorem is that if $\nabla_\theta \mathcal{L}(\hat{\theta}, \mathcal{D})$ is $\mathcal{C}^k$ differentiable at $\mathcal{D}$, for some k, then so too is $\psi(\mathcal{D})$, and by extension $\hat{\theta}(\mathcal{D})$. □

Equation (6.6) is a key assumption of the previous theorem. However, it is not restrictive. One can consider a line of reasoning analogous to that of Remark 6.1.

Consider the Lebesgue measure of the following set:

$$A = \{\mathcal{D} \in \mathbb{R}^{r \times n} : |\nabla^2_\theta \mathcal{L}(\hat{\theta}(\mathcal{D}), \mathcal{D})| = 0\}.$$

We conjecture that for non-pathological models, $\lambda(A) = 0$, and hence the probability of drawing data not satisfying (6.6) is also zero. Intuitively, this follows as a consequence of the set of singular matrices being low-dimensional in the space of nonsingular matrices. More formally, if $\lambda(A)$ were positive, we would require the existence of a $\mathcal{D} \in \mathbb{R}^n$, and $\epsilon > 0$ satisfying

$$|\nabla^2_\theta \mathcal{L}(\hat{\theta}(z), z)| = 0 \ \forall z : \|z - \mathcal{D}\|_2^2 < \epsilon;$$

Since the mapping $z \to |\nabla^2_\theta \mathcal{L}(\hat{\theta}(z), z)|$ goes from $\mathbb{R}^n$ to $\mathbb{R}$, and depends in a complicated way on the characteristics of the model output $\mathcal{Y}(\theta)$, one would not expect it to be identically 0 on an open set.

6.4 Covariance Approximation of the Maximum Likelihood Estimator

We will approximate the covariance of $\hat{\theta}(\mathcal{D}^{[E]})$ with the following function:

$$\mathfrak{C}(\mathcal{D}^{[E]}) = \nabla\hat{\theta}(\mathcal{D}^{[E]})\frac{\Sigma}{E}\nabla\hat{\theta}(\mathcal{D}^{[E]})^T. \tag{6.8}$$

The motivation for this approximation is given by Lemma 6.1, and the fact that

$$\lim_{E\to\infty} \hat{\theta}(\mathcal{D}^{[E]}) = \hat{\theta}(\mathcal{G})$$

by consistency. The expression for $\mathfrak{C}(\mathcal{D}^{[E]})$ can be calculated by using Theorem 6.1. In this section we consider the properties of this approximation, and compare it to the traditional covariance estimate predicated on the inverse Fisher Information Matrix.

Definition 6.1.

The Fisher Information Matrix of a parameterised Likelihood function $\mathcal{L}(\theta;\mathcal{D})$, at θ^*, is defined as

$$\mathcal{I}(\theta^*) = \mathbb{E}_{\theta^*}\Big[\big(\nabla_\theta \log \mathcal{L}(\theta;\mathcal{D}(\theta^*))\big)\big(\nabla_\theta \log \mathcal{L}(\theta;\mathcal{D}(\theta^*))\big)^T\Big], \tag{6.9}$$

where $\mathcal{D}(\theta^*)$ is the random variable representing the output of the experimental model at θ^*.

As mentioned in the preliminaries, the log-likelihood is additive with respect to experimental replicate number. So the Fisher Information Matrix of E-replicated data is

$$\mathcal{I}^E(\theta^*) = E\mathcal{I}(\theta^*).$$

In this chapter we have taken $\mathcal{D}(\theta)$ as the corruption of $\mathcal{Y}(\theta)$ with Gaussian measurement noise, i.e.

$$\mathcal{D}(\theta) \sim \mathcal{N}\big(\mathcal{Y}(\theta), \Sigma\big).$$

In this case, the Fisher Information Matrix takes the following form:

$$\mathcal{I}(\theta) = \nabla_\theta \mathcal{Y}(\theta)^T \Sigma^{-1} \nabla_\theta \mathcal{Y}(\theta) \tag{6.10}$$

Recall that the Cost Hessian was previously referred to as the Fisher Information Matrix in the discussion accompanying Lemma 4.1. This was for the case of a particular cost function given as (see (6.3))

$$C_{\mathcal{Y}(\theta^*)}(\theta) = KL\big[\mathcal{D}(\theta^*)||\mathcal{D}(\theta)\big].$$

This is consistent with Definition 6.1. The equality between the Hessian of the Kullback-Leibler Divergence and the Fisher Information Matrix is proved in e.g. [34]. The inverse Fisher Information constitutes an asymptotic lower bound on the covariance of the maximum likelihood estimator, through the Cramer-Rao inequality (see Lemma 4.1). Thus both $\mathfrak{C}(\mathcal{D}^{[E]})$ and $\mathcal{I}^E(\hat{\theta}(\mathcal{D}^{[E]})^{-1}$ approximate the same quantity, namely the covariance of $\hat{\theta}(\mathcal{D}^{[E]})$. They have complementary properties, as we now demonstrate:

Consider the following linearisation of $\mathcal{Y}(\theta)$ around some θ^*:

$$\tilde{\mathcal{Y}}(\theta) = \mathcal{Y}(\theta^*) + \langle \nabla_\theta \mathcal{Y}(\theta^*), (\theta - \theta^*) \rangle.$$

Note that the linearisation does not affect the quantity $\mathcal{I}(\theta^*)$, calculated according to (6.10). Furthermore, we now have that

$$\mathcal{I}(\theta) = \mathcal{I}(\theta^*), \quad \forall \theta \in \Theta.$$

This demonstrates how it is the nonlinearities in the dependence of $\mathcal{Y}(\theta)$ on θ that add an intractable error term $\mathcal{I}\big(\hat{\theta}(\mathcal{D})\big) - \mathcal{I}(\theta^*)$ to the estimate of the true Fisher Information on which the Cramer-Rao inequality is predicated, i.e. $\mathcal{I}(\theta^*)$. In the limit of increasing experimental replicates, one assumes that $\hat{\theta} \to_P \theta^*$, and therefore, correspondingly, that $\mathcal{I}^E\big(\hat{\theta}(\mathcal{D}^{[E]})\big) \to_P \mathcal{I}^E(\theta^*)$.

Now consider the original nonlinear model $\mathcal{Y}(\theta)$, and suppose that the parameter

estimation routine is itself linearised. Thus:

$$\begin{aligned}\tilde{\hat{\theta}}(\hat{\mathcal{G}}) &= \hat{\theta}(\mathcal{G}) + \nabla_{\hat{\mathcal{G}}}\hat{\theta}(\mathcal{G})(\hat{\mathcal{G}} - \mathcal{G}) \\ &= \theta^* + \nabla_{\hat{\mathcal{G}}}\hat{\theta}(\mathcal{G})(\hat{\mathcal{G}} - \mathcal{G}).\end{aligned}$$

Then we would have

$$\tilde{\hat{\theta}}(\hat{\mathcal{G}}) - \tilde{\hat{\theta}}(\mathcal{G}) \sim \mathcal{N}\Big(0, \mathfrak{C}(\mathcal{G})\Big)$$

In other words, the convergence in distribution relation (6.5) (the subject of Lemma 6.1) would be replaced by equality, regardless of the multiplicity E of repeated experiments. Furthermore, the covariance estimate would be immune to the error $\hat{\mathcal{G}} - \mathcal{G}$, as we would have $\mathfrak{C}(\hat{\mathcal{G}}) = \mathfrak{C}(\mathcal{G})$ regardless of $\hat{\mathcal{G}}$.

In conclusion, estimating the covariance of a parameter estimate, whether through $\mathcal{I}^E(\hat{\theta}\big(\mathcal{D}^{[E]}\big))$ or $\mathfrak{C}(\hat{\mathcal{G}})$, is made inaccurate due to the presence of nonlinearities. These inaccuracies vanish asymptotically, as the number of experimental replicates E increases. For $\mathcal{I}^E\big(\hat{\theta}(\mathcal{D}^{[E]})\big)$, the nonlinearities occur in the dependence of the deterministic model on its parameters, and couple with an inaccurate parameter estimate. For $\mathfrak{C}(\hat{\mathcal{G}})$, the nonlinearities occur in the dependence of the maximum likelihood estimate on the data, and couple with an inaccurate estimate of the noise-free data $\mathcal{G}$.

The covariance approximation $\mathfrak{C}(\hat{\mathcal{G}})$ was formulated without the perfect model assumption (i.e. $\mathcal{G} = \mathcal{Y}(\theta^*)$), while the inverse Fisher Information relies on it. Therefore, it is sensible to check the properties of $\mathfrak{C}(\hat{\mathcal{G}})$ when the perfect model assumption holds, in order to be assured that it yields similar conclusions to the Fisher Information. In fact $\mathfrak{C}(\hat{\mathcal{G}})$ corresponds asymptotically to the inverse Fisher Information Matrix in the limit of increasing experimental replicates and given the perfect model assumption, as we now show.

Lemma 6.2.

Suppose that $\mathcal{G} = \mathcal{Y}(\theta^*)$ and assumptions A1 to A3 hold. Then

$$\mathfrak{C}(\mathcal{D}^{[E]}) \to_D \mathcal{I}^E\big(\hat{\theta}(\mathcal{D}^{[E]})\big)^{-1} \tag{6.11}$$

Proof. By assumption, the MLE $\hat{\theta}(\mathcal{D})$ is both consistent and differentiable in an open set surrounding θ^*. So we have

$$\lim_{E\to\infty} \hat{\theta}(\hat{\mathcal{G}}) = \theta^* \ \mathcal{P} \ a.s.,$$

and differentiation gives us

$$\frac{\mathrm{d}}{\mathrm{d}\theta^*} \lim_{E\to\infty} \hat{\theta}(\hat{\mathcal{G}}) = I_q \ \mathcal{P} \ a.s., \tag{6.12}$$

where $I_q \in \mathbb{R}^{q\times q}$ is the identity matrix. Application of the chain rule to the LHS of (6.12) results in:

$$\lim_{E\to\infty} \nabla_{\hat{\mathcal{G}}}\hat{\theta}(\hat{\mathcal{G}})^T \nabla_{\theta^*}\hat{\mathcal{G}} = I_q. \tag{6.13}$$

By the perfect model assumption and consistency, we have that

$$\lim_{E\to\infty} \hat{\mathcal{G}} = \mathcal{G} = \mathcal{Y}(\theta^*).$$

So we can rewrite (6.13) as

$$\lim_{E\to\infty} \nabla_{\hat{\mathcal{G}}}\hat{\theta}(\hat{\mathcal{G}})^T \nabla_{\theta^*}\mathcal{Y}(\theta^*) = I_q.$$

So, by expansion and cancellation

$$\begin{aligned}\lim_{E\to\infty} \mathfrak{C}(\hat{\mathcal{G}})\mathcal{I}^E(\hat{\theta}(\hat{\mathcal{G}})) &= \nabla\hat{\theta}(\hat{\mathcal{G}})\frac{\Sigma}{E}E\nabla\hat{\theta}(\hat{\mathcal{G}})^T\nabla_\theta\mathcal{Y}(\theta^*)^T\Sigma^{-1}\nabla_\theta\mathcal{Y}(\theta)\\ &= I_q\end{aligned}$$

as required. □

Remark 6.2. *The Likelihood $\mathcal{L}(\theta, \mathcal{D})$ is likely to possess multiple local minima in θ, given data $\mathcal{D}$ for models exhibiting complex behaviours and many parameters [80, 130]. The form of the measurement noise may influence which of these is globally minimising. As experimental data tends to infinity, consistency (Assumption A1) implies that eventually a single, particular local minimum will become globally minimising, with a probability tending to one. This is the asymptotic regime under which both $\mathfrak{C}(\hat{\mathcal{G}})$, and $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$, become valid as approximations to the covariance of the maximum likelihood estimator. It is important to note that both covariance estimates therefore discount the effect of the parameter estimation protocol jumping between different local minima of the cost function due to the vagaries of the data. This phenomenon is likely to occur not only due to the form of the measurement noise, but also due to deficiencies in the parameter estimation protocol. After all, global optimisation is an NP-hard problem [80].*

6.5 Examples

6.5.1 Structurally Identifiable Example for which Parameter Estimation is Impossible

Consistency, as defined in Assumption A1, is a prerequisite to meaningful parameter estimation. In the case that the perfect model assumption holds, structural identifi-

ability of the model (along with other easily satisfied conditions, see equations (2.4) of the Preliminaries) is sufficient for consistency. As discussed in previous chapters, structural identifiability itself may be a hard property to verify. Even then, it ceases to be a sufficient condition for consistency in the case that the perfect model assumption does not hold. This is now demonstrated by example. We provide an example of a structurally identifiable model of an (artificially constructed) process for which parameter estimation to any degree of confidence is impossible, due to the small discrepancy between the dynamics of the process and the model. This is of course both an academic example and an extreme case, but it illustrates the principle of structural identifiability not guaranteeing consistency.

We wish to estimate a process that we will endow with the following dynamics:

$$\dot{g}(t) = Ag(t) \qquad\qquad g(0) = [4,5]^T$$

$$A = \begin{bmatrix} -3 & 1 \\ 2 & 1 \end{bmatrix}$$

Data consists of noisy process observations taken at arbitrary timepoints $\{t_j\}_{j=1}^r$. So we have $\mathcal{G} = \{g(t_i)\}_{i=1}^r$ and

$$\mathcal{D} \sim \mathcal{N}(\mathcal{G}, \Sigma),$$

with the covariance matrix Σ of measurement noise taken as the identity. Note that this choice of covariance matrix is not necessary for the methods of the chapter to be applicable, but it improves the clarity of exposition. Typically, the process being estimated is unknown, and the objective is to accurately model it. However, model structures cannot always capture the true dynamics of the modelled process. Of course in this example we know the process dynamics *a priori* by construction. We will deliberately provide a parameterised model structure that undermodels process

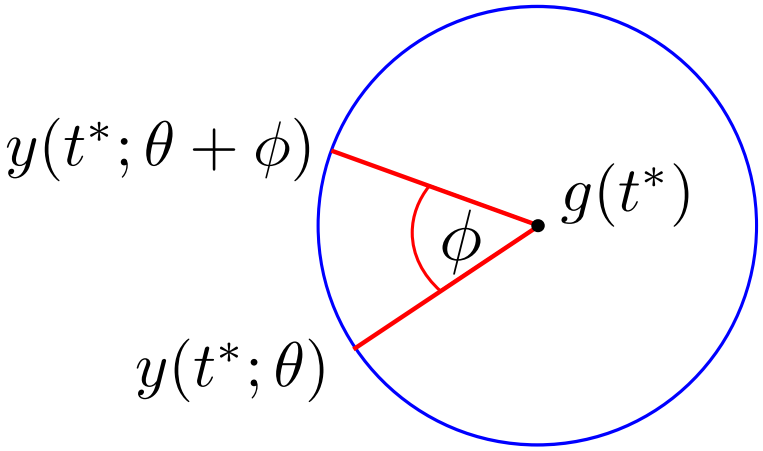

Figure 6.1: For any $t^* > 0$, the set of model outputs $y(t^*, \theta)$, for $\theta \in (-\pi, \pi]$ forms a circle with centre $g(t^*)$.

dynamics: no parameter vectors of the model can perfectly reproduce the noise-free process dynamics, even if some can closely approximate them. This is done in order to observe the consequent pathologies in the parameter estimation routine. Model dynamics are given as:

$$\dot{x}(t) = \tilde{A}x(t)$$

$$y(t, \theta) = C(\theta)x(t)$$

where

$$\tilde{A} = \begin{bmatrix} A & 0 \\ 0 & G \end{bmatrix} \qquad C(\theta) = \begin{bmatrix} 1 & 0 & \sin(\theta) & -\sin(\theta) \\ 0 & 1 & \cos(\theta) & -\cos(\theta) \end{bmatrix}$$

$$G = \begin{bmatrix} -3.5 & 1.5 \\ 2 & -1 \end{bmatrix} \qquad x(0) = [g(0), 1, 1]^T.$$

Note that the quantity

$$\|y(t, \theta) - g(t)\|_2^2 \tag{6.14}$$

is invariant to change in θ, for all $t > 0$. In particular, $y(t, \theta')$ can be obtained from $y(t, \theta)$, for $\theta' > \theta$, by a clockwise rotation through the angle $\theta' - \theta$ around $g(t)$ (see Figure 6.1).

Since measurement noise is Gaussian, the maximum likelihood can be calculated from the mean data

$$\hat{\mathcal{G}} = \frac{1}{E} \sum_{i=1} \mathcal{D}^i,$$

which is distributed as

$$\hat{\mathcal{G}} \sim \mathcal{N}\left(\mathcal{G}, \frac{\Sigma}{E}\right).$$

Note that, since Σ is the identity matrix, the probability density of $\hat{\mathcal{G}}$ is a radial function centred at $\mathcal{G}$. Therefore, $\hat{\mathcal{G}}$ is invariant over sets of the form:

$$\{x : \|x - \mathcal{G}\|_2 = c\}.$$

This observation, combined with the invariance of (6.14) with respect to θ, ensures that the likelihood of any parameter being the MLE is constant over parameter space. Since the distribution of the MLE is uniform over $(-\pi, \pi]$, we can calculate its true covariance as

$$\mathrm{Cov}\big(\hat{\theta}(\hat{\mathcal{G}})\big) = \frac{1}{3}\pi^2.$$

Note that this covariance is independent of E, the number of experimental iterations taken.

Since the covariance does not decay with increasing E, Assumption A1 (consistency) does not hold. Thus asymptotic covariance quantification of any parameter estimate,

either through $\mathfrak{C}(\hat{\mathcal{G}})$ or $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$, is mathematically invalid. Nevertheless the model is structurally identifiable. In general, verification/invalidation of Assumption A1 without the perfect model assumption is impossible in any real-life context despite being key to much of the theory associated with the field. In this case it is possible only because we know the process dynamics a priori. Therefore it is worthwhile to test the fidelity of the covariance approximations $\mathfrak{C}(\hat{\mathcal{G}})$ and $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$ regardless.

The trajectory of $g(t)$ was simulated multiple times, and data $\mathcal{D}^{[E]}$ was collected by adding measurement noise and storing outputs at the measurement timepoints. At the K^{th} iteration (i.e. after K experimental replicates), an MLE parameter estimate $\hat{\theta}(\mathcal{D}^{[K]})$ was calculated, taking into account data from all K replicates. A plot of parameter estimate against replicate number K is provided in Figure 6.2. Although the distribution of parameter estimates given K replicates is uniform for any K, as proved previously, this is not true given knowledge of the parameter estimate over $K-1$ iterations. Thus significant autocorrelation is clearly visible in Figure 6.2.

At each iteration K of the experiment, both $\mathfrak{C}(\hat{\mathcal{G}})$ and $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$ were calculated, and graphs of the estimated covariance for both methods are provided respectively in Figure 6.3.

The mean data $\hat{\mathcal{G}}$ converges to $\mathcal{G}$ as the number of replicates increase. One can easily see that the maximum likelihood estimator $\hat{\theta}$ possesses an asymptote at $\hat{\theta}(\mathcal{G})$, at which point all parameters have equal likelihood. The magnitude of the gradient $\nabla\hat{\theta}(\hat{\mathcal{G}})$ correspondingly tends to infinity as this asymptote is approached. The estimate $\mathfrak{C}(\hat{\mathcal{G}})$, despite being blind to the underlying process dynamics, is predicated on this gradient, and thus flags the estimation problem as ill-conditioned. This stands in contrast to $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$, which incorrectly predicts a swiftly decaying covariance as the number of replicates increases. A comparison is provided in Figure 6.3.

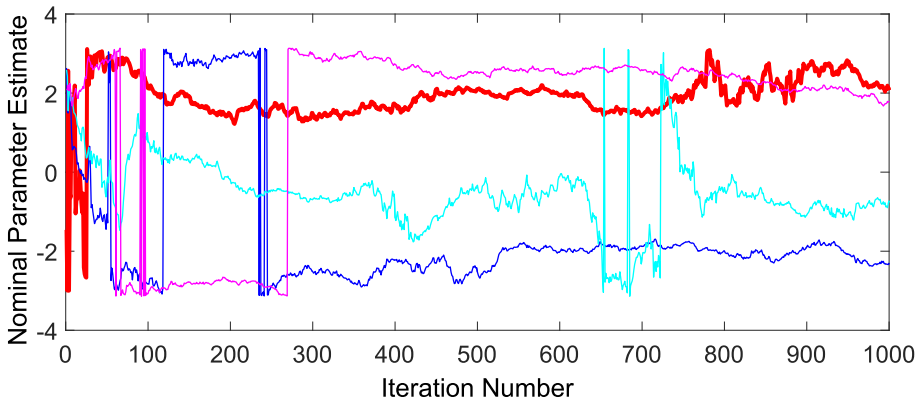

Figure 6.2: Multiple parameter estimates of an ill-posed model. *Plot of* $\hat{\theta}(\hat{\mathcal{G}})$. *Each iteration alters the sample mean, changing the nominal parameter estimate. This scheme is carried out four times, to emphasise unpredictability. Figure 6.3 uses the red set of iterations (depicted in bold).*

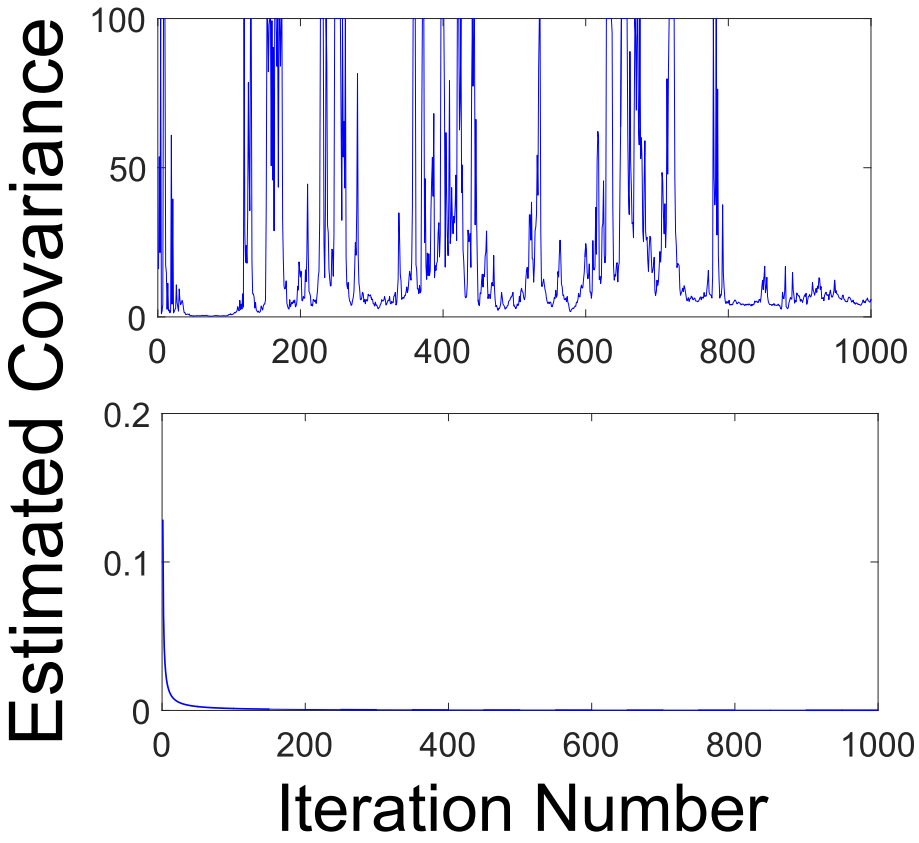

Figure 6.3: Comparison of covariance estimation routines on an ill-posed model. *Estimate of covariance using* (6.8)*, and the sample mean after each iteration. Any estimates over* 100 *were taken to be* 100 *to maintain the scaling of the graph.*

6.5.2 Chemical Reaction Network Example

We first supply a structurally identifiable ODE system corresponding to a chemical reaction network, with parameters representing unknown rate constants. We simulate data by picking an 'optimal' parameter vector θ^*, and adding measurement noise. Therefore the perfect model assumption holds, as $\mathcal{G}$, the noise free data, is equal to $\mathcal{Y}(\theta^*)$. We compare our method of covariance quantification with the existing one as increasing numbers of experiments replicates are taken. Results suggest that the methods converge, as proven mathematically in Lemma 6.2. We further produce a numerically synthesised estimate of parameter covariance, by running 1000 parameter estimations, and calculating their sample covariance. The ODE, and corresponding chemical reaction network, are described below:

$$\begin{aligned}
&\dot{x}(t,\theta) = A(\theta)x(t) && y(t,\theta) = x(t,\theta) \in \mathbb{R}^3 \\
&A(\theta) = \begin{bmatrix} -0.7-\theta_1 & \theta_2 & 0 \\ \theta_1 & -\theta_2 & 2 \\ 0.7 & 0 & -3 \end{bmatrix} && \theta \in \mathbb{R}^2 \qquad (6.15) \\
&x(0,\theta) = [50,50,50]^T
\end{aligned}$$

$$x_1 \underset{\theta_2}{\overset{\theta_1}{\rightleftharpoons}} x_2 \quad x_1 \overset{0.7}{\to} x_3 \quad x_3 \overset{2}{\to} x_2 \quad x_3 \overset{1}{\to} \emptyset$$

Global structural unidentifiability can be proved by taking a Laplace transform of the dynamics. If we denote by $\hat{x}(s,\theta)$ the component of the complex frequency s in the trajectory $x(t,\theta)$, then we have that

$$\hat{x}(s,\theta) = (s\mathbb{I} - A(\theta))^{-1} x(0,\theta).$$

For two values of θ to induce identical model dynamics, $\hat{x}(s,\theta)$ must be preserved for all values of s. Therefore one need only show that $s\mathbb{I} - A(\theta)$ is injective as a function of θ, in order to prove global structural identifiability. This can be performed using the symbolic toolbox in MATLAB, for example. Note too that a similar computation verifies nonsingularity of $A(\theta)$ for all values of θ.

We set $\theta^* = [3,4]^T$. Output is measured at a set of timepoints $T = [0,1,2,\ldots,10]$. We add Gaussian noise with identity covariance matrix to $\mathcal{Y}(\theta^*)$ to generate data. We also varied E, the number of experimental replicates, since the results of this chapter are asymptotic in the limit of increasing E. Figure 6.4 (top half) shows that the Frobenius distance between the covariance estimates produced by $\mathfrak{C}(\hat{\mathcal{G}})$ and $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$ respectively, tends to numerical zero as E increases (note that the y-axis of the graph is scaled by E, for reasons outlined in the accompanying caption). This serves as numerical validation of Lemma 6.2. The figure also suggests that the rate of convergence of either method to the sample covariance, as E increases, is sublinear. This is likely either due to the inaccuracies in both methods highlighted in Remark 6.2, or to error in the sample covariance of parameter estimates (taken, for each E, by running 1000 parameter estimation routines).

We now add an extra depletion term to one of the state derivatives of the ODE generating synthetic data. However, we do not alter the ODE model (6.15) used in the parameter estimation protocol: this is a 'hidden' depletion term. Thus the perfect model assumption does not hold. The deterministic trajectory generating the synthetic data, and the corresponding chemical reaction, are described below:

$$\dot{g}(t) = A(\theta^*)g(t) - [0,0,g_3(t)]^T \qquad\qquad \theta^* = [3,4]^T, \tag{6.16}$$

$$x_1 \underset{\theta_2}{\overset{\theta_1}{\rightleftharpoons}} x_2 \quad x_1 \xrightarrow{0.7} x_3 \quad x_3 \xrightarrow{2} x_2 \quad x_3 \xrightarrow{2} \emptyset$$

Note the change of notation from (6.15) to (6.16): the latter describes a process, whereas the former describes a model of the process.

Here A is as in (6.15), and $g_3(t)$ represents the third component of the vector $g(t) \in \mathbb{R}^3$. Figure 6.4 (bottom half) depicts the results. The figure demonstrates the superior convergence properties of (6.8) towards the asymptotic sample covariance in the case that the perfect model assumption does not hold.

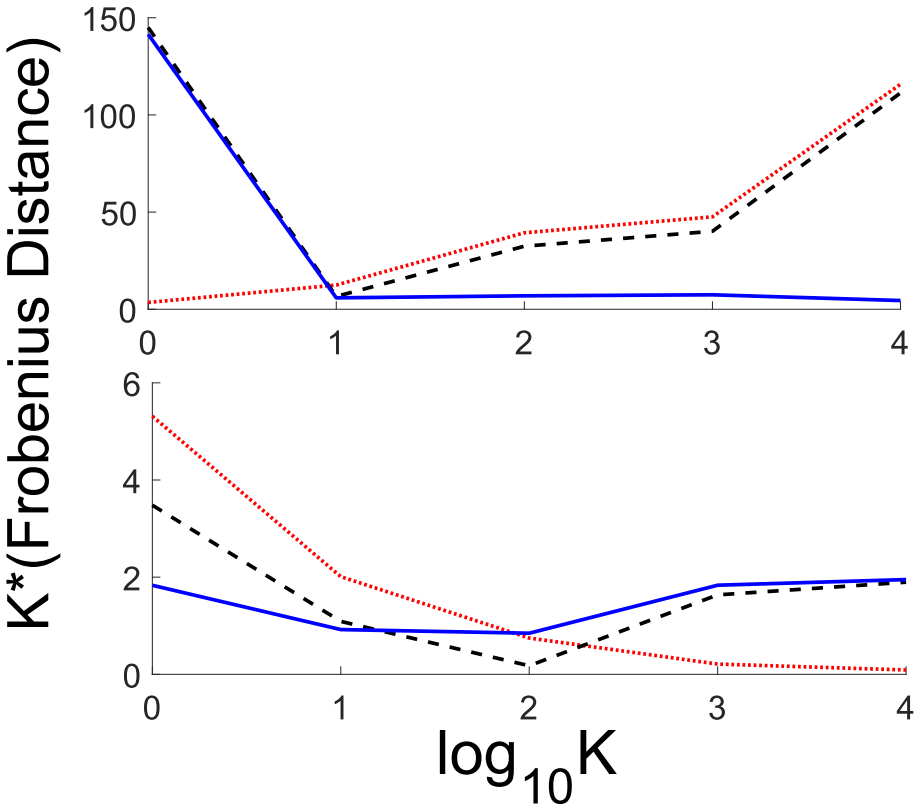

Figure 6.4: Comparison of covariance estimation routines on a chemical reaction network model. *Sample covariance matrix of* $\hat{\theta}(\hat{\mathcal{G}})$ *taken by running* 1000 *identifications. Graphs show Frobenius distance between sample covariance and covariance estimators* $\mathfrak{C}(\hat{\mathcal{G}})$ *(red, dotted) and* $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$ *(black, dashed). The blue, solid line represents Frobenius distance between* $\mathfrak{C}(\hat{\mathcal{G}})$ *and* $\mathcal{I}\big(\hat{\theta}(\hat{\mathcal{G}})\big)^{-1}$*. The* Y*-axis is scaled by* K *to cancel the effect of absolute decrease in covariance with increasing* K*. All graphed distances are monotonically decreasing with* K *in the absence of this scaling. Covariance estimates are taken at increasing integer values of* $\log_{10} K$*, starting from* 0*. Model dynamics given by* (6.15) *for both figures. TOP: Process dynamics are model dynamics at* $\theta^* = [3, 4]$*. BOTTOM: Undermodelled case: process dynamics given by* (6.16)*.*

6.6 Summary

We have considered the problem of parameter estimation in models with finite dimensional output fitted to data corrupted by Gaussian measurement noise. In particular, we have focused on the sensitivity of the estimated parameters to measurement noise. Existing methods of quantification assume that the model structure can perfectly interpolate the data, for some (unknown) parameter vector. We have produced a new method of quantification that agrees asymptotically with the existing method, given this assumption. However, it remains mathematically valid when such interpolation is not possible, and no parameter vector can recreate the noise-free process dynamics.

We have produced an engineered example in which parameter estimation is not well-posed, due to the form of the synthetic data generated. However, the model used in the estimation is structurally identifiable, and, as such, a modeller would not realise this lack of well-posedness without knowing the form of the data-generating process itself. This highlights that even structural identifiability is not sufficient for a well-posed parameter estimation routine when the process is undermodelled. In the example, our covariance quantification method detected the ill-posedness of the problem, whereas the existing method did not. We have also produced a more natural example, demonstrating agreement of our method with the existing one in the absence of undermodelling, and asymptotic superiority of our method in its presence.

Chapter 7

Conclusions

An overarching theme of this thesis has been the analysis of mathematical models through consideration of their behaviour as parameters are varied. This approach was used to undertake tasks including model simplification (Chapter 5), assessment of model suitability in the parameter estimation problem (Chapter 4), and quantification of model performance over regions of parameter space (Chapter 3). Such challenges are of increasing importance as computational advances allow for ever-more detailed models, containing many parameters, whose high-level properties may not be discernable solely from intuition and simulation. Multiple approaches were employed, which nevertheless shared several flavours:

- We emphasised the importance of non-locality in any quantification of parameter sensitivity, and provided mathematical analyses, and algorithms, reflecting this. We demonstrated that linearised analyses of parametric sensitivity can yield erroneous conclusions on models as simple as a damped harmonic oscillator.

- We analysed model behaviour over continuous ranges of parameter space, rather than relying on methods based on repeated simulation at discrete parameter

points. This allows for recovery of low-dimensional correlations in parameter space that might otherwise be missed, and provides model guarantees valid *for every* parameter within a region.

- The algorithms given in this thesis are all applicable to models containing non-linearities. During the period of this thesis, we have consistently found linear models to be a fruitful testing ground for formulating hypotheses and developing intuition. Nevertheless, the analysis of linear systems is a mature field, and the increasingly detailed mathematical models now entering the literature frequently possess nonlinearities integral to their functioning. Tools of model analysis must reflect these developments.

It should be noted that sacrifices had to be made in order to allow our analyses to enjoy the benefits previously listed. The scalability of those techniques presented in Chapter 3 that exploited sum-of-squares (SOS) programming is limited. Hopefully recent and future improvements in the computational complexity of SOS programming, such as those based on relaxation [1], or sparsity exploitation [142], will improve scalability in the future. A more fundamental limitation of the algorithms of Chapter 3 is the requirement of all amenable models to be asymptotically stable to an equilibrium point. Thus the methods of the chapter are inapplicable to common classes of model such as oscillators. The algorithm introduced in Chapter 5, by contrast, enjoys a high of scalability and general applicability. However it is only suitable for regional analyses of a model in parameter space, in that it cannot discriminate between locally and globally identifiable parameter vectors.

7.1 Summary of Contributions

We summarise the most important contributions of the thesis.

- The usage of storage functions as a tool for upper bounding model-related performance criteria dates back to the work of Willems in [135]. The nature of the convex optimisation routines nowadays used to construct storage functions is such that they will fail due to unboundedness if modified to provide lower bounds. In Chapter 3, we develop an approach that can create two-sided bounds, which, moreoever, are parameter-dependent. Importantly, our approach has novel theoretical guarantees on tightness of the bounds. Consequently, we can provide upper and lower bounds on the discrepancy between two models of wind turbines that have a near-zero residual. We believe that the method will prove useful not only in the model discrimination problem, but also in the parameter estimation problem, where it can be used to algebraically describe the set of parameters inducing model outputs close to some reference output.

- The duality between parametric sensitivity and the well-posedness of parameter estimation from data has been appreciated ever since the seminal work of Rao [103]. Many existing approaches to exploration of this duality rely on the assumption of infinitesimal measurement uncertainty (see for example the concept of model sloppiness introduced in [21]). In Chapter 4 we provide a rigorous theoretical framework through which the duality can be explored when measurement uncertainty is significant, and, correspondingly parametric sensitivity is formulated with respect to large perturbations. We show that inferences drawn under the assumption of infinitesimal measurement uncertainty can be very misleading when extrapolated to the non-infinitesimal case, even for simple classes of models such as those with linear dynamics and linear parameter dependence. The suitability of high-dimensional, complex models to parameter estimation is commonly assessed through an infinitesimal analysis. As such, the results of the chapter suggest that a more careful analysis is needed for conclusions made

on such a basis to be valid.

- Calculation of the nonlocal parametric sensitivity properties of high-dimensional models is a computationally challenging task. This motivates the current preponderance of local analyses, which, as stated previously, may lead to flawed conclusions. Chapter 5 introduces a scalable algorithm for finding curves in parameter space over which model output varies minimally. This allows for the delineation of functional relations in parameter space over which parameter estimation is impossible, or ill-posed. Crucially, quantification of the well-posedeness of parameter estimation is based on the framework of Chapter 4, and thus valid for non-infinitesimal measurement uncertainty. The algorithm represents a new approach to finding structurally unidentifiable combinations of parameters, a problem that has a long history in the literature (see e.g. [10]). It is distinguished by its scalability, applicability to very general model classes, amenability to explicit incorporation of data, and provision of unidentifiable functional relations on parameter space, as opposed to a list of structurally unidentifiable parameters.

- Much of the statistical theory underlying parameter estimation from noisy data makes use of the perfect model assumption. This requires the existence of an (unknown) parameter vector inducing model output that perfectly recreates the noise-free behaviour of the modelled process. The assumption is unlikely to be valid in reduced models of complex processes that willingly omit some degree of mechanistic detail in order to maintain computational tractability. It is also unlikely to be valid where some details of the process are unknown and/or mismodelled. Even where the perfect model assumption does not hold, model analysis can be used to draw meaningful conclusions about the process under consideration. However, as we show in Chapter 6, standard quantifications of

the well-posedness of parameter estimation lose their validity. We introduced a new approximation of the uncertainty associated with parameter estimation that holds independently of the perfect model assumption. When the latter does hold, our approximation corresponds asymptotically with the standard approximation: the inverse Fisher Information Matrix. We believe that this contribution will be useful in assessing the accuracy of parameter estimates in simplified mathematical models of complex processes.

7.2 Future Research Directions

We now propose several avenues of research based upon the results of this thesis.

- Recall the convex optimisation routine used in Chapter 3 to give two-sided, parameter-dependent bounds on the discrepancy between the outputs of two models over parameter space. Novelty lay in both the tractability of the lower bound, and the tightness of the bounds, which was demonstrated both practically and theoretically. We believe that modifications of the optimisation routine could produce tighter, two-sided bounds on a variety of model performance criteria for which only loose, one-sided bounds are currently available. These criteria could include the input-to-state and state-to-output gain of controlled systems. This would allow for tighter guarantees of controller performance over regions of parameter space.

- Gain scheduling is the technique of approximating a nonlinear model by a family of linear models, linearised around a grid of operating points in parameter space. It is widely used in the control of nonlinear models, where a controller is synthesised for each of the linear models, and a switching algorithm chooses the appropriate operating point, based on measured parameter values. Choos-

ing where to place the operating points in parameter space, and how to switch between them, is an open problem [13]. A grid of operating points that is too dense leaves the switching algorithm prone to chattering, i.e. constant switching between linear controllers as the parameters traverse the boundary of the respective operating regions [108]. In the worst case this can induce instability of the controlled system, and in the best case it degrades controller performance. A grid that is too sparse means that the linear model being used is often a poor approximation of the true dynamics, and so a conservative controller with poor performance properties is necessary. The algorithm of Chapter 3 makes it possible to express bounds on the discrepancy between the nonlinear model and its linearisation around a particular operating point as an algebraic function of the parameters. Therefore a scheme for operating point placement and switching could be based on the values attained by this algebraic function.

- Chapter 5 presented an algorithm producing curves in parameter space over which model output varied little. Quantification of the variation in model output was given by a cost function. We considered the construction of cost functions for models of processes that were assumed to be deterministic, albeit with data obscured by measurement noise. An interesting topic of research would be quantification of the variability in output between models whose dynamics were driven by stochasticity. Such models include those described by Markov processes, mixed-effects models, or stochastic differential equations. We would expect that a cost function, in these cases, would be based on the difference between moments of the probability distribution of model output at different parameter vectors. If such a cost function could be constructed, and tractably calculated, then the algorithm of Chapter 5 could be extended to account for stochastic systems. The motivation for doing so would be the same as that for deterministic systems. However, theory relating to the parametric analysis of

stochastic systems is less well-developed than the deterministic case.

- The standard protocol for establishing the uncertainty region around a parameter estimate is to find the set of parameter vectors inducing model output 'reasonably' close to that of the parameter estimate. The implicit question is 'what perturbations of the estimated parameters do not perturb model outputs much?'. The results of Chapter 6 modified this problem, in that we searched for the set of parameter vectors inducing model output 'reasonably' close to that of the experimental data. A question that builds on this would be: 'what plausible perturbations of the experimental data have the capacity to highly change the set of well-fitting parameters?'. The estimated uncertainty of a parameter estimate was given as a direct, tractable function of the experimental data in Chapter 6. Therefore we believe that further development of this work could go some way to answering the previous question. The result would be confidence regions in parameter space that were acknowledged as robust (or not) to fluctuations in the data.

Appendix A

Convergence of the Integrated Square Error between ODE systems Sharing an Equilibrium Point

In Chapter 3, we considered pairs of dynamical systems sharing an equilibrium point. We claimed that the integral over time of the square error between the two systems trajectories was finite, under weak regularity conditions. In this chapter we formalise the claim, and provide a proof.

Take two dynamical systems of the form considered in equation (3.1) of Chapter 3, namely

$$\dot{x}(t) = f(x(t), \theta), \qquad y(t) = h_1(x(t), \theta) \tag{A.1}$$

$$\dot{z}(t) = g(z(t), \theta), \qquad r(t) = h_2(z(t), \theta). \tag{A.2}$$

Let us fix θ at an arbitrary value such that $x^*(\theta)$ and $z^*(\theta)$ exist, and $y^*(\theta) = r^*(\theta)$. Assume that there exists a compact domain R^x such that both the state trajectories are asymptotically stable to the value $x^*(\theta) \in R^x$, for $x_0 \in R^x$. Assume further that f, g, h_1, h_2 are $\mathcal{C}^2$ functions. Finally, take A_x as the linearisation of the trajectory $x(t;\theta;x_0)$ around its equilibrium point $x^*(\theta)$, i.e.

$$A_x = \frac{\partial f}{\partial x}\Big|_{x=x^*(\theta)}. \tag{A.3}$$

Assume that A_x and A_z are nonsingular. Then

$$\int_0^\infty \|y(t,\theta;x_0) - r(t,\theta;z_0)\|_2^2 \, dt \tag{A.4}$$

converges to a finite value, for all $x_0 \in R^x$.

Proof. Our proof will take several steps. We first prove that:

$$\begin{aligned} \int_0^\infty \|x(t;\theta,x_0) - x^*(\theta)\|_2^2 \, dt < \infty \\ \int_0^\infty \|z(t;\theta,z_0) - z^*(\theta)\|_2^2 \, dt < \infty \end{aligned} \tag{A.5}$$

holds for the system (A.1), and the analogous condition is true for system (A.2). Since continuous differentiability of h_1 and h_2 implies that they are Lipschitz continuous on any closed set within R^x containing the origin, finiteness of (A.5) implies:

$$\begin{aligned} \int_0^\infty \|h_1(x(t;\theta,x_0),\theta) - h_1(x^*(\theta),\theta)\|_2^2 \, dt < \infty \\ \int_0^\infty \|h_2(z(t;\theta,z_0),\theta) - h_2(z^*(\theta),\theta)\|_2^2 \, dt < \infty \end{aligned} \tag{A.6}$$

The sum of the two finite quantities in (A.6) provide an upper bound to (A.4) through the triangle inequality, implying the result.

We first show that

$$\int_0^\infty \|x(t,\theta;x_0) - x^*(\theta)\|_2^2 \, dt \tag{A.7}$$

is finite. Without loss of generality, we take $x^*(\theta) = 0$. Since A_x is nonsingular, the Hartmann Theorem (Perko; Differential Equations and Dynamical Systems, Chapter 2, Page 127, 3rd Edition 2001), implies the existence of a $T \in \mathbb{R}^+$ and open neighbourhood $N \ni 0$, such that, for all $t > T$, $x(t,\theta;x_0) \in N$. Furthermore, it guarantees existence of a continuously differentiable homeomorphism $h : N \to \mathbb{R}^n$ such that

$$h(0) = 0 \tag{A.8}$$

$$\forall\ t^* > 0,\ a \in N : h\big(x(t^*;\theta,a)\big) = e^{A_x t^*} h(a). \tag{A.9}$$

Note that the Hartmann theorem relies on nonsingularity of A_x. By the definition of open sets, there exists $\epsilon > 0$ such that the the closed ball of radius $\frac{\epsilon}{2}$, centred at 0 is strictly contained within N. We shall denote this set B. Note that, by asymptotic stability, there exists $T^\dagger$ such that $t > T^\dagger \Rightarrow x(t,\theta;x_0) \in B$. Closedness of B and continuous differentiability of h on B together imply that there exists K_l, K_u such that

$$K_l\|x - x^*(\theta)\| \leq \|h(x)\| \leq K_u\|x - x^*(\theta)\| \tag{A.10}$$

The result follows as, for all $a \in B$:

$$\int_0^\infty \|h(x(t,\theta;a))\| \, dt = \int_0^\infty \|h(\phi_t(a))\| \, dt = \int_0^\infty \|e^{A_x t} h(a)\| \, dt < \infty, \tag{A.11}$$

and (A.10) implies that

$$\int_0^\infty \|x(t,\theta;a) - x^*(\theta)\|\ dt \leq \frac{1}{K_l}\int_0^\infty \|h(x(t,\theta;a))\|\ dt \tag{A.12}$$

We chose a arbitrarily within N, and know that we can therefore pick a such that it is a point on the trajectory of $x(t;\theta,x_0)$ after finite time, as asymptotic stability ensures that the trajectory eventually, irreversibly, enters N. Equation (A.5) follows. □

Appendix B

Optimal Control

The field of optimal control is concerned with choosing (time-varying) inputs to a dynamical system that allow it to optimise desired performance criteria. We are interested in the case for which the system and inputs are treated as continuous functions of time. Optimal control problems are a form of optimisation problem: one searches for an input that minimises some cost (equivalently, optimises some objective) subject to constraints. They are distinguished from the class of optimisation problems considered in Section 2.3 of the Preliminaries in that the objective and constraints are infinite-dimensional. In Section 2.3, the cost (equivalently, objective) and constraint functions previously considered were functions of a finite-dimensional variable $x \in \mathbb{R}^n$. Here, they are **functionals** on the space of system trajectories, which are indexed by time. For instance, we may want to control the acceleration/braking of a car so as to minimise the total work done in achieving a particular driving maneouvre. The cost functional is then the integral of power applied over time (equal to total work done), and the trajectory of power output, as a time-varying function, is infinite-dimensional.

There are many classes of optimal control problem, and a large body of literature

devoted to the theoretical and numerical analysis of their solutions (see e.g. [22],[67] for reviews). We can broadly divide between solution styles based on the Hamilton-Jacobi-Bellman (HJB) equations, and those based on Pontryagin's Minimum Principle. Solving the HJB equation associated with a given optimal control problem provides a globally optimal control strategy. However, the equations take the form of nonlinear partial-differential-equations, and are generally both impossible to solve analytically, and hard to approximate numerically. A solution approach based on Pontryagin's Minimum Principle formulates conditions on an optimal control that are necessary but not (usually) sufficient. However, such an approach is generally more tractable, as the conditions only need be satisfied over a single trajectory of the dynamical system. The HJB conditions, by contrast, apply to all possible system trajectories. We only consider optimal control strategies based on Pontryagin's Minimum Principle in this thesis. Our choice is based on the fact that for the particular problem we consider, a finite (and usually small) number of candidate control strategies synthesised using Pontryagin's Minimum Principle will be guaranteed to contain the globally optimal strategy.

This chapter provides the steps necessary to solve a particular class of control problem. We eventually arrive at a set of necessary conditions on the optimal control. These necessary conditions are further analysed in the main body of the thesis. We will deal with dynamical systems of the following form:

$$\dot{x}(t) = f(x(t), u(t)) \qquad x(0) = x_0 \tag{B.1a}$$

$$x(t) \in \mathbb{R}^n \qquad u(t) \in \mathbb{R}^p. \tag{B.1b}$$

Here $u(t)$ is the system **input** at time t. We can choose the dynamics of $u(t)$, and thereby manipulate the trajectory of the system state $x(t)$. Note that $u(t)$ may be subject to constraints. For instance, the acceleration we can demand of a car has

certain physical limits. Note that the main body of the thesis does not require (or allow for) the dynamical system (B.1) to correspond to the mathematical model under consideration. In fact, the state space of the system (B.1) will correspond to the parameter space of the analysed model. We can now define the particular class of optimal control problem that this chapter is devoted to solving.

Problem B.1. ***Standard optimal control problem***

$$\begin{aligned} &\textit{Minimise} \quad J(u) = \int_0^{t_F} L\big(x(t), u(t)\big)\ dt \\ &\textit{subject to} \quad K_1\big(x(t), u(t)\big) \leq 0; \qquad\qquad K_2\big(x(t), u(t)\big) = 0. \end{aligned} \tag{B.2}$$

Here t_F is a fixed end-time. $J(u)$ is referred to as the ***cost functional****, while the scalar integrand $L\big(x(t), u(t)\big) \in \mathbb{R}$ is referred to as the* ***Lagrangian****. K_1 and K_2 represent coupled (in)equality constraints on the input and state trajectories. We assume that the Lagrangian is differentiable, and that both the input and state trajectories are continuous.*

We first consider the effect of applying infinitesimal input perturbations on the cost functional. Let us take an arbitrary function $\delta u : [0, t_F] \to \mathbb{R}^p$ such that $J\big(u + \epsilon(\delta u)\big)$ is defined for any sufficiently small $\epsilon > 0$. In perturbing the input by $\epsilon(\delta u)$, we induce a change in the state at any time t. We therefore define

$$\delta x(t) := \lim_{\epsilon \to 0} \frac{1}{\epsilon} \int_0^t f\big(x(t'), u(t') + \epsilon(\delta u(t'))\big) - f\big(x(t'), u(t')\big)\ dt' \tag{B.3a}$$

$$= \int_0^t \left\langle \frac{\partial f}{\partial u}\big(x(t'), u(t')\big), \delta u(t') \right\rangle\ dt', \tag{B.3b}$$

where the final equality follows by application of the chain rule. So δx is a first order approximation of the change in the state $x(t)$ induced by infinitesimally perturbing the input. This approximation is valid in the limit of decreasing perturbation magnitude.

We also see that δx is linear in δu, by inspection of (B.3b). Note that the functions δx and δu are distinct from the partial derivatives ∂x and ∂u.

We can now define the **first variation** of J at u as

$$\delta J(\delta u) := \lim_{\epsilon \to 0} \frac{J\big(u + \epsilon(\delta u)\big) - J(u)}{\epsilon} \tag{B.4a}$$

$$= \int_0^{t_F} \left[\lim_{\epsilon \to 0} \frac{L\big(x(t), u(t) + \epsilon(\delta u(t))\big) - L\big(x(t), u(t)\big)}{\epsilon} \right] \, dt \tag{B.4b}$$

$$= \int_0^{t_F} \left[\left\langle \frac{\partial L}{\partial x}\big(x(t), u(t)\big), \delta x(t) \right\rangle + \left\langle \frac{\partial L}{\partial u}\big(x(t), u(t)\big), \delta u(t) \right\rangle \right] \, dt. \tag{B.4c}$$

We see that δJ is a functional in δu whose output quantifies the effect on $J(u)$ of infinitesimally perturbing the input $u(t)$ in the direction $\delta u(t)$. Furthermore, equation (B.4c), which follows by direct application of the chain rule, shows us that δJ is a linear functional of δu (since δx is linear in δu). Note that in deriving (B.4c), we have assumed both $x(0)$ and t_F are fixed. The first variation δJ has an alternative, more general form (derived using the Leibniz integral rule) when the latter quantities are allowed to vary. However this generalisation is unnecessary for our purposes.

We now reformulate our expression for δJ so as to remove the explicit dependence on δx. First, pick any function $\lambda : [0, t_F] \to \mathbb{R}^n$. We will refer to λ as the **costate** function, and impose constraints on its behaviour subsequently. For now, regardless of λ, we have that

$$J(u) = \int_0^{t_F} L(x(t), u(t)) + \big\langle \lambda(t), \big[\dot{x}(t) - f(x(t), u(t))\big] \big\rangle \, dt, \tag{B.5}$$

as $\dot{x}(t) - f(x(t), u(t)) = 0$ by definition.

We then define the **Hamiltonian** function as

$$H(x, u, \lambda) = L(x, u) + \langle \lambda, f(x, u) \rangle, \tag{B.6}$$

which allows for reformulation of (B.5) as

$$J(u) = \int_0^{t_F} H\big(x(t), u(t), \lambda(t)\big) - \big\langle \lambda(t), f\big(x(t), u(t)\big)\big\rangle \; dt.$$

Integrating by parts with respect to time t, we get

$$\begin{aligned} J(u) = &\int_0^{t_F} H\big(x(t), u(t), \lambda(t)\big) + \langle \dot{\lambda}(t), x(t)\rangle \; dt \\ &+\langle \lambda(0), x_0\rangle - \langle \lambda(t_F), x(t_F)\rangle. \end{aligned}$$

From this, the first variation of J, given fixed x_0 and free $x(t_F)$, is given by

$$\begin{aligned} \delta J(\delta u) = &\int_0^{t_F} \left\langle \left[\frac{\partial H}{\partial x}\big(x(t), u(t), \lambda(t)\big) + \dot{\lambda}(t)\right], \delta x(t)\right\rangle \; dt \\ &+ \int_0^{t_F} \left\langle \frac{\partial H}{\partial u}\big(x(t), u(t), \lambda(t)\big), \delta u(t)\right\rangle \; dt \\ &- \lambda(t_F)\delta x(t_F). \end{aligned} \tag{B.7}$$

Since our choice of costate $\lambda(t)$ has so far been arbitrary, we now impose the following constraints:

$$\dot{\lambda}(t) = -\frac{\partial H}{\partial x}\big(x(t), u(t), \lambda(t)\big); \qquad\qquad \lambda(t_F) = 0.$$

This means that (B.7) becomes

$$\delta J(\delta u) = \int_0^{t_F} \left\langle \frac{\partial H}{\partial u}\big(x(t), u(t), \lambda(t)\big), \delta u(t)\right\rangle \; dt. \tag{B.8}$$

We have finally arrived at an expression for δJ that is explicitly a linear functional in δu, with no explicit dependence on δx. The first stage in most optimal control solutions relying on Pontryagin's Minimum Principle consists of formulating an analogous expression of δJ in terms of δu. The mathematical techniques used to derive

the expression may however, vary, dependent on the specific problem data.

Proposition B.1. ***Pontryagin's Minimum Principle***

Suppose $u^*(t)$ is an optimal control, in that it is a minimiser of $J(u)$. Denote by $x^*(t)$ and $\lambda^*(t)$ the associated state and costate trajectories. Then, at all timepoints $t \in [0, t_F]$, we must have that

$$H\big(x^*(t), u^*(t), \lambda^*(t)\big) \leq H\big(x^*(t), u(t), \lambda^*(t)\big),$$

for any admissible control $u(t)$.

A rigorous derivation of the principle is provided in [17]. A weaker proposition that is sufficient for our purposes, and implied by Pontryagin's Minimum Principle, is now provided.

Proposition B.2. ***Local minimality of the Hamiltonian***

Suppose $u^*(t)$ is an optimal control, in that it is a minimiser of $J(u)$. Denote by $x^*(t)$ and $\lambda^*(t)$ the associated state and costate trajectories. Then, at all timepoints $t \in [0, t_F]$, for all admissible input perturbations δu, we must have

$$H\big(x^*(t), u^*(t), \lambda^*(t)\big) \leq H\big(x^*(t), u^*(t) + \epsilon(\delta u(t)), \lambda^*(t)\big),$$

for sufficiently small ϵ.

Proof. We provide an outline of the argument, while omitting some technical details.

First note that $\delta J(\delta u) \geq 0$, for any δu, by local minimality. Furthermore

$$\begin{aligned}\delta J(\delta u) &= \int_0^{t_F} \left\langle \frac{\partial H}{\partial u}\big(x(t), u(t), \lambda(t)\big), \delta u(t) \right\rangle \, dt \\ &= \lim_{\epsilon \to 0} \frac{1}{\epsilon} \int_0^{t_F} \left[H(x^*(t), u^*(t) + \epsilon(\delta u(t)), \lambda^*(t)) - H(x^*(t), u^*(t), \lambda^*(t)) \right] \, dt,\end{aligned}$$

so the integral over time of the Hamiltonian associated with the optimal control $u^*(t)$ must be locally minimal among the set of possible inputs. This in fact implies minimality of the Hamiltonian, with respect to u, at any timepoint $\hat{t}$. If not, then there would exist some admissible δu satisfying, for sufficiently small ϵ,

$$\frac{1}{\epsilon}\Big[H\big(x^*(t), u^*(t) + \epsilon(\delta u(t)), \lambda^*(t)\big) - H\big(x^*(t), u^*(t), \lambda^*(t)\big) \Big] < 0, \ \forall t \in [t_a, t_b],$$

where $[t_a, t_b]$ is some closed interval in $[0, t_F]$, containing $\hat{t}$. The validity of the inequality over an interval, rather than a point, is justified by the assumed continuity of the Hamiltonian. We could then take

$$u_\epsilon(t) = \begin{cases} u^*(t) & \text{if } t \notin [t_a, t_b]; \\ u^*(t) + \epsilon(\delta u(t)) & \text{if } t \in [t_a, t_b]. \end{cases}$$

As a consequence, we would have

$$\begin{aligned}\lim_{\epsilon \to 0} \frac{1}{\epsilon} \int_0^{t_F} \left[H\big(x^*(t), u_\epsilon(t), \lambda^*(t)\big) - H\big(x^*(t), u^*(t), \lambda^*(t)\big) \right] \, dt < 0, \\ \Rightarrow \delta J(u_1 - u^*) < 0\end{aligned}$$

contradicting local minimality. □

An immediate consequence of Proposition B.2 is that, at any point in time, the

optimal control $u^*(t)$ is a local minimum of the optimisation problem

$$\begin{aligned} &\underset{u\in\mathbb{R}^p}{\text{Minimise}} \quad H\big(x^*(t), u, \lambda^*(t)\big) \\ &\text{subject to} \quad K_1\big(x(t), u\big) \leq 0; \qquad K_2\big(x(t), u\big) = 0. \end{aligned} \tag{B.9}$$

To obtain necessary conditions on this minimisation, let us define the **augmented Hamiltonian**

$$\bar{H}(x, u, \lambda, \mu) := H(x, u, \lambda) + \mu_1 K_1(x, u) + \mu_2 K_2(x, u),$$

where μ is chosen such that $\mu_1 \geq 0$ and $\mu_1 K_1 = 0$. Note that the numerical values of $\bar{H}$ and H always coincide over feasible trajectories of the optimal control problem, since the terms $\mu_i K_i$ are necessarily zero. In fact, we can replace every previous incidence of the Hamiltonian H in this chapter with $\bar{H}$, and not affect the validity of the arguments provided, although the functional form of the modified costate derivative $\dot{\lambda} = -\frac{\partial \bar{H}}{\partial x}$ will change. Thus we get

$$\delta J(\delta u) = \int_0^{t_F} \left\langle \frac{\partial \bar{H}}{\partial u}\big(x(t), u(t), \lambda(t), \mu(t)\big), \delta u(t) \right\rangle \, dt. \tag{B.10}$$

The KKT conditions of Lemma 2.1 now imply that a necessary condition on the minimisation (B.9) is that

$$\frac{\partial \bar{H}}{\partial u}\big(x^*(t), u^*(t), \lambda^*(t), \mu^*(t)\big) = 0, \ \forall t \in [0, t_F]. \tag{B.11}$$

Note the implicit assumption that the optimisation (B.9) has sufficient regularity for the KKT conditions hold. This is shown for the specific problems encountered in the main body of the thesis.

An important point to note is that the (augmented) Hamiltonian $\bar{H}$ is time-invariant, as

$$\frac{\mathrm{d}\bar{H}}{\mathrm{d}t} = \frac{\partial \bar{H}}{\partial x}\dot{x}(t) + \frac{\partial \bar{H}}{\partial \lambda}\dot{\lambda}(t) + \frac{\partial \bar{H}}{\partial u}\dot{u}(t). \tag{B.12}$$

The first two terms of (B.12) cancel, as $\frac{\partial \bar{H}}{\partial x} = -\dot{\lambda}$, and $\frac{\partial \bar{H}}{\partial \lambda} = \dot{x}$. Finally, nullity of the third term is a necessary condition on the optimal control. Note that we need not take partial derivatives of $\bar{H}$ with respect to the multipliers $\mu_i(t)$, as the value of the Hamiltonian is independent of their values. In any case, they are not guaranteed to be differentiable functions of time.

The constancy of the value H over time adds further structure to the problem:

$$H = L\big(x(0), u(0)\big) + \langle \lambda(0), \dot{x}(0)\rangle = L\big(x(t_F), u(t_F)\big) + \langle \lambda(t_F), \dot{x}(t_F)\rangle\,,$$

where the last equality occurs due to the condition that $\lambda(t_F) = 0$. So we can set the value of the Hamiltonian as the desired final-value of L. This also imposes a constraint on the as-yet undetermined value $\lambda(0)$. In the specific problems that we encounter later in the thesis, this constraint is combined with (B.11) to completely determine $\lambda(0)$, a trick that is not possible in the general case now being presented.

In summary, we have formulated necessary conditions for a control $u(t)$ to be locally optimal with respect to the optimal control problem defined in Problem B.1. The search for a $u(t)$ satisfying these conditions takes the form of a boundary-value problem (i.e. an ODE with terminal-time constraints) in the coupled dynamical variables $\lambda(t)$ and $x(t)$. Specifically, we have

$$\begin{aligned}
&\dot{x}(t) = f\big(x(t), u(t)\big) \\
&\dot{\lambda}(t) = -\frac{\partial \bar{H}}{\partial x}\big(x(t), u(t), \lambda(t), \mu(t)\big); \qquad\qquad \lambda(t_F) = 0,
\end{aligned}$$

where

$$\bar{H}\big(x(t),u(t),\lambda(t),\mu(t)\big) = L\big(x(t),u(t)\big) + \big\langle \lambda(t), f\big(x(t),u(t)\big)\big\rangle + \mu_1(t)K_1\big(x(t),u(t)\big) + \mu_2(t)K_2\big(x(t),u(t)\big).$$

The indeterminates $u(t)$ and $\mu_1(t)$ are chosen so as to satisfy

1. $\frac{\partial \bar{H}}{\partial u}\big(x(t),u(t),\lambda(t),\mu(t)\big) = 0.$
2. $\mu_1(t)K_1\big(x(t),u(t)\big) = 0,$
3. $\bar{H}\big(x(t),u(t),\lambda(t),\mu(t)\big) = \bar{H}\big(x(0),u(0),\lambda(0),\mu(0)\big)$

while the feasibility constraints

$$K_1\big(x(t),u(t)\big) \leq 0 \qquad\qquad K_2\big(x(t),u(t)\big) = 0,$$

must also hold at all times.

Note that the fixed end-time t_F only enters the necessary conditions provided through the constraint $\lambda(t_F) = 0$. Meanwhile, $\lambda(0)$ must be determined so as to ensure that $\lambda(t_F) = 0$ holds. If, however, the other constraints are sufficient to determine $\lambda(0)$, then we need not define t_F *a priori*. In this case, numerical solution of the necessary conditions above can proceed as an explicit ODE without terminal constraints. If and when the condition $\lambda(t) = 0$ (which is equivalent to $L\big(x(t),u(t),\lambda(t)\big) = \bar{H}$) is satisfied, we can set $t = t_F$ and terminate the ODE.

Appendix C

IL13-induced JAK-STAT Pathway Model

The IL13 JAK-STAT pathway is a metabolic signalling pathway within the cell, which is constitutively activated in various types of lymphoma. It has been modelled as a system of ODE's in [99]. This system was suggested as a benchmark for identifiability analysis in [104]. The original system has a 14-dimensional vector of time-dependent state variables, which we denoted $x(t)$. Each denotes a chemical concentration, details of which are provided in [99, 104]. State evolution is dependent on the parameter vector θ, which is 23-dimensional, with each component representing the rate constant of a chemical reaction. We additionally have a constant u_1 representing the level of extracellular IL13. The 7-dimensional vector of measured variables is denoted $y(t)$. The system equations are:

$$\dot{x}_1(t) = -2.265k_1u_1x_1(t) - k_5x_1(t) + k_6x_2(t)$$

$$\dot{x}_2(t) = k_5x_1(t) - k_6x_2(t)$$

$$\dot{x}_3(t) = k_1 2.265u_1x_1(t) - k_2x_3(t)x_7(t)$$

$$\dot{x}_4(t) = k_2x_3(t)x_7(t) - k_3x_4(t)$$

$$\dot{x}_5(t) = k_3x_4(t) - k_4x_5(t)$$

$$\dot{x}_6(t) = \frac{-k_7x_6(t)\big(x_4(t) + x_3(t)\big)}{1 + k_{13}x_{15}(t)} + x_9k_8x_7(t)$$

$$\dot{x}_7(t) = \frac{k_7x_6(t)\big(x_4(t) + x_3(t)\big)}{1 + k_{13}x_{15}(t)} - x_9k_8x_7(t)$$

$$\dot{x}_8(t) = 0$$

$$\dot{x}_9(t) = -k_9x_9(t)x_7(t) + x_9k_{10}x_{10}(t)$$

$$\dot{x}_{10}(t) = k_9x_9(t)x_7(t) - x_9k_{10}x_{10}(t)$$

$$\dot{x}_{11}(t) = k_{11}x_{10}(t)$$

$$\dot{x}_{12}(t) = -2.265k_{12}u_1x_{12}(t)$$

$$\dot{x}_{13}(t) = 2.265k_{12}u_1x_{12}(t)$$

$$\dot{x}_{14}(t) = \frac{k_{14}x_{11}(t)}{k_{15} + x_{11}(t)} - k_{16}x_{14}(t)$$

$$\dot{x}_{15}(t) = k_{17}x_{10}(t).$$

$$y_1(t) = x_1(t) + x_3(t) + x_4(t)$$

$$y_2(t) = k_{18}\big(x_3(t) + x_4(t) + x_5(t) + x_{13}(t)\big)$$

$$y_3(t) = k_{19}\big(x_4(t) + x_5(t)\big)$$

$$y_4(t) = k_{20}x_7(t)$$

$$y_5(t) = k_{21}x_{11}(t)$$

$$y_6(t) = k_{22}x_{15}(t)$$

$$y_7(t) = x_{14}(t)$$

$$y_8(t) = x_{10}(t).$$

The parameter vector is taken as:

$$\theta = [k_1, k_2, k_3, k_4, k_5, k_6, k_7, k_8, k_9, k_{10}, k_{11}, k_{12}, k_{13}, k_{14}, k_{15}, k_{16}, k_{17}, k_{18}, k_{19}, k_{20}, k_{21}, k_{22}, k_{23}].$$

The published parameter estimates θ^* and initial conditions $x(0)$ are provided in [99, 104]. We assumed that measurements of $y(t;\theta)$ were taken at seventeen timepoints,

between $t = 0$ and $t = 100$ minutes, corresponding to those shown in Figure A of [99]. Model output $\mathcal{Y}(\theta)$ is taken as the concatenation of the measurements $y(t;\theta)$ over these seventeen timepoints.

Appendix D

NF-κB Regulatory Module Model

We consider the mathematical model of the NF-κB regulatory module formulated in [68]. This module forms part of numerous metabolic signalling pathways. Our notation and analysis is identical to that of the JAK-STAT model described in Appendix C. Each state $x_i(t)$ represents the concentration of a chemical, while each parameter describes the rate constant of a chemical reaction. Details of the specific chemicals

and associated reactions are provided in [68]. The system equations are

$$
\begin{aligned}
\dot{x}_1(t) &= k_{prod} - k_{deg}x_1(t) - k_{prod}x_1(t)v(t) \\
\dot{x}_2(t) &= -\big(k_3 + k_{deg} + a_2x_{10}(t) + a_3x_{13}(t)\big)x_2(t) \\
&\quad + t_1x_4(t) + t_2x_5(t) + \big(k_1x_1(t) - k_2x_2(t)x_8(t)\big)v(t) \\
\dot{x}_3(t) &= k_3x_2(t) - k_{deg}x_3(t) + k_2x_2(t)x_8(t)v(t) \\
\dot{x}_4(t) &= a_2x_2(t)x_{10}(t) - t_1x_4(t) \\
\dot{x}_5(t) &= a_3x_2(t)x_{13}(t) - t_2x_5(t) \\
\dot{x}_6(t) &= c_{6a}x_{13}(t) - a_1x_6(t)x_{10}(t) + t_2x_5(t) - i_1x_6(t) \\
\dot{x}_7(t) &= i_1k_vx_6(t) - a_1x_{11}(t)x_7(t) \\
\dot{x}_8(t) &= c_4x_9(t) - c_5x_8(t) \\
\dot{x}_9(t) &= c_2 + c_1x_7(t) - c_3x_9(t) \\
\dot{x}_{10}(t) &= -a_2x_2(t)x_{10}(t) - a_1x_{10}(t)x_6(t) + c_{4a}x_{12}(t) \\
&\quad - c_{5a}x_{10}(t) - i_{1a}x_{10}(t) + e_{1a}x_{11}(t) \\
\dot{x}_{11}(t) &= -a_1x_{11}(t)x_7(t) + i_{1a}k_vx_{10}(t) - e_{1a}k_vx_{11}(t) \\
\dot{x}_{12}(t) &= c_{2a} + c_{1a}x_7(t) - c_{3a}x_{12}(t) \\
\dot{x}_{13}(t) &= a_1x_{10}(t)x_6(t) - c_{6a}x_{13}(t) - a_3x_2(t)x_{13}(t) + e_{2a}x_{14}(t) \\
\dot{x}_{14}(t) &= a_1x_{11}(t)x_7(t) - e_{2a}k_vx_{14}(t) \\
\dot{x}_{15}(t) &= c_{2c} + c_{1c}x_7(t) - c_{3c}x_{15}(t).
\end{aligned}
$$

$$
\begin{aligned}
y_1(t) &= x_7(t) \\
y_2(t) &= x_{10}(t) + x_{13}(t) \\
y_3(t) &= x_1(t) + x_2(t) + x_3(t) \\
y_4(t) &= x_2(t) \\
y_5(t) &= x_{12}(t),
\end{aligned}
$$

while the parameter vector is

$$
\begin{aligned}
\theta = &[k_{prod}, k_{deg}, k_1, k_2, k_3, a_1, a_2, a_3, t_1, t_2, c_{6a}, i_1, k_v, \\
&c_1, c_2, c_3, c_4, c_5, c_{4a}, c_{5a}, i_{1a}, e_{1a}, c_{1a}, c_{2a}, c_{3a}, e_{2a}].
\end{aligned}
$$

Here $v(t)$ represents activity of TNFα, an extracellular signal that excites the regulatory module. We follow [68] in taking this as a Heaviside step function that is turned on 3600 seconds after time zero. Output was taken as the measurement of $y(t)$ at 100 second intervals, up to 50000 seconds. The parameter components c_{1c}, c_{2c}, and c_{3c} were fixed at their published values and omitted from the parameter vector: each is structurally unidentifiable a priori. We see this as they only affect the state x_{15}, which is neither observed nor affects dynamics of any other state.

Bibliography

[1] A. A. Ahmadi and A. Majumdar. Dsos and sdsos optimization: Lp and socp-based alternatives to sum of squares optimization. In *Information Sciences and Systems (CISS), 2014 48th Annual Conference on*, pages 1–5. IEEE, 2014.

[2] J. Anderson and A. Papachristodoulou. On validation and invalidation of biological models. *BMC bioinformatics*, 10(1):132, 2009.

[3] J. Anderson and A. Papachristodoulou. Robust nonlinear stability and performance analysis of an F/A-18 aircraft model using sum of squares programming. *International Journal of Robust and Nonlinear Control*, 23(10):1099–1114, 2013.

[4] M. Anguelova. *Observability and identifiability of nonlinear systems with applications in biology*. PhD thesis, Chalmers University of Technology, 2007.

[5] M. Anguelova, M. Jirstrand, and J. Karlsson. An Efficient Method for Structural Identiability Analysis of Large Dynamic Systems. In *16th IFAC Symposium on System Identification*, volume 16, pages 941–946, 2012.

[6] J. F. Apgar, D. K. Witmer, F. M. White, and B. Tidor. Sloppy models, parameter uncertainty, and the role of experimental design. *Molecular bioSystems*, 6(10):1890–900, 2010.

[7] W. Bateson and G. Mendel. *Mendel's principles of heredity*. Courier Corporation, 2013.

[8] D. J. Bearup, N. D. Evans, and M. J. Chappell. The input–output relationship approach to structural identifiability analysis. *Computer methods and programs in biomedicine*, 109(2):171–181, 2013.

[9] G. Becker and A. Packard. Robust performance of linear parametrically varying systems using parametrically-dependent linear feedback. *Systems & Control Letters*, 23(3):205–215, 1994.

[10] R. Bellman and Åström, K.J. On structural identifiability. *Mathematical Biosciences*, 7(3):329–339, 1970.

[11] A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. SIAM, 2001.

[12] F. D. Bianchi. Gain scheduling control of variable-speed wind energy conversion systems using quasi-LPV models. *Control Engineering Practice*, 13(2):247–255, 2005.

[13] F. D. Bianchi. *Wind turbine control systems: principles, modelling & gain scheduling design*. Advances in Industrial Control. Librairie Lavoisier, 2006.

[14] G. Blekherman, P. A. Parrilo, and R. R. Thomas. *Semidefinite optimization and convex algebraic geometry*, volume 13. SIAM, 2013.

[15] P-A. Bliman. A Convex Approach to Robust Stability for Linear Systems with Uncertain Scalar Parameters. *SIAM Journal on Control and Optimization*, 42(6):2016–2042, 2004.

[16] P-A. Bliman. An existence result for polynomial solutions of parameter-dependent LMIs. *Systems & Control Letters*, 51(3-4):165–169, 2004.

[17] V. G. Boltyanskii, R. V. Gamkrelidze, and L. S. Pontryagin. The theory of optimal processes. i. the maximum principle. Technical report, DTIC Document,

1960.

[18] G. E. P. Box. All models are wrong, but some are useful. *Robustness in Statistics*, page 202, 1979.

[19] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[20] F. Brauer and J. A. Nohel. *The qualitative theory of ordinary differential equations: an introduction*. Courier Corporation, 2012.

[21] K. S. Brown and J. P. Sethna. Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68(2):021904, 2003.

[22] A.E. Bryson. *Applied optimal control: optimization, estimation and control*. CRC Press, 1975.

[23] G. Buzzi-Ferraris and P. Forzatti. A new sequential experimental design procedure for discriminating among rival models. *Chemical Engineering Science*, 38(2):225–232, 1983.

[24] B. Calderhead and M. Girolami. Statistical analysis of nonlinear dynamical systems using differential geometric sampling methods. *Interface Focus*, 2011.

[25] G. Casella and R.L. Berger. *Statistical inference*. Duxbury, 2nd edition, 2002.

[26] E. A. Catchpole and B. J. T. Morgan. Detecting parameter redundancy. *Biometrika*, pages 187–196, 1997.

[27] M. J. Chappell and R. N. Gunn. A procedure for generating locally identifiable reparameterisations of unidentifiable non-linear systems by the similarity transformation approach. *Mathematical biosciences*, 148(1):21–41, 1998.

[28] B. H. Chen and S. P. Asprey. On the Design of Optimally Informative Dynamic Experiments for Model Discrimination in Multiresponse Nonlinear Situations. *Industrial & Engineering Chemistry Research*, 42(7):1379–1390, 2003.

[29] G. Chesi. *Domain of attraction: analysis and control via SOS programming*. Lecture notes in Control and Information Sciences. Springer, 2011.

[30] G. Chesi, A. Garulli, A. Tesi, and A. Vicino. *Homogeneous Polynomial Forms for Robustness Analysis of Uncertain Systems*. Lecture Notes in Control and Information Sciences. Springer, 2009.

[31] O-T. Chis, J.R. Banga, and E. Balsa-Canto. Structural identifiability of systems biology models: a critical comparison of methods. *PloS one*, 6(11):e27755, 2011.

[32] O-T. Chis, J.R. Banga, and E. Balsa-Canto. Sloppy models can be identifiable. *arXiv preprint: arXiv:1403.1417*, 2014.

[33] D. J. Cole, B. J. T. Morgan, and DM Titterington. Determining the parametric structure of models. *Mathematical biosciences*, 228(1):16–30, 2010.

[34] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[35] B.C. Daniels, Y-J. Chen, J.P. Sethna, R.N. Gutenkunst, and C.R. Myers. Sloppiness, robustness, and evolvability in systems biology. *Current opinion in biotechnology*, 19(4):389–95, 2008.

[36] L. Denis-Vidal and G. Joly-Blanchard. An easy to check criterion for (un)indentifiability of uncontrolled systems and its applications. *IEEE Transactions on Automatic Control*, 45(4):768–771, 2000.

[37] L. Denis-Vidal and G. Joly-Blanchard. Equivalence and identifiability analysis of uncontrolled nonlinear dynamical systems. *Automatica*, 40(2):287–292, 2004.

[38] E. Dufresne, H. A. Harrington, and D. V. Raman. The geometry of sloppiness. *arXiv preprint arXiv:1608.05679*, 2016.

[39] M.C. Eisenberg and M.A.L. Hayashi. Determining identifiable parameter combinations using subset profiling. *Mathematical biosciences*, 256:116–126, 2014.

[40] T. Eißing, F. Allgöwer, and E. Bullinger. Robustness properties of apoptosis models with respect to parameter variations and intrinsic noise. *IEE Proceedings-Systems Biology*, 152(4):221–228, 2005.

[41] N. D. Evans and M. J. Chappell. Extensions to a procedure for generating locally identifiable reparameterisations of unidentifiable systems. *Mathematical biosciences*, 168(2):137–159, 2000.

[42] N. D. Evans, M. J. Chappell, M. J. Chapman, and K. R. Godfrey. Structural indistinguishability between uncontrolled (autonomous) nonlinear analytic systems. *Automatica*, 40(11):1947–1953, 2004.

[43] N.D. Evans, M.J. Chappell, and K.R. Godfrey. Identifiability of uncontrolled nonlinear rational systems. *Automatica*, 38(10):1799–1805, 2002.

[44] D. Fernández-Slezak, C. Suárez, G.A. Cecchi, G. Marshall, and G. Stolovitzky. When the optimal is not the best: parameter estimation in complex biological models. *PloS one*, 5(10):e13283, 2010.

[45] P. Gahinet, P. Apkarian, and M. Chilali. Affine parameter-dependent Lyapunov functions and real parametric uncertainty. *IEEE Transactions on Automatic Control*, 41(3):436–442, 1996.

[46] M. Girolami. Bayesian inference for differential equations. *Theoretical Computer Science*, 408(1):4–16, 2008.

[47] M. Girolami and B. Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

[48] K. R. Godfrey, M. J. Chapman, and S. Vajda. Identifiability and indistinguishability of nonlinear pharmacokinetic models. *Journal of pharmacokinetics and biopharmaceutics*, 22(3):229–251, 1994.

[49] G. C. Goodwin and M. E. Salgado. A stochastic embedding approach for quantifying uncertainty in the estimation of restricted complexity models. *International Journal of Adaptive Control and Signal Processing*, 3(4):333–356, 1989.

[50] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS computational biology*, 3(10):1871–78, 2007.

[51] M. J. Hayes, D. G. Bates, and I. Postlethwaite. New tools for computing tight bounds on the real structured singular value. *Journal of Guidance, Control, and Dynamics*, 24(6):1204–1213, 2001.

[52] S. Hengl, C. Kreutz, J. Timmer, and T. Maiwald. Data-based identifiability analysis of non-linear dynamical models. *Bioinformatics*, 23(19):2612–8, 2007.

[53] D. Henrion, J. Lofberg, M. Kocvara, and M. Stingl. Solving polynomial static output feedback problems with penbmi. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 7581–7586. IEEE, 2005.

[54] R. Hermann and A. J. Krener. Nonlinear controllability and observability. *IEEE Transactions on automatic control*, 22(5):728–740, 1977.

[55] G. H. Hines, M. Arcak, and A. K. Packard. Equilibrium-independent passivity: A new definition and numerical certification. *Automatica*, 47(9):1949–1956, 2011.

[56] K.E. Hines, T. R. Middendorf, and R. W. Aldrich. Determination of parameter identifiability in nonlinear biophysical models: A Bayesian approach. *The Journal of General Physiology*, 143(3):401–16, 2014.

[57] D. Hinrichsen and A. J. Pritchard. *Real and complex stability radii: a survey*. Springer, 1990.

[58] W. G. Hunter and A. M. Reiner. Designs for Discriminating Between Two Rival Models. *Technometrics*, 7(3):307–323, 1965.

[59] M. Joshi, A. Seidel-Morgenstern, and A. Kremling. Exploiting the bootstrap method for quantifying parameter confidence intervals in dynamical systems. *Metabolic engineering*, 8(5):447–55, 2006.

[60] E. W. Kamen and P. P. Khargonekar. On the control of linear systems whose coefficients are functions of parameters. *IEEE Transactions on Automatic Control*, 29(1):25–33, 1984.

[61] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 302–311. ACM, 1984.

[62] J. R. Karr, J. C. Sanghvi, D. N. Macklin, M. V. Gutschow, J. M. Jacobs, B. Bolival, N. Assad-Garcia, J. I. Glass, and M. W. Covert. A whole-cell computational model predicts phenotype from genotype. *Cell*, 150(2):389–401, 2012.

[63] H Khalil. *Nonlinear Systems Third Edition*. Prentice Hall Upper Saddle River, 3rd edition, 2002.

[64] A. Kremling, S. Fischer, and K. Gadkar. A benchmark for methods in reverse engineering and model discrimination: problem formulation and solutions. *Genome Research*, 14(9):1773–1785, 2004.

[65] J. B. Lasserre and T. Netzer. Sos approximations of nonnegative polynomials via simple high degree perturbations. *Mathematische Zeitschrift*, 256(1):99–112, 2007.

[66] C. Letellier, L. A. Aguirre, and J. Maquet. Relation between observability and differential embeddings for nonlinear dynamics. *Physical Review E*, 71(6):066213, 2005.

[67] D. Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2012.

[68] T. Lipniacki, P. Paszek, and A.R. Brasier. Mathematical model of NF-κB regulatory module. *Journal of theoretical biology*, pages 195–215, 2004.

[69] Y-Y. Liu, J-J. Slotine, and A-L. Barabási. Observability of complex systems. *Proceedings of the National Academy of Sciences*, 110(7):2460–5, 2013.

[70] L. Ljung. *System identification: Theory for the User*. Birkhauser Boston, 1998.

[71] L. Ljung. Estimating Linear Time-invariant Models of Nonlinear Time-varying Systems. *European Journal of Control*, 7(2-3):203–219, 2001.

[72] L. Ljung and T. Glad. On global identifiability for arbitrary model parametrizations. *Automatica*, 30(2):265–276, 1994.

[73] L. Ljung, B. Wahlberg, and H. Hjalmarsson. Model quality: the roles of prior knowledge and data information. In *Proceedings of the 30th IEEE Conference on Decision and Control*, pages 273–278. IEEE, 1991.

[74] G. Margaria, E. Riccomagno, M. J. Chappell, and H. P. Wynn. Differential algebra methods for the study of the structural identifiability of rational function state-space models in the biosciences. *Mathematical Biosciences*, 174(1):1–26, 2001.

[75] MATLAB. version 8.2.0 (r2013b), 2013.

[76] W. Q. Meeker and L. A. Escobar. Teaching about approximate confidence regions based on maximum likelihood estimation. *The American Statistician*, 49(1):48–53, 1995.

[77] N. Meshkat, C. Anderson, and J. J. DiStefano. Finding identifiable parameter combinations in nonlinear ode models and the rational reparameterization of their input–output equations. *Mathematical biosciences*, 233(1):19–31, 2011.

[78] N. Meshkat, M. Eisenberg, and J. J. Distefano. An algorithm for finding globally identifiable parameter combinations of nonlinear ODE models using Gröbner Bases. *Mathematical biosciences*, 222(2):61–72, 2009.

[79] H. Miao, X. Xia, A.S. Perelson, and H. Wu. On identifiability of nonlinear ODE models and applications in viral dynamics. *SIAM review. Society for Industrial and Applied Mathematics*, 53(1):3–39, 2011.

[80] C.G. Moles, P. Mendes, and J.R. Banga. Parameter estimation in biochemical pathways: a comparison of global optimization methods. *Genome research*, 13(11):2467–74, 2003.

[81] M. C. Neale and M. B. Miller. The use of likelihood-based confidence intervals in genetic models. *Behavior genetics*, 27(2):113–120, 1997.

[82] Y. Nesterov and A. Nemirovskii. *Interior-point polynomial algorithms in convex programming*, volume 13. SIAM, 1994.

[83] M. P. Newlin and R. S. Smith. A generalization of the structured singular value and its application to model validation. *IEEE Transactions on Automatic Control*, 43(7):901–907, 1998.

[84] B. Ninness and G. C. Goodwin. Estimation of model quality. *Automatica*, 31(12):1771–1797, 1995.

[85] A. E. Nordsjö and T. Wigren. On estimation of errors caused by non-linear undermodelling in system identification. *International Journal of Control*, 75(14):1100–1113, 2002.

[86] A. Packard. Gain scheduling via linear fractional transformations. *Systems & Control Letters*, 22(2):79–92, 1994.

[87] A. Packard and G. Becker. Induced L2-norm control for LPV system with bounded parameter variation rates. In *Proceedings of 1995 American Control Conference*, volume 3, pages 2379–2383, 1995.

[88] A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, and P. Parrilo. Sostools version 3.00 sum of squares optimization toolbox for matlab. *arXiv preprint arXiv:1310.4716*, 2013.

[89] A. Papachristodoulou and S. Prajna. A tutorial on sum of squares techniques for systems analysis. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 2686–2700. IEEE, 2005.

[90] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.

[91] D. Peaucelle and D. Arzelier. Robust performance analysis with LMI-based methods for real parametric uncertainty via parameter-dependent Lyapunov functions. *IEEE Transactions on Automatic Control*, 46(4):624–630, 2001.

[92] M. M. Peet and A. Papachristodoulou. A Converse Sum of Squares Lyapunov Result With a Degree Bound. *IEEE Transactions on Automatic Control*, 57(9):2281–2293, 2012.

[93] T. P. Peixoto. Model Selection and Hypothesis Testing for Large-Scale Network Models with Overlapping Groups. *Physical Review X*, 4(1):011033, 2014.

[94] L. Perko. *Differential equations and dynamical systems*. Texts in Applied Mathematics. Springer-Verlag, 1991.

[95] F. Permenter and P. A. Parrilo. Selecting a monomial basis for sums of squares programming over a quotient ring. In *Proceedings of the 51st IEEE Conference on Decision and Control*, pages 1871–1876), 2012.

[96] H Pohjanpalo. System identifiability based on the power series expansion of the solution. *Mathematical biosciences*, 1978.

[97] S. Prajna and H. Sandberg. On Model Reduction of Polynomial Dynamical Systems. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 1666–1671. IEEE, 2005.

[98] T. P. Prescott and A. Papachristodoulou. Guaranteed error bounds for structured complexity reduction of biochemical networks. *Journal of Theoretical Biology*, 304:172–182, 2012.

[99] V. Raia, M. Schilling, M. Böhm, B. Hahn, A. Kowarsch, A. Raue, C. Sticht, S. Bohl, M. Saile, P. Möller, N. Gretz, J. Timmer, F. Theis, W.-D. Lehmann, P. Lichter, and U. Klingmüller. Dynamic mathematical modeling of IL13-induced signaling in Hodgkin and primary mediastinal B-cell lymphoma allows prediction of therapeutic targets. *Cancer research*, 71(3):693–704, 2011.

[100] D. V. Raman, J. Anderson, and A. Papachristodoulou. Delineating parameter unidentifiabilities in complex models. *arXiv preprint arXiv:1607.07705*, 2016.

[101] D. V. Raman, J. Anderson, and A. Papachristodoulou. A new approach for estimating the robustness of parameter estimates to measurement noise. In *2016 American Control Conference (ACC)*, pages 1820–1825. IEEE, 2016.

[102] D. V. Raman, J. Anderson, and A. Papachristodoulou. On the performance of nonlinear dynamical systems under parameter perturbation. *Automatica*, 63:265–273, 2016.

[103] C. R. Rao. Information and accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta. Mathematical Society*, 37:81–91, 1945.

[104] A. Raue, J. Karlsson, and M.P. Saccomani. Comparison of approaches for parameter identifiability analysis of biological systems. *Bioinformatics*, 2014.

[105] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–9, 2009.

[106] B. Reznick. Some concrete aspects of hilbert's 17th problem. *Contemporary Mathematics*, 253:251–272, 2000.

[107] W. Rudin. *Real and complex analysis*. Tata McGraw-Hill Education, 1987.

[108] W. J. Rugh and J. S. Shamma. Research on gain scheduling. *Automatica*, 36(10):1401–1425, 2000.

[109] P. Rumschinski, S. Borchers, S. Bosio, R. Weismantel, and R. Findeisen. Set-base dynamical parameter estimation and model invalidation for biochemical reaction networks. *BMC systems biology*, 4(1):69, 2010.

[110] P. Rumschinski, S. Streif, and R. Findeisen. Combining qualitative information and semi-quantitative data for guaranteed invalidation of biochemical network models. *International Journal of Robust and Nonlinear Control*, 22(10):1157–1173, 2012.

[111] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.

[112] A. Sedoglavic. A probabilistic algorithm to test local algebraic observability in polynomial time. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*. ACM, 2001.

[113] S. M. Shahruz and S. Behtash. Design of controllers for linear parameter-varying systems by the gain scheduling technique. *Journal of Mathematical Analysis and Applications*, 168(1):195–217, 1992.

[114] R. S. Smith and J. C. Doyle. Model validation: a connection between robust control and identification. *IEEE Transactions on Automatic Control*, 37(7):942–952, July 1992.

[115] J. F. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999.

[116] Hector J Sussmann. Existence and uniqueness of minimal realizations of nonlinear systems. *Mathematical Systems Theory*, 10(1):263–284, 1976.

[117] W. A. Sutherland. *Introduction to metric and topological spaces*. Oxford University Press, 1975.

[118] D. T. Suzuki, A. J. F. Griffiths, J. H. Miller, and R. C. Lewontin. *An introduction to genetic analysis*. WH Freeman and Company, 3 edition, 1986.

[119] W. Tan and A. Packard. Stability region analysis using polynomial and composite polynomial Lyapunov functions and sum-of-squares programming. *IEEE Transactions on Automatic Control*, 2008.

[120] U. Topcu and A. Packard. Local robust performance analysis for nonlinear dynamical systems. *IEEE Transactions on Automatic Control*, 2009.

[121] U. Topcu and A. Packard. Local Stability Analysis for Uncertain Nonlinear Systems. *IEEE Transactions on Automatic Control*, 54(5):1042–1047, 2009.

[122] U. Topcu, A. K. Packard, P. Seiler, and G. J. Balas. Robust Region-of-Attraction Estimation. *IEEE Transactions on Automatic Control*, 55(1):137–142, 2010.

[123] M. K. Transtrum, B. B. Machta, and J. P. Sethna. Why are nonlinear fits to data so challenging? *Physical Review Letters*, 104(6):060201, 2010.

[124] M. K. Transtrum, B. B. Machta, and J. P. Sethna. Geometry of nonlinear least squares with applications to sloppy models and optimization. *Physical Review E*, 83(3):036701, 2011.

[125] M. K. Transtrum and P. Qiu. Model reduction by manifold boundaries. *Physical Review Letters*, 113(9):098701, 2014.

[126] S. Vajda, K. R. Godfrey, and H. Rabitz. Similarity transformation approach to identifiability analysis of nonlinear compartmental models. *Mathematical Biosciences*, 93(2):217–248, 1989.

[127] S. Vajda, H. Rabitz, E. Walter, and Y. Lecourtier. Qualitative and quantitative identifiability analysis of nonlinear chemical kinetic models. *Chemical Engineering Communications*, 83(1):191–219, 1989.

[128] M. Vallisneri. Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Physical Review D*, 77(4):042001, 2008.

[129] G. Valmorbida and J. Anderson. Region of attraction analysis via invariant sets. In *2014 American Control Conference*, pages 3591–3596. IEEE, 2014.

[130] J. Vanlier, C.A. Tiemann, P.A.J. Hilbers, and N.A.W. van Riel. Parameter uncertainty in biochemical models described by ordinary differential equations. *Mathematical biosciences*, 246(2):305–14, 2013.

[131] E. Walter and H. Piet-Lahanier. Estimation of parameter bounds from bounded-error data: a survey. *Mathematics and Computers in simulation*, 32(5-6):449–468, 1990.

[132] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna. The sloppy model universality class and the Vandermonde matrix. *Physical Review Letters*, 97(15):150601, 2006.

[133] A. J. Whalen, S. N. Brennan, T. D. Sauer, and S. J. Schiff. Observability and Controllability of Nonlinear Networks: The Role of Symmetry. *Physical Review X*, 5(1):011005, 2015.

[134] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 1938.

[135] J. C. Willems. Dissipative dynamical systems part I: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.

[136] D. Williams. *Probability with martingales*. Cambridge university press, 1991.

[137] F. Wu and S. Prajna. SOS-based solution approach to polynomial LPV system analysis and synthesis problems. *International Journal of Control*, 78(8):600–611, 2005.

[138] X. Xia and C. H. Moog. Identifiability of nonlinear systems with application to HIV/AIDS models. *IEEE Transactions on Automatic Control*, 48(2):330–336,

2003.

[139] J. W.T. Yates, N.D. Evans, and M.J. Chappell. Structural identifiability analysis via symmetries of differential equations. *Automatica*, 45(11):2585–2591, 2009.

[140] D.E. Zak, G.E. Gonye, J.S. Schwaber, and F.J. Doyle. Importance of input perturbations and stochastic gene expression in the reverse engineering of genetic regulatory networks: insights from an identifiability analysis of an in silico network. *Genome research*, 13(11):2396–405, 2003.

[141] X. Zhang. *Parameter-Dependent Lyapunov Functions and Stability Analysis of Linear Parameter-Dependent Dynamical Systems*. PhD thesis, Georgia Institute of Technology, 2003.

[142] Yang Zheng, Giovanni Fantuzzi, Antonis Papachristodoulou, Paul Goulart, and Andrew Wynn. Fast admm for semidefinite programs with chordal sparsity. *arXiv preprint arXiv:1609.06068*, 2016.

[143] K. Zhou, J. C. Doyle, and K. Glover. *Robust and optimal control*. Macmillan, 1996.