Chapter 1
Introduction

> We must think differently about our ideas — and how we test them. We must become more comfortable with probability and uncertainty. We must think more carefully about the assumptions and beliefs that we bring to a problem.
>
> *The Signal and the Noise: The Art of Science and Prediction*
> Nate Silver

1.1 What is Uncertainty Quantification?

This book is an introduction to the mathematics of Uncertainty Quantification (UQ), but what is UQ? It is, roughly put, the coming together of probability theory and statistical practice with 'the real world'. These two anecdotes illustrate something of what is meant by this statement:

- Until the early-to-mid 1990s, risk modelling for catastrophe insurance and re-insurance (i.e. insurance for property owners against risks arising from earthquakes, hurricanes, terrorism, etc., and then insurance for the providers of such insurance) was done on a purely statistical basis. Since that time, catastrophe modellers have tried to incorporate models for the underlying physics or human behaviour, hoping to gain a more accurate predictive understanding of risks by blending the statistics and the physics, e.g. by focussing on what is both statistically *and* physically reasonable. This approach also allows risk modellers to study interesting hypothetical scenarios in a meaningful way, e.g. using a physics-based model of water drainage to assess potential damage from rainfall 10% in excess of the historical maximum.

© Springer International Publishing Switzerland 2015

T.J. Sullivan, *Introduction to Uncertainty Quantification*, Texts in Applied Mathematics 63, DOI 10.1007/978-3-319-23395-6_1

- Over roughly the same period of time, deterministic engineering models of complex physical processes began to incorporate some element of uncertainty to account for lack of knowledge about important physical parameters, random variability in operating circumstances, or outright ignorance about what the form of a 'correct' model would be. Again the aim is to provide more accurate predictions about systems' behaviour.

Thus, a 'typical' UQ problem involves one or more mathematical models for a process of interest, subject to some uncertainty about the correct form of, or parameter values for, those models. Often, though not always, these uncertainties are treated probabilistically.

Perhaps as a result of its history, there are many perspectives on what UQ is, including at the extremes assertions like "UQ is just a buzzword for statistics" or "UQ is just error analysis". These points of view are somewhat extremist, but they do contain a kernel of truth: very often, the probabilistic theory underlying UQ methods is actually quite simple, but is obscured by the details of the application. However, the complications that practical applications present are also part of the essence of UQ: it is all very well giving an accurate prediction for some insurance risk in terms of an elementary mathematical object such as an expected value, but how will you actually go about evaluating that expected value when it is an integral over a million-dimensional parameter space? Thus, it is important to appreciate both the underlying mathematics and the practicalities of implementation, and the presentation here leans towards the former while keeping the latter in mind.

Typical UQ problems of interest include certification, prediction, model and software verification and validation, parameter estimation, data assimilation, and inverse problems. At its very broadest,

> "UQ studies all sources of error and uncertainty, including the following: systematic and stochastic measurement error; ignorance; limitations of theoretical models; limitations of numerical representations of those models; limitations of the accuracy and reliability of computations, approximations, and algorithms; and human error. A more precise definition is UQ is the end-to-end study of the reliability of scientific inferences." (U.S. Department of Energy, 2009, p. 135)

It is especially important to appreciate the "end-to-end" nature of UQ studies: one is interested in *relationships between pieces of information*, not the 'truth' of those pieces of information/assumptions, bearing in mind that they are only approximations of reality. There is always going to be a risk of 'Garbage In, Garbage Out'. UQ cannot tell you that your model is 'right' or 'true', but only that, *if* you accept the validity of the model (to some quantified degree), *then* you must logically accept the validity of certain conclusions (to some quantified degree). In the author's view, this is the proper interpretation of philosophically sound but somewhat unhelpful assertions like "Verification and validation of numerical models of natural systems is impossible" and "The primary value of models is heuristic" (Oreskes et al., 1994). UQ can, however, tell you that two or more of your modelling assumptions are mutually contradictory, and hence that your model is wrong, and a complete UQ analysis will include a meta-analysis examining the sensitivity of the original analysis to perturbations of the governing assumptions.

A prototypical, if rather over-used, example for UQ is an elliptic PDE with uncertainty coefficients:

Example 1.1. Consider the following elliptic boundary value problem on a connected Lipschitz domain $\mathcal{X} \subseteq \mathbb{R}^n$ (typically $n = 2$ or 3):

$$\begin{aligned} -\nabla \cdot (\kappa \nabla u) &= f && \text{in } \mathcal{X}, \\ u &= b && \text{on } \partial\mathcal{X}. \end{aligned} \tag{1.1}$$

Problem (1.1) is a simple but not overly naïve model for the pressure field u of some fluid occupying a domain $\mathcal{X}$. The domain $\mathcal{X}$ consists of a material, and the tensor field $\kappa\colon \mathcal{X} \to \mathbb{R}^{n\times n}$ describes the permeability of this material to the fluid. There is a source term $f\colon \mathcal{X} \to \mathbb{R}$, and the boundary condition specifies the values $b\colon \partial\mathcal{X} \to \mathbb{R}$ that the pressure takes on the boundary of $\mathcal{X}$. This model is of interest in the earth sciences because Darcy's law asserts that the velocity field v of the fluid flow in this medium is related to the gradient of the pressure field by

$$v = \kappa \nabla u.$$

If the fluid contains some kind of contaminant, then it may be important to understand where fluid following the velocity field v will end up, and when.

In a course on PDE theory, you will learn that, for each given positive-definite and essentially bounded permeability field κ, problem (1.1) has a unique weak solution u in the Sobolev space $H_0^1(\mathcal{X})$ for each forcing term f in the dual Sobolev space $H^{-1}(\mathcal{X})$. This is known as the *forward problem.* One objective of this book is to tell you that this is far from the end of the story! As far as practical applications go, existence and uniqueness of solutions to the forward problem is only the beginning. For one thing, this PDE model is only an approximation of reality. Secondly, even if the PDE were a perfectly accurate model, the 'true' κ, f and b are not known precisely, so our knowledge about $u = u(\kappa, f, b)$ is also uncertain in some way. If κ, f and b are treated as random variables, then u is also a random variable, and one is naturally interested in properties of that random variable such as mean, variance, deviation probabilities, etc. This is known as the *forward propagation of uncertainty*, and to perform it we must build some theory for probability on function spaces.

Another issue is that often we want to solve an *inverse problem*: perhaps we know something about f, b and u and want to infer κ via the relationship (1.1). For example, we may observe the pressure $u(x_i)$ at finitely many points $x_i \in \mathcal{X}$; This problem is hugely underdetermined, and hence ill-posed; ill-posedness is characteristic of many inverse problems, and is only worsened by the fact that the observations may be corrupted by observational noise. Even a prototypical inverse problem such as this one is of enormous practical interest: it is by solving such inverse problems that oil companies attempt to infer the location of oil deposits in order to make a profit, and seismologists the structure of the planet in order to make earthquake predictions. Both of these problems, the forward and inverse propagation of uncertainty, fall under the very general remit of UQ. Furthermore, in practice, the domain $\mathcal{X}$ and the fields f, b, κ and u are all discretized and solved for numerically (i.e. approximately and finite-dimensionally), so it is of interest to understand the impact of these discretization errors.

Epistemic and Aleatoric Uncertainty. It is common to divide uncertainty into two types, *aleatoric* and *epistemic* uncertainty. Aleatoric uncertainty — from the Latin *alea*, meaning a die — refers to uncertainty about an inherently variable phenomenon. Epistemic uncertainty — from the Greek ἐπιστήμη, meaning knowledge — refers to uncertainty arising from lack of knowledge. If one has at hand a model for some system of interest, then epistemic uncertainty is often further subdivided into *model form* uncertainty, in which one has significant doubts that the model is even 'structurally correct', and *parametric* uncertainty, in which one believes that the form of the model reflects reality well, but one is uncertain about the correct values to use for particular parameters in the model.

To a certain extent, the distinction between epistemic and aleatoric uncertainty is an imprecise one, and repeats the old debate between frequentist and subjectivist (e.g. Bayesian) statisticians. Someone who was simultaneously a devout Newtonian physicist and a devout Bayesian might argue that the results of dice rolls are not aleatoric uncertainties — one simply doesn't have complete enough information about the initial conditions of die, the material and geometry of the die, any gusts of wind that might affect the flight of the die, and so forth. On the other hand, it is usually clear that some forms of uncertainty are epistemic rather than aleatoric: for example, when physicists say that they have yet to come up with a Theory of Everything, they are expressing a lack of knowledge about the laws of physics in our universe, and the correct mathematical description of those laws. In any case, regardless of one's favoured interpretation of probability, the language of probability theory is a powerful tool in describing uncertainty.

Some Typical UQ Objectives. Many common UQ objectives can be illustrated in the context of a system, F, that maps inputs X in some space $\mathcal{X}$ to outputs $Y = F(X)$ in some space $\mathcal{Y}$. Some common UQ objectives include:

- The *forward propagation* or *push-forward problem*. Suppose that the uncertainty about the inputs of F can be summarized in a probability distribution μ on $\mathcal{X}$. Given this, determine the induced probability distribution $F_*\mu$ on the output space $\mathcal{Y}$, as defined by

$$(F_*\mu)(E) := \mathbb{P}_\mu(\{x \in \mathcal{X} \mid F(x) \in E\}) \equiv \mathbb{P}_\mu[F(X) \in E].$$

This task is typically complicated by μ being a complicated distribution, or F being non-linear. Because $(F_*\mu)$ is a very high-dimensional object, it is often more practical to identify some specific outcomes of interest and settle for a solution of the following problem:

- The *reliability* or *certification problem*. Suppose that some set $\mathcal{Y}_{\text{fail}} \subseteq \mathcal{Y}$ is identified as a 'failure set', i.e. the outcome $F(X) \in \mathcal{Y}_{\text{fail}}$ is undesirable in some way. Given appropriate information about the inputs X and forward process F, determine the failure probability,

$$\mathbb{P}_\mu[F(X) \in \mathcal{Y}_{\text{fail}}].$$

Furthermore, in the case of a failure, how large will the deviation from acceptable performance be, and what are the consequences?

- The *prediction problem*. Dually to the reliability problem, given a maximum acceptable probability of error $\varepsilon > 0$, find a set $\mathcal{Y}_\varepsilon \subseteq \mathcal{Y}$ such that

$$\mathbb{P}_\mu[F(X) \in \mathcal{Y}_\varepsilon] \geq 1 - \varepsilon.$$

i.e. the prediction $F(X) \in \mathcal{Y}_\varepsilon$ is wrong with probability at most ε.

- An *inverse problem*, such as *state estimation* (often for a quantity that is changing in time) or *parameter identification* (usually for a quantity that is not changing, or is non-physical model parameter). Given some observations of the output, Y, which may be corrupted or unreliable in some way, attempt to determine the corresponding inputs X such that $F(X) = Y$. In what sense are some estimates for X more or less reliable than others?
- The *model reduction* or *model calibration problem*. Construct another function F_h (perhaps a numerical model with certain numerical parameters to be *calibrated*, or one involving far fewer input or output variables) such that $F_h \approx F$ in an appropriate sense. Quantifying the accuracy of the approximation may itself be a certification or prediction problem.

Sometimes a UQ problem consists of several of these problems coupled together: for example, one might have to solve an inverse problem to produce or improve some model parameters, and then use those parameters to propagate some other uncertainties forwards, and hence produce a prediction that can be used for decision support in some certification problem.

Typical issues to be confronted in addressing these problems include the high dimension of the parameter spaces associated with practical problems; the approximation of integrals (expected values) by numerical quadrature; the cost of evaluating functions that often correspond to expensive computer simulations or physical experiments; and non-negligible epistemic uncertainty about the correct form of vital ingredients in the analysis, such as the functions and probability measures in key integrals.

The aim of this book is to provide an introduction to the fundamental mathematical ideas underlying the basic approaches to these types of problems. Practical UQ applications almost always require some ad hoc combination of multiple techniques, adapted and specialized to suit the circumstances, but the emphasis here is on basic ideas, with simple illustrative examples. The hope is that interested students or practitioners will be able to generalize from the topics covered here to their particular problems of interest, with the help of additional resources cited in the bibliographic discussions at the end of each chapter. So, for example, while Chapter 12 discusses intrusive (Galerkin) methods for UQ with an implicit assumption that the basis is a polynomial chaos basis, one should be able to adapt these ideas to non-polynomial bases.

A Word of Warning. UQ is not a mature field like linear algebra or single-variable complex analysis, with stately textbooks containing well-polished presentations of classical theorems bearing August names like Cauchy, Gauss and Hamilton. Both because of its youth as a field and its very close engagement with applications, UQ is much more about problems, methods and 'good enough for the job'. There are some very elegant approaches *within* UQ, but as yet no single, general, over-arching theory *of* UQ.

1.2 Mathematical Prerequisites

Like any course or text, this book has some prerequisites. The perspective on UQ that runs through this book is strongly (but not exclusively) grounded in probability theory and Hilbert spaces, so the main prerequisite is familiarity with the language of linear functional analysis and measure/probability theory. As a crude diagnostic test, read the following sentence:

> Given any σ-finite measure space $(\mathcal{X}, \mathscr{F}, \mu)$, the set of all $\mathscr{F}$-measurable functions $f \colon \mathcal{X} \to \mathbb{C}$ for which $\int_{\mathcal{X}} |f|^2 \, \mathrm{d}\mu$ is finite, modulo equality μ-almost everywhere, is a Hilbert space with respect to the inner product $\langle f, g \rangle := \int_{\mathcal{X}} \bar{f} g \, \mathrm{d}\mu$.

None of the symbols, concepts or terms used or implicit in that sentence should give prospective students or readers any serious problems. Chapters 2 and 3 give a recap, without proof, of the necessary concepts and results, and most of the material therein should be familiar territory. In addition, Chapters 4 and 5 provide additional mathematical background on optimization and information theory respectively. It is assumed that readers have greater prior familiarity with the material in Chapters 2 and 3 than the material in Chapters 4 and 5; this is reflected in the way that results are presented mostly without proof in Chapters 2 and 3, but with proof in Chapters 4 and 5.

If, in addition, students or readers have some familiarity with topics such as numerical analysis, ordinary and partial differential equations, and stochastic analysis, then certain techniques, examples and remarks will make more sense. None of these are essential prerequisites, but, some ability and willingness to implement UQ methods — even in simple settings — in, e.g., C/C++, Mathematica, Matlab, or Python is highly desirable. (Some of the concepts

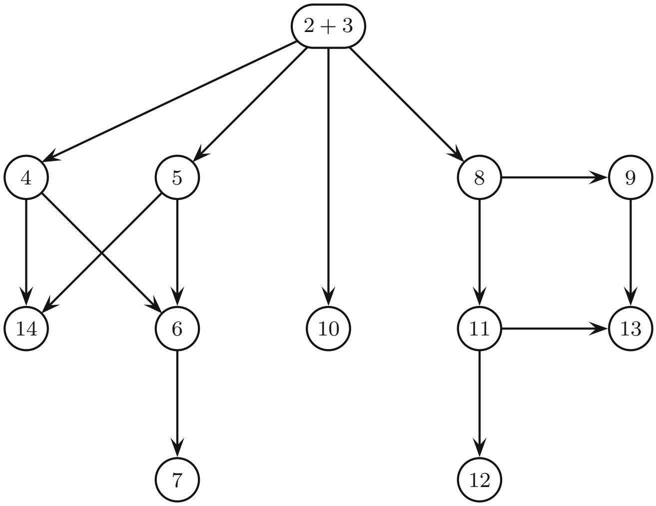

Fig. 1.1: Outline of the book (Leitfaden). An arrow from m to n indicates that Chapter n substantially depends upon material in Chapter m.

covered in the book will be given example numerical implementations in Python.) Although the aim of this book is to give an overview of the mathematical elements of UQ, this is a topic best learned in the doing, not through pure theory. However, in the interests of accessibility and pedagogy, none of the examples or exercises in this book will involve serious programming legerdemain.

1.3 Outline of the Book

The first part of this book lays out basic and general mathematical tools for the later discussion of UQ. Chapter 2 covers measure and probability theory, which are essential tools given the probabilistic description of many UQ problems. Chapter 3 covers some elements of linear functional analysis on Banach and Hilbert spaces, and constructions such as tensor products, all of which are natural spaces for the representation of random quantities and fields. Many UQ problems involve a notion of 'best fit', and so Chapter 4 provides a brief introduction to optimization theory in general, with particular attention to linear programming and least squares. Finally, although much of the UQ theory in this book is probabilistic, and is furthermore an L^2 theory, Chapter 5 covers more general notions of information and uncertainty.

The second part of the book is concerned with mathematical tools that are much closer to the practice of UQ. We begin in Chapter 6 with a mathematical treatment of inverse problems, and specifically their Bayesian interpretation; we take advantage of the tools developed in Chapters 2 and 3 to discuss Bayesian inverse problems on function spaces, which are especially important in PDE applications. In Chapter 7, this leads to a specific class of inverse problems, filtering and data assimilation problems, in which data and unknowns are decomposed in a sequential manner. Chapter 8 introduces orthogonal polynomial theory, a classical area of mathematics that has a double application in UQ: orthogonal polynomials are useful basis functions for the representation of random processes, and form the basis of powerful numerical integration (quadrature) algorithms. Chapter 9 discusses these quadrature methods in more detail, along with other methods such as Monte Carlo. Chapter 10 covers one aspect of forward uncertainty propagation, namely sensitivity analysis and model reduction, i.e. finding out which input parameters are influential in determining the values of some output process. Chapter 11 introduces spectral decompositions of random variables and other random quantities, including but not limited to polynomial chaos methods. Chapter 12 covers the intrusive (or Galerkin) approach to the determination of coefficients in spectral expansions; Chapter 13 covers the alternative non-intrusive (sample-based) paradigm. Finally, Chapter 14 discusses approaches to probability-based UQ that apply when even the probability distributions of interest are uncertain in some way.

The conceptual relationships among the chapters are summarized in Figure 1.1.

1.4 The Road Not Taken

There are many topics relevant to UQ that are either not covered or discussed only briefly here, including: detailed treatment of data assimilation beyond the confines of the Kálmán filter and its variations; accuracy, stability and computational cost of numerical methods; details of numerical implementation of optimization methods; stochastic homogenization and other multiscale methods; optimal control and robust optimization; machine learning; issues related to ‘big data’; and the visualization of uncertainty.