

Chapter 13

Non-Intrusive Methods

[W]hen people thought the Earth was flat, they were wrong. When people thought the Earth was spherical, they were wrong. But if *you* think that thinking the Earth is spherical is *just as wrong* as thinking the Earth is flat, then your view is wronger than both of them put together.

The Relativity of Wrong
ISAAC ASIMOV

Chapter 12 considers a spectral approach to UQ, namely Galerkin expansion, that is mathematically very attractive in that it is a natural extension of the Galerkin methods that are commonly used for deterministic PDEs and (up to a constant) minimizes the stochastic residual, but has the severe disadvantage that the stochastic modes of the solution are coupled together by a large system such as (12.15). Hence, the Galerkin formalism is not suitable for situations in which deterministic solutions are slow and expensive to obtain, and the deterministic solution method cannot be modified. Many so-called *legacy codes* are not amenable to such *intrusive* methods of UQ.

In contrast, this chapter considers *non-intrusive* spectral methods for UQ. These are characterized by the feature that the solution $U(\theta)$ of the deterministic problem is a ‘black box’ that does not need to be modified for use in the spectral method, beyond being able to be evaluated at any desired point θ of the probability space $(\Theta, \mathcal{F}, \mu)$. Indeed, sometimes, it is necessary to go one step further than this and consider the case of *legacy data*, i.e. an archive or data set of past input-output pairs $\{(\theta_n, U(\theta_n)) \mid n = 1, \dots, N\}$, sampled according to a possibly unknown or sub-optimal strategy, that is provided ‘as is’ and that cannot be modified or extended at all: the reasons for such restrictions may range from financial or practical difficulties to legal and ethical concerns.

There is a substantial overlap between non-intrusive methods for UQ and deterministic methods for interpolation and approximation as discussed in Chapter 8. However, this chapter additionally considers the method of *Gaussian process regression* (also known as *kriging*), which produces a probabilistic prediction of $U(\theta)$ away from the data set, including a variance-based measure of uncertainty in that prediction.

13.1 Non-Intrusive Spectral Methods

One class of non-intrusive UQ methods is the family of *non-intrusive spectral methods*, namely the determination of approximate spectral coefficients, e.g. polynomial chaos coefficients, of an uncertain quantity U . The distinguishing feature here, in contrast to the approximate spectral coefficients calculated in Chapter 12, is that realizations of U are used directly. A good mental model is that the realizations of U will be used as evaluations in a quadrature rule, to determine an approximate orthogonal projection onto a finite-dimensional subspace of the stochastic solution space. For this reason, these methods are sometimes called *non-intrusive spectral projection* (NISP).

Consider a square-integrable stochastic process $U: \Theta \rightarrow \mathcal{U}$ taking values in a separable Hilbert space¹ \mathcal{U} , with a spectral expansion

$$U = \sum_{k \in \mathbb{N}_0} u_k \Psi_k$$

of $U \in L^2(\Theta, \mu; \mathcal{U}) \cong \mathcal{U} \otimes L^2(\Theta, \mu; \mathbb{R})$ in terms of coefficients (stochastic modes) $u_k \in \mathcal{U}$ and an orthogonal basis $\{\Psi_k \mid k \in \mathbb{N}_0\}$ of $L^2(\Theta, \mu; \mathbb{R})$. As usual, the stochastic modes are given by

$$u_k = \frac{\langle U \Psi_k \rangle}{\langle \Psi_k^2 \rangle} = \frac{1}{\gamma_k} \int_{\Theta} U(\theta) \Psi_k(\theta) \, d\mu(\theta). \quad (13.1)$$

If the normalization constants $\gamma_k := \langle \Psi_k^2 \rangle \equiv \|\Psi_k\|_{L^2(\mu)}^2$ are known ahead of time, then it remains only to approximate the integral with respect to μ of the product of U with each basis function Ψ_k ; in some cases, the normalization constants must also be approximated. In any case, the aim is to use realizations of U to determine approximate stochastic modes $\tilde{u}_k \in \mathcal{U}$, with $\tilde{u}_k \approx u_k$, and hence an approximation

$$\tilde{U} := \sum_{k \in \mathbb{N}_0} \tilde{u}_k \Psi_k \approx U.$$

Such a stochastic process \tilde{U} is sometimes called a *surrogate* or *emulator* for the original process U .

¹ As usual, readers will lose little by assuming that $\mathcal{U} = \mathbb{R}$ on a first reading. Later, \mathcal{U} should be thought of as a non-trivial space of time- and space-dependent fields, so that $U(t, x; \theta) = \sum_{k \in \mathbb{N}_0} (t, x) \Psi_k(\theta)$.

Deterministic Quadrature. If the dimension of Θ is low and $U(\theta)$ is relatively smooth as a function of θ , then an appealing approach to the estimation of $\langle U\Psi_k \rangle$ is deterministic quadrature. For optimal polynomial accuracy, Gaussian quadrature (i.e. nodes at the roots of μ -orthogonal polynomials) may be used. In practice, nested quadrature rules such as Clenshaw–Curtis may be preferable since one does not wish to have to discard past solutions of U upon passing to a more accurate quadrature rule. For multi-dimensional domains of integration Θ , sparse quadrature rules may be used to partially alleviate the curse of dimension.

Note that, if the basis elements Ψ_k are polynomials, then the normalization constant $\gamma_k := \langle \Psi_k^2 \rangle$ can be evaluated numerically but with zero quadrature error by Gaussian quadrature with at least $(k + 1)/2$ nodes.

Monte Carlo and Quasi-Monte Carlo Integration. If the dimension of Θ is high, or $U(\theta)$ is a non-smooth function of θ , then it is tempting to resort to Monte Carlo approximation of $\langle U\Psi_k \rangle$. This approach is also appealing because the calculation of the stochastic nodes u_k can be written as a straightforward (but often large) matrix-matrix multiplication. The problem with Monte Carlo methods, as ever, is the slow convergence rate of $\sim (\text{number of samples})^{-1/2}$; quasi-Monte Carlo quadrature may be used to improve the convergence rate for smoother integrands.

Connection with Linear Least Squares. There is a close connection between least-squares minimization and the determination of approximate spectral coefficients via quadrature (be it deterministic or stochastic). Let basis functions Ψ_0, \dots, Ψ_K and nodes $\theta_1, \dots, \theta_N$ be given, and let

$$V := \begin{bmatrix} \Psi_0(\theta_1) & \cdots & \Psi_K(\theta_1) \\ \vdots & \ddots & \vdots \\ \Psi_0(\theta_N) & \cdots & \Psi_K(\theta_N) \end{bmatrix} \in \mathbb{R}^{N \times (K+1)} \tag{13.2}$$

be the associated Vandermonde-like matrix. Also, let $Q(f) := \sum_{n=1}^N w_n f(\theta_n)$ be an N -point quadrature rule using the nodes $\theta_1, \dots, \theta_N$, and let $W := \text{diag}(w_1, \dots, w_N) \in \mathbb{R}^{N \times N}$. For example, if the θ_n are i.i.d. draws from the measure μ on Θ , then

$$w_1 = \cdots = w_N = \frac{1}{N}$$

corresponds to the ‘vanilla’ Monte Carlo quadrature rule Q .

Theorem 13.1. *Given observed data $y_n := U(\theta_n)$ for $n = 1, \dots, N$, and $\mathbf{y} = [y_1, \dots, y_N]$, the following statements about approximate spectral coefficients $\tilde{\mathbf{u}} = (\tilde{u}_0, \dots, \tilde{u}_K)$ for $\tilde{U} := \sum_{k=0}^K \tilde{u}_k \Psi_k$ are equivalent:*

(a) \tilde{U} minimizes the weighted sum of squared residuals

$$R^2 := \sum_{n=1}^N w_n |\tilde{U}(\theta_n) - y_n|^2;$$

(b) $\tilde{\mathbf{u}}$ satisfies

$$V^T W V \tilde{\mathbf{u}} = V^T W \mathbf{y}^T; \quad (13.3)$$

(c) $\tilde{U} = U$ in the weak sense, tested against Ψ_0, \dots, Ψ_K using the quadrature rule Q , i.e., for $k = 0, \dots, K$,

$$Q(\Psi_k \tilde{U}) = Q(\Psi_k U).$$

Proof. Since

$$V \tilde{\mathbf{u}} = \begin{bmatrix} \tilde{U}(\theta_1) \\ \vdots \\ \tilde{U}(\theta_N) \end{bmatrix},$$

the weighted sum of squared residuals $\sum_{n=1}^N w_n |\tilde{U}(\theta_n) - y_n|^2$ for approximate model \tilde{U} equals $\|V \tilde{\mathbf{u}} - \mathbf{y}^T\|_W^2$. By Theorem 4.28, this function of $\tilde{\mathbf{u}}$ is minimized if and only if $\tilde{\mathbf{u}}$ satisfies the normal equations (13.3), which shows that (a) \iff (b). Explicit calculation of the left- and right-hand sides of (13.3) yields

$$\sum_{n=1}^N w_n \begin{bmatrix} \Psi_0(\theta_n) \tilde{U}(\theta_n) \\ \vdots \\ \Psi_K(\theta_n) \tilde{U}(\theta_n) \end{bmatrix} = \sum_{n=1}^N w_n \begin{bmatrix} \Psi_0(\theta_n) y_n \\ \vdots \\ \Psi_K(\theta_n) y_n \end{bmatrix},$$

which shows that (b) \iff (c) □

Note that the matrix $V^T W V$ on the left-hand side of (13.3) is

$$V^T W V = \begin{bmatrix} Q(\Psi_0 \Psi_0) & \cdots & Q(\Psi_0 \Psi_K) \\ \vdots & \ddots & \vdots \\ Q(\Psi_K \Psi_0) & \cdots & Q(\Psi_K \Psi_K) \end{bmatrix} \in \mathbb{R}^{(K+1) \times (K+1)},$$

i.e. is the Gram matrix of the basis functions Ψ_0, \dots, Ψ_K with respect to the quadrature rule Q 's associated inner product. Therefore, if the quadrature rule Q is one associated to μ (e.g. a Gaussian quadrature formula for μ , or a Monte Carlo quadrature with i.i.d. $\theta_n \sim \mu$), then $V^T W V$ will be an approximation to the Gram matrix of the basis functions Ψ_0, \dots, Ψ_K in the $L^2(\mu)$ inner product. In particular, dependent upon the accuracy of the quadrature rule Q , we will have $V^T W V \approx \text{diag}(\gamma_0, \dots, \gamma_K)$, and then

$$\tilde{u}_k \approx \frac{Q(\Psi_k U)}{\gamma_k},$$

i.e. \tilde{u}_k approximately satisfies the orthogonal projection condition (13.1) satisfied by u_k .

In practice, when given $\{\theta_n\}_{n=1}^N$ that are not necessarily associated with some quadrature rule for μ , along with corresponding output values $\{y_n := U(\theta_n)\}_{n=1}^N$, it is common to construct approximate stochastic modes and hence an approximate spectral expansion \tilde{U} by choosing $\tilde{u}_0, \dots, \tilde{u}_k$ to minimize the some weighted sum of squared residuals, i.e. according to (13.3).

Conversely, one can engage in the *design of experiments* — i.e. the selection of $\{\theta_n\}_{n=1}^N$ — to optimize some derived quantity of the matrix V ; common choices include

- A-optimality, in which the trace of $(V^T V)^{-1}$ is minimized;
- D-optimality, in which the determinant of $V^T V$ is maximized;
- E-optimality, in which the least singular value of $V^T V$ is maximized; and
- G-optimality, in which the largest diagonal term in the orthogonal projection $V(V^T V)^{-1} V^T \in \mathbb{R}^{N \times N}$ is minimized.

Remark 13.2. The Vandermonde-like matrix V from (13.2) is often ill-conditioned, i.e. has singular values of hugely different magnitudes. Often, this is a property of the normalization constants of the basis functions $\{\Psi_k\}_{k=0}^K$. As can be seen from Table 8.2, many of the standard families of orthogonal polynomials have normalization constants $\|\psi_k\|_{L^2}$ that tend to 0 or to ∞ as $k \rightarrow \infty$. A tensor product system $\{\psi_\alpha\}_{\alpha \in \mathbb{N}_0^g}$ of multivariate orthogonal polynomials, as in Theorem 8.25, might well have

$$\liminf_{|\alpha| \rightarrow \infty} \|\psi_\alpha\|_{L^2} = 0 \quad \text{and} \quad \limsup_{|\alpha| \rightarrow \infty} \|\psi_\alpha\|_{L^2} = \infty;$$

this phenomenon arises in, for example, the products of the Legendre and Hermite, or the Legendre and Charlier, bases. Working with orthonormal bases, or using preconditioners, alleviates the difficulties caused by such ill-conditioned matrices V .

Remark 13.3. In practice, the following sources of error arise when computing non-intrusive approximate spectral expansions in the fashion outlined in this section:

- (a) *discretization error* comes about through the approximation of \mathcal{U} by a finite-dimensional subspace \mathcal{U}_M , i.e. the approximation the stochastic modes u_k by a finite sum $u_k \approx \sum_{m=1}^M u_{km} \phi_m$, where $\{\phi_m \mid m \in \mathbb{N}\}$ is some basis for \mathcal{U} ;
- (b) *truncation error* comes about through the truncation of the spectral expansion for U after finitely many terms, i.e. $U \approx \sum_{k=0}^K u_k \Psi_k$;
- (c) *quadrature error* comes about through the approximate nature of the numerical integration scheme used to find the stochastic modes; classical statistical concerns about the unbiasedness of estimators for expected values fall into this category. The choice of integration nodes contributes greatly to this source of error.

A complete quantification of the uncertainty associated with predictions of U made using a truncated non-intrusively constructed spectral stochastic model $\tilde{U} := \sum_{k=0}^K \tilde{u}_k \Psi_k$ requires an understanding of all three of these sources of

error, and there is necessarily some tradeoff among them when trying to give ‘optimal’ predictions for a given level of computational and experimental cost.

Remark 13.4. It often happens in practice that the process U is not initially defined on the same probability space as the gPC basis functions, in which case some appropriate changes of variables must be used. In particular, this situation can arise if we are given an archive of legacy data values of U without the corresponding inputs. See Exercise 13.5 for a discussion of these issues in the example setting of Gaussian mixtures.

Example 13.5. Consider again the simple harmonic oscillator


$$\ddot{U}(t) = -\Omega^2 U(t)$$

with the initial conditions $U(0) = 1$, $\dot{U}(0) = 0$. Suppose that $\Omega \sim \text{Unif}([0.8, 1.2])$, so that $\Omega = 1.0 + 0.2\varepsilon$, where $\varepsilon \sim \text{Unif}([-1, 1])$ is the stochastic germ, with its associated Legendre basis polynomials. Figure 13.1 shows the evolution of the approximate stochastic modes for U , calculated using $N = 1000$ i.i.d. samples of ε and the least squares approach of Theorem 13.1. As in previous examples of this type, the forward solution of the ODE is performed using a symplectic integrator with time step 0.01.


Note that many standard computational algebra routines, such as Python’s `numpy.linalg.lstsq`, will solve all the least squares problems of finding $\{\tilde{u}_k(t_i)\}_{k=0}^K$ for all time points t_i in a vectorized manner. That is, it is not necessary to call `numpy.linalg.lstsq` with matrix V and data $\{U(t_0, \omega_n)\}_{n=1}^N$ to obtain $\{\tilde{u}_k(t_0)\}_{k=0}^K$, and then do the same for t_1 , etc. Instead, all the data $\{U(t_i, \omega_n) \mid n = 1, \dots, N; i \in \mathbb{N}_0\}$ can be supplied at once as a matrix, yielding a matrix $\{\tilde{u}_k(t_i) \mid k = 0, \dots, K; i \in \mathbb{N}_0\}$.

13.2 Stochastic Collocation

Collocation methods for ordinary and partial differential equations are another form of interpolation. The idea is to find a low-dimensional object — usually a polynomial — that approximates the true solution to the differential equation by means of *exactly* satisfying the differential equation at a selected set of points, called *collocation points* or *collocation nodes*. An important feature of the collocation approach is that an approximation is constructed not on a pre-defined stochastic subspace, but instead uses interpolation, and hence both the approximation and the approximation space are implicitly prescribed by the collocation nodes. As the number of collocation nodes increases, the space in which the solution is sought becomes correspondingly larger.



The solid curve shows the mean $u_0(t)$ of the solution, the dashed curves show the higher-degree Legendre coefficients $u_k(t)$, and the grey envelope shows the mean \pm one standard deviation. Over the time interval shown, $|u_k(t)| \leq 10^{-2}$ for $k \geq 5$.



Contour plot of the truncated NISP model $\tilde{U}(t) = \sum_{k=0}^{10} \tilde{u}_k(t)Le_k$.

Fig. 13.1: The degree-10 Legendre PC NISP solution to the simple harmonic oscillator equation of Example 13.5 with $\Omega \sim \text{Unif}([0.8, 1.2])$.

Example 13.6 (Collocation for an ODE). Consider, for example, the initial value problem

$$\begin{aligned} \dot{u}(t) &= f(t, u(t)), & \text{for } t \in [a, b] \\ u(a) &= u_a, \end{aligned}$$

to be solved on an interval of time $[a, b]$. Choose n points

$$a \leq t_1 < t_2 < \cdots < t_n \leq b,$$

called *collocation nodes*. Now find a polynomial $p(t) \in \mathbb{R}_{\leq n}[t]$ so that the ODE

$$\dot{p}(t_k) = f(t_k, p(t_k))$$

is satisfied for $k = 1, \dots, n$, as is the initial condition $p(a) = u_a$. For example, if $n = 2$, $t_1 = a$ and $t_2 = b$, then the coefficients $c_2, c_1, c_0 \in \mathbb{R}$ of the polynomial approximation

$$p(t) = \sum_{k=0}^2 c_k (t-a)^k,$$

which has derivative $\dot{p}(t) = 2c_2(t-a) + c_1$, are required to satisfy

$$\begin{aligned} \dot{p}(a) &= c_1 = f(a, p(a)) \\ \dot{p}(b) &= 2c_2(b-a) + c_1 = f(b, p(b)) \\ p(a) &= c_0 = u_a \end{aligned}$$

i.e.

$$p(t) = \frac{f(b, p(b)) - f(a, u_a)}{2(b-a)}(t-a)^2 + f(a, u_a)(t-a) + u_a.$$

The above equation implicitly defines the final value $p(b)$ of the collocation solution. This method is also known as the *trapezoidal rule* for ODEs, since the same solution is obtained by rewriting the differential equation as

$$u(t) = u(a) + \int_a^t f(s, u(s)) ds$$

and approximating the integral on the right-hand side by the trapezoidal quadrature rule for integrals.

It should be made clear at the outset that there is nothing stochastic about ‘stochastic collocation’, just as there is nothing chaotic about ‘polynomial chaos’. The meaning of the term ‘stochastic’ in this case is that the collocation principle is being applied across the ‘stochastic space’ (i.e. the probability space) of a stochastic process, rather than the space/time/space-time

domain. That is, for a stochastic process U with known values $U(\theta_n)$ at known collocation points $\theta_1, \dots, \theta_N \in \Theta$, we seek an approximation \tilde{U} such that

$$\tilde{U}(\theta_n) = U(\theta_n) \quad \text{for } n = 1, \dots, N.$$

There is, however, some flexibility in how to approximate $U(\theta)$ for $\theta \neq \theta_1, \dots, \theta_N$.

Example 13.7. Consider, for example, the random PDE

$$\begin{aligned} \mathcal{L}_\theta[U(x, \theta)] &= 0 & \text{for } x \in \mathcal{X}, \theta \in \Theta, \\ \mathcal{B}_\theta[U(x, \theta)] &= 0 & \text{for } x \in \partial\mathcal{X}, \theta \in \Theta, \end{aligned}$$

where, for μ -a.e. θ in some probability space $(\Theta, \mathcal{F}, \mu)$, the differential operator \mathcal{L}_θ and boundary operator \mathcal{B}_θ are well defined and the PDE admits a unique solution $U(\cdot, \theta): \mathcal{X} \rightarrow \mathbb{R}$. The solution $U: \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ is then a stochastic process. We now let $\Theta_M := \{\theta_1, \dots, \theta_M\} \subseteq \Theta$ be a finite set of prescribed collocation nodes. The collocation problem is to find a *collocation solution* \tilde{U} , an approximation to the exact solution U , that satisfies

$$\begin{aligned} \mathcal{L}_{\theta_m}[\tilde{U}(x, \theta_m)] &= 0 & \text{for } x \in \mathcal{X}, \\ \mathcal{B}_{\theta_m}[\tilde{U}(x, \theta_m)] &= 0 & \text{for } x \in \partial\mathcal{X}, \end{aligned}$$

for $m = 1, \dots, M$.

Interpolation Approach. An obvious first approach is to use interpolating polynomials when they are available. This is easiest when the stochastic space Θ is one-dimensional, in which case the Lagrange basis polynomials of a given nodal set are an attractive choice of interpolation basis. As always, though, care must be taken to use nodal sets that will not lead to Runge oscillations; if there is very little a priori information about the process U , then constructing a ‘good’ nodal set may be a matter of trial and error. In general, the choice of collocation nodes is a significant contributor to the error and uncertainty in the resulting predictions.

Given the values $U(\theta_1), \dots, U(\theta_N)$ of U at nodes $\theta_1, \dots, \theta_N$ in a one-dimensional space Θ , the (Lagrange-form polynomial interpolation) collocation approximation \tilde{U} to U is given by

$$\tilde{U}(\theta) = \sum_{n=1}^N U(\theta_n) \ell_n(\theta) = \sum_{n=1}^N U(\theta_n) \prod_{\substack{1 \leq k \leq N \\ k \neq n}} \frac{\theta - \theta_k}{\theta_n - \theta_k}.$$

Example 13.8. Figure 13.2 shows the results of the interpolation-collocation approach for the simple harmonic oscillator equation considered earlier, again for $\omega \in [0.8, 1.2]$. Two nodal sets $\omega_1, \dots, \omega_N \in \mathbb{R}$ are considered: uniform

nodes, and Chebyshev nodes. In order to make the differences between the two solutions more easily visible, only $N = 4$ nodes are used.

The collocation solution $\tilde{U}(\cdot, \omega_n)$ at each of the collocation nodes ω_n is the solution of the deterministic problem

$$\begin{aligned}\frac{d^2}{dt^2}\tilde{U}(t, \omega_n) &= -\omega_n^2 U(t, \omega_n), \\ \tilde{U}(0, \omega_n) &= 1, \\ \frac{d}{dt}\tilde{U}(0, \omega_n) &= 0.\end{aligned}$$

Away from the collocation nodes, \tilde{U} is defined by polynomial interpolation: for each t , $\tilde{U}(t, \omega)$ is a polynomial in ω of degree at most N with prescribed values at the collocation nodes. Writing this interpolation in terms of Lagrange basis polynomials

$$\ell_n(\omega; \omega_1, \dots, \omega_N) := \prod_{\substack{1 \leq k \leq N \\ k \neq n}} \frac{\omega - \omega_k}{\omega_n - \omega_k}$$


yields

$$\tilde{U}(t, \omega) = \sum_{n=1}^N U(t, \omega_n) \ell_n(\omega).$$

As can be seen in Figure 13.2(c-d), both nodal sets have the undesirable property that the approximate solution $\tilde{U}(t, \omega)$ has with the undesirable property that $|\tilde{U}(t, \omega)| > |\tilde{U}(0, \omega)| = 1$ for some $t > 0$ and $\omega \in [0.8, 1.2]$. Therefore, for general ω , $\tilde{U}(t, \omega)$ is *not* a solution of the original ODE. However, as the discussion around Runge's phenomenon in Section 8.5 would lead us to expect, the regions in (t, ω) -space where such unphysical values are attained are smaller with the Chebyshev nodes than the uniformly distributed ones.


The extension of one-dimensional interpolation methods to the multi-dimensional case can be handled in a theoretically straightforward manner using tensor product grids, similar to the constructions used in quadrature. In tensor product constructions, both the grid of interpolation points and the interpolation polynomials are products of the associated one-dimensional objects. Thus, in a product space $\Theta = \Theta_1 \times \dots \times \Theta_d$, we take nodes

$$\begin{aligned}\theta_1^1, \dots, \theta_{N_1}^1 &\in \Theta_1 \\ &\vdots \\ \theta_1^d, \dots, \theta_{N_d}^d &\in \Theta_d\end{aligned}$$




$U(t, \omega)$ with uniform (Newton–Cotes) nodes for $\omega \in [0.8, 1.2]$.

$U(t, \omega)$ with Chebyshev nodes for $\omega \in [0.8, 1.2]$.



Interpolation of (a) with respect to ω .



Interpolation of (b) with respect to ω .

Fig. 13.2: Interpolation solutions for a simple harmonic oscillator with uncertain natural frequency ω , $U(0, \omega) = 1$, $\dot{U}(0, \omega) = 0$. Both cases use four interpolation nodes. Note that the Chebyshev nodes produce smaller regions in (t, ω) -space with unphysical values $|\tilde{U}(t, \omega)| > 1$.

and construct a product grid of nodes $\theta_{\mathbf{n}} := (\theta_{n_1}^1, \dots, \theta_{n_d}^d) \in \Theta$, where the multi-index $\mathbf{n} = (n_1, \dots, n_d)$ runs over $\{1, \dots, N_1\} \times \dots \times \{1, \dots, N_d\}$. The corresponding interpolation formula, in terms of Lagrange basis polynomials, is then

$$\tilde{U}(\theta) = \sum_{\mathbf{n}=(1,\dots,1)}^{(N_1,\dots,N_d)} U(\theta_{\mathbf{n}}) \prod_{i=1}^d \ell_{n_i}(\theta^i; \theta_1^i, \dots, \theta_{N_i}^i).$$

The problem with tensor product grids for interpolative collocation is the same as for tensor product quadrature: the curse of dimension, i.e. the large number of nodes needed to adequately resolve features of functions on high-dimensional spaces. The curse of dimension can be partially circumvented by using interpolation through sparse grids, e.g. those of Smolyak type.

Collocation for arbitrary unstructured sets of nodes — such as those that arise when inheriting an archive of ‘legacy’ data that cannot be modified or extended for whatever reason — is a notably tricky subject, essentially because it boils down to polynomial interpolation through an unstructured set of nodes. Even the existence of interpolating polynomials such as analogues of the Lagrange basis polynomials is not, in general, guaranteed.

Other Approximation Strategies. There are many other strategies for the construction of collocation solutions, especially in high dimension, besides polynomial bases. Common choices include splines and radial basis functions; see the bibliographic notes at the end of the chapter for references. Another popular method is Gaussian process regression, which is the topic of the next section.

13.3 Gaussian Process Regression

The interpolation approaches of the previous section were all deterministic in two senses: they assume that the values $U(\theta_{\mathbf{n}})$ are observed exactly, without error and with perfect reproducibility; they also assume that the correct form for an interpolated value $\tilde{U}(\theta)$ away from the nodal set is a deterministic function of the nodes and observed values. In many situations in the natural sciences and commerce, these assumptions are not appropriate. Instead, it is appropriate to incorporate an estimate of the observational uncertainties, and to produce probabilistic predictions; this is another area in which the Bayesian perspective is quite natural.

This section surveys one such method of stochastic interpolation, known as *Gaussian process regression* or *kriging*; as ever, the quite rigid properties of Gaussian measures hugely simplify the presentation. The essential idea is that we will model U as a Gaussian random field; the prior information on U consists of a mean field and a covariance operator, the latter often being given in practice by a correlation length; the observations of U at discrete points

are then used to condition the prior Gaussian using Schur complementation, and thereby produce a posterior Gaussian prediction for the value of U at any other point.

Noise-Free Observations. Suppose for simplicity that we observe the values $y_n := U(\theta_n)$ exactly, without any observational error. We wish to use the data $\{(\theta_n, y_n) \mid n = 1, \dots, N\}$ to make a prediction for the values of U at other points in the domain Θ . To save space, we will refer to $\theta^\circ = (\theta_1, \dots, \theta_N)$ as the *observed points* and $y^\circ = (y_1, \dots, y_N)$ as the *observed values*; together, (θ°, y°) constitute the *observed data* or *training set*. By way of contrast, we wish to predict the value(s) y^p of U at point(s) θ^p , referred to as the *prediction points* or *test points*. We will abuse notation and write $m(\theta^\circ)$ for $(m(\theta_0), \dots, m(\theta_N))$, and so on.

Under the prior assumption that U is a Gaussian random field with known mean $m: \Theta \rightarrow \mathbb{R}$ and known covariance function $C: \Theta \times \Theta \rightarrow \mathbb{R}$, the random vector (y°, y^p) is a draw from a multivariate Gaussian distribution with mean $(m(\theta^\circ), m(\theta^p))$ and covariance matrix

$$\begin{bmatrix} C(\theta^\circ, \theta^\circ) & C(\theta^\circ, \theta^p) \\ C(\theta^\circ, \theta^p) & C(\theta^p, \theta^p) \end{bmatrix}$$

(Note that in the case of N observed data points and one new value to be predicted, $C(\theta^\circ, \theta^\circ)$ is an $N \times N$ block, $C(\theta^p, \theta^p)$ is 1×1 , and $C(\theta^\circ, \theta^p)$ is a $1 \times N$ ‘row vector’.) By Theorem 2.54, the conditional distribution of $U(\theta^p)$ given the observations $U(\theta^\circ) = y^\circ$ is Gaussian, with its mean and variance given in terms of the Schur complement

$$S := C(\theta^p, \theta^p) - C(\theta^p, \theta^\circ)^\top C(\theta^\circ, \theta^\circ)^{-1} C(\theta^\circ, \theta^p)$$

by

$$U(\theta^p) \mid \theta^\circ, y^\circ \sim \mathcal{N}(m^p + C(\theta^p, \theta^\circ) C(\theta^\circ, \theta^\circ)^{-1} (y^\circ - m(\theta^\circ)), S).$$

This means that, in practice, a draw $\tilde{U}(\theta^p)$ from this conditioned Gaussian measure would be used as a proxy/prediction for the value $U(\theta^p)$. Note that S depends only upon the locations of the interpolation nodes θ° and θ^p . Thus, if variance is to be used as a measure of the precision of the estimate $\tilde{U}(\theta^p)$, then it will be independent of the observed data y° .

Noisy Observations. The above derivation is very easily adapted to the case of noisy observations, i.e. $y^\circ = U(\theta^\circ) + \eta$, where η is some random noise vector. As usual, the Gaussian case is the simplest, and if $\eta \sim \mathcal{N}(0, \Gamma)$, then the net effect is to replace each occurrence of “ $C(\theta^\circ, \theta^\circ)$ ” above by “ $\Gamma + C(\theta^\circ, \theta^\circ)$ ”. In terms of regularization, this is nothing other than quadratic regularization using the norm $\|\cdot\|_{\Gamma^{1/2}} = \|\Gamma^{-1/2} \cdot\|$ on \mathbb{R}^N .

One advantage of regularization, as ever, is that it sacrifices the interpolation property (exactly fitting the data) for better-conditioned solutions and even the ability to assimilate ‘contradictory’ observed data, i.e. $\theta_n = \theta_m$ but $y_n \neq y_m$. See Figure 13.3 for simple examples.

Example 13.9. Consider $\Theta = [0, 1]$, and suppose that the prior description of U is as a zero-mean Gaussian process with Gaussian covariance kernel

$$C(\theta, \theta') := \exp\left(-\frac{|\theta - \theta'|^2}{2\ell^2}\right);$$

$\ell > 0$ is the correlation length of the process, and the numerical results illustrated in Figure 13.3 use $\ell = \frac{1}{4}$.

- (a) Suppose that values $y^\circ = 0.1, 0.8$ and 0.5 are observed for U at $\theta^\circ = 0.1, 0.5, 0.9$ respectively. In this case, the matrix $C(\theta^\circ, \theta^\circ)$ and its inverse are approximately

$$C(\theta^\circ, \theta^\circ) = \begin{bmatrix} 1.000 & 0.278 & 0.006 \\ 0.278 & 1.000 & 0.278 \\ 0.006 & 0.278 & 1.000 \end{bmatrix}$$

$$C(\theta^\circ, \theta^\circ)^{-1} = \begin{bmatrix} 1.090 & -0.327 & 0.084 \\ -0.327 & 1.182 & -0.327 \\ 0.084 & -0.327 & 1.090 \end{bmatrix}.$$

Figure 13.3(a) shows the posterior mean field and posterior variance: note that the posterior mean interpolates the given data.

- (b) Now suppose that values $y^\circ = 0.1, 0.8, 0.9$, and 0.5 are observed for U at $\theta^\circ = 0.1, 0.5, 0.5, 0.9$ respectively. In this case, because there are two contradictory values for U at $\theta = 0.5$, we do not expect the posterior mean to be a function that interpolates the data. Indeed, the matrix $C(\theta^\circ, \theta^\circ)$ has a repeated row and column:

$$C(\theta^\circ, \theta^\circ) = \begin{bmatrix} 1.000 & 0.278 & 0.278 & 0.006 \\ 0.278 & 1.000 & 1.000 & 0.278 \\ 0.278 & 1.000 & 1.000 & 0.278 \\ 0.006 & 0.278 & 0.278 & 1.000 \end{bmatrix},$$

and hence $C(\theta^\circ, \theta^\circ)$ is not invertible. However, assuming that $y^\circ = U(\theta^\circ) + \mathcal{N}(0, \eta^2)$, with $\eta > 0$, restores well-posedness to the problem. Figure 13.3(b) shows the posterior mean and covariance field with the regularization $\eta = 0.1$.

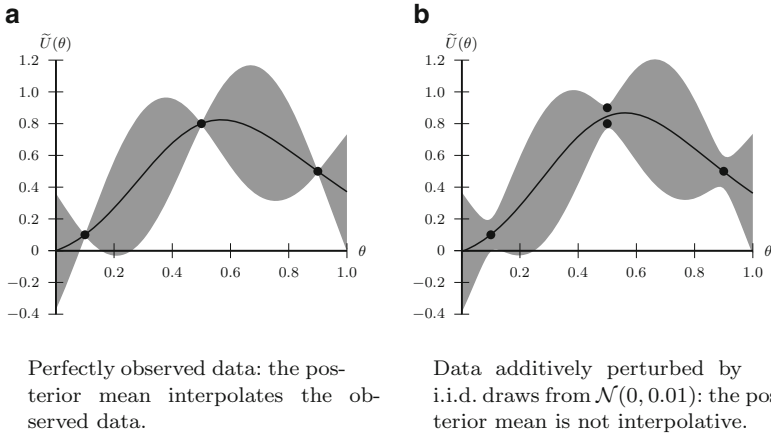


Fig. 13.3: A simple example of Gaussian process regression/kriging in one dimension. The dots show the observed data points, the black curve the posterior mean of the Gaussian process \tilde{U} , and the shaded region the posterior mean \pm one posterior standard deviation.

Variations. There are many ‘flavours’ of the kriging method, essentially determined by the choice of the prior, and in particular the choice of the prior mean. For example, *simple kriging* assumes a known spatially constant mean field, i.e. $\mathbb{E}[U(\theta)] = m$ for all θ .

A mild generalization is *ordinary kriging*, in which it is again assumed that $\mathbb{E}[U(\theta)] = m$ for all θ , but m is not assumed to be known. This underdetermined situation can be rendered tractable by including additional assumptions on the form of $\tilde{U}(\theta^p)$ as a function of the data (θ^o, y^o) : one simple assumption of this type is a linear model of the form $\tilde{U}(\theta^p) = \sum_{n=1}^N w_n y_n$ for some weights $w = (w_1, \dots, w_N) \in \mathbb{R}^N$ — note well that this is *not* the same as linearly interpolating the observed data.

In this situation, as in the Gauss–Markov theorem (Theorem 6.2), the natural criteria of zero mean error (unbiasedness) and minimal squared error are used to determine the estimate of $U(\theta^p)$: writing $\tilde{U}(\theta^p) = \sum_{n=1}^N w_n y_n$, the unbiasedness requirement that $\mathbb{E}[\tilde{U}(\theta^p) - U(\theta^p)] = 0$ implies that the weights w_n sum to 1, and minimizing $\mathbb{E}[(\tilde{U}(\theta^p) - U(\theta^p))^2]$ becomes the constrained optimization problem

$$\begin{aligned}
 &\text{minimize: } C(\theta^p, \theta^p) - 2w^\top C(\theta^p, \theta^o) + w^\top C(\theta^o, \theta^o)w \\
 &\text{among: } w \in \mathbb{R}^N \\
 &\text{subject to: } \sum_{n=1}^N w_n = 1.
 \end{aligned}$$



By the method of Lagrange multipliers, the weight vector w and the Lagrange multiplier $\lambda \in \mathbb{R}$ are given jointly as the solutions of

$$\begin{bmatrix} C(\theta^\circ, \theta^\circ) & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} w \\ \lambda \end{bmatrix} = \begin{bmatrix} C(\theta^p, \theta^\circ) \\ 1 \end{bmatrix}. \quad (13.4)$$

Even when $C(\theta^\circ, \theta^\circ)$ is positive-definite, the matrix on the left-hand side is not invertible: however, the column vector on the right-hand side does lie in the range, and so it is possible² to solve for (w, λ) .

13.4 Bibliography

Non-intrusive methods for UQ, including non-intrusive spectral projection and stochastic collocation, are covered by Le Maître and Knio (2010, Chapter 3) and Xiu (2010, Chapter 7). A classic paper on interpolation using sparse grids is that of Barthelmann et al. (2000), and applications to UQ for PDEs with random input data have been explored by, e.g., Nobile et al. (2008a,b). Narayan and Xiu (2012) give a method for stochastic collocation on arbitrary sets of nodes using the framework of least orthogonal interpolation, following an earlier Gaussian construction of de Boor and Ron (1990). Yan et al. (2012) consider stochastic collocation algorithms with sparsity-promoting ℓ^1 regularizations. Buhmann (2003) provides a general introduction to the theory and practical usage of radial basis functions. A comprehensive introduction to splines is the book of de Boor (2001); for a more statistical interpretation, see, e.g., Smith (1979).

Kriging was introduced by Krige (1951) and popularized in geostatistics by Matheron (1963). See, e.g., Conti et al. (2009) for applications to the interpolation of results from slow or expensive computational methods. Rasmussen and Williams (2006) cover the theory and application of Gaussian processes to machine learning; their text also gives a good overview of the relationships between Gaussian processes and other modelling perspectives, including regularization, reproducing kernel Hilbert spaces, and support vector machines.

13.5 Exercises

Exercise 13.1. Choose distinct nodes $\theta_1, \dots, \theta_N \in \Theta = [0, 1]$ and corresponding values $y_1, \dots, y_N \in \mathbb{R}$. Interpolate these data points in all the ways discussed so far in the text. In particular, interpolate the data using

² Indeed, many standard numerical linear algebra packages will readily solve the system (13.4) without throwing any error whatsoever.

apiecewise linear interpolation, using a polynomial of degree $N - 1$, and using Gaussian processes with various choices of covariance kernel. Plot the interpolants on the same axes to get an idea of their qualitative features.

Exercise 13.2. Extend the analysis of the simple harmonic oscillator from Examples 13.5 and 13.8 to incorporate uncertainty in the initial condition, and calculate sensitivity indices with respect to the various uncertainties. Perform the same analyses with an alternative uncertainty model, e.g. the log-normal model of Example 12.6.

Exercise 13.3. Perform the analogue of Exercise 13.2 for the Van der Pol oscillator

$$\ddot{u}(t) - \mu(1 - u(t)^2)\dot{u}(t) + \omega^2 u(t) = 0.$$

Compare your results with those of the active subspace method (Example 10.20 and Figure 10.1).

Exercise 13.4. Extend the analysis of Exercises 13.2 and 13.3 by treating the time step $h > 0$ of the numerical ODE solver as an additional source of uncertainty and error. Suppose that the numerical integration scheme for the ODE has a global truncation error at most Ch^r for some $C, r > 0$, and so model the exact solution to the ODE as the computed solution plus a draw from $\text{Unif}(-Ch^r, Ch^r)$. Using this randomly perturbed observational data, calculate approximate spectral coefficients for the process using the NISP scheme. (For more sophisticated randomized numerical schemes for ODEs and PDEs, see, e.g., Schober et al. (2014) and the works listed as part of the Probabilistic Numerics project <http://www.probablistic-numeric.org>.)

Exercise 13.5. It often happens that the process U is not initially defined on the same probability space as the gPC basis functions: in particular, this situation can arise if we are given an archive of legacy data values of U without corresponding inputs. In this situation, it is necessary to transform both sets of random variables to a common probability space. This exercise concerns an example implementation of this procedure in the case that U is a real-valued *Gaussian mixture*: for some weights $w_1, \dots, w_J \geq 0$ summing to 1, means $m_1, \dots, m_J \in \mathbb{R}$, and variances $\sigma_1^2, \dots, \sigma_J^2 > 0$, the Lebesgue probability density $f_U: \mathbb{R} \rightarrow [0, \infty)$ of U is given as the following convex combination of Gaussian densities:

$$f_U(x) := \sum_{j=1}^J \frac{w_j}{\sqrt{2\pi\sigma_j^2}} \exp\left(-\frac{(x - m_j)^2}{2\sigma_j^2}\right). \quad (13.5)$$

Suppose that we wish to perform a Hermite expansion of U , i.e. to write $U = \sum_{k \in \mathbb{N}_0} u_k \text{He}_k(Z)$, where $Z \sim \gamma = \mathcal{N}(0, 1)$. The immediate problem is that U is defined as a function of θ in some abstract probability space $(\Theta, \mathcal{F}, \mu)$, not as a function of z in the concrete probability space $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \gamma)$.

(a) Let $\Theta = \{1, \dots, J\} \times \mathbb{R}$, and define a probability measure μ on Θ by

$$\mu := \sum_{j=1}^J w_j \delta_j \otimes \mathcal{N}(m_j, \sigma_j^2).$$

(In terms of sampling, this means that draws (j, y) from μ are performed by first choosing $j \in \{1, \dots, J\}$ at random according to the weights w_1, \dots, w_J , and then drawing a Gaussian sample $y \sim \mathcal{N}(m_j, \sigma_j^2)$.) Let $P: \Theta \rightarrow \mathbb{R}$ denote projection onto the second component, i.e. $P(j, y) := y$. Show that the push-forward measure $P_*\mu$ on \mathbb{R} is the Gaussian mixture (13.5).

- (b) Let $F_U: \mathbb{R} \rightarrow [0, 1]$ denote the cumulative distribution function (CDF) of U , i.e.

$$F_U(x) := \mathbb{P}_\mu[U \leq x] = \int_{-\infty}^x f_U(s) ds.$$

Show that F_U is invertible, and that if $V \sim \text{Unif}([0, 1])$, then $F_U^{-1}(V)$ has the same distribution as U .

- (c) Let Φ denote the CDF of the standard normal distribution γ . Show, by change of integration variables, that

$$\langle U, \text{He}_k \rangle_{L^2(\gamma)} = \int_0^1 F_U^{-1}(v) \text{He}_k(\Phi^{-1}(v)) dv. \quad (13.6)$$

- (d) Use your favourite quadrature rule for uniform measure on $[0, 1]$ to approximately evaluate (13.6), and hence calculate approximate Hermite PC coefficients \tilde{u}_k for U .
- (e) Choose some m_j and σ_j^2 , and generate N i.i.d. sample realizations y_1, \dots, y_N of U using the observation of part (a). Approximate F_U by the empirical CDF of the data, i.e.

$$F_U(x) \approx \widehat{F}_{\mathbf{y}}(x) := \frac{|\{1 \leq n \leq N \mid y_n \leq x\}|}{N}.$$

Use this approximation and your favourite quadrature rule for uniform measure on $[0, 1]$ to approximately evaluate (13.6), and hence calculate approximate Hermite PC coefficients \tilde{u}_k for U . (This procedure, using the empirical CDF, is essentially the one that we must use if we are given only the data \mathbf{y} and no functional relationship of the form $y_n = U(\theta_n)$.)

- (f) Compare the results of parts (d) and (e).

Exercise 13.6. Choose nodes in the square $[0, 1]^2$ and corresponding data values, and interpolate them using Gaussian process regression with a radial covariance function such as $C(x, x') = \exp(-\|x - x'\|^2/r^2)$, with $r > 0$ being a correlation length parameter. Produce accompanying plots of the posterior variance field.