1
Introduction

This introduction reviews, summarizes, and illustrates fundamental connections between Bayesian inference, numerical quadrature, Gausssian process regression, polyharmonic splines, information-based complexity, optimal recovery, and game theory that form the basis of the book. This is followed by describing a sample of the results derived from these interplays, including those in numerical homogenization, operator-adapted wavelets, fast solvers, and Gaussian process regression. It finishes with an outline of the structure of the book.

1.1 Statistical Numerical Approximation

Although numerical approximation and statistical inference are traditionally seen as entirely separate subjects, they are intimately connected through the common purpose of making estimations with partial information. In [95], Diaconis presents a simple but compelling example of this connection, which we now present.

1.1.1 Bayesian Numerical Approximation

Consider the problem of computing

$$\int_0^1 u(t)\,dt \tag{1.1}$$

for a given function u, e.g., with

$$u(t) = \sin(t)e^{t^3+\cos(t+t^4)}\,. \tag{1.2}$$

Although u is explicitly known, it does not have a trivial primitive and (1.1) must be approximated by evaluating u at a finite number of points (e.g., $t_i = \frac{i}{N}$, $i \in \{0, 1, \ldots, N\}$; see Figure 1.1) and using a quadrature formula, e.g.,

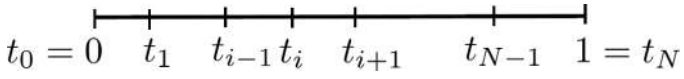

Figure 1.1 Quadrature points.

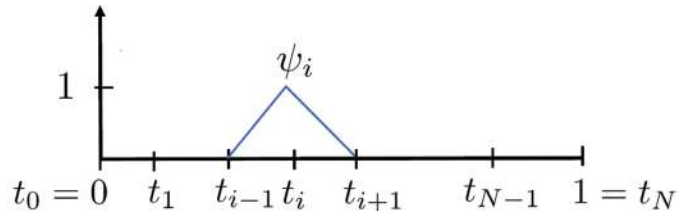

Figure 1.2 Ψ_i.

$$\int_0^1 u(t)\,dt \approx \sum_{i=1}^{N} \frac{u(t_i)+u(t_{i-1})}{2}\Delta t\,, \tag{1.3}$$

with $\Delta t = 1/N$.

Surprisingly, if we instead assume u to be generated by a Brownian motion B_t and approximate $\int_0^1 u(t)\,dt$ with the conditional expectation

$$\mathbb{E}[\int_0^1 B_t\,dt \mid B_{t_i} = u(t_i)\,, \forall i]\,, \tag{1.4}$$

we rediscover the trapezoidal quadrature rule (1.3). To obtain this, switch the integral with the expected value in (1.4), i.e.

$$\mathbb{E}\Big[\int_0^1 B_t\,dt\Big|B_{t_i} = u(t_i) \text{ for all } i\Big] = \int_0^1 \mathbb{E}\Big[B_t\Big|B_{t_i} = u(t_i) \text{ for all } i\Big]\,dt\,, \tag{1.5}$$

and observe that

$$\mathbb{E}\Big[B_t\Big|B_{t_i} = u(t_i) \text{ for all } i\Big] = \sum_i u(t_i)\psi_i(t) \tag{1.6}$$

is the piecewise linear interpolation of u between the points $t_0, \ldots, t_N$ represented by the piecewise linear tent basis functions ψ_i illustrated in Figure 1.2. Moreover, assuming u to be generated by integrals of Brownian motion yields higher-order quadrature rules, i.e., replacing B_t by $\int_0^t B_s\,ds$ yields cubic splines in (1.6) and cubic splines quadrature rules in (1.4). Integrating the Brownian motion k times yields splines and quadrature rules of order $2k+1$.

1.1.2 Bayesian Numerical Homogenization

The Bayesian approach to the discoveries of old and new quadrature rules presented in Section 1.1.1 and in the pioneering works of Diaconis [95], Shaw [278], O'Hagan [232, 233], and Skilling [281] has a natural generalization to partial differential equations (PDEs) [238]. Consider, for instance, the problem of identifying accurate basis functions for the PDE

$$\begin{cases} -\operatorname{div}\big(a(x)\nabla u(x)\big) = f(x) & x \in \Omega; \\ u = 0 \quad \text{on} \quad \partial\Omega, \end{cases} \tag{1.7}$$

where div is the divergence operator, ∇ is the gradient, $\partial\Omega$ is the boundary of a regular subset $\Omega \subset \mathbb{R}^d$, $d \leq 3$, and a is a uniformly elliptic symmetric matrix with entries in $L^\infty(\Omega)$.

We will now consider assuming white noise as a prior for the function f on the right-hand side of (1.7). White noise is a type of weak Gaussian random variable called a *Gaussian field* that we will describe in Definition 7.18, and in full in Section 17.2. In this case, f being white noise amounts to the fact that for each $\phi \in L^2(\Omega)$, the spatial integral $\int_\Omega f\phi$ is a Gaussian random variable with a mean 0 and a variance $\|\phi\|^2_{L^2(\Omega)}$.

Assuming white noise as a prior on the right-hand side of (1.7) and conditioning the solution u to (1.7) on $(u(x_i))_{i\in\{1,\dots,m\}}$ (see Figure 1.3) leads to

$$\mathbb{E}\big[u(x)\big|u(x_i),\, i \in \{1,\dots,m\}\big] = \sum_i u(x_i)\psi_i(x), \tag{1.8}$$

where the ψ_i are deterministic functions. When $a(x) = I_d$, the d-dimensional identity matrix, these ψ_i are the polyharmonic splines of Harder, Desmarais, and Duchon [104, 105, 106, 156] that were originally discovered [156], thanks to the insight of aircraft engineers seeking basis functions adapted to the bending energy of airplane wings.

When the conductivity a is arbitrary in $L^\infty(\Omega)$, then these ψ_i are basis functions that are adapted to the irregularities (microstructure) of the conductivity and

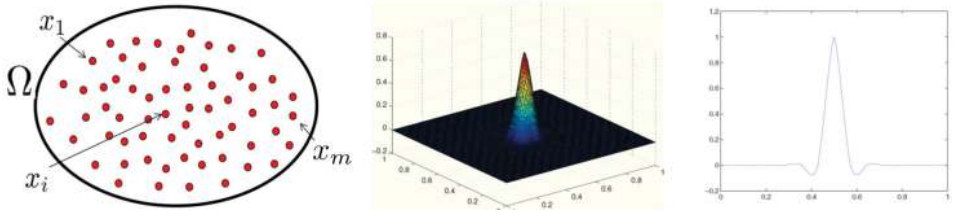

Figure 1.3 Left: the interpolation points. Center: a rough polyharmonic spline ψ_i. Right: x-slice of ψ_i. Reproduced with permission from [249].

provide a generalization of polyharmonic splines to PDEs with rough coefficients, i.e., rough polyharmonic splines [249] (see Figure 1.3), which were discovered after a laborious process of scientific investigation.

1.1.3 Information-Based Complexity

Although this process of randomizing a perfectly known function may seem counterintuitive, a natural framework for understanding it can be found in the pioneering works of Kadane, Traub, Wasilkowsi, and Woźniakowski on information-based complexity (IBC) [175, 306, 334], the branch of computational complexity founded on the observation that numerical implementation requires computation with partial information and limited resources. In IBC, the performance of an algorithm operating on incomplete information can be analyzed in the worst-case or the average-case (randomized) setting with respect to the missing information and as observed by Packel [252], the average case setting could be interpreted as a possible mixed strategy in an adversarial game obtained by lifting a (worst-case) minmax problem to a minmax problem over mixed (randomized) strategies.[1] This observation initiates [239] a natural connection between numerical approximation and Wald's decision theory [322], evidently influenced by von Neumann's theory of games [320].

1.2 The Game Theoretic Perspective

1.2.1 Optimal Recovery

The framework of optimal recovery of Micchelli and Rivlin [218] provides a natural setting for presenting the correspondence between numerical approximation (NA) and Gaussian process regression (GPR) from a game theoretic perspective. Consider a Banach space $\mathcal{B}$ and write $[\cdot,\cdot]$ for the duality product between $\mathcal{B}$ and its dual space $\mathcal{B}^*$. When $\mathcal{B}$ is infinite- or high-dimensional, one cannot directly compute with $u \in \mathcal{B}$ but only with a finite number of *features* of u. The types of features we consider here are represented as a vector

$$\Phi(u) := \big([\phi_1, u], \ldots, [\phi_m, u]\big)$$

corresponding to m linearly independent measurements $\phi_1, \ldots, \phi_m \in \mathcal{B}^*$. The objective is to recover/approximate u from the partial information contained in the feature vector $\Phi(u)$. To quantify errors in the recovery, let

$$Q : \mathcal{B}^* \to \mathcal{B}$$

[1] Such results for certain minmax statistical estimators have also been presented in Li [203] and Sacks and Ylvisaker [264].

be a bijection that is symmetric and positive, in that $[\phi, Q\varphi] = [\varphi, Q\phi]$ and $[\phi, Q\phi] \geq 0$ for $\phi, \varphi \in \mathcal{B}^*$, and endow $\mathcal{B}$ with the quadratic norm $\|\cdot\|$ defined by

$$\|u\|^2 := [Q^{-1}u, u].$$

Then, using the relative error in $\|\cdot\|$-norm as a loss, the classical numerical analysis approach is to approximate u with a minimizer $v^\dagger$ of

$$\min_v \max_u \frac{\|u - v(\Phi(u))\|}{\|u\|}. \tag{1.9}$$

A minimum over all possible functions of the m linear measurements is

$$v^\dagger = \sum_{i=1}^{m} [\phi_i, u]\psi_i, \tag{1.10}$$

where the elements

$$\psi_i := \sum_{j=1}^{m} \Theta^{-1}_{i,j} Q\phi_j, \quad i \in \{1, \ldots, m\}, \tag{1.11}$$

of $\mathcal{B}$, known as optimal recovery splines, are defined using the components $\Theta^{-1}_{i,j}$ of the inverse Θ^{-1} of the Gram matrix Θ defined by $\Theta_{i,j} := [\phi_i, Q\phi_j]$.

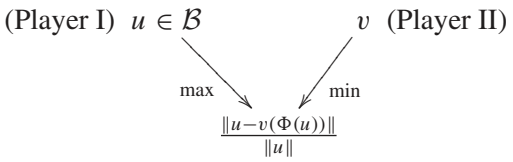

(1.12)

1.2.2 Randomized Strategies and Gamblets

The minmax problem (1.9) can be viewed as the adversarial zero sum game (1.12) in which Player I chooses an element u of the linear space $\mathcal{B}$ and Player II (who does not see u) must approximate Player I's choice based on seeing the finite number of linear measurements $\Phi(u)$ of u.

The function $(u, v) \mapsto \frac{\|u-v(\Phi(u))\|}{\|u\|}$ has no saddle points, so to identify a minmax solution as a saddle point one can proceed, as in von Neumann's game theory [320], by introducing mixed/randomized strategies and lift the problem to probability measures over all possible choices for Players I and II. To articulate optimal strategies, observe that a centered Gaussian field ξ with covariance operator Q, denoted $\xi \sim \mathcal{N}(0, Q)$, is an isometry mapping $\mathcal{B}^*$ to a space of centered Gaussian random variables such that $[\phi, \xi] \sim \mathcal{N}(0, \|\phi\|_*^2)$, $\phi \in \mathcal{B}^*$, where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$ defined by $\|\phi\|_* = \sup_{v \in \mathcal{B}} [\phi, v]/\|v\| = [\phi, Q\phi]^{\frac{1}{2}}$; see, e.g., Janson [171].

For the lifted version of the game (1.12), an optimal strategy of Player I is the centered Gaussian field $\xi \sim \mathcal{N}(0, Q)$ and an optimal strategy of Player II is the pure (deterministic) strategy defined by its conditional expectation

$$v^\dagger = \mathbb{E}\big[\xi \mid [\phi_i, \xi] = [\phi_i, u] \text{ for all } i\big], \tag{1.13}$$

which is equal to the optimal recovery solution (1.10). The optimal recovery splines (1.11) can also be interpreted as elementary gambles/bets

$$\psi_i = \mathbb{E}\big[\xi \mid [\phi_j, \xi] = \delta_{i,j} \text{ for all } j\big], \tag{1.14}$$

which we call *gamblets*, for playing the game. Here the optimal strategy of Player II is a pure strategy because $\|\cdot\|$ is convex and the optimal strategy of Player I is Gaussian because $\|\cdot\|$ is quadratic.

1.2.3 Illustrations

As an illustration of this approach, consider again the numerical quadrature problem associated with computing $\int_0^1 u(t)\,dt$. Take $\mathcal{B} = \mathcal{H}^1[0,1]$ to be the Sobolev space of functions whose first derivatives are square-integrable endowed with the quadratic norm $\|u\|^2 := (u(0))^2 + \int_0^1 (\frac{du(t)}{dt})^2\,dt$ and consider the problem of recovering $u \in \mathcal{B}$ from the incomplete measurements $u(t_i)$ $(= \int_0^1 u\phi_i$ with $\phi_i = \delta(\cdot - t_i))$ using the relative error in $\|\cdot\|$-norm as a loss. Then the Gaussian field ξ defined by the norm $\|\cdot\|$ is a scaled and shifted Brownian motion and (1.13) leads to an approximation that is optimal in both the optimal recovery (worst-case) sense and game theoretic sense, identifying [218] the optimal recovery estimate of the integral with the integral of the optimally estimated u. This recovers the trapezoidal rule with $\int_0^1 u(t)\,dt \approx \int_0^1 \mathbb{E}\big[\xi \mid [\phi_i, \xi] = [\phi_i, u] \text{ for all } i\big]$ by observing that the splines (1.14) are the usual piecewise linear tent basis functions and (1.13) is the piecewise linear interpolation of u.

As another illustration, consider the problem of identifying accurate basis functions for (1.7). Take

$$\mathcal{B} = \big\{u \in \mathcal{H}_0^1(\Omega) \big| \, \|\operatorname{div}(a\nabla u)\|_{L^2(\Omega)} < \infty\big\} \tag{1.15}$$

endowed with the quadratic norm

$$\|u\| := \|\operatorname{div}(a\nabla u)\|_{L^2(\Omega)}\,. \tag{1.16}$$

Consider the problem of recovering $u \in \mathcal{B}$ from the incomplete measurements $u(x_i)$ using the relative error in $\|\cdot\|$ as a loss. Then the Gaussian field ξ defined by the norm $\|\cdot\|$ is white noise and its expectation with respect to the measurements $u(x_i)$ (1.8) leads to an approximation that is optimal in both the optimal recovery (worst-case) sense and game theoretic sense.

1.3 In the Setting of Sobolev Spaces

These interplays provide simple solutions to classical problems in numerical approximation and Gaussian process regression, and we will illustrate this in the setting of a linear operator

$$\mathcal{L} : \mathcal{H}_0^s(\Omega) \to \mathcal{H}^{-s}(\Omega) \tag{1.17}$$

mapping the Sobolev space $\mathcal{H}_0^s(\Omega)$, of functions possessing s square-integrable derivatives that vanish on the boundary (see Section 2.1), to its dual space $\mathcal{H}^{-s}(\Omega)$, where $s, d \in \mathbb{N}^*$ and $\Omega \subset \mathbb{R}^d$ is a regular bounded domain. Assume $\mathcal{L}$ to be an arbitrary symmetric ($\int_\Omega u\mathcal{L}v = \int_\Omega v\mathcal{L}u$), positive ($\int_\Omega u\mathcal{L}u \geq 0$), and local ($\int_\Omega u\mathcal{L}v = 0$ if u and v have disjoint supports) linear bijection. Write $[\phi, u] := \int_\Omega \phi u$ for the duality product between $\phi \in \mathcal{H}^{-s}(\Omega)$ and $u \in \mathcal{H}_0^s(\Omega)$. Let $\mathcal{B}$ be the Sobolev space $\mathcal{H}_0^s(\Omega)$ endowed with the quadratic energy norm $\|u\|^2 := [\mathcal{L}u, u]$.

When $s > d/2$, the Green's function G of $\mathcal{L}$ is a well-defined continuous symmetric positive definite kernel, and one can consider the centered Gaussian process ξ with covariance function G (see Figure 1.4). Consider the problem of finding an approximation of an unknown element $u \in \mathcal{H}_0^s(\Omega)$ given its values at the points $x_1, \ldots, x_m$ (see Figure 1.5). Then, using the relative error in $\|\cdot\|$, as in (1.9), as a loss, the minmax recovery of u is obtained in (1.13) by conditioning the Gaussian process ξ on the values of u at the points $x_1, \ldots, x_m$, and the optimal solution (1.10) corresponds to the formula

$$v^\dagger(x) = \sum_{i,j=1}^{m} u(x_i)\Theta_{i,j}^{-1}G(x_j, x), \tag{1.18}$$

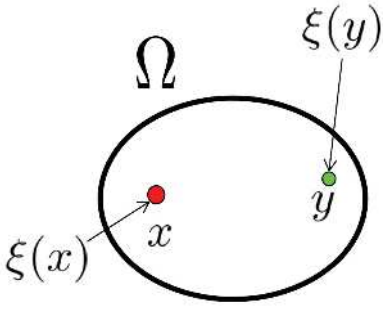

Figure 1.4 For $s > d/2$, ξ is a centered Gaussian process on Ω with covariance function $\mathbb{E}[\xi(x)\xi(y)] = G(x, y)$, where G is the Green's function of the operator $\mathcal{L}$.

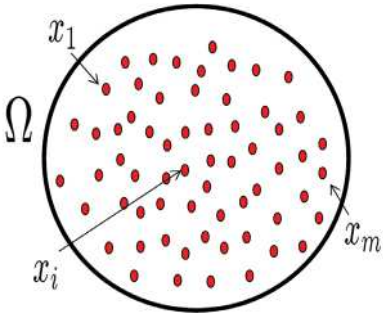

Figure 1.5 Ω and $x_1, \ldots, x_m$.

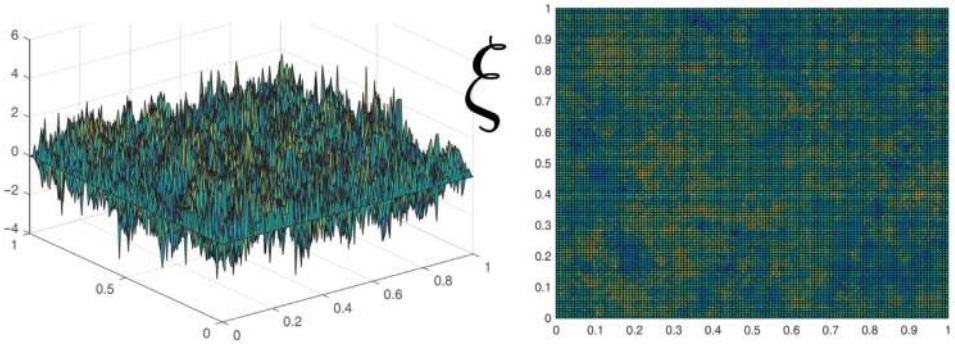

Figure 1.6 Simulation of the Gaussian field ξ.

where $\Theta^{-1}_{i,j}$ is the (i,j)th entry of the inverse Θ^{-1} of the kernel matrix Θ defined by $\Theta_{i,j} := G(x_i, x_j)$, obtained by Kriging u with the kernel G in the reproducing kernel Hilbert space $(\mathcal{H}^s_0(\Omega), \|\cdot\|)$ with reproducing kernel G.

When $s \leq d/2$, the Green's function G of $\mathcal{L}$ exists in the sense of distributions and $\xi \sim \mathcal{N}(0, \mathcal{L}^{-1})$ is defined in a weak sense as a Gaussian field, that is, after integration against a test function $\phi \in \mathcal{H}^{-s}(\Omega)$,

$$\int_\Omega \xi\phi \sim \mathcal{N}(0, \int_{\Omega^2} \phi(x) G(x,y) \phi(y)\, dx\, dy).$$

Figure 1.6 shows an instantiation of ξ for the divergence form elliptic operator $\mathcal{L} := -\operatorname{div}(a\nabla\cdot)$ with a uniformly elliptic, bounded, and rough conductivity $a(x)$.

1.3.1 Numerical Homogenization

Consider the problem of identifying m basis functions that are (1) as accurate as possible in approximating the solution space $\mathcal{L}^{-1}(L^2(\Omega))$ of $\mathcal{L}$ and (2) as localized

as possible. This problem, known as numerical homogenization, is nontrivial because requirements (1) and (2) are conflicting. Indeed, the optimal basis functions for accuracy are the eigenfunctions associated with the lowest eigenvalues of $\mathcal{L}$, which are nonlocalized. As a consequence, this problem is also related to that of identifying Wannier functions for $\mathcal{L}$, i.e., linear combinations of eigenfunctions associated with eigenvalues concentrated around a given eigenvalue such that the resulting linear combinations are also concentrated in space, enabling a low complexity approximation of the eigensubspaces of $\mathcal{L}$.

Conditioning the Gaussian process ξ in (1.14) provides a simple solution [239, 242] to this problem, along with a generalization [238] of polyharmonic splines [156, 249] and of variational multiscale/ localized orthogonal decomposition (LOD basis functions [169, 208]. Given $h > 0$ and $\delta \in (0,1)$, partition Ω into subsets $\tau_1, \ldots, \tau_m$ such that each τ_i is contained in a ball of center x_i and radius $\delta^{-1}h$ and contains a ball of radius δh (see Figure 1.7). Let $\phi_i := 1_{\tau_i}/\sqrt{|\tau_i|}$ be the weighted indicator function of τ_i, where $|\tau_i|$ is the volume of τ_i, or, for $s > d/2$, let $\phi_i := h^{d/2}\delta(\cdot - x_i)$ be the scaled Dirac delta function located at x_i. Then, the splines ψ_i, defined in (1.11) and (1.14) and illustrated in Figure 1.8, achieve the same

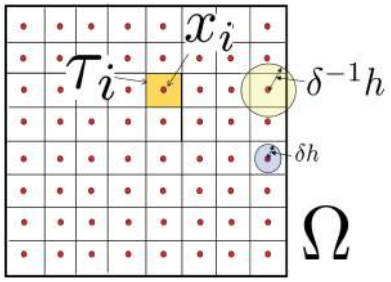

Figure 1.7 τ_i and x_i. h relates to the size of the τ_i and δ^{-2} to their aspect ratios.

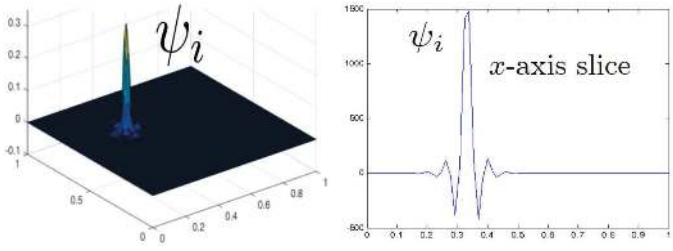

Figure 1.8 Left: ψ_i. Right: x-axis slice of ψ_i. Reproduced from [239] (copyright ©2017 Society for Industrial and Applied Mathematics, reprinted with permission, all rights reserved).

accuracy as the eigenfunctions of $\mathcal{L}$ associated with the m lowest eigenvalues up to a multiplicative constant[2], in that

$$\inf_{v\in\operatorname{span}\{\psi_1,\dots,\psi_m\}} \|\mathcal{L}^{-1}f - v\|_{\mathcal{H}_0^s(\Omega)} \leq Ch^s \|f\|_{L^2(\Omega)},$$

for $f \in L^2(\Omega)$ ($h \approx m^{-\frac{1}{d}}$), *and* they are exponentially localized, in that

$$\|\psi_i\|_{\mathcal{H}^s(\Omega\setminus B(x_i,nh))} \leq Ch^{-s}e^{-n/C}. \tag{1.19}$$

1.3.2 Screening Effect

The preceding results on exponential decay also provide a proof of a version of the phenomenon known, in Kriging and geostatistics, as the *screening effect* [288]. The heuristic idea (for $s > d/2$) is that although $\xi(x)$ and $\xi(y)$ are significantly correlated due to the slow decay of the Green's function $G(x,y)$ in the distance between x and y (see Figure 1.4), they become nearly independent after conditioning on the values of the field at the points in between. For homogeneously spaced points, this effect is obtained from the exponential decay of the gamblets as follows. Write $\operatorname{Cor}(X,Y|\cdot)$ for the conditional correlation between random variables X and Y, and $\langle u,v\rangle := \int_\Omega u\mathcal{L}v$ for the energy scalar product. Then the general identity

$$\operatorname{Cor}\left([\phi_i,\xi],[\phi_j,\xi]\middle|[\phi_l,\xi] \text{ for } l\neq i,j\right) = -\frac{\langle\psi_i,\psi_j\rangle}{\|\psi_i\|\|\psi_j\|},$$

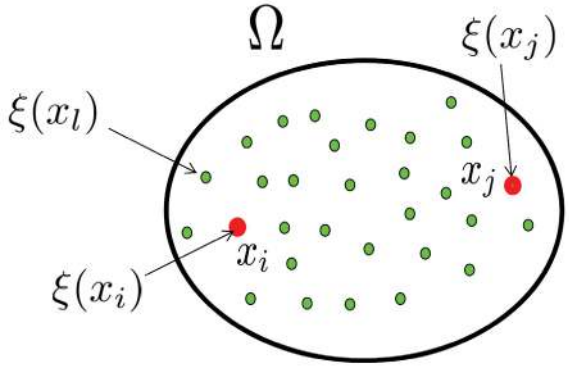

Figure 1.9 Consider the correlation between $\xi(x_i)$ and $\xi(x_j)$ given $\xi(x_l)$ for all $l \neq i,j$.

[2] Throughout, write C for a constant depending only on $\Omega, s, d, \delta, h, \|\mathcal{L}\|$, and $\|\mathcal{L}^{-1}\|$.