


Statistical Numerical Approximation



Houman Owhadi, Clint Scovel, and Florian Schäfer

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. This shared purpose is currently stimulating a growing interest in statistical inference/machine-learning approaches to solving PDEs [8, 13], in the use of randomized algorithms in linear algebra [2], and in the merging of numerical errors with modeling errors in uncertainty quantification [3]. While this interest might be perceived as a recent phenomenon, interplays between numerical approximation and statistical inference are not new. Indeed, they can be traced back to Poincaré’s course in probability theory (1896) and to the pioneering investigations of Sul’din [19], Palasti and Renyi [12], Sard [14], Kimeldorf and Wahba [4] (on the

correspondence between Bayesian estimation and spline smoothing/interpolation), and Larkin [5] (on the correspondence between Gaussian process regression and numerical approximation). Although their study initially “attracted little attention among numerical analysts” [5], it was revived in information-based complexity (IBC) [20], Bayesian numerical analysis [1], and more recently in probabilistic numerics [3]. This short review is an invitation to explore these connections from the consolidating perspective of game/decision theory as presented in [10]. It is motivated by the suggestion that these confluences might not just be objects of curiosity but constitute a pathway to simple solutions to fundamental problems in both areas.

Modeling a Known Function as the Instantiation of a Random Process

In [1], Diaconis presents a simple but compelling connection between numerical analysis and Bayesian inference: Consider the problem of computing

$$\int_0^1 u(t) dt \quad (1)$$

for a given function u (e.g., $u(t) = te^{\sin\sqrt{t}}$). Although u is perfectly known, it does not have a trivial primitive, and its integral must be numerically approximated by evaluating u at a finite number of points (e.g., $t_i = \frac{i}{N}$, $i \in$

Houman Owhadi is a professor of applied and computational mathematics and control and dynamical systems at the California Institute of Technology. His email address is owhadi@caltech.edu.

Clint Scovel is a research associate at the California Institute of Technology. His email address is clintscovel@gmail.com.

Florian Schäfer is a graduate student in applied and computational mathematics at the California Institute of Technology. His email address is Florian.Schaefer@caltech.edu.

Communicated by Notices Associate Editor Reza Malek-Madani.

For permission to reprint this article, please contact:

reprint-permission@ams.org.

DOI: <https://doi.org/10.1090/noti1963>

$\{0, 1, \dots, N\}$) and using a quadrature formula (e.g., $\int_0^1 u(t) dt \approx \sum_{i=1}^N \frac{u(t_i) + u(t_{i-1})}{2} \Delta t$ 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 $\mathbb{E}[\int_0^1 B_t dt \mid B_{t_i} = u(t_i) \forall i]$, we rediscover the trapezoidal quadrature rule. Moreover, assuming u to be generated by integrals of Brownian motion yields higher-order quadrature rules.

Although this approach of modelling a perfectly known function as a sample from a random process may seem counterintuitive, a natural framework for understanding it can be found in information-based complexity (IBC) [20], 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. Moreover, as observed by Packer [11], the average case setting could be interpreted as a mixed strategy in an adversarial game obtained by lifting a (worst case) minmax problem to a minmax problem over mixed (randomized) strategies. This observation initiates [9, 10] a natural connection between numerical approximation and Wald's decision theory, evidently influenced by von Neumann's theory of games.

Optimal Recovery and Gaussian Process Regression

The framework of optimal recovery of Micchelli and Rivlin [7] 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 type of features we consider here are represented as a vector $\Phi(u) := ([\phi_1, u], \dots, [\phi_m, u])$ corresponding to m linearly independent measurements $\phi_1, \dots, \phi_m \in \mathcal{B}^*$. 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 $Q : \mathcal{B}^* \rightarrow \mathcal{B}$ 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 the minimizer v^\dagger of

$$\min_v \max_u \frac{\|u - v(\Phi(u))\|}{\|u\|}. \quad (2)$$

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

$$v^\dagger = \sum_{i=1}^m [\phi_i, u] \psi_i, \quad (3)$$

where the elements

$$\psi_i := \sum_{j=1}^m \Theta_{i,j}^{-1} Q \phi_j, \quad i \in \{1, \dots, m\}, \quad (4)$$

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

The minmax problem (2) can be viewed as the adversarial zero sum game

$$\begin{array}{ccc} \text{(Player I)} & u \in \mathcal{B} & \\ & \searrow & \swarrow \\ & \max & \min & v \text{ (Player II)} \\ & & & \\ & & & \frac{\|u - v(\Phi(u))\|}{\|u\|} \end{array} \quad (5)$$

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, by introducing mixed/randomized strategies and lifting the problem to probability measures over all possible choices for Players I and II. To articulate the 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

$$[\phi, \xi] \sim \mathcal{N}(0, \|\phi\|_*^2), \quad \phi \in \mathcal{B}^*,$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$ defined by $\|\phi\|_* = \sup_{v \in \mathcal{B}} [\phi, v] / \|v\| = [\phi, Q\phi]^{1/2}$. For the lifted version of the game (5), the optimal strategy of Player I is the centered Gaussian field $\xi \sim \mathcal{N}(0, Q)$, and the optimal strategy of Player II is the pure (deterministic) strategy defined by its conditional expectation

$$v^\dagger = \mathbb{E}[\xi \mid [\phi_i, \xi] = [\phi_i, u] \text{ for all } i], \quad (6)$$

which is equal to the optimal recovery solution (3). The optimal recovery splines (4) can also be interpreted as elementary gambles/bets

$$\psi_i = \mathbb{E}[\xi \mid [\phi_j, \xi] = \delta_{i,j} \text{ for all } j], \quad (7)$$

which we call *gamblers*, 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.

As an illustration of this approach, consider again the numerical quadrature problem associated with computing

$\int_0^1 u(t) dt$. Take $\mathcal{B} = \mathcal{H}^1[0, 1]$ 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 (6) leads to an approximation that is optimal in both the optimal recovery (worst case) sense and the game theoretic sense, identifying [7] the optimal recovery estimate of the integral with the integral of the optimally estimated u . This recovers the trapezoidal rule with

$$\int_0^1 u(t) dt \approx \int_0^1 \mathbb{E}[\xi \mid [\phi_i, \xi] = [\phi_i, u] \text{ for all } i]$$

by observing that the splines (7) are the usual piecewise linear tent basis functions and (6) is the piecewise linear interpolation of u .

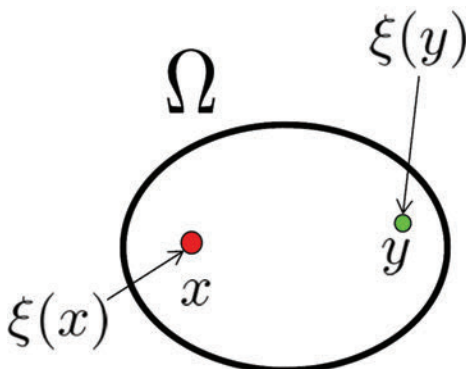


Figure 1. For $s > d/2$, ξ is a centered Gaussian process on Ω with covariance function $\mathbb{E}[\xi(x)\xi(y)] = G(x, y)$, where G is Green's function of the operator \mathcal{L} [10].

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

$$\mathcal{L} : \mathcal{H}_0^s(\Omega) \rightarrow \mathcal{H}^{-s}(\Omega) \quad (8)$$

mapping the Sobolev space $\mathcal{H}_0^s(\Omega)$ 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 $s > d/2$, 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).




Figure 2. Ω and x_1, \dots, x_m [10].

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, \dots, x_m (see Figure 2). Then, using the relative error in $\|\cdot\|$, as in (2), as a loss, the minmax recovery of u is obtained in (6) by conditioning the Gaussian process ξ on the values of u at the points x_1, \dots, x_m , and the optimal solution (3) corresponds to the formula

$$v^\dagger(x) = \sum_{i,j=1}^m u(x_i) \Theta_{i,j}^{-1} G(x_j, x), \quad (9)$$

where $\Theta_{i,j}^{-1}$ is the (i, j) th entry of the inverse Θ^{-1} of the kernel matrix Θ defined by $\Theta_{i,j} := G(x_i, x_j)$. Surprisingly, the standard deviation $\sigma(x)$ of $\xi(x)$ conditioned on the values of ξ at the points x_i (represented by $\sigma^2(x) := G(x, x) - \sum_{i,j=1}^m \Theta_{i,j}^{-1} G(x, x_j) G(x, x_i)$) also bounds the deterministic interpolation error via

$$|u(x) - v^\dagger(x)| \leq \sigma(x) \|u\| \quad \forall x. \quad (10)$$

Estimates such as (10) are obtained in radial basis function interpolation (where the conditional variance $\sigma^2(x)$ is known as the *power function*) by identifying $(\mathcal{H}_0^s(\Omega), \|\cdot\|)$ as a reproducing kernel Hilbert space (RKHS) with reproducing kernel G . In the RKHS framework v^\dagger is identified as the minimizer of $\|v\|$ over $v \in \mathcal{H}_0^s(\Omega)$ subject to $v(x_i) = u(x_i)$ for all i , which, by the representer theorem, can be expressed as (3) (i.e., $v^\dagger(x) = \sum_{j=1}^m c_j G(x_j, x)$, and the c_i are obtained by enforcing the interpolation constraints $v^\dagger(x_i) = u(x_i)$ for all i). This recovery approach is naturally generalized, in regularization and learning theory, to noisy observations y_i of the data $u(x_i)$ by minimizing $\sum_{i=1}^m (y_i - v(x_i))^2 + \lambda \|v\|^2$ over $v \in \mathcal{H}_0^s(\Omega)$ (which corresponds to conditioning ξ on $y_i = u(x_i) + Z_i$, where the Z_i are i.i.d. centered Gaussian random variables with variance λ).

When $s \leq d/2$, 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 3 shows an instan-




Figure 3. Simulation of the Gaussian field ξ [10].

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

Numerical Homogenization

Consider the problem of identifying m basis functions that are (i) as accurate as possible in approximating the solution space $\mathcal{L}^{-1}(L^2(\Omega))$ of \mathcal{L} and (ii) as localized as possible. This problem, known as numerical homogenization, is nontrivial because requirements (i) and (ii) are conflicting. Indeed the optimal basis functions for accuracy are the eigenfunctions associated with the lowest eigenvalues of \mathcal{L} , which are nonlocalized. Conditioning the Gaussian process ξ in (7) provides a simple solution [9, 10] to this problem, along with a generalization [8] of rough polyharmonic splines and of variational multiscale/LOD basis functions [6].

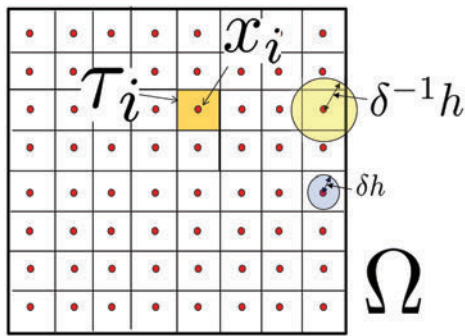


Figure 4. τ_i and x_i . h relates to the size of the τ_i and δ^{-2} to their aspect ratios [10].

Given $h > 0$ and $\delta \in (0, 1)$, partition Ω into subsets τ_1, \dots, τ_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 4). 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 (4) and (7) and illustrated in Figure 5, achieve the same accuracy as the eigenfunctions of \mathcal{L} associated with the m

lowest eigenvalues up to a multiplicative constant,¹ 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

$$\|\psi_i\|_{\mathcal{H}^s(\Omega \setminus B(x_i, nh))} \leq Ch^{-s} e^{-n/C}. \quad (11)$$




Figure 5. Left: ψ_i . Right: x -axis slice of ψ_i [10].

When \mathcal{L} is a power of the Laplacian (i.e., $\mathcal{L} = \Delta^s$) and the ϕ_i are (unscaled) Dirac delta functions, then the ψ_i provide a generalization of the polyharmonic splines of Harder and Desmarais (1972) and Duchon (1977) and of the cardinal splines of Schoenberg [16]. Indeed, when $d = 1$ and the x_i are located on the grid \mathbb{Z} of the real line, the representation (4) identifies the ψ_i with cardinal splines [16] (characterized by $\psi_i(x_j) = \delta_{i,j}$, $\frac{d^{2s}\psi_i}{dx^{2s}} = 0$ on the complement of the points x_i and the continuity of derivatives of order $2s - 2$ of ψ_i). Furthermore, when $d \geq 1$, the general variational formulation

$$\psi_i = \operatorname{argmin} \begin{cases} \text{Minimize } \|\psi\|, \\ \text{Subject to } \psi \in \mathcal{B}, \\ \text{and } [\phi_j, \psi] = \delta_{i,j} \quad \forall j \end{cases} \quad (12)$$

of the optimal recovery splines (4) identifies the ψ_i with polyharmonic splines and, as observed by Madych and Nelson (1990), also identifies polyharmonic splines as a multivariate generalization of cardinal splines.

Screening Effect

The above results on exponential decay also provide proof of a version of the phenomenon known, in Kriging and geostatistics, as the *screening effect* [17]. The heuristic idea (for $s > d/2$) is that although $\xi(x)$ and $\xi(y)$ are significantly correlated due to the slow decay of Green's function $G(x, y)$ in the distance between x and y (see Figure 1), 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

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

$\text{Cor}(X, Y|\cdot)$ 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

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

employed with $\phi_i := \delta(\cdot - x_i)$, equates the correlation of the random variables $\xi(x_i)$ and $\xi(x_j)$ conditioned on the values of $\xi(x_l)$ for all $l \neq i, j$ (see Figure 6) with the negative cosine of the angle between the gamblets ψ_i and ψ_j . Combined with the exponential decay (11), this leads to

$$|\text{Cor}(\xi(x_i), \xi(x_j) | \xi(x_l) \text{ for } l \neq i, j)| \leq C e^{-C^{-1} \frac{|x_i - x_j|}{h}}.$$




Figure 6. Consider the correlation between $\xi(x_i)$ and $\xi(x_j)$ given $\xi(x_l)$ for all $l \neq i, j$ [10].

Operator Adapted Wavelets

Consider the problem of identifying wavelets adapted to the operator \mathcal{L} in the sense that the matrix representation of \mathcal{L} in the basis formed by these wavelets is block-diagonal with uniformly well-conditioned and sparse blocks (Figure 7). The three corresponding properties for these wavelets are (i) orthogonality across scales in the energy scalar product; (ii) uniform boundedness of the condition numbers of the operator within each subband, i.e., uniform Riesz stability in the energy norm; and (iii) exponential decay. As reviewed in [18], although adapted wavelets achieving two of these properties have been constructed, it was not known “if there is a practical technique for ensuring all the three properties simultaneously in general” [18, p. 83].

We now present a solution to this problem [9, 10, 15] using the construction of the elementary bets of the game (5), as illustrated in Figure 8. First, construct a hierarchy

$$\phi_i^{(k)} = \sum_{j \in \mathcal{I}^{(k+1)}} \pi_{i,j}^{(k,k+1)} \phi_j^{(k+1)} \quad (13)$$

of linearly nested elements of $\mathcal{H}^{-s}(\Omega)$, employed to represent the process of computing over a hierarchy of levels of complexity. Figure 8 displays this solution when these




Figure 7. Matrix representation of \mathcal{L} in a finite element basis of fully adapted wavelets [10].

elements are chosen to be the Haar pre-wavelets $\phi_i^{(k)} := \mathbf{1}_{\tau_i^{(k)}} / \sqrt{|\tau_i^{(k)}|}$ of Figure 9.

To construct these Haar pre-wavelets, set $q \in \mathbb{N}^* \cup \{\infty\}$ and $h, \delta \in (0, 1)$. Use k as an index for scale and $i \in \mathcal{I}^{(k)}$ as an index for location, and let the $\tau_i^{(k)}$ be subsets of Ω such that (a) each $\tau_i^{(k)}$ contains a ball of radius δh^k and is contained in a ball of radius $\delta^{-1} h^k$, (b) $(\tau_i^{(k)})_{i \in \mathcal{I}^{(k)}}$ forms a partition of Ω , and (c) $(\tau_i^{(k+1)})_{i \in \mathcal{I}^{(k+1)}}$ forms a subpartition of $(\tau_i^{(k)})_{i \in \mathcal{I}^{(k)}}$.² Now consider the down-scaling game where Player I chooses an unknown element $u \in \mathcal{H}_0^s(\Omega)$ and Player II must approximate u after seeing level k measurements $([\phi_i^{(k)}, u])_{i \in \mathcal{I}^{(k)}}$. Using relative error in $\|\cdot\|$ -norm as a loss, the sequence of optimal bets of Player II, $u^{(k)} = \mathbb{E}[\xi | [\phi_i^{(k)}, \xi] = [\phi_i^{(k)}, u] \forall i \in \mathcal{I}^{(k)}]$, is obtained by conditioning $\xi \sim \mathcal{N}(0, \mathcal{L}^{-1})$ and forming a martingale under the filtration induced by these hierarchical measurements. Conditioning ξ on the elementary measurements $[\phi_i^{(k)}, \xi] = \delta_{i,j}$, represented in the transition from left to right in the upper half of Figure 8, produces the hierarchy of elementary bets, or gamblets,

$$\psi_i^{(k)} := \sum_{j \in \mathcal{I}^{(k)}} \Theta_{i,j}^{(k,-1)} \mathcal{L}^{-1} \phi_j^{(k)}, \quad i \in \mathcal{I}^{(k)}, \quad (14)$$

acting as \mathcal{L} -adapted pre-wavelets, displayed in more detail in Figure 10. In (14) $\Theta_{i,j}^{(k,-1)}$ is the (i, j) th coefficient of the inverse $\Theta^{(k,-1)}$ of $\Theta^{(k)}$ with entries $\Theta_{i,j}^{(k)} = \int \phi_i^{(k)} \mathcal{L}^{-1} \phi_j^{(k)}$.

The nesting of the $\phi_i^{(k)}$ implies that of the $\psi_i^{(k)}$. More specifically, we obtain that

$$\psi_i^{(k)} = \sum_{j \in \mathcal{I}^{(k+1)}} R_{i,j}^{(k,k+1)} \psi_j^{(k+1)}, \quad i \in \mathcal{I}^{(k)}, \quad (15)$$

where the interpolation matrix $R^{(k,k+1)}$ has the entries

$$R_{i,j}^{(k,k+1)} = \mathbb{E}[[\phi_j^{(k)}, \xi] | [\phi_l^{(k-1)}, \xi] = \delta_{i,l}, \quad l \in \mathbb{R}^{\mathcal{I}^{(k-1)}}],$$

²Let the set of labels $\mathcal{I}^{(k)}$ used to label the nested subsets $(\tau_i^{(k)})_{i \in \mathcal{I}^{(k)}, k \in \{1, \dots, q\}}$ be chosen to be a finite set of k -tuples of the form $i = (i_1, \dots, i_k)$.

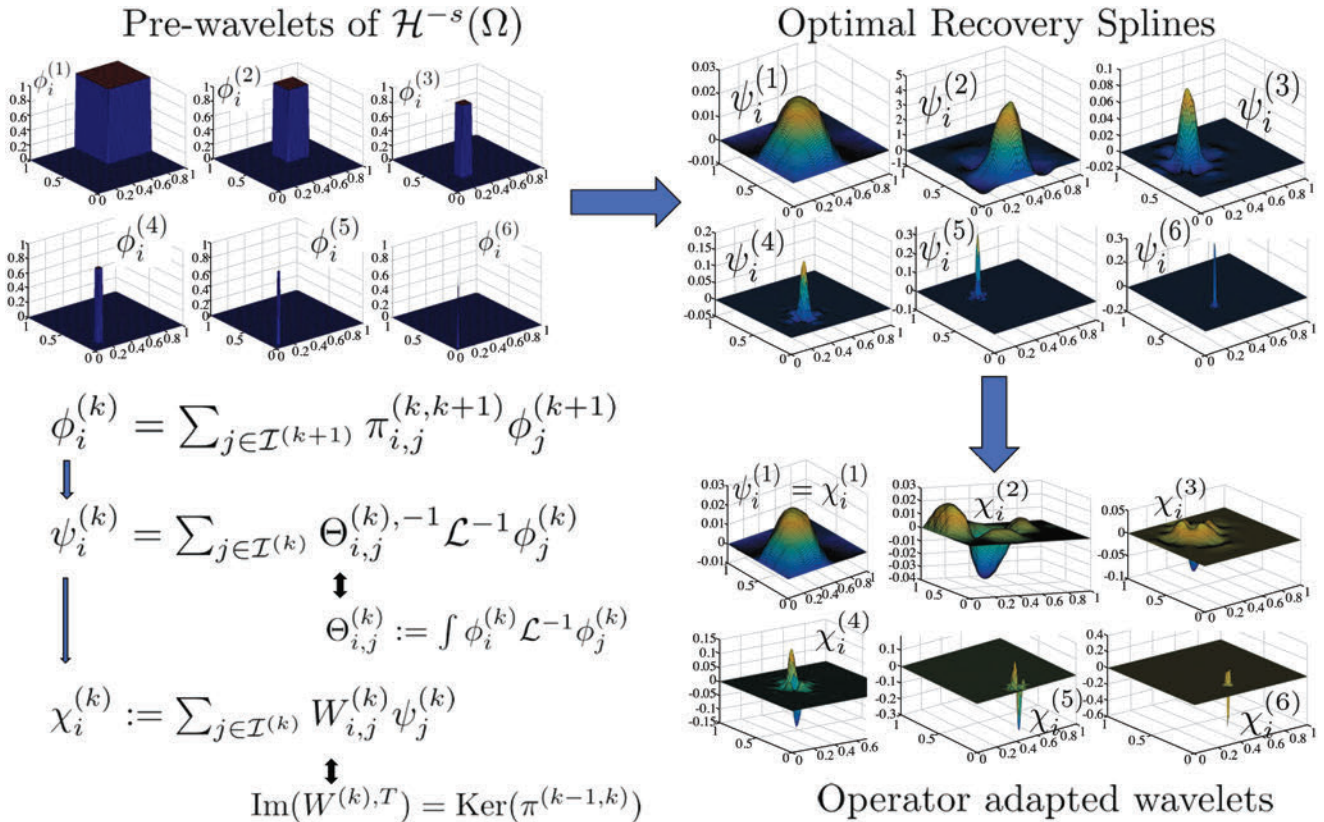


Figure 8. Overview of the construction of operator adapted wavelets [10].

identified, in the upscaling game, as the optimal bet of Player II on $[\phi_j^{(k)}, u]$ given that $[\phi_l^{(k-1)}, u] = \delta_{i,l}$ for $l \in \mathcal{I}^{(k-1)}$. Since the corresponding linear spaces

$$\Psi^{(k)} := \text{span}\{\psi_i^{(k)} \mid i \in \mathcal{I}^{(k)}\}$$

satisfy the nesting relation $\Psi^{(k)} \subset \Psi^{(k+1)}$, the $\langle \cdot, \cdot \rangle$ -orthogonal complement

$$\Psi^{(k)} = \Psi^{(k-1)} \oplus \mathfrak{X}^{(k)}$$

of $\Psi^{(k-1)}$ in $\Psi^{(k)}$ is, under conditions to be discussed, identical to the span

$$\mathfrak{X}^{(k)} = \text{span}\{\chi_i^{(k)} \mid i \in \mathcal{J}^{(k)}\}$$

of the operator adapted wavelets

$$\chi_i^{(k)} := \sum_{j \in \mathcal{J}^{(k)}} W_{i,j}^{(k)} \psi_j^{(k)}, \quad i \in \mathcal{J}^{(k)}, \quad (16)$$

obtained by taking what amounts to local differences of the pre-wavelets $\psi_i^{(k)}$. This transition from the optimal recovery splines $\psi_i^{(k)}$ to the operator adapted wavelets $\chi_i^{(k)}$ is represented in the transition from top to bottom in the right half of Figure 8 and is illustrated in more detail in Figure 11. Here, using $\mathcal{J}^{(k)}$ to label³ the elements $\chi_i^{(k)}$,

³Let $(\mathcal{J}^{(k)})_{2 \leq k \leq q}$ be a finite set of k -tuples of the form $j = (j_1, \dots, j_k)$ such that $(j_1, \dots, j_{k-1}) \in \mathcal{J}^{(k-1)}$ and $|\mathcal{J}^{(k)}| = |\mathcal{J}^{(k)}| - |\mathcal{J}^{(k-1)}|$.

the $W^{(k)}$ are $\mathcal{J}^{(k)} \times \mathcal{I}^{(k)}$ matrices (with orthonormal rows) such that $\text{Ker}(\pi^{(k-1,k)}) = \text{Im}((W^{(k)})^T)$ and $W_{i,j}^{(k)} = 0$ for $(i_1, \dots, i_{k-1}) \neq (j_1, \dots, j_{k-1})$. For simplicity write $\chi_i^{(1)}$ for $\psi_i^{(1)}$ and $\mathfrak{X}^{(1)}$ for $\Psi^{(1)}$. Then (i) the $\chi_i^{(k)}$ are scale-orthogonal across k with respect to the energy scalar product $\langle \cdot, \cdot \rangle$, (ii) they are exponentially localized, and (iii) \mathcal{L} is well conditioned in each subband $\mathfrak{X}^{(k)}$, in the sense that the condition number of the stiffness matrix $B^{(k)}$ with entries

$$B_{i,j}^{(k)} := \langle \chi_i^{(k)}, \chi_j^{(k)} \rangle \quad (17)$$

is uniformly bounded across k by

$$\text{Cond}(B^{(k)}) \leq Ch^{-2} \quad \forall k.$$




Figure 9. The subsets $\tau_i^{(k)}$. Selecting $h := \frac{1}{2}$ and $\delta := \frac{1}{2}$ [10].

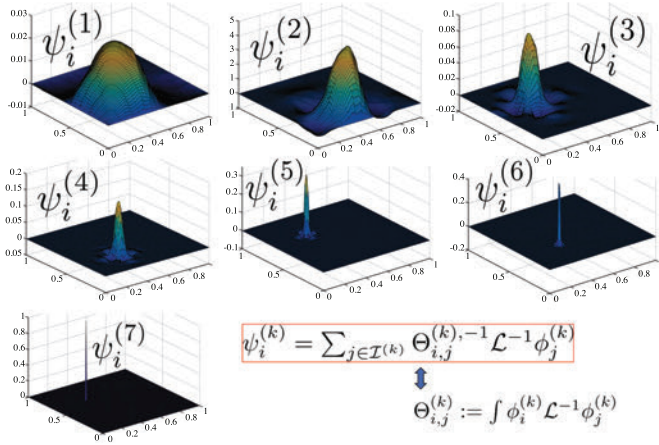


Figure 10. Gamblets $\psi_i^{(k)}$ for $1 \leq k \leq 7$ [10].

Fast Solvers

Since the $\mathfrak{X}^{(k)}$ are scale-orthogonal, the linear system

$$\mathcal{L}u = f \quad (18)$$

can be solved independently in each subband $\mathfrak{X}^{(k)}$ (using the $\chi_i^{(k)}$ as finite elements; see Figure 12 for an illustration of the corresponding subband solutions $u^{(k)} - u^{(k-1)}$), and this gamblet transform has turned (18) into a set of uniformly well-conditioned and sparse linear systems that can be solved independently. The gamblets $\chi_i^{(k)}$ can be computed in $\mathcal{O}(N \log^{2d+1} N)$ complexity based on three properties: (i) the nesting $\psi_i^{(k)} = \sum_j R_{i,j}^{(k,k+1)} \psi_j^{(k+1)}$ enables the hierarchical computation of the $\psi_i^{(k)}$, (ii) the exponential decay of the $\psi_i^{(k)}$ implies the near-sparsity of the interpolation matrices $R^{(k-1,k)}$ and stiffness matrices $B^{(k)}$, and (iii) the uniform bound on $\text{Cond}(B^{(k)})$. Once these gamblets have been computed the linear system (18) can be solved in $\mathcal{O}(N \log^{d+1} N)$ complexity.

Sparse and Rank-revealing Multiresolution Representation of Green's Function

For $q = \infty$ we have the following multiresolution orthogonal direct sum decomposition of the solution space,

$$\mathcal{H}_0^s(\Omega) = \bigoplus_{k=1}^{\infty} \mathfrak{X}^{(k)}, \quad (19)$$

and the multiresolution decomposition

$$G(x, y) = \sum_{k=1}^{\infty} \sum_{i, j \in \mathcal{J}^{(k)}} B_{i,j}^{(k),-1} \chi_i^{(k)}(x) \chi_j^{(k)}(y) \quad (20)$$

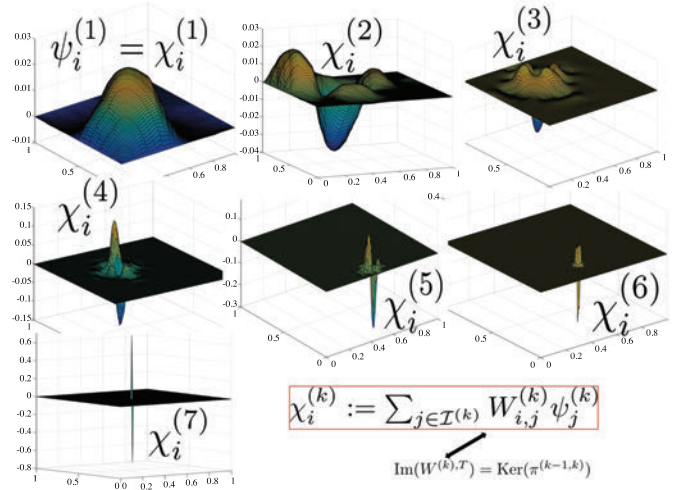


Figure 11. $\psi_i^{(1)}$ and scale-orthogonalized gamblets $\chi_i^{(k)}$ for fixed i and $2 \leq k \leq 7$ [10].

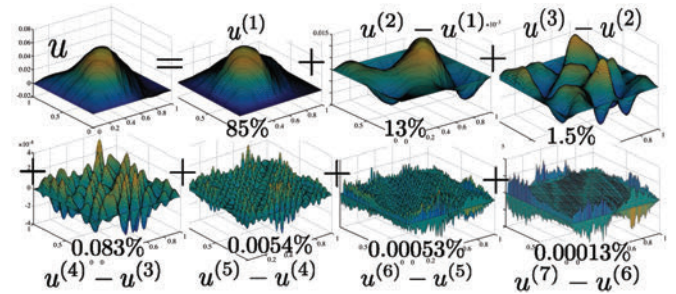


Figure 12. Multiresolution decomposition of the solution u of $\mathcal{L}u = f$ with $f \in L^2(\Omega)$ [10]. The % numbers below the subband projections correspond to the relative energy content of that subband.

of Green's function G of \mathcal{L} . The estimates⁴

$$C^{-1} h^{2sk} J^{(k)} \leq B^{(k),-1} \leq C h^{2sk} J^{(k)} \quad (21)$$

and

$$B_{i,j}^{(k),-1} \leq C h^{2sk} e^{-\frac{d_{i,j}}{cn^k}} \quad (22)$$

imply that the representation (20) is (i) **rank-revealing**, in the sense that the principal submatrix truncation

$$G^{(k)}(x, y) = \sum_{k'=1}^k \sum_{i, j \in \mathcal{J}^{(k')}} B_{i,j}^{(k'),-1} \chi_i^{(k')}(x) \chi_j^{(k')}(y)$$

of G is a low rank approximation of G that is optimal up to a multiplicative constant, i.e., $\|Gf - G^{(k)}f\|_{\mathcal{H}_0^s(\Omega)} \leq C h^{ks} \|f\|_{L^2(\Omega)}$, and (ii) **sparse**, in the sense that for $i, j \in$

⁴Write $J^{(k)}$ for the $\mathcal{J}^{(k)} \times \mathcal{J}^{(k)}$ identity matrix. For $i, j \in \mathcal{J}^{(1)}$ write $d_{i,j}$ for the distance between the support of $\phi_i^{(1)}$ and that of $\phi_j^{(1)}$. For $k \geq 2$ and $i, j \in \mathcal{J}^{(k)}$ write $d_{i,j}$ for the distance between the support of $\phi_i^{(k-1)}$ and that of $\phi_j^{(k-1)}$.

$\mathcal{J}^{(k)}$ and $f \in \mathcal{H}^{-s}(\Omega)$, $\|B_{i,j}^{(k),-1} \chi_i^{(k)} [f, \chi_j^{(k)}]\|_{\mathcal{H}_0^s(\Omega)} \leq C e^{-\frac{d_{i,j}}{ch^k}} \|f\|_{\mathcal{H}^{-s}(\Omega)}$.

Multiresolution Representation of the Gaussian Field ξ

The multiresolution decomposition (20) of Green's function naturally corresponds to a multiresolution decomposition of the Gaussian field $\xi \sim \mathcal{N}(0, \mathcal{L}^{-1})$. Letting $(Y^{(k)})_{k \geq 1}$ be independent Gaussian vectors $Y^{(k)} \sim \mathcal{N}(0, B^{(k),-1})$ with $B^{(k)}$ defined in (17), we have the decomposition

$$\xi = \sum_{k=1}^{\infty} \sum_{i \in \mathcal{J}^{(k)}} Y_i^{(k)} \chi_i^{(k)} \quad (23)$$

of the Gaussian field ξ into modes $\xi^{(k)} - \xi^{(k-1)} = \sum_{i \in \mathcal{J}^{(k)}} Y_i^{(k)} \chi_i^{(k)}$ oscillating at different scales (see Figure 13). Interpreting the Gaussian field ξ as a randomization of u , it follows that the subband components $u^{(k)} - u^{(k-1)}$ of a solution u to $\mathcal{L}u = f$, illustrated in Figure 12, are simply conditional realizations of the modes $\xi^{(k)} - \xi^{(k-1)}$ on the information contained in the difference between level k and level $k - 1$ measurements.




Figure 13. Simulation of the Gaussian fields $\xi^{(1)}$ and $(\xi^{(k)} - \xi^{(k-1)})_{k \geq 2}$ [10].

The Cholesky Factorization of Dense Kernel Matrices

Now, let us demonstrate the confluence of NA and GPR in elementary linear algebra operations. Let x_1, \dots, x_N be N homogeneously distributed points of Ω . For $s > d/2$ consider the $N \times N$ symmetric positive definite dense kernel matrix Θ defined by

$$\Theta_{i,j} = G(x_i, x_j). \quad (24)$$

Since Θ is dense the computational complexities of simple linear algebra operations (with vanilla methods) are as follows: (i) storage: $\mathcal{O}(N^2)$, (ii) Θv : $\mathcal{O}(N^2)$, (iii) $\Theta^{-1} v$: $\mathcal{O}(N^3)$, (iv) $\det(\Theta)$: $\mathcal{O}(N^3)$, and (v) principal component analysis of Θ : $\mathcal{O}(N^3)$.

Opening these complexity bottlenecks is of practical importance in (i) computational physics, (ii) Gaussian process statistics (where Θ is the covariance matrix of the vector $(\xi(x_1), \dots, \xi(x_N))$ corresponding to the centered Gaussian process ξ with covariance function G), and (iii) kernel methods for machine learning (e.g., support vector machines).




Figure 14. Decomposition of x_1, \dots, x_N into a nested hierarchy [15].

Here, the link between NA and GPR can be used to derive a simple incomplete Cholesky factorization algorithm enabling the operations mentioned above in near-linear complexity [15]. This algorithm (i) selects $\mathcal{O}(N \log N \log^d(N/\epsilon))$ entries of Θ and an ordering x_1, \dots, x_N , represented by a permutation matrix P ; (ii) from these entries and permutation, in complexity $\mathcal{O}(N \log N \log^d(N/\epsilon))$ in space and $\mathcal{O}(N \log^2 N \log^{2d}(N/\epsilon))$ in time, it computes a sparse lower triangular matrix L such that the number of nonzero entries of L is $\mathcal{O}(N \log N \log^d(N/\epsilon))$ and

$$\|P^T \Theta P - LL^T\|_F \leq \epsilon, \quad (25)$$

where $\|\cdot\|_F$ is the Frobenius norm.

The idea of the algorithm is to first decompose x_1, \dots, x_N into a nested hierarchy $\{x_i\}_{i \in \mathcal{J}^{(1)}} \subset \{x_i\}_{i \in \mathcal{J}^{(2)}} \subset \dots \subset \{x_i\}_{i \in \mathcal{J}^{(q)}}$ (see Figure 14) and order the points from $\mathcal{J}^{(1)} = \mathcal{J}^{(1)}$, $\mathcal{J}^{(2)} = \mathcal{J}^{(2)}/\mathcal{J}^{(1)}$ to $\mathcal{J}^{(q)} = \mathcal{J}^{(q)}/\mathcal{J}^{(q-1)}$. We obtain $P^T \Theta P$ by ordering the rows and columns of Θ accordingly.

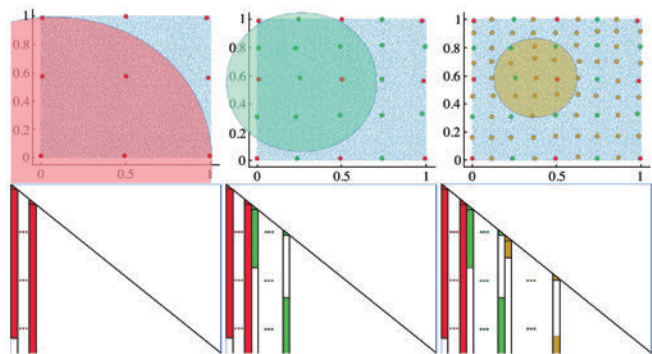


Figure 15. The sparsity pattern S [15]. For an index i at level k , e.g., a red point, indices j corresponding to points within a ball of radius $\sim 2^{-k} \ln \frac{1}{\epsilon}$ comprise the i th column in red.

The next step is to use this hierarchical ordering to define the sparsity pattern $S := \{i, j \in \mathcal{I}^{(q)} \mid i \in \mathcal{J}^{(k)}, j \in \mathcal{J}^{(l)}, |x_i - x_j| \leq \ln \frac{N}{\epsilon} 2^{-\min(k,l)}\}$ corresponding to a sparse subset of the entries of the matrix Θ (see Figure 15). The incomplete Cholesky factorization LL^T of $P^T\Theta P$ is then computed in Algorithm 1 as one small tweak from the classical algorithm: skip all operations for which (k, i) , (k, j) , or (i, j) are outside the sparsity pattern S .

Algorithm 1 Incomplete *sparse* Cholesky factorization of [15]. Takes the entries $(i, j) \in S$ of the $N \times N$ symmetric positive definite matrix $A \leftarrow P^T\Theta P$ as input and returns the $N \times N$ lower triangular matrix L as output.

```

1: for  $i = 1$  to  $N$  do
2:    $L_{:,i} \leftarrow A_{:,i} / \sqrt{A_{i,i}}$ 
3:   for  $j \in \{i + 1, \dots, N\} : (i, j) \in S$  do
4:     for  $k \in \{j, \dots, N\} : (k, i), (k, j) \in S$  do
5:        $A_{k,j} \leftarrow A_{k,j} - \frac{A_{k,i}A_{i,j}}{A_{i,i}}$ 
6:     end for
7:   end for
8: end for

```

Why does it work? The Cholesky factorization computes a lower triangular matrix L that satisfies $P^T\Theta P = LL^T$. The resulting factor L depends on the ordering of the points x_1, \dots, x_N (as represented by P). Although the factors L are dense under the lexicographic ordering, they are near sparse, that is, approximately sparse, under the hierarchical ordering of Figure 14. While this observation is new, it has a simple explanation from a GPR perspective. To simplify notation, assume that the rows and columns of Θ are ordered in the hierarchical ordering, and write $\Theta_{1:l,1:l}$ for the $\{1, \dots, l\} \times \{1, \dots, l\}$ submatrix of Θ . The Cholesky factorization algorithm iteratively computes the Schur complements $S^{(l)}$ (from $l = 1$ to $l = N$) of $\Theta_{1:l,1:l}$ in Θ , which have the probabilistic interpretation as conditional covariance matrices

$$S_{i,j}^{(l)} = \text{Cov}(\xi(x_i), \xi(x_j) \mid \xi(x_k), k \leq l),$$

where ξ is the centered Gaussian process with covariance function G . The screening effect implies that as l reaches finer levels of the hierarchy, the conditional correlation length of ξ decreases and the Schur complements $S^{(l)}$ become increasingly sparse, leading to near sparsity of the Cholesky factors L .

The near sparsity of L is also related to the exponential decay of the gamblets. Choose the hierarchy (13) of linearly nested measurement functions to be subsampled Diracs $\phi_i^{(k)} := \delta(\cdot - x_i)$ for $i \in \mathcal{I}^{(k)}$. Let $\chi_i^{(k)}$ be the corresponding operator adapted wavelets obtained through (16) with $W_{i,j}^{(k)} = \delta_{i,j}$ and let $B^{(k)}$ be as in (17). Then Θ

admits the block-Cholesky factorization

$$\Theta = \bar{L}D\bar{L}^T,$$

where D is a block diagonal matrix with the $\mathcal{J}^{(k)} \times \mathcal{J}^{(k)}$ diagonal block equal to $B^{(k),-1}$, and, for $i \in \mathcal{J}^{(k)}, j \in \mathcal{J}^{(k')}$, $\bar{L}_{i,j}$ is defined by

$$\bar{L}_{i,j} = \begin{cases} 0, & k < k', \\ \delta_{i,j}, & k = k', \\ [\phi_i^{(k)}, \chi_j^{(k')}], & k > k'. \end{cases}$$

The near sparsity of \bar{L} follows from the exponential decay of the $\chi_j^{(k')}$ and the fact that the measurement functions $\phi_i^{(k)}$ are Dirac measures. The estimates (21) and (22) imply the near sparsity of the Cholesky factors \hat{L} of D and hence of $L = \bar{L}\hat{L}$. The approximation error estimate (25) is then obtained by matching the sparsity set S with the near sparsity structure of L . Furthermore, the rank revealing property (20) of gamblets implies that the Cholesky decomposition is not only sparse but also rank-revealing, i.e.,

$$\|P^T\Theta P - L^{(k)}L^{(k),T}\| \leq C\|\Theta\|k^{-\frac{2s}{d}},$$

where $L^{(k)}$ is the rank- k matrix defined by the first k columns of the Cholesky factor L of Θ and $\|\Theta\|$ is the operator norm of Θ .

ACKNOWLEDGMENTS. The authors gratefully acknowledge that this work was supported by the Air Force Office of Scientific Research and the DARPA EQUIPS Program under award number FA9550-16-1-0054 (Computational Information Games), the Air Force Office of Scientific Research under award number FA9550-18-1-0271 (Games for Computation and Learning), and the Office of Naval Research under award N00014-18-1-2363 (Toward scalable universal solvers for linear systems). Figures 1 to 13 have been used from the forthcoming book [10] with permission from Cambridge University Press.

References

- [1] Diaconis P. Bayesian numerical analysis. *Statistical Decision Theory and Related Topics, IV, Vol. 1* (West Lafayette, Ind., 1986), Springer, New York, 1988, 163–175. MR927099
- [2] Halko N, Martinsson PG, Tropp JA. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.*, no. 2 (53):217–288, 2011, dx.doi.org/10.1137/090771806. MR2806637
- [3] Hennig P, Osborne MA, Girolami M. Probabilistic numerics and uncertainty in computations, *Proc. Roy. Soc. A.*, no. 2179 (471):20150142, 2015, dx.doi.org/10.1098/rspa.2015.0142. MR3378744

- [4] Kimeldorf GS, Wahba G. A correspondence between Bayesian estimation on stochastic processes and smoothing by splines, *Ann. Math. Statist.* (41):495–502, 1970. MR254999
- [5] Larkin FM. Gaussian measure in Hilbert space and applications in numerical analysis, *J. Math.*, no. 3 (2), 1972. MR0303193
- [6] Målqvist A, Peterseim D. Localization of elliptic multiscale problems, *Math. of Comp.*, no. 290 (83):2583–2603, 2014. MR3246801
- [7] Micchelli CA, Rivlin TJ. A survey of optimal recovery. *Optimal Estimation in Approximation Theory*, Springer, 1977, 1–54. MR0617931
- [8] Owhadi H. Bayesian numerical homogenization, *Multiscale Model. Simul.*, no. 3 (13):812–828, 2015, [dx.doi.org/10.1137/140974596](https://doi.org/10.1137/140974596). MR3369060
- [9] Owhadi H. Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games, *SIAM Rev.*, no. 1 (59):99–149, 2017. MR3605827
- [10] Owhadi H, Scovel C. *Operator Adapted Wavelets, Fast Solvers, and Numerical Homogenization from a Game Theoretic Approach to Numerical Approximation and Algorithm Design*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, 2019.
- [11] Packer EW. The algorithm designer versus nature: a game-theoretic approach to information-based complexity, *J. Complexity*, no. 3 (3):244–257, 1987, [dx.doi.org/10.1016/0885-064X\(87\)90014-8](https://doi.org/10.1016/0885-064X(87)90014-8). MR919675
- [12] Palasti I, Renyi A. On interpolation theory and the theory of games, *MTA Mat. Kat. Int. Kozl* (1):529–540, 1956.
- [13] Raissi M, Perdikaris P, Karniadakis GE. Inferring solutions of differential equations using noisy multi-fidelity data, *J. Comput. Phys.* (335):736–746, 2017. MR3612520
- [14] Sard A. *Linear Approximation*, Vol. 9, Amer. Math. Soc., 1963. MR0158203
- [15] Schäfer F, Sullivan TJ, Owhadi H. Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity, arXiv:1706.02205, 2017.
- [16] Schoenberg IJ. *Cardinal Spline Interpolation*, Vol. 12, SIAM, 1973. MR0420078
- [17] Stein ML. The screening effect in Kriging, *Ann. Statist.*:298–323, 2002. MR1892665
- [18] Sudarshan R. *Operator-adapted finite element wavelets: Theory and applications to a posteriori error estimation and adaptive computational modeling*. Thesis (Ph.D.)—Massachusetts Institute of Technology, 2005, ProQuest LLC, Ann Arbor, MI. MR2717264
- [19] Sul'din AV. Wiener measure and its applications to approximation methods. I, *Iz. Vysš. Učebn. Zaved. Matematika*. (6):145–158, 1959. MR0157489
- [20] Traub JF, Wasilkowski GW, Woźniakowski H. *Information-Based Complexity*, with contributions by A. G. Werschulz and T. Boult, Computer Science and Scientific Computing, Academic Press, Inc., Boston, MA, 1988. MR958691



Houman Owhadi



Clint Scovel



Florian Schäfer

Credits

All figures and author photos are courtesy of the authors.

2019-2020 MEMBERSHIP

INSTITUTE FOR
ADVANCED STUDY

MEMBERSHIPS

THE IAS SCHOOL OF MATHEMATICS welcomes applications from mathematicians and theoretical computer scientists at all career levels, and strongly encourages applications from women, minorities, and mid-career scientists (5-15 years from Ph.D.). Competitive salaries, on-campus housing, and other resources are available for periods of 4-11 months for researchers in all mathematical subject areas. The School supports approximately 40 post-docs per year.

In 2020-2021, there will be a special-year program, **GEOMETRIC AND MODULAR REPRESENTATION THEORY**, led by Geordie Williamson of the University of Sydney; however, Membership will not be limited to mathematicians in this field.

To apply, submit your application at [mathjobs](https://mathjobs.org).
For more info, please visit: math.ias.edu

MID-CAREER

Are you 5-15 years from your Ph.D.? Ask us about funding!

PROGRAMS

SUMMER COLLABORATORS

math.ias.edu/summercollaborators

WOMEN & MATHEMATICS

math.ias.edu/wam/2020

EMERGING TOPICS

math.ias.edu/

DEADLINE: DECEMBER 1, 2019 • [MATHJOBS.ORG](https://mathjobs.org)