Acta Numerica (2015), pp. 1–159
doi:10.1017/S0962492915000033

Approximation of high-dimensional parametric PDEs*

Albert Cohen
Laboratoire Jacques-Louis Lions,
Université Pierre et Marie Curie,
75005 Paris, France
E-mail: cohen@ann.jussieu.fr

Ronald DeVore
Department of Mathematics,
Texas A&M University,
College Station, TX 77840, USA
E-mail: ronald.a.devore@gmail.com

Parametrized families of PDEs arise in various contexts such as inverse problems, control and optimization, risk assessment, and uncertainty quantification. In most of these applications, the number of parameters is large or perhaps even infinite. Thus, the development of numerical methods for these parametric problems is faced with the possible curse of dimensionality. This article is directed at (i) identifying and understanding which properties of parametric equations allow one to avoid this curse and (ii) developing and analysing effective numerical methods which fully exploit these properties and, in turn, are immune to the growth in dimensionality.

Part I of this article studies the smoothness and approximability of the solution map, that is, the map $a \mapsto u(a)$, where a is the parameter value and $u(a)$ is the corresponding solution to the PDE. It is shown that for many relevant parametric PDEs, the parametric smoothness of this map is typically holomorphic and also highly anisotropic, in that the relevant parameters are of widely varying importance in describing the solution. These two properties are then exploited to establish convergence rates of n-term approximations to the solution map, for which each term is separable in the parametric and physical variables. These results reveal that, at least on a theoretical level, the solution map can be well approximated by discretizations of moderate complexity, thereby showing how the curse of dimensionality is broken. This theoretical analysis is carried out through concepts of approximation theory such as best n-term approximation, sparsity, and n-widths. These notions

* This research was supported by ONR contracts N00014-11-1-0712 and N00014-12-1-0561, NSF grant DMS 1222715, the Institut Universitaire de France and ERC advanced grant BREAD.

determine *a priori* the best possible performance of numerical methods and thus serve as a benchmark for concrete algorithms.

Part II of this article turns to the development of numerical algorithms based on the theoretically established sparse separable approximations. The numerical methods studied fall into two general categories. The first uses polynomial expansions in terms of the parameters to approximate the solution map. The second one searches for suitable low-dimensional spaces for simultaneously approximating all members of the parametric family. The numerical implementation of these approaches is carried out through adaptive and greedy algorithms. An *a priori* analysis of the performance of these algorithms establishes how well they meet the theoretical benchmarks.

CONTENTS

1. Overview

1.1. Parametric and stochastic PDEs

Partial differential equations (PDEs) are commonly used to model complex systems in a variety of physical contexts. When solving a given PDE, one typically fixes certain *parameters*: the shape of the physical domain, the diffusion or velocity field, the source term, the flux or reaction law, *etc.* We use the terminology *parametric PDEs* when some of these parameters are allowed to vary over a certain range of interest. When treating such parametric PDEs, one is interested in finding the solution for all parameters in the range of interest.

To describe such problems in their full generality, we adopt the formulation

$$\mathcal{P}(u, a) = 0, \tag{1.1}$$

where a denotes the parameters, u is the unknown of the problem, and

$$\mathcal{P} : V \times X \to W \tag{1.2}$$

is a linear or nonlinear partial differential operator, with (X, V, W) a triplet of Banach spaces. We assume that the parameter a ranges over a compact set $\mathcal{A} \subset X$, and that for any a in this range there exists a unique solution $u = u(a) \in V$ to (1.1). This allows us to define the *solution map*

$$u : a \mapsto u(a), \tag{1.3}$$

which acts from X onto V and is well defined over $\mathcal{A}$. We also define the *solution manifold* as the family

$$\mathcal{M} = u(\mathcal{A}) = \{u(a) : a \in \mathcal{A}\}, \tag{1.4}$$

which gathers together all solutions as the parameter varies within its range.

One simple guiding example, which will be often used throughout this article, is the linear elliptic equation

$$\begin{aligned} -\operatorname{div}(a\nabla u) &= f \quad \text{on } D, \\ u &= 0 \quad \text{on } \partial D, \end{aligned} \tag{1.5}$$

set on a Lipschitz domain $D \subset \mathbb{R}^m$. Here, we fix the right-hand side f as a real-valued function and consider real-valued diffusion coefficients a as the parameter. The corresponding operator $\mathcal{P}$ is therefore given by

$$\mathcal{P}(u, a) = f + \operatorname{div}(a\nabla u). \tag{1.6}$$

A possible choice for the triplet of spaces is then

$$(X, V, W) = (L^\infty(D), H_0^1(D), H^{-1}(D)). \tag{1.7}$$

Indeed, if $u \in V$, $a \in X$ and $f \in W$, one then defines $\mathcal{P}(u, a)$ as an element of W by requiring that

$$\langle \mathcal{P}(u, a), v\rangle = \langle f, v\rangle + \int_D a\nabla u \cdot \nabla v, \quad v \in V, \tag{1.8}$$

where $\langle \cdot, \cdot \rangle$ is the duality bracket between $V' = W$ and V. Lax–Milgram theory ensures the existence and uniqueness of a solution $u(a)$ to (1.1) from V if, for some $r > 0$, the diffusion a satisfies the ellipticity condition

$$a(x) \geq r, \quad x \in D. \tag{1.9}$$

Therefore, a typical parameter range is a set $\mathcal{A} \subset \{a \in L^\infty(D) : a \geq r\}$, which in addition is assumed to be compact in L^∞.

Although elementary, the above example gathers important features that are present in other relevant examples of parametric PDEs. In particular, the solution map $a \mapsto u(a)$ acts from an infinite-dimensional space into another infinite-dimensional space. Also note that, while the operator $\mathcal{P}$

of (1.6) is linear both in a and u (up to the constant additive term f), the solution map is *nonlinear*. Because of the high-dimensionality of the parameter space X, such problems represent a significant challenge when trying to capture this map numerically. One objective of this article is to understand which properties of this map allow for a successful numerical treatment. Concepts such as holomorphy, sparsity, and adaptivity are at the heart of our development.

The solution map may also be viewed as a function

$$(x, a) \mapsto u(x, a) \tag{1.10}$$

of both the physical variable $x \in D$ and the parametric variable $a \in \mathcal{A}$. The parametric variable has a particular status because its different instances are uncoupled: for any fixed instance $a = a_0$, we may solve the PDE exactly or approximately and therefore compute $u(x, a_0)$ for all values of x, while ignoring all other values $a \neq a_0$. This plays an important role for certain numerical methods which are based on solving the parametric PDE for different particular values of a, since this task can be parallelized.

Parametric PDEs occur in a variety of modelling contexts. We draw the following major distinctions in their setting.

- *Deterministic modelling.* The parameters are deterministic *design or control* variables, which may be tuned by the user so that the solution u, or some quantity of interest $Q(u)$, has prescribed properties. For instance, if the elliptic equation (1.5) is used to model the heat conduction in a component produced by an industrial process, one may want to design the material in order to minimize the heat flux on a certain part on the boundary $\Gamma \subset \partial D$, in which case the quantity of interest to be optimized is

$$Q(u) = \int_\Gamma a \nabla u \cdot \mathbf{n}, \tag{1.11}$$

  where $\mathbf{n}$ is the outer normal. This amounts to optimizing the function

$$a \mapsto F(a) := Q(u(a)) \tag{1.12}$$

  over $\mathcal{A}$.

- *Stochastic modelling.* The parameters are random variables with prescribed probability laws, which account for *uncertainties* in the model. Therefore a has a certain probability distribution μ supported on $\mathcal{A}$. One is then typically interested in the resulting probabilistic properties of the solution u, which is itself a random variable over $\mathcal{A}$ with values in V, or in the probabilistic properties of a quantity of interest $Q(u)$. For instance, if the elliptic equation (1.5) is used to model oil or ground water diffusion, a common way to deal with the uncertainty of the

underground porous media is to define a as a random field with some prescribed law. Then one might want to estimate the mean solution

$$\overline{u} := \mathbb{E}(u), \tag{1.13}$$

or its variance

$$V(u) := \mathbb{E}(\|u - \overline{u}\|_V^2), \tag{1.14}$$

or the average flux through a certain interface Γ, that is, $\mathbb{E}(Q(u))$ with $Q(u)$ as in (1.11), or the probability that this flux exceeds a certain quantity, *etc.*

In both deterministic and stochastic settings, the given application may require evaluation of $u(a^i)$ for a large number $n \gg 1$ of values $\{a^1, \dots, a^n\}$ of the parameter a. This is the case, for instance, when using a descent method for optimizing a quantity of interest in the deterministic framework, or a Monte Carlo method for evaluating an expectation in the stochastic framework. Since each individual instance $u(a)$ is the solution of a PDE, its exact evaluation is typically out of reach. Instead, each query for $u(a^i)$ is approximately evaluated by a numerical solver, which may itself be computationally intensive for high accuracy.

In order to significantly reduce the number of computations needed for attaining a prescribed accuracy, alternative strategies, commonly referred to as *reduced modelling*, have been developed. Understanding which of these strategies are effective in the case where the parameter has large or infinite dimension, and why, is the subject of this article.

1.2. Affine representation of the parameters

So far our description of the set $\mathcal{A}$ of parameters allows it to be any compact subset of X. An important ingredient in both our theoretical and numerical developments is to identify any $a \in \mathcal{A}$ through a sequence of real numbers. We are especially interested in affine representations of $\mathcal{A}$. We say that a sequence $(\psi_j)_{j\geq 1}$ of functions $\psi_j \in X$ is an *affine representer* for $\mathcal{A}$, or *representer* for short, if we can write each a as

$$a = a(y) = \overline{a} + \sum_{j\geq 1} y_j \psi_j, \quad y := (y_j)_{j\geq 1}, \quad y_j = y_j(a), \tag{1.15}$$

where the y_j are real numbers, $\overline{a}$ is a fixed function from X, and the series converges in the norm of X for each $a \in \mathcal{A}$. We are making a slight abuse of notation here since we use a to represent a general element of $\mathcal{A}$ and also use a to represent the map

$$a : y \mapsto a(y), \tag{1.16}$$

from $\mathbb{R}^{\mathbb{N}}$ to X. However, the meaning will always be clear from the context.

It is easy to see that for any compact set $\mathcal{A}$ in a Banach space X, affine representers exist. For example, if X has a Schauder basis then any such basis will be a representer. Even if X does not admit a Schauder basis, as is the case for our example $X = L^\infty(D)$, we can still find representers as follows. Choose any $\bar{a} \in \mathcal{A}$. Since $\mathcal{K} := \mathcal{A} - \bar{a}$ is compact, there exist finite-dimensional spaces $(X_n)_{n\geq 0}$, with $\dim(X_n) = n$, such that

$$\operatorname{dist}(\mathcal{K}, X_n)_X := \sup_{a\in\mathcal{K}} \min_{b\in X_n} \|a - b\|_X \to 0 \quad \text{as } n \to \infty. \tag{1.17}$$

We can also take the spaces X_n to be nested, that is

$$X_n \subset X_{n+1}, \quad n \geq 0. \tag{1.18}$$

Let $(\phi_{j,n})_{j=1,\dots,n}$ be any basis for X_n and define $N_n := n(n-1)/2$. The sequence

$$\psi_j := \phi_{j-N_n,n}, \quad j = N_n, \dots, N_{n+1}, \tag{1.19}$$

contains each of the bases $(\phi_{j,n})_{j=1,\dots,n}$ for all $n \geq 1$. Given any $a \in \mathcal{A}$, let a_n be a best approximation in X to $a - \bar{a}$ from X_n, with $a_0 := 0$. Then we can write

$$a = \bar{a} + \sum_{n=1}^{\infty} (a_n - a_{n-1}). \tag{1.20}$$

Each term $a_n - a_{n-1}$ is in X_n and hence can be written as a linear combination of the ψ_j. Therefore, $(\psi_j)_{j\geq 1}$ is a representer for $\mathcal{A}$.

Affine representations (1.15) often occur in the natural formulation of the parametric problem. For instance, if the diffusion coefficient a in (1.5) is piecewise constant over a fixed partition $\{D_j\}_{j=1,\dots,d}$ of the physical domain D, then it is natural to set

$$a(y) = \overline{a} + \sum_{j=1}^{d} y_j \chi_{D_j}, \tag{1.21}$$

where $\overline{a}$ is a constant and the χ_{D_j} are the characteristic functions of the subdomains D_j. Similarly, if the parameter a describes the shape of the boundary of the physical domain in a computer-aided design setting, a typical format is

$$a(y) = \overline{a} + \sum_{j=1}^{d} y_j B_j, \tag{1.22}$$

where $\overline{a}$ represents a nominal shape and the B_j are B-spline functions associated to control points. In these two examples, d is finite, yet possibly very large.

In the statistical context, if a is a second-order random field over a domain D, a frequently used choice in (1.15) is $\overline{a} := \mathbb{E}(a)$, the average field,

and $(\psi_j)_{j\geq 0}$, the Karhunen–Loève basis, that is, the eigenfunctions of the covariance operator

$$v \mapsto R_a v := \int_D C_a(\cdot, x)v(x)\,\mathrm{d}x, \quad C_a(z,x) := \mathbb{E}((a(z)-\overline{a}(z))(a(x)-\overline{a}(x)). \tag{1.23}$$

Then the resulting scalar variables are centred and uncorrelated, that is, $\mathbb{E}(y_i) = 0$ and $\mathbb{E}(y_i y_j) = 0$ when $i \neq j$.

Even if an affine representation of the form (1.15) is not given in the formulation of the problem, one can be derived by taking any representation system $(\psi_j)_{j\geq 1}$ in the Banach space X. For example, if X admits a Schauder basis, then one can take any such basis $(\psi_j)_{j\geq 1}$ for X and arrive at such an expansion. In classical spaces X, such as L^p or Sobolev spaces, standard systems of approximation, such as Fourier series, splines, or wavelets can be used.

The advantage of the representation (1.15) is that a can now be identified through the sequence $(y_j)_{j\geq 1}$. When considering all $a \in \mathcal{A}$, we obtain a family of such sequences. Note that this family can be quite complicated. In order to simplify matters, we normalize the ψ_j, so that for any j,

$$\sup_{a\in\mathcal{A}} |y_j(a)| = 1. \tag{1.24}$$

Such a renormalization is usually possible because $\mathcal{A}$ is compact and $y_j(a)$ depends continuously on a. After this normalization, for each $a \in \mathcal{A}$, the sequence $(y_j(a))_{j\geq 1}$ belongs to the infinite-dimensional cube

$$U := [-1,1]^{\mathbb{N}}. \tag{1.25}$$

Note that taking a general sequence $(y_j)_{j\geq 1}$ from this cube, there may not be an $a \in \mathcal{A}$ with $y_j = y_j(a)$, $j \geq 1$. Also, if $\{\psi_j\}_{j\geq 1}$ is not a basis, the representation (1.15) may not be unique. We define

$$U_{\mathcal{A}} := \left\{(y_j)_{j\geq 1} \in U : \sum_{j\geq 1} y_j\psi_j \in \mathcal{A}\right\}. \tag{1.26}$$

We are mainly interested in representers $\bar{a}$, $(\psi_j)_{j\geq 1}$ for which

$$\bar{a} + \sum_{j\geq 1} y_j\psi_j \tag{1.27}$$

converges in X for each $(y_j) \in U$. We call such representers *complete*. In this case, we may define

$$a(U) := \left\{a = a(y) = \overline{a} + \sum_{j\geq 1} y_j\psi_j : (y_j)_{j\geq 1} \in U\right\}, \tag{1.28}$$

so that

$$\mathcal{A} \subset a(U). \tag{1.29}$$

A typical case of a complete representer is when $(\|\psi_j\|_X)$ is a sequence in $\ell^1(\mathbb{N})$.

Once an affine representation has been chosen, the initial solution map $a \mapsto u(a)$ becomes equivalent to the map $y \mapsto u(a(y))$ which is defined on $U_{\mathcal{A}}$. With an abuse of notation, we write this new solution map as

$$y \mapsto u(y) := u(a(y)). \tag{1.30}$$

This is a Banach-space-valued function of an infinite number of variables. Note that in the case where the affine representation has a finite number d of terms, the range of y is $[-1,1]^d$. However, the infinite-dimensional case subsumes the finite-dimensional case, since the latter may be viewed as a particular case with $\psi_j = 0$ for $j > d$.

In the case of a complete representer, $a(y)$ is defined on all of U. However, we do not know whether the solution map u is defined on all of U. To guarantee this, the following assumption will be used often.

Assumption A. The parameter set $\mathcal{A}$ has a complete representer $(\psi_j)_{j\geq 1}$ and the solution map $a \mapsto u(a)$ is well defined on the whole set $a(U)$, or equivalently the solution map $y \mapsto u(y)$ is well defined on the whole set U.

This assumption naturally holds when the set $\mathcal{A}$ is exactly defined as $a(U)$.

1.3. Smoothness of the solution map

One objective of this article is to develop efficient numerical approximations to the solution maps of (1.3) or (1.30). One of the main difficulties is that these maps are high- or infinite-dimensional, in the sense that the dimension of the variable a or y is high or infinite. In order to understand what might be good strategies for constructing such approximations, we need first to understand the inherent properties of these maps that might allow us to circumvent this difficulty.

We initiate such a program in Section 2, where we first analyse the smoothness of the solution map $a \mapsto u(a)$. In the case of the elliptic equation (1.5), it is easily seen that this map is not only infinitely differentiable, but also admits a holomorphic extension to certain subdomains of the complex-valued $X = L^\infty(D)$. We propose two general approaches which allow us to establish similar holomorphy properties for other relevant instances of linear and nonlinear parametric PDEs. The first approach is based on the Ladyzhenskaya–Babuška–Brezzi (LBB) theory. It applies to a range of linear PDEs where the operator and the right-hand side have holomorphic dependence in a. These include parabolic and saddle-point problems, such

as the heat equations or the Stokes problem, with parameter a in the diffusion term, similar to (1.5). The second approach is based on the implicit function theorem in complex-valued Banach spaces. In contrast to the first approach, it can be applied to certain nonlinear PDEs.

Using the affine representation (1.15) of a, we then study the solution map $y \mapsto u(y)$ under Assumption A, which means that it is defined on the whole of U. In addition to holomorphy, an important property of u can be extracted from the affine representation (1.15). The functions ψ_j appearing in (1.15) have norms $\|\psi_j\|_X$ of varying size. Since the variable y_j is scaled to be in $[-1,1]$, when $\|\psi_j\|_X$ is small, this variable has a reduced effect on the variations of $u(y)$. Thus the variables $(y_j)_{j\geq 1}$ are not democratic, but rather they have varying importance. In other words, the map $y \mapsto u(y)$ is highly anisotropic. More specifically, we derive holomorphic extension results for this map on certain multivariate complex domains of tensor product type. In particular, we consider polydiscs of the general form

$$\mathcal{U}_\rho := \otimes\{|z_j| \leq \rho_j\} = \{z = (z_j)_{j\geq 1} \in \mathbb{C}^{\mathbb{N}} : |z_j| \leq \rho_j\}, \tag{1.31}$$

where $\rho = (\rho_j)_{j\geq 1}$ is a positive sequence which serves to describe the anisotropy of the solution map. We also consider polyellipses, which deviate less far from the real axis. These holomorphy domains play a key role in the derivation of approximation results.

Remark 1.1. While we are generally interested in real-valued solutions u to the parametric PDE (1.1), corresponding to real-valued parameters a or y, our analysis of holomorphic smoothness leads us naturally to complex-valued solutions, corresponding to complex-valued parameters. For this reason, the spaces X, V, W are always assumed to be complex-valued Banach spaces throughout this paper.

1.4. Approximation of the solution map

Reduced modelling methods seek to take advantage of the properties of the solution maps $a \mapsto u(a)$ or $y \mapsto u(y)$ such as holomorphy and anisotropy mentioned above. These properties suggest strategies for approximating these maps u by simple functions u_n, in which the physical variables x and the parametric variables a or y are separated and hence take the form

$$(x,a) \mapsto u_n(x,a) := \sum_{i=1}^{n} v_i(x)\phi_i(a), \tag{1.32}$$

or

$$(x,y) \mapsto u_n(x,y) := \sum_{i=1}^{n} v_i(x)\phi_i(y), \tag{1.33}$$

where $\{v_1, \dots, v_n\}$ are functions of x living in the solution space V and $\{\phi_1, \dots, \phi_n\}$ are functions of a or y with values in $\mathbb{R}$ or $\mathbb{C}$.

We may view u_n as a rank n approximation to u, in analogy with low-rank approximation of matrices. We adopt the notations

$$a \mapsto u_n(a) = u_n(\cdot, a) \quad \text{and} \quad y \mapsto u_n(y) = u_n(\cdot, y), \tag{1.34}$$

for the above approximations.

Let us discuss the potential accuracy of separable approximations of the form (1.32). If our objective is to capture $u(a)$ for all $a \in \mathcal{A}$ with a prescribed accuracy $\varepsilon(n)$, this means that we search for an error bound in the uniform sense, that is, of the form

$$\|u - u_n\|_{L^\infty(\mathcal{A},V)} := \sup_{a\in\mathcal{A}} \|u(a) - u_n(a)\|_V \le \varepsilon(n). \tag{1.35}$$

For certain applications, in particular in the stochastic framework, we may instead decide to measure the error on average, for instance by searching for an error bound in the mean square sense,

$$\mathbb{E}(\|u - u_n\|_V^2) := \|u - u_n\|^2_{L^2(\mathcal{A},V,\mu)} = \int_{\mathcal{A}} \|u(a) - u_n(a)\|_V^2 \, \mathrm{d}\mu(a) \le \varepsilon(n)^2, \tag{1.36}$$

where μ is the probability measure for the distribution of a over $\mathcal{A}$. Since μ is a probability measure, we have for any v

$$\|v\|_{L^2(\mathcal{A},V,\mu)} \le \|v\|_{L^\infty(\mathcal{A},V)}. \tag{1.37}$$

Therefore the uniform bound is stronger than the average bound, in the sense that (1.35) implies (1.36) with the same value of $\varepsilon(n)$.

Likewise, for the approximation of the map $y \mapsto u(y)$, we may search for a uniform bound

$$\|u - u_n\|_{L^\infty(U_\mathcal{A},V)} := \sup_{y\in U_\mathcal{A}} \|u(y) - u_n(y)\|_V \le \varepsilon(n) \tag{1.38}$$

or a mean-square bound

$$\mathbb{E}(\|u - u_n\|_V^2) := \|u - u_n\|^2_{L^2(U_\mathcal{A},V,\mu)} = \int_{U_\mathcal{A}} \|u(y) - u_n(y)\|_V^2 \, \mathrm{d}\mu(y) \le \varepsilon(n)^2, \tag{1.39}$$

where μ is the probability measure for the distribution of y over $U_\mathcal{A}$.

Remark 1.2. We do not indicate the measure μ in our notation $L^\infty(\mathcal{A}, V)$ or $L^\infty(U_\mathcal{A}, V)$, since in all relevant examples considered in this article we always consider the exact supremum over a in $\mathcal{A}$ or over y in $U_\mathcal{A}$, rather than the essential supremum.

For any $a \in \mathcal{A}$, the approximation $u_n(a)$ belongs to

$$V_n := \operatorname{span}\{v_1, \dots, v_n\}, \tag{1.40}$$