

State Estimation—The Role of Reduced Models



Albert Cohen, Wolfgang Dahmen, and Ron DeVore

Abstract The exploration of complex physical or technological processes usually requires exploiting available information from different sources: (i) *physical laws* often represented as a family of *parameter dependent partial differential equations* and (ii) *data* provided by *measurement devices* or *sensors*. The amount of sensors is typically limited and data acquisition may be expensive and in some cases even harmful. This article reviews some recent developments for this “small-data” scenario where inversion is strongly aggravated by the typically large parametric dimensionality. The proposed concepts may be viewed as exploring alternatives to *Bayesian inversion* in favor of more deterministic accuracy quantification related to the required computational complexity. We discuss *optimality criteria* which delineate intrinsic information limits, and highlight the role of *reduced models* for developing efficient computational strategies. In particular, the need to *adapt* the reduced models—not to a specific (possibly noisy) data set but rather to the sensor system—is a central theme. This, in turn, is facilitated by exploiting *geometric* perspectives based on proper *stable variational formulations* of the continuous model.

1 Introduction

Modern sensor technology and data acquisition capabilities generate an ever increasing wealth of data about virtually every branch of science and social life. Machine learning offers novel techniques for extracting quantifiable information from such large data sets. While machine learning has already had a transformative impact on

A. Cohen

Laboratoire Jacques-Louis Lions, Sorbonne Université, 4, Place Jussieu, 75005 Paris, France
e-mail: cohen@ann.jussieu.fr

W. Dahmen (✉)

Mathematics Department, University of South Carolina, 1523 Greene Street, Columbia, SC 29208, USA
e-mail: dahmen@math.sc.edu

R. DeVore

Department of Mathematics, Texas A & M University, College Station, TX 77843-3368, USA

a diversity of application areas in the “big-data” regime, particularly in image classification and artificial intelligence, it is yet to have a similar impact in many other areas of science.

Utilizing data observations in the analysis of scientific processes differs from traditional learning in that one has the additional information that these processes are described by mathematical models—systems of partial differential equations (PDE) or integral equations—that encode the physical laws that govern the process. Such models, however, are often deficient, inaccurate, incomplete or need to be further calibrated by determining a large number of *parameters* in order to accurately represent an observed process. Typical guiding examples are Darcy’s equation for the pressure in ground-water flow or electron impedance tomography. Both are based on second order elliptic equations as core models. The diffusion coefficients in these examples describe permeability or conductivity, respectively. The parametric representations of the coefficients could arise, for instance, from Karhunen-Loève expansions of a random field that represent “unresolvable” features to be captured by the model. In this case the number of parameters could actually be *infinite*.

The use of machine learning to describe complex states of interest or even the underlying laws, solely through data, seems to bear little hope. In fact, data acquisition is often expensive or even harmful as in applications involving radiation. Thus, a severe undersampling poses principal obstructions to *state* or *parameter estimation* by solely processing observational data through standard machine learning techniques. It is therefore more natural to try to effectively combine the data information with the knowledge of the underlying physical laws represented by parameter dependent families of PDEs.

Methods that fuse together data-driven and model-based approaches fall roughly into two categories. One prototype of a *data assimilation* scenario arises in meteorology where data are used to stabilize otherwise chaotic dynamical systems, typically with the aid of (stochastic) filtering techniques. A second setting, in line with the above examples, uses an underlying *stable* continuous model to *regularize* otherwise ill-posed estimation tasks in a “small-data” scenario. *Bayesian inversion* is a prominent way of regularizing such problems. It relaxes the estimation task to asking only for *posterior probabilities* of states or parameters to explain given observations.

The present article reviews some recent developments on data driven state and parameter estimation that can be viewed as seeking alternatives to Bayesian inversion by placing a stronger focus on deterministic uncertainty quantification and its relation to *computational complexity*. The emphasis is on foundational aspects such as the optimality of algorithms (formulated in an appropriate sense) when treating estimation tasks for “small-data” problems in *high-dimensional parameter* regimes. Central issues concern the role of *reduced modeling* and the exploitation of intrinsic problem metrics provided by the *variational formulation* of the underlying continuous family of PDEs. This is used by the so called *Parametrized Background Data-Weak* (PBDW) framework, introduced in [20] and further analyzed in [4], to identify a suitable trial (Hilbert) space \mathbb{U} that accommodates the states and eventually also the data. An important point is to distinguish between the *data* and corresponding *sensors*—here linear functionals in the dual \mathbb{U}' of \mathbb{U} —from which the data are

generated. This will be seen to actually open a *geometric* perspective that sheds light on intrinsic estimation limits. Moreover, in the deterministic setting, a pivotal role is played by the so called *solution manifold*, which is the set of all states that can be attained when the parameters in the PDE traverse the whole parameter domain.

Even with full knowledge of a state in the solution manifold, to infer from it a corresponding parameter is a *nonlinear severely ill-posed* problem typically formulated as a *non-convex* optimization problem. On the other hand, state estimation from data is a *linear*, and hence a more benign inversion task mainly suffering under the current premises from a severe undersampling. We will, however, indicate how to reduce, under certain circumstances, the latter to the former problem so as to end up with a *convex* optimization problem. This motivates focusing in what follows mainly on state estimation. A central question then becomes how to best invoke knowledge on the solution manifold to regularize the estimation problem without introducing unnecessarily ambiguous bias. Our principal viewpoint is to recast state estimation as an *optimal recovery* problem which then naturally leads one to explore the role and potential of *reduced modeling*.

The layout of the paper is as follows. Section 2 describes the conceptual framework for *state estimation* as an *optimal recovery task*. This formulation allows the identification of lower bounds for the best achievable recovery accuracy.

Section 3 reviews recent developments concerning a certain *affine* recovery scheme and highlights the role of *reduced models* adapted to the recovery task. The overarching theme is to establish certified recovery bounds. When striving for optimality of such affine recovery maps, high parameter dimensionality is identified as a major challenge. We outline a recent remedy that avoids the *Curse of Dimensionality* by trading deterministic accuracy guarantees against analogs that hold with quantifiable high probability.

Even optimal affine reduced models can, in general, not be expected to realize the benchmarks identified in Sect. 2. To put the results in Sect. 3 in proper perspective, we comment in Sect. 4 on ongoing work that uses the results on affine reduced models and corresponding estimators as a central building block for nonlinear estimators. We also indicate briefly some ramifications on parameter estimation.

2 Models and Data

2.1 The Model

Technological design or simulating physical processes is often based on continuum models given by a family

$$\mathcal{R}(u, y) = 0, \quad y \in \mathcal{Y}, \quad (2.1)$$

of partial differential Equations (PDEs) that depend on parameters y ranging over a parameter domain $\mathcal{Y} \subset \mathbb{R}^{d_y}$. We will always assume *uniform well-posedness* of

(2.1): for each $y \in \mathcal{Y}$, there exists a unique solution $u = u(y)$ in some trial Hilbert space \mathbb{U} which satisfies $\mathcal{R}(u(y), y) = 0$.

Specifically, we consider only linear problems of the form $\mathcal{B}_y u = f$, that is,

$$\mathcal{R}(u, y) = f - \mathcal{B}_y u. \quad (2.2)$$

Here f belongs to the dual \mathbb{V}' of a suitable *test space* \mathbb{V} and \mathcal{B}_y is a linear operator acting from \mathbb{U} to \mathbb{V}' that depends on $y \in \mathcal{Y}$. Here, uniform well-posedness means then that \mathcal{B}_y is boundedly invertible with bounds independent of y . By the Babuška-Banach-Nečas Theorem, this is equivalent to saying that the bilinear form

$$(u, v) \mapsto b_y(u, v) := (\mathcal{B}_y u)(v) \quad (2.3)$$

satisfies the following *continuity* and *inf-sup conditions*

$$\sup_{u \in \mathbb{U}} \sup_{v \in \mathbb{V}} \frac{b_y(u, v)}{\|u\|_{\mathbb{U}} \|v\|_{\mathbb{V}}} \leq C_b \quad \text{and} \quad \inf_{u \in \mathbb{U}} \sup_{v \in \mathbb{V}} \frac{b_y(u, v)}{\|u\|_{\mathbb{U}} \|v\|_{\mathbb{V}}} \geq c_b > 0, \quad y \in \mathcal{Y}, \quad (2.4)$$

together with the property that $b_y(u, v) = 0$, $u \in \mathbb{U}$, implies $v = 0$ (injectivity of \mathcal{B}_y^*). The relevance of this stability notion lies in the entailed validity of the *error-residual relation*

$$C_b^{-1} \|f - \mathcal{B}_y v\|_{\mathbb{V}'} \leq \|u(y) - v\|_{\mathbb{U}} \leq c_b^{-1} \|f - \mathcal{B}_y v\|_{\mathbb{V}'}, \quad v \in \mathbb{U}, y \in \mathcal{Y}, \quad (2.5)$$

where $\|g\|_{\mathbb{V}'} := \sup\{g(v) : \|v\|_{\mathbb{V}} = 1\}$. Thus, errors in the trial norm are equivalent to residuals in the dual test norm which will be exploited in what follows.

For a wide range of problems such as space-time variational formulations, e.g. of parabolic or convection-diffusion problems, indefinite or singularly perturbed problems, the identification of a suitable pair \mathbb{U}, \mathbb{V} that guarantees stability in the above sense is not entirely straightforward. In particular, trial and test space may have to differ from each other, see e.g. [6, 11, 17, 23] for examples as well as some general principles.

The simplest example, used for illustration purposes, is the *elliptic* family

$$\mathcal{R}(u, y) = f + \operatorname{div}(a(y) \nabla u), \quad (2.6)$$

set in $\Omega \subset \mathbb{R}^{d_x}$ where $d_x \in \{1, 2, 3\}$, with boundary conditions $u|_{\partial\Omega} = 0$. Uniform well-posedness follows then for $\mathbb{U} = \mathbb{V} = H_0^1(\Omega)$ if we have for some fixed constants $0 < r \leq R < \infty$ the bounds

$$r \leq a(x, y) \leq R, \quad (x, y) \in \Omega \times \mathcal{Y}, \quad (2.7)$$

readily implying (2.4).

Aside from well-posedness, a second important structural property of the model (2.1) is *affine parameter dependence*. By this we mean that

$$\mathcal{B}_y u = \mathcal{B}_0 u + \sum_{j=1}^{d_y} y_j \mathcal{B}_j u, \quad y = (y_j)_{j=1, \dots, d_y} \in \mathcal{Y}, \quad (2.8)$$

where the operators $\mathcal{B}_j : \mathbb{U} \rightarrow \mathbb{V}'$ are *independent* of y . In turn, the residual has a similar affine dependence structure

$$\mathcal{R}(u, y) = \mathcal{R}_0(u) + \sum_{j=1}^{d_y} y_j \mathcal{R}_j u, \quad \mathcal{R}_0(u) := f - \mathcal{B}_0 u, \quad \mathcal{R}_j = -\mathcal{B}_j. \quad (2.9)$$

For the example (2.6) such a structure is encountered for *affine* parametric representations of the diffusion coefficients

$$a(x, y) = a_0(x) + \sum_{j=1}^{d_y} y_j \theta_j(x), \quad (x, y) \in \Omega \times \mathcal{Y}, \quad (2.10)$$

i.e., the field a is expanded in terms of some given spatial basis functions θ_j . As indicated earlier, the pressure equation in Darcy's law for porous media flow is an example for (2.6) where the diffusion coefficient $a(y)$ of the form (2.10) may arise from a stochastic model for permeability via a Karhunen-Loëve expansion. In this case (upon proper normalization) $y \in [-1, 1]^{\mathbb{N}}$ has, in principle, *infinitely* many entries, that is $d_y = \infty$. However, due to (2.7), the θ_j should then have some decay as $j \rightarrow \infty$ which means that the parameters become less and less important when j increases. Another example is electron impedance tomography involving the same type of elliptic operator where parametric expansions represent possible variations of conductivity often modeled as piecewise constants, i.e., the θ_j could be characteristic functions subordinate to a partition of Ω . In this case data are acquired through sensors that act through trace functionals greatly adding to ill-posedness.

A central role in the subsequent discussion is played by the solution manifold

$$\mathcal{M} = u(\mathcal{Y}) := \{u(y) : y \in \mathcal{Y}\} \quad (2.11)$$

which is then the range of the *parameter-to-solution map* $u : y \mapsto u(y)$ comprised of all states that can be attained when y traverses \mathcal{Y} . Without further mention, \mathcal{M} will be assumed to be compact which actually follows under standard assumptions met in all above mentioned examples.

Estimating states in \mathcal{M} or corresponding parameters from measurements requires the efficient approximation of elements in \mathcal{M} . A common challenge encountered in all such models lies in the inherent *high-dimensionality* of the states $u = u(\cdot, y)$ as functions of d_x spatial variables $x \in \Omega$ and $d_y \gg 1$ parametric variables $y \in \mathcal{Y}$. In particular, when $d_y = \infty$ any calculation, of course, has to work with finitely many "activated" parameters whose number, however, has to be coordinated with the spatial resolution of a numerical scheme to retain *model-consistency*. It is especially

this issue that hinders standard approaches based on *first discretizing* the parametric model because rigorously balancing spatial and parametric uncertainties becomes then difficult.

What renders such problem scenarios nevertheless numerically tractable is a further property that will be implicitly assumed in what follows, namely that the *Kolmogorov n-widths* of the solution manifold

$$d_n(\mathcal{M})_{\mathbb{U}} := \inf_{\dim \mathbb{U}_n = n} \sup_{u \in \mathcal{M}} \inf_{v \in \mathbb{U}_n} \|u - v\|_{\mathbb{U}} \quad (2.12)$$

exhibits at least some algebraic decay

$$d_n(\mathcal{M})_{\mathbb{U}} \lesssim n^{-s} \quad (2.13)$$

for some $s > 0$, see [13] for a comprehensive account.

For instance, this is known to be the case for elliptic models (2.6) with (2.7), as a consequence of the results of sparse polynomial approximation of the parameter to solution map $y \mapsto u(y)$ established e.g. in [15]. More generally, (2.13) can be established under a general holomorphy property of the parameter to solution map, as a consequence of a similar algebraic decay assumed on the n -widths of the parameter set, see [14]. For a fixed finite number $d_y < \infty$ of parameters, under certain structural assumptions on the parameter representations (e.g. piecewise constants on checkerboard partitions) one can even establish (sub-) exponential decay rates, see [2] for more details. Assuming s in (2.13) to have a “substantial” size for any range of d_y , is therefore justified.

In summary, the results discussed below are valid and practically feasible for well posed linear models (2.4) with affine parameter dependence (2.9) whose solution manifolds have rapidly decaying n -widths (2.13).

2.2 The Data

Suppose we are given data $\mathbf{w} = (w_1, \dots, w_m)^\top \in \mathbb{R}^m$ representing observations of an unknown state $u \in \mathbb{U}$ obtained through m linearly independent linear functionals $\ell_i \in \mathbb{U}'$, i.e.,

$$w_i = \ell_i(u), \quad i = 1, \dots, m. \quad (2.14)$$

Since in real applications data acquisition may be costly or harmful we assume that m is *fixed*. The central task to be discussed in what follows is to recover from this information an estimate for the observed unknown state u , based on the prior assumption that u belongs to \mathcal{M} or is close to \mathcal{M} . Moreover, to bring out the essence of this estimation task we assume for the moment that the data are noise-free.

Following [4, 20], we first recast the data in a “compliant” metric, by introducing the Riesz representers $\psi_i \in \mathbb{U}$, defined by

$$(\psi_i, v)_{\mathbb{U}} = \ell_i(v), \quad v \in \mathbb{U}, \quad i = 1, \dots, m,$$

The ψ_i now span the m -dimensional subspace $\mathbb{W} \subset \mathbb{U}$ which we refer to as *measurement space*, and the information carried by the $\ell_i(u)$ is equivalent to that of the orthogonal projection $P_{\mathbb{W}}u$ of u to \mathbb{W} . The decomposition

$$u = P_{\mathbb{W}}u + P_{\mathbb{W}^\perp}u, \quad u \in \mathbb{U}, \quad (2.15)$$

thus contains a first term that is “seen” by the sensors and a second (infinite-dimensional) term which cannot be detected. The decomposition (2.15) may be seen as a sensor-induced “coordinate system” thereby opening up a *geometric perspective* that will prove very useful in what follows. State estimation can then be viewed as learning from samples $w := P_{\mathbb{W}}u$ the unknown “labels” $P_{\mathbb{W}^\perp}u \in \mathbb{W}^\perp$.

In this article, we are interested in how well we can approximate u from the information that $u \in \mathcal{M}$ and $P_{\mathbb{W}}u = w$ with w given to us. Any such approximation is given by a mapping $A : w \rightarrow A(w) \in \mathbb{U}$. The overall performance of recovery on all of \mathcal{M} by the mapping A is typically measured in the worst case setting, that is,

$$E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}) = \sup_{u \in \mathcal{M}} \|u - A(P_{\mathbb{W}}u)\|_{\mathbb{U}}. \quad (2.16)$$

The optimal recovery error on \mathcal{M} is then defined as

$$E_{\text{wc}}(\mathcal{M}, \mathbb{W}) := \inf_A E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}), \quad (2.17)$$

where the infimum is over all possible recovery maps. Let us observe that the construction of recovery maps can be restricted to be of the form

$$A : w \rightarrow A(w), \quad A(w) = w + B(w), \quad \text{with } B : \mathbb{W} \rightarrow \mathbb{W}^\perp. \quad (2.18)$$

Indeed, given any recovery mapping A , we can write $A(w) = P_{\mathbb{W}}A(w) + P_{\mathbb{W}^\perp}A(w)$ and the performance of the recovery can only be improved if we replace the first term by w . In other words, $A(w)$ should belong to the affine space

$$\mathbb{U}_w := w + \mathbb{W}^\perp, \quad (2.19)$$

that contains u . The mappings B are commonly referred to as *liftings* into \mathbb{W}^\perp .

2.3 Optimality Criteria and Numerical Recovery

Finding a best recovery map A attaining (2.17) is known as *optimal recovery*. The best mapping has a well-known simple theoretical description, see e.g. [21], that we now describe. Note first that a precise recovery of the unknown state u from the given

information is generally impossible. Indeed, the best we can say about u is that it lies in the *manifold slice*

$$\mathcal{M}_w := \{u \in \mathcal{M} : P_{\mathbb{W}}u = w\} = \mathcal{M} \cap \mathbb{U}_w, \quad (2.20)$$

which is comprised of all elements in \mathcal{M} sharing the same measurement $w \in \mathbb{W}$. The Chebyshev ball $B(\mathcal{M}_w)$ is the smallest ball in \mathbb{U} that contains \mathcal{M}_w . The best recovery algorithm is then given by the mapping

$$A^*(w) := \text{cen}(\mathcal{M}_w), \quad (2.21)$$

that assigns to each $w \in \mathbb{W}$ the center $\text{cen}(\mathcal{M}_w)$ of $B(\mathcal{M}_w)$, called the Chebyshev center of \mathcal{M}_w . Then, the radius $\text{rad}(\mathcal{M}_w)$ of $B(\mathcal{M}_w)$ is the best worst case error over the class \mathcal{M}_w . The best worst case error over \mathcal{M} , which is achieved by A^* , is thus given by

$$E_{\text{wc}}(\mathcal{M}, W) = E_{\text{wc}}(A^*, \mathcal{M}, \mathbb{W}) = \max_{w \in \mathbb{W}} \text{rad}(\mathcal{M}_w). \quad (2.22)$$

While the above mapping A^* gives a nice theoretical description of the optimal recovery algorithm, it is typically not numerically implementable since the Chebyshev center $\text{cen}(\mathcal{M}_w)$ is not easily found. Moreover, such an optimal algorithm is highly nonlinear and possibly discontinuous. The purpose of this section is to formulate a more modest goal for the performance of a recovery algorithm with the hope that this more modest goal can be met with a numerically realizable algorithm. The remaining sections of the paper introduce numerically implementable recovery mappings, analyze their performance, and evaluate the numerical cost in constructing these mappings.

The search for a numerically realizable algorithm must out of necessity lessen the performance criteria. A first possibility is to weaken the performance criteria to *near best* algorithms. This means that we search for an algorithm A such that

$$E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}) \leq C_0 E_{\text{wc}}(\mathcal{M}, \mathbb{W}), \quad (2.23)$$

with a reasonable value of $C_0 > 1$. For example, any mapping A which takes w into an element in the Chebyshev ball of \mathcal{M}_w is near best with constant $C_0 = 2$. However, finding near best mappings A also seems to be numerically out of reach.

In order to formulate a more attainable performance criterion, we return to our earlier observations about uncertainty in both the model class \mathcal{M} and in the measurements w . The former is a modeling error while the latter is an inherent measurement error. Both of these uncertainties can be quantified by introducing for each $\varepsilon > 0$, the ε -neighborhood of the manifold

$$\mathcal{M}^\varepsilon := \{v \in \mathbb{U} : \text{dist}(v, \mathcal{M})_{\mathbb{U}} \leq \varepsilon\}. \quad (2.24)$$

The uncertainty in the model can be thought of as saying the sought after u is in \mathcal{M}^ε rather than $u \in \mathcal{M}$. Also, we may formulate uncertainty (noise) in the measurements as saying that they are not measurements of a $u \in \mathcal{M}$ but rather some $u \in \mathcal{M}^\varepsilon$. Here the value of ε quantifies these uncertainties.

Our new goal is to numerically construct a recovery map A that is near-optimal on \mathcal{M}^ε , for some given $\varepsilon > 0$. Let us note that \mathcal{M}^ε is not compact. An algorithm A is worst-case near optimal for \mathcal{M}^ε if and only if its performance is bounded by a constant multiple of the diameter

$$\delta_\varepsilon(\mathcal{M}, \mathbb{W}) := \max \{ \|u - v\|_{\mathbb{U}} : u, v \in \mathcal{M}^\varepsilon, P_{\mathbb{W}}(u - v) = 0 \}. \quad (2.25)$$

Notice that $\varepsilon = 0$ gives the performance criterion for near optimal recovery over \mathcal{M} . One can show that the function $\varepsilon \mapsto \delta_\varepsilon(\mathcal{M}, \mathbb{W})$ is monotone non-decreasing in ε , continuous from the right, and $\lim_{\varepsilon \rightarrow 0^+} \delta_\varepsilon(\mathcal{M}, \mathbb{W}) = \delta_0(\mathcal{M}, \mathbb{W})$. The speed at which $\delta_\varepsilon(\mathcal{M}, \mathbb{W})$ approaches $\delta_0(\mathcal{M}, \mathbb{W})$ reflects the “condition” of the estimation problem depending on \mathcal{M} and \mathbb{W} . While the practical realization of worst-case near-optimality for \mathcal{M}^ε is already a challenge, quantifying corresponding computational cost would require assumptions on the condition of the problem.

One central theme, guiding subsequent discussions, is therefore to find recovery maps A_ε that realize an error bound of the form

$$E_{wc}(A_\varepsilon, \mathcal{M}, \mathbb{W}) \leq C_0 \delta_\varepsilon(\mathcal{M}, \mathbb{W}). \quad (2.26)$$

Any a priori information on measurement accuracy and model bias might be used to choose a viable tolerance ε .

High parametric dimensionality poses particular challenges to estimation tasks when the targeted error bounds are in the above worst case sense. These challenges can be somewhat mitigated when adopting a Bayesian point of view [24]. The prior information on u is then described by a probability distribution p on \mathbb{U} , which is supported on \mathcal{M} . Such a measure is typically induced by a probability distribution on \mathcal{Y} that may or may not be known. In the latter case, sampling \mathcal{M} , i.e., computing snapshots $u(y^i)$, $i = 1, \dots, N$, for i.i.d. samples $y^i \in \mathcal{Y}$, provides labeled data $(w_i, w_i^\perp) = (P_{\mathbb{W}}u(y^i), P_{\mathbb{W}^\perp}u(y^i))$ according to the sensor-based decomposition (2.15). This puts us into the setting of *regression* in machine learning asking for an estimator that predicts for any new measurement $w \in \mathbb{W}$ its lifting $w^\perp = B(w)$. It is then natural to measure the performance of an algorithm in an averaged sense. The best estimator A that minimizes the mean-square risk

$$E_{ms}(A, p, \mathbb{W}) = \mathbb{E}(\|u - A(P_{\mathbb{W}}u)\|^2) = \int_{\mathbb{U}} \|u - A(P_{\mathbb{W}}u)\|^2 dp(u) \quad (2.27)$$

is given by the conditional expectation

$$A(w) = \mathbb{E}(u | P_{\mathbb{W}}u = w). \quad (2.28)$$

Since always $E_{\text{ms}}(A, p, \mathbb{W}) \leq E_{\text{wc}}(A, \mathcal{M}, \mathbb{W})$, the optimality benchmarks are somewhat weaker. In the rest of this paper, we adhere to the worst case error in the deterministic setting that only assumes membership of u to \mathcal{M} or \mathcal{M}^ε .

The following section is concerned with an important *building block* on a pathway towards achieving (2.26) at quantifiable computational cost. This building block, referred to as *one-space method* is a linear (affine) scheme which is, in principle, simple and easy to numerically implement. It depends on suitably chosen subspaces. We highlight the *regularizing property* of these subspaces as well as ways to *optimize* them. This will reveal certain intrinsic obstructions caused by *parameter dimensionality*. The one-space method by itself will generally not achieve (2.26) but, as indicated earlier, can be used as a building block in a *nonlinear* recovery scheme that may indeed meet the goal (2.26).

3 The One-Space Method

3.1 Subspace Regularization

The one space method can be viewed as a simple regularizer for state estimation. The resulting recovery map is induced by an n -dimensional subspace \mathbb{U}_n of \mathbb{U} for $n \leq m$. Assume that, for each $n \geq 0$, we are given a subspace $\mathbb{U}_n \subset \mathbb{U}$ of dimension n whose distance from \mathcal{M} can be assessed

$$\text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}} := \max_{u \in \mathcal{M}} \text{dist}(u, \mathbb{U}_n)_{\mathbb{U}} \leq \varepsilon_n. \quad (3.1)$$

Then the cylinder

$$\mathcal{K}(\mathbb{U}_n, \varepsilon_n) := \{u \in \mathbb{U} : \text{dist}(u, \mathbb{U}_n)_{\mathbb{U}} \leq \varepsilon_n\} \quad (3.2)$$

contains \mathcal{M} and likewise the cylinder $\mathcal{K}(\mathbb{U}_n, \varepsilon_n + \varepsilon)$ contains \mathcal{M}^ε . Our prior assumption that the observed state belongs to \mathcal{M} or \mathcal{M}^ε can then be relaxed by assuming membership to these larger but simpler sets.

Remarkably, one can now realize an optimal recovery map quite easily that meets the relaxed benchmark $E_{\text{wc}}(\mathcal{K}(\mathbb{U}_n, \varepsilon_n), \mathbb{W})$: in [4] it was shown that the Chebyshev center of the slice

$$\mathcal{K}_w(\mathbb{U}_n, \varepsilon_n) := \mathcal{K}(\mathbb{U}_n, \varepsilon_n) \cap \mathbb{U}_w, \quad (3.3)$$

is exactly given by the state in \mathbb{U}_w that is closest to \mathbb{U}_n , that is

$$u^* = u^*(w) := \underset{u \in \mathbb{U}_w}{\text{argmin}} \|u - P_{\mathbb{U}_n} u\|_{\mathbb{U}}. \quad (3.4)$$

This minimizer exists and can be shown to be unique as long as $\mathbb{U}_n \cap \mathbb{W}^\perp = \{0\}$. The corresponding optimal recovery map

$$A_{\mathbb{U}_n} : w \mapsto u^*(w) \quad (3.5)$$

was first introduced in [20] as the Parametrized Background Data Weak (PBDW) algorithm, and is referred to as the *one-space* method in [4]. Due to its above minimizing property, it is readily checked that this map is linear and can be determined with the aid of the singular value decomposition of the cross-Gramian between any pair of orthonormal basis for \mathbb{U}_n and \mathbb{W} .

The worst case error $E_{\text{wc}}(\mathcal{K}(\mathbb{U}_n, \varepsilon_n), \mathbb{W})$ can be described more precisely by introducing

$$\mu(\mathbb{U}_n, \mathbb{W}) := \sup_{v \in \mathbb{U}_n} \frac{\|v\|_{\mathbb{U}}}{\|P_{\mathbb{W}}v\|_{\mathbb{U}}} \quad (3.6)$$

which is finite if and only if $\mathbb{U}_n \cap \mathbb{W}^\perp = \{0\}$. This quantity, also introduced in a related but slightly different context in [1], is therefore related to the angle between the spaces \mathbb{U}_n and \mathbb{W} . It becomes large when \mathbb{U}_n contains elements that are nearly perpendicular to \mathbb{W} . It is actually computable: one has $\mu(\mathbb{U}_n, \mathbb{W}) = \beta(\mathbb{U}_n, \mathbb{W})^{-1}$ where

$$\beta(\mathbb{U}_n, \mathbb{W}) := \inf_{v \in \mathbb{U}_n} \sup_{w \in \mathbb{W}} \frac{\langle v, w \rangle_{\mathbb{U}}}{\|v\|_{\mathbb{U}} \|w\|_{\mathbb{U}}}, \quad (3.7)$$

and $\beta(\mathbb{U}_n, \mathbb{W})$ is the smallest singular value of the cross-Gramian between any pair of orthonormal bases for \mathbb{W} and \mathbb{U}_n . It has been shown in [4, 20] that the worst case error bound over $\mathcal{K}(\mathbb{U}_n, \varepsilon_n)$ is given by

$$E_{\text{wc}}(A_{\mathbb{U}_n}, \mathcal{K}(\mathbb{U}_n, \varepsilon_n), \mathbb{W}) = E_{\text{wc}}(\mathcal{K}(\mathbb{U}_n, \varepsilon_n), \mathbb{W}) = \mu(\mathbb{U}_n, \mathbb{W})\varepsilon_n. \quad (3.8)$$

The quantity $\mu(\mathbb{U}_n, \mathbb{W})$ also coincides with the norm of the linear recovery map $A_{\mathbb{U}_n}$. Relaxing the prior $u \in \mathcal{M}$ by exploiting information on \mathcal{M} solely through approximability of \mathcal{M} by \mathbb{U}_n , thus implicitly *regularizes* the estimation task: whenever $\mu(\mathbb{U}_n, \mathbb{W})$ is finite, the optimal recovery map $A_{\mathbb{U}_n}$ is bounded and hence Lipschitz.

One important observation is that the map $A_{\mathbb{U}_n}$ is actually independent of ε_n . In particular, it achieves optimality for the smallest possible containment cylinder

$$\mathcal{K}(\mathbb{U}_n) := \mathcal{K}(\mathbb{U}_n, \text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}}), \quad (3.9)$$

and therefore, since $E_{\text{wc}}(A_{\mathbb{U}_n}, \mathcal{M}, \mathbb{W}) \leq E_{\text{wc}}(A_{\mathbb{U}_n}, \mathcal{K}(\mathbb{U}_n), \mathbb{W}) = E_{\text{wc}}(\mathcal{K}(\mathbb{U}_n), \mathbb{W})$,

$$E_{\text{wc}}(A_{\mathbb{U}_n}, \mathcal{M}, \mathbb{W}) \leq \mu(\mathbb{U}_n, \mathbb{W})\text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}}. \quad (3.10)$$

Likewise, the containment $\mathcal{M}^\varepsilon \subset \mathcal{K}(\mathbb{U}_n, \text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}} + \varepsilon)$ implies that

$$E_{\text{wc}}(A_{\mathbb{U}_n}, \mathcal{M}^\varepsilon, \mathbb{W}) \leq \mu(\mathbb{U}_n, \mathbb{W})(\text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}} + \varepsilon). \quad (3.11)$$

On the other hand, the recovery map $A_{\mathbb{U}_n}$ may be far from optimal over the sets \mathcal{M} or \mathcal{M}^e . This is due to the fact that the cylinders $\mathcal{K}(\mathbb{U}_n, \varepsilon_n)$ and $\mathcal{K}(\mathbb{U}_n, \varepsilon_n + \varepsilon)$ may be much larger than \mathcal{M} or \mathcal{M}^e . In particular, it is quite possible that for a particular observation w , one has $\text{rad}(\mathcal{M}_w) \ll \text{rad}(\mathcal{K}_w(\mathbb{U}_n, \varepsilon_n))$. Therefore, we cannot generally expect that the one space method achieves our goal (2.26). In particular, the condition $n \leq m$, which is necessary to avoid that $\mu(\mathbb{U}_n, \mathbb{W}) = \infty$, limits the dimension of an approximating subspace \mathbb{U}_n and therefore ε_n itself is inherently bounded from below. The “dimension budget” has therefore to be used wisely in order to obtain good performance bounds. This typically rules out “generic approximation spaces” such as finite element spaces, and raises the question which subspace \mathbb{U}_n yields the best estimator when applying the above method.

3.2 Optimal Affine Recovery

The results of the previous section bring forward the question as to what is the best choice of the space \mathbb{U}_n for the given \mathcal{M} . On the one hand, proximity to \mathcal{M} is desirable since $\text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}}$ enters the error bound. However, favoring proximity, may increase $\mu(\mathbb{U}_n, \mathbb{W})$. Before addressing this question systematically, it is important to note that the above results carry over verbatim when \mathbb{U}_n is replaced by an *affine space* $\mathbb{U}_n = \bar{u} + \tilde{\mathbb{U}}_n$ where $\tilde{\mathbb{U}}_n \subset \mathbb{U}$ is a linear space. This means the reduced model $\mathcal{K}(\mathbb{U}_n, \varepsilon_n)$ is of the form

$$\mathcal{K}(\mathbb{U}_n, \varepsilon_n) := \bar{u} + \mathcal{K}(\tilde{\mathbb{U}}_n, \varepsilon_n).$$

The best worst-case recovery bound is now given by

$$E_{\text{wc}}(\mathcal{K}(\mathbb{U}_n, \varepsilon_n), \mathbb{W}) = \mu(\tilde{\mathbb{U}}_n, \mathbb{W})\varepsilon_n. \quad (3.12)$$

Intuitively, this may help to better control the angle between \mathbb{W} and \mathbb{U}_n by anchoring the affine space at a suitable location (typically near or on \mathcal{M}). More importantly, it helps in *localizing* models via parameter domain decompositions that will be discussed later.

The one-space algorithm discussed in the previous section confines the “dimensionality” budget of the approximation spaces \mathbb{U}_n to $n \leq m$. In view of (3.10), to obtain an overall good estimation accuracy, this space can clearly not be chosen arbitrarily but should be well adapted both to the solution manifold \mathcal{M} and to measurement space W , that is, to the given observation functionals giving rise to the data.

A simple way of *adapting* a recovery space to \mathbb{W} is as follows: suppose for a moment that we were able to construct for $n = 1, \dots, m$, a hierarchy of spaces $\mathbb{U}_1^{\text{nb}} \subset \mathbb{U}_2^{\text{nb}} \subset \dots \subset \mathbb{U}_m^{\text{nb}}$, that approximate \mathcal{M} in a *near-best* way, namely

$$\text{dist}(\mathcal{M}, \mathbb{U}_n^{\text{nb}})_{\mathbb{U}} \leq Cd_n(\mathcal{M})_{\mathbb{U}}. \quad (3.13)$$

We may compute along the way the quantities $\mu(\mathbb{U}_j^{\text{nb}}, \mathbb{W})$, then choose

$$n^* = \operatorname{argmin}_{n \leq m} \mu(\mathbb{U}_n^{\text{nb}}, \mathbb{W}) \operatorname{dist}(\mathcal{M}, \mathbb{U}_n^{\text{nb}})_{\mathbb{U}}, \quad (3.14)$$

and take the map $A_{\mathbb{U}_{n^*}^{\text{nb}}}$. We sometimes refer to this choice as “*poor man’s algorithm*”. It is not clear though whether $\mathbb{U}_{n^*}^{\text{nb}}$ is indeed a near-best choice for state recovery by the one-space method. In other words, one may question whether

$$E_{\text{wc}}(A_{\mathbb{U}_{n^*}^{\text{nb}}}, \mathcal{M}, \mathbb{W}) \leq C \inf_{\dim \tilde{\mathbb{U}} \leq m} E_{\text{wc}}(A_{\tilde{\mathbb{U}}}, \mathcal{M}, \mathbb{W}), \quad (3.15)$$

holds with a uniform constant $C < \infty$. In fact, numerical tests strongly suggest otherwise, which motivated in [12] the following alternative to the poor man’s algorithm.

Recall that a given linear space \mathbb{U}_n determines the linear recovery map $A_{\mathbb{U}_n}$. Likewise a given affine space \mathbb{U}_n determines an affine recovery map $A_{\mathbb{U}_n}$. Conversely, it can be checked that an affine recovery map A determines an affine space \mathbb{U}_n that allows one to interpret the recovery scheme as a one-space method in the sense that $A = A_{\mathbb{U}_n}$. Denoting by \mathcal{A} the class of all affine mappings of the form

$$A(w) = w + z + Bw, \quad (3.16)$$

where $z \in \mathbb{W}^{\perp}$ and $B \in \mathcal{L}(\mathbb{W}, \mathbb{W}^{\perp})$ is linear, we might thus as well directly look for a mapping that minimizes

$$E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}) := \sup_{u \in \mathcal{M}} \|u - A(P_{\mathbb{W}}u)\|_{\mathbb{U}} = \sup_{u \in \mathcal{M}} \|P_{\mathbb{W}^{\perp}}u - z - BP_{\mathbb{W}}u\|_{\mathbb{U}} =: \mathcal{E}(z, B) \quad (3.17)$$

over \mathcal{A} , i.e., over all $(z, B) \in \mathbb{W}^{\perp} \times \mathcal{L}(\mathbb{W}, \mathbb{W}^{\perp})$. It can be shown that indeed a minimizing pair (z^*, B^*) exists, i.e.,

$$\mathcal{E}(z^*, B^*) = \min_{A \in \mathcal{A}} E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}) =: E_{\text{wc}, \mathcal{A}}(\mathcal{M}, \mathbb{W}),$$

see [12]. However, the minimization of $E_{\text{wc}}(A, \mathcal{M}, \mathbb{W})$ over $(z, B) \in \mathbb{W}^{\perp} \times \mathcal{L}(\mathbb{W}, \mathbb{W}^{\perp})$ is far from practically feasible. In fact, each evaluation of $E_{\text{wc}}(A, \mathcal{M}, \mathbb{W})$ requires exploring \mathcal{M} and B can have a range in the infinite dimensional space \mathbb{W}^{\perp} . In order to arrive at a computationally tractable problem, one needs to

- (i) Replace \mathcal{M} by a finite set $\tilde{\mathcal{M}} \subset \mathcal{M}$, that should be sufficiently dense. Denseness can be quantified by requiring that $\tilde{\mathcal{M}} = \tilde{\mathcal{M}}^{\delta}$ is a δ -net for \mathcal{M} for some $\delta > 0$, i.e., for any $u \in \mathcal{M}$, there exists $\tilde{u} \in \tilde{\mathcal{M}}^{\delta}$ such that $\|u - \tilde{u}\|_{\mathbb{U}} \leq \delta$.
- (ii) Choose a finite dimensional space $\mathbb{U}_L \subset \mathbb{U}$ that approximates \mathcal{M} to a desired precision $\operatorname{dist}(\mathcal{M}, \mathbb{U}_L)_{\mathbb{U}} \leq \eta$, and replace \mathbb{W}^{\perp} by the finite dimensional complement

$$\tilde{\mathbb{W}}^\perp := \mathbb{U}_L \ominus \mathbb{W} \quad (3.18)$$

of \mathbb{W} in \mathbb{U}_L .

The resulting optimization problem

$$(\tilde{z}, \tilde{B}) = \operatorname{argmin}_{(z, B) \in \tilde{\mathbb{W}}^\perp \times \mathcal{L}(\mathbb{W}, \tilde{\mathbb{W}}^\perp)} \sup_{u \in \tilde{\mathcal{M}}^\delta} \|P_{\mathbb{W}^\perp} u - z - B P_{\mathbb{W}} u\|_{\mathbb{U}}. \quad (3.19)$$

can be solved by primal-dual splitting methods providing a $O(1/k)$ convergence rate, [12].

Due to the perturbations (i) and (ii) of the ideal minimization problem, the resulting (\tilde{z}, \tilde{B}) is no longer optimal. However, one can show that

$$E_{\text{wc}}(\tilde{A}, \mathcal{M}, \mathbb{W}) \leq E_{\text{wc}, \mathcal{A}}(\mathcal{M}, \mathbb{W}) + \eta + C\delta, \quad (3.20)$$

where the constant C is the operator norm of B minimizing (3.17). On the other hand, since the range of any affine mapping A is an affine space of dimension at most m , therefore contained in a linear space of dimension at most $m+1$, one always has $E_{\text{wc}, \mathcal{A}}(\mathcal{M}, \mathbb{W}) \geq d_{m+1}(\mathcal{M})_{\mathbb{U}}$. Therefore (\tilde{z}, \tilde{B}) satisfies a near-optimal bound

$$E_{\text{wc}}(\tilde{A}, \mathcal{M}, \mathbb{W}) \lesssim E_{\text{wc}, \mathcal{A}}(\mathcal{M}, \mathbb{W}), \quad (3.21)$$

whenever η and δ are picked such that

$$\eta \lesssim d_{m+1}(\mathcal{M})_{\mathbb{U}}, \quad \text{and} \quad \delta \lesssim d_{m+1}(\mathcal{M})_{\mathbb{U}}. \quad (3.22)$$

The numerical tests in [12] for a model problem of the type (2.6) with piecewise constant checkerboard diffusion coefficients and d_y up to $d_y = 64$ show that this recovery map exhibits significantly better accuracy than the method based on (3.14). It even yields smaller error bounds than the affine mean square estimator (2.27). The following section discusses the numerical cost entailed by conditions like (3.22).

3.3 Rate-Optimal Reduced Bases

To keep the dimension L of the space \mathbb{U}_L in (3.18) small, a near-best subspace \mathbb{U}_L^{nb} in the sense of (3.13) would be highly desirable. Likewise the poor man's scheme (3.14) would benefit from such subspaces. Unfortunately, such near-best subspaces are not practically accessible. The *reduced basis method* aims to construct subspaces which come close to near-optimality in a sense that we further explain next. The main idea is to generate theses subspaces by a sequence of elements picked from the manifold \mathcal{M} itself, by means of a *weak-greedy algorithm* introduced and studied in [8]. In an idealized form, this algorithm proceeds as follows: given a current

space $\mathbb{U}_n^{\text{wg}} = \text{span}\{u_1, \dots, u_n\}$, one takes $u_{n+1} = u(y_{n+1})$ such that, for some fixed $\gamma \in]0, 1]$, $\|u_{n+1} - P_{\mathbb{U}_n} u_{n+1}\|_{\mathbb{U}} \geq \gamma \max_{u \in \mathcal{M}} \|u - P_{\mathbb{U}_n} u\|_{\mathbb{U}}$, or equivalently

$$\|u(y_{n+1}) - P_{\mathbb{U}_n} u(y_{n+1})\|_{\mathbb{U}} \geq \gamma \max_{y \in \mathcal{Y}} \|u(y) - P_{\mathbb{U}_n} u(y)\|_{\mathbb{U}}. \quad (3.23)$$

Then, one defines $\mathbb{U}_{n+1}^{\text{wg}} = \text{span}\{u_1, \dots, u_{n+1}\}$. While unfortunately, the weak greedy algorithm does in general not produce spaces satisfying (3.13), it does come close. Namely, it has been shown in [3, 19] that the spaces \mathbb{U}_n^{wg} are *rate-optimal* in the following sense:

(i) For any $s > 0$ one has

$$d_n(\mathcal{M})_{\mathbb{U}} \leq C(n+1)^{-s}, \quad n \geq 0 \implies \text{dist}(\mathcal{M}, \mathbb{U}_n^{\text{wg}})_{\mathbb{U}} \leq \tilde{C}(n+1)^{-s}, \quad n \geq 0, \quad (3.24)$$

where \tilde{C} depends on C, s, γ .

(ii) For any $\beta > 0$, one has

$$d_n(\mathcal{M})_{\mathbb{U}} \leq C e^{-cn^\beta}, \quad n \geq 0 \implies \text{dist}(\mathcal{M}, \mathbb{U}_n^{\text{wg}})_{\mathbb{U}} \leq \tilde{C} e^{-\tilde{c}n^\beta}, \quad n \geq 0, \quad (3.25)$$

where the constants \tilde{c}, \tilde{C} depend on c, C, β, γ .

In the form described above, the weak-greedy concept seems infeasible since it would, in principle, require computing the solution $u(y)$ for all values of $y \in \mathcal{Y}$ exactly, exploring the whole exact solution manifold. However, its practical applicability is facilitated when there exists a *tight* surrogate $R(y, \mathbb{U}_n)$, satisfying

$$c_R R(y, \mathbb{U}_n) \leq \|u(y) - P_{\mathbb{U}_n} u(y)\|_{\mathbb{U}} = \text{dist}(u(y), \mathbb{U}_n) \leq C_R R(y, \mathbb{U}_n), \quad y \in \mathcal{Y}, \quad (3.26)$$

for uniform constants $0 < c_R \leq C_R < \infty$, which can be evaluated at affordable cost. Then, maximization of $R(y, \mathbb{U}_n)$ over \mathcal{Y} amounts to the weak-greedy step (3.23) with $\gamma := \frac{c_R}{C_R}$. According to [18], the validity of the following two conditions indeed allows one to derive computable surrogates that satisfy (3.26):

- (i) The underlying parametric family of PDEs (2.1) permits a uniformly *stable variational formulation* (2.4), and one has *affine parameter dependence* (2.9);
- (ii) The discrete projection $\Pi_{\mathbb{U}_n}$ (of Galerkin or Petrov-Galerkin type) has the *best approximation property*, i.e., resulting errors are uniformly comparable to the best approximation error.

Conditions (i) and (ii) ensure, in view of (2.5), that $\|u(y) - P_{\mathbb{U}_n} u(y)\|_{\mathbb{U}} \sim \|\mathcal{R}(y, \Pi_{\mathbb{U}_n} u(y))\|_{\mathbb{V}}$ holds uniformly in $y \in \mathcal{Y}$. Thus,

$$R(y, \mathbb{U}_n) := \|\mathcal{R}(y, \Pi_{\mathbb{U}_n} u(y))\|_{\mathbb{V}} = \sup_{v \in \mathbb{V}} \frac{\mathcal{R}(y, \Pi_{\mathbb{U}_n} u(y))(v)}{\|v\|_{\mathbb{V}}} \quad (3.27)$$

satisfies (3.26) and is therefore a tight surrogate for $\text{dist}(\mathcal{M}, \mathbb{U}_n)_{\mathbb{U}}$. In the elliptic case (2.6) under assumption (2.7), (i) and (ii) hold and the above comments reflect standard

practice. For the wider scope of stable but *unsymmetric* variational formulations [6, 16, 23] the inf-sup conditions (2.4) imply (i), but the Galerkin projection in (ii) needs to be replaced by a stable *Petrov-Galerkin* projection with respect to suitable test spaces \mathbb{V}_n accompanying the reduced trial spaces \mathbb{U}_n . It has been shown in [18] how to generate such test spaces with the aid of a *double-greedy* strategy, see also [16].

The main pay-off of using the surrogate $R(y, \mathbb{U}_n)$ is that one no longer needs to compute $u(y)$ but only the low-dimensional projection $\Pi_{\mathbb{U}_n} u(y)$ by solving for each y an $n \times n$ system, which itself can be rapidly assembled thanks to the affine parameter dependence [22]. However, one still faces the problem of its exact maximization over $y \in \mathcal{Y}$. A standard approach is to maximize instead over a *discrete* training set $\tilde{\mathcal{Y}}_n \subset \mathcal{Y}$, which in turn induces a discretization of the solution manifold

$$\tilde{\mathcal{M}}_n = \{u(y) : y \in \tilde{\mathcal{Y}}_n\}. \quad (3.28)$$

The resulting weak-greedy algorithm can be shown to remain rate optimal in the sense of (3.24) and (3.25) if the discretization is fine enough so that $\tilde{\mathcal{M}}_n$ constitutes an ε_n -approximation net of \mathcal{M} where ε_n does not exceed $\text{cdist}(\mathcal{M}, \mathbb{U}_n^{\text{wg}})_{\mathbb{U}}$ for a suitable constant $0 < c < 1$. In the current regime of *large or even infinite parameter dimensionality*, this becomes prohibitive because $\#\tilde{\mathcal{Y}}_n$ would then typically scale like $O(\varepsilon_n^{-cd_y})$, [10].

As a remedy it has been proposed in [10] to use training sets $\tilde{\mathcal{Y}}_n$ that are generated by *randomly sampling* \mathcal{Y} , and ask that the objective of rate optimality is met with high probability. This turns out to be achievable with training sets of much less prohibitive size. In an informal and simplified manner the main result can be stated as follows.

Theorem 1 *Given any target accuracy $\varepsilon > 0$ and some $0 < \eta < 1$, then the weak greedy reduced basis algorithm based on choosing at each step $N = N(\varepsilon, \eta) \sim |\ln \eta| + |\ln \varepsilon|$ randomly chosen training points in \mathcal{Y} has the following properties with probability at least $1 - \eta$: it terminates with $\text{dist}(\mathcal{M}, \mathbb{U}_{n(\varepsilon)})_{\mathbb{U}} \leq \varepsilon$ as soon as the maximum of the surrogate over the current training set falls below $c\varepsilon^{1+a}$ for some $c, a > 0$. Moreover, if $d_n(\mathcal{M})_{\mathbb{U}} \leq Cn^{-s}$, then $n(\varepsilon) \lesssim \varepsilon^{-\frac{1}{s}-b}$. The constants c, a, b depend on the constants in (3.26), as well as on the rate r of polynomial approximability of the parameter to solution map $y \mapsto u(y)$. The larger s and r , the smaller a and b , and the closer the performance becomes to the ideal one.*

4 Nonlinear Models

4.1 Piecewise Affine Reduced Models

As already noted, schemes based on linear or affine reduced models of the form $\mathcal{K}(\mathbb{U}_n, \varepsilon)$ can, in general, not be expected to realize the benchmark (2.26), discussed earlier in Sect. 2. The convexity of the containment set $\mathcal{K}(\mathbb{U}_n, \varepsilon)$ may cause the

reconstruction error to be significantly larger than $\delta_\varepsilon(\mathcal{M}, \mathbb{W})$. Another way of understanding this limitation is that in order to make ε small, one is enforced to raise the dimension n of \mathbb{U}_n , making the quantity $\mu(\mathbb{U}_n, \mathbb{W})$ larger and eventually infinite if $n > m$.

To overcome this principal limitation one needs to resort to *nonlinear* models that better capture the non-convex geometry of \mathcal{M} . One natural approach consists in replacing the single space \mathbb{U}_n by a family $(\mathbb{U}^k)_{k=1,\dots,K}$ of affine spaces

$$\mathbb{U}^k = \bar{u}_k + \tilde{\mathbb{U}}^k, \quad \dim(\tilde{\mathbb{U}}^k) = n_k \leq m, \quad (4.1)$$

each of which aims to approximate a *portion* \mathcal{M}_k of \mathcal{M} to a prescribed target accuracy simultaneously controlling $\mu(\mathbb{U}^k, \mathbb{W})$: fixing $\varepsilon > 0$, we assume that we have at hand a partition of \mathcal{M} into portions

$$\mathcal{M} = \bigcup_{k=1}^K \mathcal{M}_k \quad (4.2)$$

such that

$$\text{dist}(\mathcal{M}_k, \mathbb{U}^k)_{\mathbb{U}} \leq \varepsilon_k, \quad \text{and} \quad \mu(\tilde{\mathbb{U}}^k, \mathbb{W}) \varepsilon_k \leq \varepsilon, \quad k = 1, \dots, K. \quad (4.3)$$

One way of obtaining such a partition is through a greedy splitting procedure of the domain $\mathcal{Y} = [-1, 1]^{d_y}$ which is detailed in [9]. The procedure terminates when for each cell \mathcal{Y}_k the corresponding portion of the manifold \mathcal{M}_k can be associated to an affine \mathbb{U}_k satisfying these properties. We are ensured that this eventually occurs since for a sufficiently fine cell \mathcal{Y}_k one has $\text{rad}(\mathcal{M}_k) \leq \varepsilon$ which means that we could then use a zero dimensional affine space $\mathbb{U}_k = \{\bar{u}_k\}$ for which we know that $\mu(\tilde{\mathbb{U}}^k, \mathbb{W}) = 1$. In this piecewise affine model, the containment property is now

$$\mathcal{M} \subset \bigcup_{k=1}^K \mathcal{K}(\mathbb{U}_k, \varepsilon_k). \quad (4.4)$$

and the cardinality K of the partition depends on the prescribed ε .

For a given measurement $w \in \mathbb{W}$, we may now compute the state estimates

$$u_k^*(w) = A_{\mathbb{U}^k}(w), \quad k = 1, \dots, K, \quad (4.5)$$

by the affine variant of the one-space method from (3.4). Since $u \in \mathcal{M}_{k_0}$ for some value k_0 , we are ensured that

$$\|u - u_{k_0}^*(w)\|_{\mathbb{U}} \leq \varepsilon, \quad (4.6)$$

for this particular choice. However k_0 is unknown to us and one has to rely on the data w in order to decide which one among the affine models is most appropriate for

the recovery. One natural *model selection* criterion can be derived if for any $\bar{u} \in \mathbb{U}$ we have at our disposal a computable surrogate $S(\bar{u})$ that is equivalent to the distance from \bar{u} to \mathcal{M} , that is

$$cS(\bar{u}) \leq \text{dist}(\bar{u}, \mathcal{M})_{\mathbb{U}} \leq CS(\bar{u}), \quad \text{dist}(\bar{u}, \mathcal{M})_{\mathbb{U}} = \min_{y \in \mathcal{Y}} \|\bar{u} - u(y)\|_{\mathbb{U}}, \quad (4.7)$$

for some fixed $0 < c \leq C$. We give an instance of such a computable surrogate in Sect. 4.2 below. The selection criterion then consists in picking k^* minimizing this surrogate between the different available state estimates, that is,

$$u^*(w) := u_{k^*}^*(w) = \operatorname{argmin} \{S(u_k^*(w)) : k = 1, \dots, K\}. \quad (4.8)$$

The following result, established in [9], shows that this estimator now realizes the benchmark (2.26) up to a multiplication of ε by $\kappa := C/c$, where c, C are the constants from (4.7).

Theorem 2 *Assume that (4.2) and (4.3) hold. For any $u \in \mathcal{M}$, if $w = P_{\mathbb{W}}u$, one has*

$$\|u - u^*(w)\| \leq \delta_{\kappa\varepsilon}(\mathcal{M}, \mathbb{W}), \quad (4.9)$$

where $\delta_\varepsilon(\mathcal{M}, \mathbb{W})$ is given by (2.25).

4.2 Approximate Metric Projection and Parameter Estimation

A practically affordable realization of the surrogate $S(\bar{u})$, providing a *near-metric projection distance* to \mathcal{M} , is a key ingredient of the above nonlinear recovery scheme. Since it has further useful implications we add a few comments on that matter.

As already observed in (2.5), whenever (2.1) admits a stable variational formulation with respect to a suitable pair (\mathbb{U}, \mathbb{V}) of trial and test spaces, the distance of any $\bar{u} \in \mathbb{U}$ to any $u(y) \in \mathcal{M}$ is uniformly equivalent to the residual of the PDE in \mathbb{V}'

$$c\|\mathcal{R}(\bar{u}, y)\|_{\mathbb{V}'} \leq \|u(y) - \bar{u}\|_{\mathbb{U}} \leq C\|\mathcal{R}(\bar{u}, y)\|_{\mathbb{V}'}, \quad (4.10)$$

with $c = C_b^{-1}$, $C = c_b^{-1}$ from (2.5). Assume in addition that $\mathcal{R}(u, y)$ depends *affinely* on $y \in \mathcal{Y}$, according to (2.9). Then, minimizing $\|\mathcal{R}(\bar{u}, y)\|_{\mathbb{V}'}$ over y is equivalent to solving a *constrained least squares* problem

$$\bar{y} = \operatorname{argmin}_{y \in \mathcal{Y}} \|\mathbf{g} - \mathbf{M}y\|_2, \quad (4.11)$$

where \mathbf{M} is a matrix of size $d_y \times d_y$ resulting from Riesz-lifts of the functionals $\mathcal{R}_j(\bar{u})$.

The solution to this problem therefore satisfies

$$\|\bar{u} - u(\bar{y})\|_{\mathbb{U}} \leq \kappa \inf_{y \in \mathcal{Y}} \|\bar{u} - u(y)\|_{\mathbb{U}} = \kappa \text{dist}(\bar{u}, \mathcal{M})_{\mathbb{U}}. \quad (4.12)$$

where $\kappa = C/c = C_b/c_b$ is the quotient between the equivalence constants in (4.10). The surrogate

$$S(\bar{u}) := \|\mathcal{R}(\bar{u}, y)\|_{\mathbb{V}'} \quad (4.13)$$

for the metric projection distance of \bar{u} onto \mathcal{M} obviously satisfies (4.7). It is indeed computable at affordable cost using (an approximation to) its Riesz-lifted version $\|e(\bar{u}, y)\|_{\mathbb{V}'} = \|\mathcal{R}(\bar{u}, y)\|_{\mathbb{V}'}$ (in $\mathbb{V}_h \subset \mathbb{V}$) assembled from the Riesz-lifts of the components $\mathcal{R}_j(\bar{u})$, see [9] for details in the affine expansion (2.9).

Since solving the above problem provides an admissible parameter value $\bar{y} \in \mathcal{Y}$, this also has some immediate bearing on *parameter estimation*. Suppose we wish to estimate from $w = P_{\mathbb{W}}u(y^*)$ the unknown parameter $y^* \in \mathcal{Y}$. Assume further that A is any given linear or nonlinear recovery map. Computing along the above lines

$$\bar{y}_w = \underset{y \in \mathcal{Y}}{\operatorname{argmin}} \|\mathcal{R}(A(w), y)\|_{\mathbb{V}'}$$

we have

$$\begin{aligned} \|u(y^*) - u(\bar{y}_w)\|_{\mathbb{U}} &\leq \|u(y^*) - A(w)\|_{\mathbb{U}} + \|A(w) - u(\bar{y}_w)\|_{\mathbb{U}} \\ &\leq E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}) + \kappa \text{dist}(A(w), \mathcal{M})_{\mathbb{U}} \leq (1 + \kappa)E_{\text{wc}}(A, \mathcal{M}, \mathbb{W}). \end{aligned} \quad (4.14)$$

We consider now the specific elliptic model (2.6) with affine diffusion coefficients $a(y)$ given by (2.10). For this model, it was established in [5] that for strictly positive f and certain regularity assumptions on $a(y)$ as functions of $x \in \Omega$, parameters may be estimated by states. Specifically, when $a(y) \in H^1(\Omega)$ uniformly in $y \in \mathcal{Y}$, one has an inverse stability estimate of the form

$$\|a(y) - a(\tilde{y})\|_{L_2(\Omega)} \leq C \|u(y) - u(\tilde{y})\|_{\mathbb{U}}^{1/6}. \quad (4.15)$$

Thus, whenever the recovery map A satisfies (4.9) for some prescribed $\varepsilon > 0$, we obtain a parameter estimation bound of the form

$$\|a(y^*) - a(\bar{y}_w)\|_{L_2(\Omega)} \leq C \delta_{\kappa\varepsilon}(\mathcal{M}, \mathbb{W})^{1/6},$$

Note that when the basis functions θ_j are L_2 -orthogonal, $\|a(y^*) - a(\bar{y}_w)\|_{L_2(\Omega)}$ is equivalent to a (weighted) ℓ_2 norm of $y^* - \bar{y}_w$.

4.3 Concluding Remarks

The affine or piecewise affine recovery scheme hinges on the ability to approximate a solution manifold effectively by linear or affine spaces, globally or locally. As explained earlier this is true for problems of elliptic or parabolic type that may include convective terms as long as they are dominated by diffusion. This may however no longer be the case when dealing with pure transport equations or models involving strongly dominating convection.

An interesting alternative would then be to adopt a stochastic model according to (2.27) and (2.28) that allows one to view the construction of the recovery map as a regression problem. In particular, when dealing with transport models, a natural candidate for parametrizing a reduced model are *deep neural networks*. However, properly adapting the architecture, regularization and training principles pose wide open questions addressed in current work in progress.

Acknowledgements A. C. was supported by ERC Adv Grant BREAD. W. D. was supported in part by the NSF-grants DMS 17-20297, DMS-2012469, by the Smart State Program of the State of South Carolina, and the Williams-Hedberg Foundation. R. D. was supported by the NSF grant DMS 18-17603. A portion of this work was completed when the authors were supported as visitors to the Isaac Newton Institute of Cambridge University.

References

1. Adcock, B., Hansen, A.C., Poon, C.: Beyond consistent reconstructions: optimality and sharp bounds for generalized sampling, and application to the uniform resampling problem. *SIAM J. Math. Anal.* **45**, 3132–3167 (2013)
2. Bachmayr, M., Cohen, A.: Kolmogorov widths and low-rank approximations of parametric elliptic PDEs. *Math. Comp.* **86**, 701–724 (2017)
3. Binev, P., Cohen, A., Dahmen, W., DeVore, R., Petrova, G., Wojtaszczyk, P.: Convergence rates for greedy algorithms in reduced basis methods. *SIAM J. Math. Anal.* **43**, 1457–1472 (2011)
4. Binev, P., Cohen, A., Dahmen, W., DeVore, R., Petrova, G., Wojtaszczyk, P.: Data assimilation in reduced modeling. *SIAM J. Uncert Quantif* **5**(1), 1–29 (2017)
5. Bonito, A., Cohen, A., DeVore, R., Petrova, G., Welper, G.: Diffusion coefficients estimation for elliptic partial differential equations. *SIAM J. Math. Anal.* **49**(2), 1570–1592 (2017)
6. Broersen, D., Stevenson, R.: A robust Petrov-Galerkin discretisation of convection-diffusions equations. *Comput. Math. Appl.* **68**(11), 1605–1618 (2014)
7. Broersen, D., Dahmen, W., Stevenson, R.: On the stability of DPG formulations of transport equations. *Math. Comp.* **87**(311), 1051–1082 (2018)
8. Buffa, A., Maday, Y., Patera, A.T., Prud'homme, C., Turicini, G.: A priori convergence of the greedy algorithm for the parametrized reduced bases. *ESAIM Math. Model. Numer. Anal.* **46**(03), 595–603 (2012)
9. Cohen, A., Dahmen, W., Mula, O., Nichols, J.: Nonlinear reduced models for state and parameter estimation (to appear in *SIAM J. Uncert. Quantif.*). <https://arxiv.org/submit/3356244>
10. Cohen, A., Dahmen, W., DeVore, R., Nichols, J.: Reduced basis greedy selection using random training sets. *M2AN* **54**(5), 1509–1524. <https://doi.org/10.1051/m2an/2020004>, <http://arxiv.org/abs/1810.09344> [math.NA]
11. Cohen, A., Dahmen, W., Welper, G.: Adaptivity and variational stabilization for convection-diffusion equations. *ESAIM: Math. Model. Numer. Anal.* **46**(5), 1247–1273 (2012)

12. Cohen, A., Dahmen, W., DeVore, R., Fadili, J., Mula, O., Nichols, J.: Optimal reduced model algorithms for data-based state estimation. *SIAM J. Numer. Anal.* **58**(6), 3355–3381 (2020). <http://arxiv.org/abs/1903.07938>
13. Cohen, A., DeVore, R.: Approximation of high-dimensional PDEs. *Acta Numer.* **24**, 1–159 (2015)
14. Cohen, A., DeVore, R.: Kolmogorov widths under holomorphic mappings. *IMA J. Numer. Anal.* **36**(1), 1–12 (2016)
15. Cohen, A., DeVore, R., Schwab, C.: Convergence rates of best N -term Galerkin approximations for a class of elliptic sPDEs. *Found. Comput. Math.* **10**(6), 615–646 (2010)
16. Dahmen W.: How to best sample a solution manifold? In: Pfander, G.E. (eds.) *Sampling Theory, a Renaissance, Applied and Numerical Harmonic Analysis*. Birkhäuser. ISBN 978-3-319-19748-7. https://doi.org/10.1007/978-3-319-19749-4_11, <http://arxiv.org/abs/1503.00307> [math.NA]
17. Dahmen, W., Huang, C., Schwab, C., Welper, G.: Adaptive Petrov-Galerkin methods for first order transport equations. *SIAM J. Numer. Anal.* **50**(5), 2420–2445 (2012)
18. Dahmen, W., Plesken, C., Welper, G.: Double greedy algorithms: reduced basis methods for transport dominated problems. *ESAIM: Math. Model. Numer. Anal.* **48**(3), 623–663 (2014)
19. DeVore R., Petrova, G., Wojtaszczyk, P.: Greedy algorithms for reduced bases in Banach spaces. *Constr. Approx.* **37**(3), 455–466 (2013)
20. Maday, Y., Patera, A.T., Penn, J.D., Yano, M.: A parametrized-background data-weak approach to variational data assimilation: formulation, analysis, and application to acoustics. *Int. J. Numer. Methods Eng. Spec. Issue Model. Reduct.* **102**(5), 931–1292 (2015)
21. Micchelli, C.A., Rivlin, T.J.: A survey of optimal recovery. In: Micchelli, C.A., Rivlin, T.J. (eds.) *Optimal Estimation in Approximation Theory*, pp. 1–54. Plenum, NY (1977)
22. Rozza, G., Huynh, D.B.P., Patera, A.T.: Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics. *Arch. Comput. Methods Eng.* **15**(3), 229–275 (2008). <https://doi.org/10.1007/s11831-008-9019-9>
23. Stevenson, R., Westerdiep, J.: Stability of Galerkin discretizations of a mixed space-time variational formulation of parabolic evolution equations. *IMA J. Numer. Anal.* **41**(1) (2021). <https://doi.org/10.1093/imanum/drz069>
24. Stuart, A.M.: Inverse problems: a Bayesian perspective. *Acta Numer.* **19**, 451–559 (2010). <https://doi.org/10.1017/S0962492910000061>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

