Approximation of multivariate functions : reduced modeling and recovery from uncomplete measurements

Albert Cohen

Laboratoire Jacques-Louis Lions
Sorbonne Université
Paris

Nantes 16-05-2022

An ubiquitous numerical problem

Reconstruct an unknown multivariate function

$$u : x \mapsto u(x), \quad x = (x_1, \dots, x_d) \in D \subset \mathbb{R}^d,$$

from (possibly noisy) observations $y^i \approx \ell_i(u) \in \mathbb{R}$ for $i = 1, \dots, m$.

Here the ℓ_i are linear forms.

An important case : evaluation $y^i \approx u(x^i)$ at sample points $x^i \in D$ for $i = 1, \dots, m$.

Distinction between two data acquisition settings :

Passive setting : we do not choose the x^i (or the ℓ_i).

Active setting : we choose the x^i (or the ℓ_i).

General questions

In the active setting :

How should we sample ?

How should we reconstruct ?

Due to limited amount of data, even if noiseless, we will reconstruct in some simpler class of function V_n represented by $n \leq m$ parameters, typically a n-dimensional space.

Thus, we reconstruct some approximation $\tilde{u} \in V_n$ of $u \notin V_n$.

Measuring performance : error $\|u - \tilde{u}\|$ versus number of measurements m ?

But also : computational complexity of the reconstruction method (offline/online).

Here $\|\cdot\| = \|\cdot\|_V$ is a norm of a Banach space V containing u.

Benchmark : best approximation error $e_n(u) = e_n(u)_V = \min_{v \in V_n} \|u - v\|$.

How should we pick good approximation spaces V_n ?

Agenda

Day 1

1. Recovery problems : applicative settings and objectives.
2. Tools from linear and nonlinear approximation theory : n-widths.
3. Reduced bases and PCA.
4. Breaking the curse of dimensionality : anisotropy and sparse approximation.

Day 2

5. Recovery from point evaluation : weighted least-squares methods.
6. Optimal sampling measure : theory and practical aspects.
7. More general measurements : the PBDW method.

Passive aquisition setting

Input-output modeled by $(x, y) \in D \times \mathbb{R}$ is a random variable of unknown joint law.

We observe independant realizations (x^i, y^i) for $i = 1, \dots, m$. We search for a function that best explains y from x.

Applicative context : regression, machine learning, denoising...

The quadratic risk $\mathbb{E}(|y - v(x)|^2)$ is minimized among all functions v by $u(x) := \mathbb{E}(y|x)$ which is unknown.

For $\tilde{u} \neq u$, one has

$$\mathbb{E}(|y - \tilde{u}(x)|^2) = \mathbb{E}(|y - u(x)|^2) + \mathbb{E}(|\tilde{u}(x) - u(x)|^2) = \sigma^2 + \int_D |u(x) - \tilde{u}(x)|^2 d\mu,$$

where $d\mu$ is the unknown probability measure of x.

We thus measure performance of a reconstruction $\tilde{u}$ by $\|u - \tilde{u}\|_{L^2(D,\mu)}$.

Inherently noisy setting : $y^i = u(x^i) + \eta^i$, where η^i is a noise $\mathbb{E}(\eta|x) = 0$.

Active aquisition setting

We are allowed to query an unknown map $x \mapsto u(x)$, typically by running an experiment or a numerical simulation.

Each (offline) query $x^i \mapsto y^i = u(x^i)$ is costly (and could be noisy).

We want to compute an approximation map $x \mapsto \tilde{u}(x)$ that is much cheaper to evaluate (online) than u.

Applicative context : model reduction, data aquisition, inverse problems, design of computer experiments.

We measure performance in some Banach space norm $\|\cdot\| = \|\cdot\|_V$ of interest, for example $\|u - \tilde{u}\|_{L^2(D,\mu)}$ where μ can be chosen by us, for example the Lebesgue measure.

Is there an optimal choice of the sample $(x^1, \dots, x^m)$? Easy to construct ?

We can invest some offline time designing the sample (prune from a larger sample).

When $d >> 1$ we want to avoid uniform grids (curse of dimensionality).

The function u may take its value in $\mathbb{R}$, or $\mathbb{R}^k$, or in an infinite dimensional space.

Example 1 : recover a physical phenomenon from pointwise sensing

An acoustic pressure field $p(x,t)$ generated by a source is measured by n microphones at positions $x^1,\dots,x^m \in D \subset \mathbb{R}^2$ or $\mathbb{R}^3$, for $t \in [0,T]$.

Fourier analysis in time $p(x^i,t) \mapsto \hat{p}(x^i,\omega)$ and focus at a frequency ω of interest.

One wants to reconstruct the function $u(x) := \hat{p}(x,\omega)$ on Y, from the observed data $u(x^i),\ i = 1,\dots,m$.

Example 2 : Fast approximate solutions to high dimensional parametric PDE's

Partial differential equation $\mathcal{P}(u, y) = 0$ depending on a parameter vector $y \in Y \subset \mathbb{R}^d$, with $d >> 1$. For each $y \in Y$, the PDE is well posed in some Hilbert space V : parameter to solution map $y \in Y \mapsto u(y) \in V$ (can be queried by a numerical solver).

Example : steady state diffusion equation

$$-\mathrm{div}(a(x)\nabla u(x)) = f(x), \quad x \in D \subset \mathbb{R}^2 \text{ or } \mathbb{R}^3 \quad + \text{ boundary conditions}$$

where diffusion function a is piecewise constant on subdomains $D_1, \dots, D_d$, with values $y_1, \dots, y_d$, which define the parameter vector $y = (y_1, \dots, y_d) \in Y = [y_{\min}, y_{\max}]^d$.

Surrogate/reduced model : from snapshots, i.e. particular instances of solutions $u(y^i)$, $i = 1, \dots, m$ computed by the numerical solver, we want to reconstruct an approximation $y \mapsto \tilde{u}(y)$ to the parameter to solution map, much cheaper to evaluate.

Example 2 : Fast approximate solutions to high dimensional parametric PDE's

Partial differential equation $\mathcal{P}(u,y)=0$ depending on a parameter vector $y\in Y\subset\mathbb{R}^d$, with $d>>1$. For each $y\in Y$, the PDE is well posed in some Hilbert space V : parameter to solution map $y\in Y\mapsto u(y)\in V$ (can be queried by a numerical solver).

Example : steady state diffusion equation

$$-\mathrm{div}(a(x)\nabla u(x))=f(x),\quad x\in D\subset\mathbb{R}^2 \text{ or } \mathbb{R}^3 \quad + \text{ boundary conditions}$$

where diffusion function a is piecewise constant on subdomains $D_1,\dots,D_d$, with values $y_1,\dots,y_d$, which define the parameter vector $y=(y_1,\dots,y_d)\in Y=[y_{\min},y_{\max}]^d$.

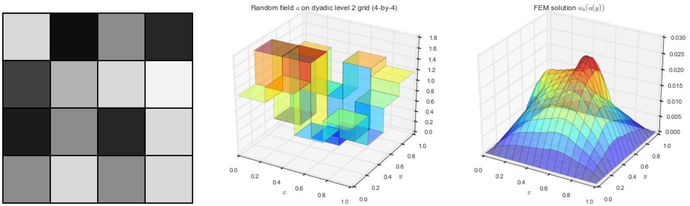

Surrogate/reduced model : from snapshots, i.e. particular instances of solutions $u(y^i)$, $i=1,\dots,m$ computed by the numerical solver, we want to reconstruct an approximation $y\mapsto\tilde{u}(y)$ to the parameter to solution map, much cheaper to evaluate.

Example 2 : Fast approximate solutions to high dimensional parametric PDE's

Partial differential equation $\mathcal{P}(u, y) = 0$ depending on a parameter vector $y \in Y \subset \mathbb{R}^d$, with $d >> 1$. For each $y \in Y$, the PDE is well posed in some Hilbert space V : parameter to solution map $y \in Y \mapsto u(y) \in V$ (can be queried by a numerical solver).

Example : steady state diffusion equation

$$-\mathrm{div}(a(x)\nabla u(x)) = f(x), \quad x \in D \subset \mathbb{R}^2 \text{ or } \mathbb{R}^3 \quad + \text{ boundary conditions}$$

where diffusion function a is piecewise constant on subdomains $D_1, \dots, D_d$, with values $y_1, \dots, y_d$, which define the parameter vector $y = (y_1, \dots, y_d) \in Y = [y_{\min}, y_{\max}]^d$.

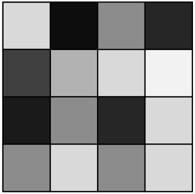

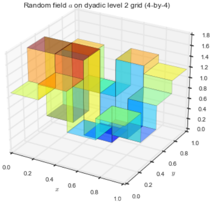

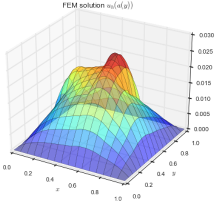

Surrogate/reduced model : from snapshots, i.e. particular instances of solutions $u(y^i)$, $i = 1, \dots, m$ computed by the numerical solver, we want to reconstruct an approximation $y \mapsto \tilde{u}(y)$ to the parameter to solution map, much cheaper to evaluate.

Example 3 : inverse problems in parametric PDE's

Partial differential equation $\mathcal{P}(u, y) = 0$ depending on a parameter vector $y \in Y \subset \mathbb{R}^d$ with $d >> 1$.

In certain settings, we know the governing PDE but do not know the parameters y of the particular solution solution $u(y)$ that we are trying to capture.

For instance, the diffusion properties of the underground could be unknown to us. So we make some local measurements by "drilling" inside the domain, or on the boundary of the domain : we are given m measurements

$$\ell_i(u), \quad i = 1, \dots, m, \quad u = u(y)$$

where the ℓ_i are linear forms on V (that could be point evaluation, local averages...).

How can we best combine these measurements with the model to recontruct the state u (state estimation) or parameter y (parameter estimation) ?

$$y \in Y \quad \mapsto \quad u = u(y) \in V \quad \mapsto \quad \ell(u) = \ell(u(y)) = (\ell_i(u(y)))_{i=1,\dots,m} \in \mathbb{R}^m.$$

The 2d Radon transform, a simplified model for tomography.

The unknown function $x \mapsto u(x)$ represent the density of tissues in a 2d slice human body. The detector at angle α and offset r measures

$$Ru(r,\alpha) = \int_{-\infty}^{+\infty} u(re_\alpha + se_\alpha^\perp)ds,$$

the attenuation along the line ray $L(r,\alpha) := \{re_\alpha + se_\alpha^\perp \ : \ s \in \mathbb{R}\}$.

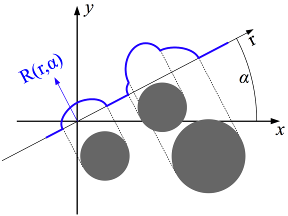

These measurements are sampled, so we observe limited data

$$\ell_i(u) = Ru(r^i,\alpha^i), \quad i = 1,\dots,m.$$

How should we sample / reconstruct ?

Approximation

Error measure : $\|u-\tilde{u}\|_V$, where $V := L^2(D,\mu)$, or other Banach space of interest.

Most often, the reconstruction $\tilde{u}$ takes place within a family $V_n \subset V$ that can be parametrized by $n \leq m$ numbers.

So it is relevant to compare $\|u-\tilde{u}\|_V$ with

$$e_n(u)_V = \min_{v \in V_n} \|u-v\|_V.$$

We restrict our attention to linear families : V_n is a linear space with $n = \dim(V_n)$.

If V is a Hilbert space, $e_n(u) = \|u - P_{V_n} u\|_V$ with P_{V_n} the V-orthogonal projection.

Classical choices : algebraic polynomials, spline spaces, trigonometric polynomials, piecewise constant functions on a given partition of D.

If our prior information is that $u \in \mathcal{K}$ where $\mathcal{K} \subset V$ is a certain class of functions, we could search for other choices of spaces V_n that are better fitted to the class $\mathcal{K}$.

Kolmogorov linear n-width

We are interested in approximating general functions $u \in V$, where V is a Banach space, by simpler functions v picked from a linear subspace $V_n \subset V$ of finite dimension n.

Classical Banach spaces : Lebesgue $L^p(D)$, Sobolev $W^{m,p}(D)$ for $D \subset \mathbb{R}^d$.

Classical linear subspaces : algebraic or trigonometric polynomials of some prescribed degree, splines or finite elements on some given mesh, span of the n first elements $\{e_1, \dots, e_n\}$ from a given basis $(e_k)_{k\geq 1}$ of V.

Model class reflecting the properties the target function : $u \in \mathcal{K}$, where $\mathcal{K}$ is a compact set of V. Best choice of approximation spaces for this model class ?

The space V_n approximate $\mathcal{K}$ with uniform accuracy

$$\mathrm{dist}(\mathcal{K}, V_n)_V := \max_{u\in\mathcal{K}} \min_{v\in V_n} \|u - v\|_V$$

A.N. Kolmogorov (1936) defines the linear n-width of $\mathcal{K}$ in the metric V as

$$d_n = d_n(\mathcal{K})_V := \inf_{\dim(V_n)=n} \mathrm{dist}(\mathcal{K}, V_n)_V$$

Intuition

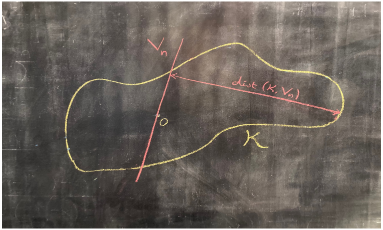

The optimal space achieving the infimum in

$$d_n(\mathcal{K})_V = \inf_{\dim(V_n)=n} \max_{u\in\mathcal{K}} \min_{v\in V_n} \|u-v\|_V.$$

may not exist. One often assumes it exists in order to avoid limiting arguments. It is generally not easy to identify or characterize.

The quantity $d_n(\mathcal{K})_V$ can be viewed as a benchmark/bottleneck for numerical methods applied to the elements from $\mathcal{K}$ that create approximations from linear spaces : interpolation, projection, least squares, Galerkin methods for solving PDEs...

Intuition

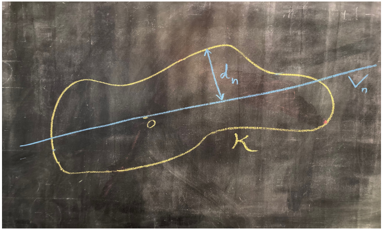

The optimal space achieving the infimum in

$$d_n(\mathcal{K})_V = \inf_{\dim(V_n)=n} \max_{u\in\mathcal{K}} \min_{v\in V_n} \|u-v\|_V.$$

may not exist. One often assumes it exists in order to avoid limiting arguments. It is generally not easy to identify or characterize.

The quantity $d_n(\mathcal{K})_V$ can be viewed as a benchmark/bottleneck for numerical methods applied to the elements from $\mathcal{K}$ that create approximations from linear spaces : interpolation, projection, least squares, Galerkin methods for solving PDEs...

An analog concept in the stochastic framework : PCA

Assume that V is a Hilbert space and u is a random variable taking its value in V.

Optimal spaces in the mean-square sense.

$$\kappa_n^2 = \kappa_n(u)_V^2 := \min_{\dim(V_n)=n} \mathbb{E}\Big(\|u - P_{V_n} u\|_V^2\Big).$$

The space achieving the minimum is easily characterized by principal component analysis : consider the covariance operator

$$v \mapsto Rv = \mathbb{E}(\langle u, v\rangle u),$$

which is compact, when assuming that $\mathbb{E}(\|u\|_V^2) < \infty$. Diagonalized in the Karhunen-Loeve basis $(\varphi_k)_{k\geq 1}$ with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \to 0$.

Then $V_n := \mathrm{span}\{\varphi_1, \dots, \varphi_n\}$ and $\kappa_n^2 = \sum_{k>n} \lambda_k$.

Note that $\kappa_n(u)_V^2 \leq d_n(\mathcal{K})_V^2$ when u is supported in $\mathcal{K}$.

Variants to n-width : realization of the approximation

The best approximation $u_n = \operatorname{argmin}\{\|u - v\|_V \ : \ v \in V_n\}$ is the orthogonal projection $u_n = P_{V_n} u$ if V is a Hilbert space.

For a general Banach space, the map $u \mapsto u_n$ is not linear, and may even not be continuous (non-uniqueness of best approximation).

This motivates alternate definitions of widths where we impose linearity or continuity of the approximation process.

Approximation numbers are defined as

$$a_n(\mathcal{K})_V := \inf_L \max_{u \in \mathcal{K}} \|u - Lu\|_V,$$

with infimum taken over all linear maps L such that $\operatorname{rank}(L) \leq n$.

In a general Banach space $d_n \leq a_n \leq \sqrt{n} d_n$ and right equality may hold.

On the other hand one can prove that

$$d_n(\mathcal{K})_V := \inf_F \max_{u \in \mathcal{K}} \|u - F(u)\|_V,$$

with infimum taken over all continuous maps F such that $\operatorname{rank}(F) \leq n$.

Behaviour of n-widths of smoothness classes

Typical compact sets in $V = L^p(D)$ are balls of smoothness spaces. The behaviour of n-width is well understood for such sets. Example : $V = L^\infty(I)$ where $I = [0,1] \subset \mathbb{R}$

and

$$\mathcal{K} = \mathcal{U}(\mathrm{Lip}(I)) = \{u \ : \ \max\{\|u\|_{L^\infty}, \|u'\|_{L^\infty}\} \leq 1\},$$

Then one can prove

$$d_n(\mathcal{K})_V = \frac{1}{2n}, \quad n \geq 1, .$$

More generally when $V = W^{t,p}(D)$ for some bounded Lipschitz domain $D \subset \mathbb{R}^d$ and $\mathcal{K}$ is the unit ball of $W^{s,p}(D)$ with $s > t$, one can prove

$$cn^{-(s-t)/d} \leq d_n(\mathcal{K})_V \leq Cn^{-(s-t)/d}, \quad n \geq 1.$$

Curse of dimensionality : exponential growth in d of the needed n to reach accuracy ε.

Proof of upper bound : use a standard approximation method (piecewise polynomials, finite elements, or splines, on uniform partitions of D)

Proof of lower bound ? Two systematic approaches.

Lemma : let $B_W = \{u \in W \ : \ \|u\|_V \leq 1\}$ be the unit ball of a subspace $W \subset V$ of dimension $n+1$, then $d_n(B_W)_V = 1$.

Proof : trivial if V is a Hilbert space. Follows from Borsuk-Ulam antipodality theorem in the Banach space case : for any continuous application F from an n-sphere $S_n = \partial B_W$ to an n dimensional space V_n, there exists $x \in S_n$ such that $F(x) = F(-x)$.

It follows that $d_n(\mathcal{K})_V \geq r$ if $\mathcal{K}$ contains the rescaled ball rB_W of an $n+1$-dimensional space W. In other words

$$d_n(\mathcal{K})_V \geq b_n(\mathcal{K})_V, \quad n \geq 1,$$

where the Bernstein n-width $b_n(\mathcal{K})_V$ is defined as the largest $r \geq 0$ such that there exists $W \subset V$ of dimension $n+1$ with $rB_W \subset \mathcal{K}$.

Lemma : let $B_W = \{u \in W \; : \; \|u\|_V \leq 1\}$ be the unit ball of a subspace $W \subset V$ of dimension $n+1$, then $d_n(B_W)_V = 1$.

Proof : trivial if V is a Hilbert space. Follows from Borsuk-Ulam antipodality theorem in the Banach space case : for any continuous application F from an n-sphere $S_n = \partial B_W$ to an n dimensional space V_n, there exists $x \in S_n$ such that $F(x) = F(-x)$.

It follows that $d_n(\mathcal{K})_V \geq r$ if $\mathcal{K}$ contains the rescaled ball rB_W of an $n+1$-dimensional space W. In other words

$$d_n(\mathcal{K})_V \geq b_n(\mathcal{K})_V, \quad n \geq 1,$$

where the Bernstein n-width $b_n(\mathcal{K})_V$ is defined as the largest $r \geq 0$ such that there exists $W \subset V$ of dimension $n+1$ with $rB_W \subset \mathcal{K}$.

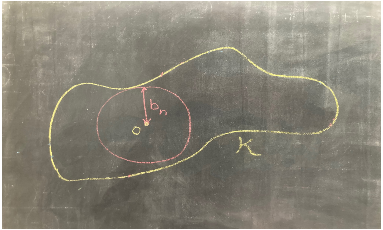

Entropy numbers

Define $\varepsilon_n(\mathcal{K})_V$ as the smallest ε such that $\mathcal{K}$ can be covered by 2^n balls of radius ε :

$$\mathcal{K} \subset \bigcup_{i=1,\dots,2^n} B(u^i,\varepsilon), \qquad B(u^i,\varepsilon) := \{u \,:\, \|u-u^i\|_V \le \varepsilon\}.$$

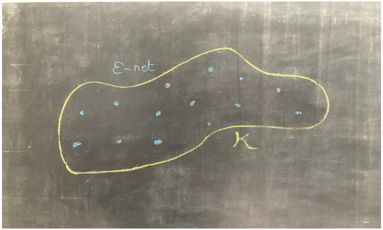

Related to lossy coding : Elements of $\mathcal{K}$ can be encoded with n bits up to precision ε_n.

Carl's inequality : for all $s>0$ one has

$$(n+1)^s \varepsilon_n(\mathcal{K})_V \le C_s \sup_{m=0,\dots,n} (m+1)^s d_m(\mathcal{K})_V, \quad n \ge 0$$

In particular

$$d_n(\mathcal{K})_V \lesssim n^{-s}, \quad n\ge 0 \implies \varepsilon_n(\mathcal{K})_V \lesssim n^{-s}, \quad n \ge 0.$$

Complex problems are often modelled by PDEs involving several physical parameters $y = (y^1, \dots, y^d) \in Y \subset \mathbb{R}^d$.

$$\mathcal{P}(u, y) = 0,$$

For each $y \in Y$, we assume well-posedness and therefore existence of a unique solution $u(y) \in V$.

In certain applications (optimization, inverse problems, uncertainty quantification), we may need to solve $y \mapsto u(y)$ for many instances of $y \in Y$: requires computational methods that are uniformly cheap and efficient, uniformly over $y \in Y$.

We are interested in well approximating the solution manifold

$$\mathcal{K} := \{u(y) \ : \ y \in Y\} \subset V,$$

which we assume to be compact.

Reduced modeling usually involves two steps :

1. In a (costly) offline stage, we search for spaces V_n of dimension n that approximate as best as possible the set $\mathcal{K}$ (benchmark $d_n(\mathcal{K})_V$). These spaces are quite different from classical finite element spaces.

2. In a (cheap) online stage, for any required $y \in Y$ we may compute an accurate approximation $u_n(y) \in V_n$ of $u(y)$, for example by the Galerkin method.

Estimating n-width of solution manifolds

An instructive example : consider the steady-state diffusion equation

$$-\mathrm{div}(a\nabla u) = f,$$

on a $2d$ domain D ($+$ boundary conditions), with piecewise constant diffusion function $a = a(y)$ having value $\overline{a} + y_j$ on subdomain D_j, where $y = (y_1, \dots, y_d) \in Y = [-c, c]^d$.

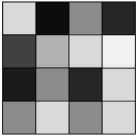

How large is the n-width of $\mathcal{K} = \{u(y) \;:\; y \in Y\} \subset V = H^1(D)$?

Solutions $u(y)$ are bounded in H^s iff $s < 3/2$ and $d_n(\mathcal{U}(H^s))_{H^1} \sim n^{-(s-1)/2} \gtrsim n^{-1/4}$.

In fact $d_n(\mathcal{K})_{H^1}$ decreases faster than $\mathcal{O}(\exp(-cn^{1/d}))$: approximate by power series

$$\max_{y\in Y} \Big\| u(y) - \sum_{|\nu|\le k} u_\nu y^\nu \Big\|_{H^1} \le C\exp(-ck), \quad y^\nu = y_1^{\nu_1} \dots y_d^{\nu_d},$$

and use $V_n = \mathrm{span}\{u_\nu \;:\; |\nu| \le k\}$ of dimension $n = \binom{k+d}{k}$.

A general result for infinite dimensional parameter dependence

Theorem (Cohen-DeVore, 2016) : Let V_1 and V_2 be two complex valued Banach spaces and $\mathcal{K}_1 \subset V_1$ be a compact set. Let

$$F : V_1 \to V_2,$$

be a map that is holomorphic on an open neighbourhood of $\mathcal{K}_1$. Then, with $\mathcal{K}_2 := F(\mathcal{K}_1)$, one has for all $s > 1$

$$\sup_{n\geq 0} n^s d_n(\mathcal{K}_1)_{V_1} < \infty \implies \sup_{n\geq 0} n^t d_n(\mathcal{K}_2)_{V_2} < \infty, \quad t < s - 1.$$

Note that if F was a continuous linear map, one would simply have

$$d_n(\mathcal{K}_2)_{V_2} \leq C d_n(\mathcal{K}_1)_{V_1}, \quad C = \|F\|_{V_1 \to V_2}.$$

The proof goes by expanding $a \in \mathcal{K}_1$ in a suitable basis $a = a(y) = \sum_{j\geq 1} y_j \psi_j$ with decay properties on the $\|\psi_j\|_{V_1}$ and then approximate $F(a(y))$ by polynomials in y. This induces a loss of 1 in the rate of decay. Open problem : same rate $t = s$?

This result applies to elliptic equations such as $-\mathrm{div}(a\nabla u) = f$ for the map $F : a \to u$ with $V_1 = L^\infty$ and $V_2 = H^1$. Also applies to parabolic equations, nonlinear problems such as Navier-Stokes equations, and to these problems set on parametrized domains. It does not apply to hyperbolic equations.

The reduced basis algorithm

Idea : use particular instances $u_i = u(y^i) \in \mathcal{K}$ for generating $V_n = \mathrm{span}\{u_1, \dots, u_n\}$.

Greedy selection in offline stage : having generated $u_1, \dots, u_{k-1}$, select next instance

$$\|u_k - P_{V_{k-1}} u_k\|_V = \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V,$$

where $P_{V_{k-1}}$ is the orthogonal projection. Here we assume V to be a Hilbert space.

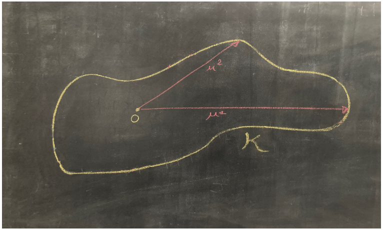

In practice, u_k and $\|u_k - P_{V_{k-1}} u_k\|_V$ approximated by solver and a-posteriori analysis.

Weak selection $\|u_k - P_{V_{k-1}} u_k\|_V \geq \gamma \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V$, for fixed $\gamma \in]0, 1[$.

In the maximization, $\mathcal{K}$ is often replaced by a large but finite training set $\widetilde{\mathcal{K}}$.

Variant : POD bases diagonalizing $v \mapsto \#(\widetilde{\mathcal{K}})^{-1} \sum_{u \in \widetilde{\mathcal{K}}} \langle u, v \rangle u$.

Approximation performances

For the greedily generated spaces V_n, we would like to compare

$$\sigma_n(\mathcal{K})_V = \mathrm{dist}(\mathcal{K}, V_n)_V = \max_{u\in\mathcal{K}} \|u - P_{V_n}u\|_V,$$

with the n-widths $d_n(\mathcal{K})_V$ that correspond to the optimal spaces.

Direct comparison is deceiving.

Buffa-Maday-Patera-Turinici (2010) : $\sigma_n \le n2^n d_n$.

For all $n \ge 0$ and $\varepsilon > 0$, there exists $\mathcal{K}$ such that $\sigma_n(\mathcal{K})_V \ge (1-\varepsilon)2^n d_n(\mathcal{K})_V$.

Comparison is much more favorable in terms of convergence rate.

Theorem (Binev-Cohen-Dahmen-DeVore-Petrova-Wojtaszczyk, 2013) : For any $s > 0$,

$$\sup_{n\ge1} n^s d_n(\mathcal{K})_V < \infty \Rightarrow \sup_{n\ge1} n^s \sigma_n(\mathcal{K})_V < \infty,$$

and

$$\sup_{n\ge1} e^{cn^s} d_n(\mathcal{K})_V < \infty \Rightarrow \sup_{n\ge1} e^{\tilde{c}n^s} \sigma_n(\mathcal{K})_V < \infty,$$

A matrix reformulation

In order to prove the theorem, we introduce the functions $\{u_0^*, u_1^*, \cdots\}$ obtained by applying Gram-Schmidt orthonormalization algorithm to the sequence $\{u_0, u_1, \cdots\}$. We consider the lower triangular matrix $A = (a_{i,j})_{i,j\geq 0}$ defined by

$$u_i = \sum_{j=0}^{i} a_{i,j} u_j^*.$$

This matrix satisfies two fundamental properties. Since $a_{n,n} = \langle u_n, u_n^* \rangle = \|u_n - P_{V_n} u_n\|_V$, we have

$$\gamma\sigma_n \leq |a_{n,n}| \leq \sigma_n \quad (P1),$$

where $\sigma_n := \sigma_n(\mathcal{M})_V$. Since for $m \geq n$ we have $\|u_m - P_{V_n} u_m\|_V \leq \sigma_n$, we have

$$\sum_{j=n}^{m} a_{m,j}^2 \leq \sigma_n^2 \qquad (P2)$$

Conversly, for any matrix satisfying these two properties with $(\sigma_n)_{n\geq 0}$ a non-increasing sequence going to 0, there exists a compact set $\mathcal{K}$ in $\ell^2(\mathbb{N})$ (the lines of the matrix) such that a realization the weak-greedy algorithm exactly leads to this matrix.

Note that since $u_i \in \mathcal{K}$ for all i, there exists a m dimensional space W of $\ell^2(\mathbb{N})$ such that each row of A is approximated by W with accuracy $d_m := d_m(\mathcal{K})_V$ in $\ell^2(\mathbb{N})$.

The same holds for any submatrix of A by restriction of W.

Lemma : let $G = (g_{i,j})$ be a $K \times K$ matrix with rows $\mathbf{g}_1, \dots, \mathbf{g}_K$. If W is any m dimensional subspace of $\mathbb{R}^K$ for some $0 < m \leq K$, and P is the orthogonal projection from $\mathbb{R}^K$ onto W, then

$$\det(G)^2 \leq \Big(\frac{1}{m}\sum_{i=1}^{K} \|P\mathbf{g}_i\|_{\ell^2}^2\Big)^m \Big(\frac{1}{K-m}\sum_{i=1}^{K} \|\mathbf{g}_i - P\mathbf{g}_i\|_{\ell^2}^2\Big)^{K-m}.$$

We apply this lemma to $K \times K$ matrix $G = (g_{i,j})$ which is formed by the rows and columns of A with indices $N+1, \dots, N+K$. By Property $(P2)$, we obtain

$$\|P\mathbf{g}_i\|_{\ell^2} \leq \|\mathbf{g}_i\|_{\ell^2} \leq \sigma_{N+1}, \quad i = 1, \dots, K,$$

We also have,

$$\|\mathbf{g}_i - P\mathbf{g}_i\|_{\ell^2} \leq d_m, \quad i = 1, \dots, K.$$

It follows that

$$\gamma^{2K} \prod_{i=1}^{K} \sigma_{N+i}^2 \leq \prod_{i=1}^{K} a_{N+i,N+i}^2 = \det(G)^2 \leq \Big(\frac{K}{m}\Big)^m \Big(\frac{K}{K-m}\Big)^{K-m} \sigma_{N+1}^{2m} d_m^{2K-2m}.$$

Application : exponential rates

We take $N = 0$, $K = n$ and any $1 \le m < n$. Using the monotonicity of $(\sigma_n)_{n \ge 0}$ and $\sigma_1 \le \sigma_0 \le d_0$, we obtain

$$\sigma_n^{2n} \le \prod_{j=1}^{n} \sigma_j^2 \le \gamma^{-2n} \left(\frac{n}{m}\right)^m \left(\frac{n}{n-m}\right)^{n-m} d_m^{2n-2m} d_0^{2m}.$$

Since $x^{-x}(1-x)^{x-1} \le 2$ for $0 < x < 1$, it follows that

$$\sigma_n \le \sqrt{2}\gamma^{-1} d_0^{\frac{m}{n}} \min_{1 \le m < n} d_m^{\frac{n-m}{n}}, \quad n \ge 1,$$

and particular

$$\sigma_{2n} \le \gamma^{-1}\sqrt{2 d_0 d_n}.$$

From this, one easily derive

$$\sup_{n \ge 1} e^{an^s} d_n(\mathcal{K})_V < \infty \Rightarrow \sup_{n \ge 1} e^{bn^s} \sigma_n(\mathcal{K})_V < \infty.$$

We take $N = K = n$ and any $1 \leq m < n$. Using the monotonicity of $(\sigma_n)_{n\geq 0}$, we obtain

$$\sigma_{2n}^{2n} \leq \prod_{j=n+1}^{2n} \sigma_j^2 \leq \gamma^{-2n} \left(\frac{n}{m}\right)^m \left(\frac{n}{n-m}\right)^{n-m} \sigma_n^{2m} d_m^{2n-2m}.$$

In the case $n = 2k$ and $m = k$ we have for any positive integer k,

$$\sigma_{4k} \leq \sqrt{2}\gamma^{-1}\sqrt{\sigma_{2k} d_k}.$$

Assuming that $d_n \leq C_0 n^{-s}$ for all $n \geq 1$ and $d_0 \leq C_0$, we obtain by induction that for all $j \geq 0$ and $n = 2^j$,

$$\sigma_n = \sigma_{2^j} \leq C 2^{-sj} \leq n^{-s}, \quad C := 2^{3s+1}\gamma^{-2} C_0.$$

Indeed, the above obviously holds for $j = 0$ or 1 since for these values, we have $\sigma_{2^j} \leq \sigma_0 = d_0 \leq C_0 \leq C2^{-sj}$. Assuming its validity for some $j \geq 1$, we find that

$$\begin{aligned}
\sigma_{2^{j+1}} &\leq \sqrt{2}\gamma^{-1}\sqrt{\sigma_{2^j} d_{2^{j-1}}} \\
&\leq \gamma^{-1} 2^{\frac{3s}{2}} \sqrt{2CC_0} 2^{-s(j+1)} \\
&= \sqrt{C}\sqrt{2^{3s+1} C_0 \gamma^{-2}} 2^{-s(j+1)} = C 2^{-s(j+1)},
\end{aligned}$$

where we have used the definition of C. For values $2^j < n < 2^{j+1}$, we obtain the general result by writing

$$\sigma_n \leq \sigma_{2^j} \leq C 2^{-sj} \leq 2^s C n^{-s} = C_1 n^{-s}.$$

Proof of the key lemma

Let $G = (g_{i,j})$ be a $K \times K$ matrix with rows $\mathbf{g}_1, \dots, \mathbf{g}_K$, and let W be any m dimensional subspace of $\mathbb{R}^K$ for some $0 < m \leq K$ with projector P. Take $\varphi_1, \dots, \varphi_m$ any orthonormal basis for the space W and complete it into an orthonormal basis $\varphi_1, \dots, \varphi_K$ for $\mathbb{R}^K$.

We denote by Φ the $K \times K$ orthogonal matrix whose j-th column is φ_j, then the matrix $C := G\Phi$ has entries $c_{i,j} = \langle \mathbf{g}_i, \varphi_j \rangle$. We have

$$\det(G)^2 = \det(C)^2.$$

With $\mathbf{c}_j$ the j-th column of C, the arithmetic-geometric mean inequality yields

$$\prod_{j=1}^{m} \|\mathbf{c}_j\|_{\ell^2}^2 \leq \Big(\frac{1}{m}\sum_{j=1}^{m} \|\mathbf{c}_j\|_{\ell^2}^2\Big)^m = \Big(\frac{1}{m}\sum_{j=1}^{m}\sum_{i=1}^{K} \langle \mathbf{g}_i, \varphi_j \rangle^2\Big)^m = \Big(\frac{1}{m}\sum_{i=1}^{K} \|P\mathbf{g}_i\|_{\ell^2}^2\Big)^m.$$

Likewise, since φ_j is orthogonal to W when $j > m$,

$$\prod_{j=m+1}^{K} \|\mathbf{c}_j\|_{\ell^2}^2 \leq \Big(\frac{1}{K-m}\sum_{j=m+1}^{K} \|\mathbf{c}_j\|_{\ell^2}^2\Big)^{K-m} = \Big(\frac{1}{K-m}\sum_{i=1}^{K} \|\mathbf{g}_i - P\mathbf{g}_i\|_{\ell^2}^2\Big)^{K-m}.$$

We conclude by using Hadamard's inequality, which gives

$$\det(C)^2 \leq \prod_{j=1}^{K} \|\mathbf{c}_j\|_{\ell^2}^2 \leq \Big(\frac{1}{m}\sum_{i=1}^{K} \|P\mathbf{g}_i\|_{\ell^2}^2\Big)^m \Big(\frac{1}{K-m}\sum_{i=1}^{K} \|\mathbf{g}_i - P\mathbf{g}_i\|_{\ell^2}^2\Big)^{K-m}.$$

Failure of linear reduced modeling

Linear reduced modeling for parametrized hyperbolic PDEs suffers from a slow decay of Kolmogorov n-width.

Simple example : consider the univariate linear transport equation

$$\partial_t u + a\partial_x u = 0,$$

with constant velocity $a \in \mathbb{R}$ and initial condition $u_0 = u(x, 0) = \chi_{[0,1]}(x)$.

Parametrize the solution by the velocity $a \in [a_{\min}, a_{\max}]$ and consider the solution manifold at final time $T = 1$,

$$\mathcal{H} = \{\chi_{[a,a+1]} \, : \, a \in [a_{\min}, a_{\max}]\}.$$

It can be proved that for $1 \leq p < \infty$,

$$d_n(\mathcal{H})_{L^p} \sim n^{-1/p}.$$

In particular, we cannot hope for a good performance of reduced basis methods (not better than piecewise constant approximation on uniform meshes).

Nonlinear approximation

For such problems, one expects improved performance by nonlinear methods.

Non-linear approximation : the function u is approximated by simpler function $v \in \Sigma_n$ that can be described by $\mathcal{O}(n)$ parameters, however Σ_n is not a linear space.

- Rational fractions : $\Sigma_n = \left\{ \frac{p}{q} \,;\, p, q \in \mathbb{P}_n \right\}$.

- Best n-term / sparse approximation in a basis $(e_k)_{k \geq 1}$: pick approximation from the set $\Sigma_n = \{\sum_{k \in E} c_k e_k \;:\; \#(E) \leq n\}$.

- Piecewise polynomials, splines, finite elements on meshes generated after n step of adaptive refinement (select and split an element in the current partition).

- Neural networks : functions $v : \mathbb{R}^d \to \mathbb{R}^m$ of the form

$$v = A_k \circ \sigma \circ A_{k-1} \circ \sigma \circ A_{k-2} \circ \cdots \circ \sigma \circ A_1,$$

where $A_j : \mathbb{R}^{d_j} \to \mathbb{R}^{d_{j+1}}$ is affine and σ is a nonlinear (rectifier) function applied componentwise, for example $\sigma(x) = RELU(x) = \max\{x, 0\}$. Here Σ_n is the set of such functions when the total number of parameters does not exceed n.

Is there a natural notion of width describing optimal nonlinear approximation ?

Library widths

A library $\mathcal{L}_n$ is a finite collection of linear spaces $V_n \subset V$ of dimension at most n.

We approximate u by picking a space from $\mathcal{L}_n$, resulting in the error

$$e(u, \mathcal{L}_n)_V = \min_{V_n \in \mathcal{L}_n} \min_{v \in V_n} \|u - v\|_V.$$

Temlyakov (1998) defines the library width

$$d_{N,n}(\mathcal{K})_V := \inf_{\#(\mathcal{L}_n) \leq N} \max_{u \in \mathcal{K}} e(u, \mathcal{L}_n)_V.$$

Note that $d_{1,n} = d_n$.

The interesting regime is when $N >> n$. Typical choices that have been studied are $N = A^n$ or $N = n^{an}$ for some $A > 1$ or $a > 0$.

This type of width is well adapted to describe optimality for best n-term approximation or adaptive refinements, but not for neural networks or rational fractions.

Manifold widths

Naive idea : replace linear spaces V_n of dimension n by smooth manifolds $\mathcal{M}_n$ of dimension n in the definition of d_n.

This would lead to the quantity

$$\inf_{\dim(\mathcal{M}_n)=n} \max_{u\in\mathcal{K}} \min_{v\in V} \|u-v\|_V,$$

However its value is 0 even for $n=1$: space filling curves !

DeVore-Howard-Michelli (1989) : impose continuous selection by defining

$$\delta_n(\mathcal{K})_V := \inf_{D,E} \max_{u\in\mathcal{K}} \|u-D(E(u))\|_V,$$

where infimum is taken on all continuous pairs $E: V \to \mathbb{R}^n$ and $D: \mathbb{R}^n \to V$.

Naive idea : replace linear spaces V_n of dimension n by smooth manifolds $\mathcal{M}_n$ of dimension n in the definition of d_n.

This would lead to the quantity

$$\inf_{\dim(\mathcal{M}_n)=n} \max_{u\in\mathcal{K}} \min_{v\in V} \|u-v\|_V,$$

However its value is 0 even for $n=1$: space filling curves !

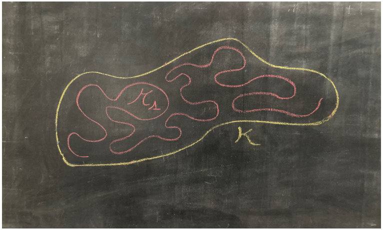

DeVore-Howard-Michelli (1989) : impose continuous selection by defining

$$\delta_n(\mathcal{K})_V := \inf_{D,E} \max_{u\in\mathcal{K}} \|u-D(E(u))\|_V,$$

where infimum is taken on all continuous pairs $E : V \to \mathbb{R}^n$ and $D : \mathbb{R}^n \to V$.

Naive idea : replace linear spaces V_n of dimension n by smooth manifolds $\mathcal{M}_n$ of dimension n in the definition of d_n.

This would lead to the quantity

$$\inf_{\dim(\mathcal{M}_n)=n} \max_{u\in\mathcal{K}} \min_{v\in V} \|u-v\|_V,$$

However its value is 0 even for $n=1$: space filling curves !

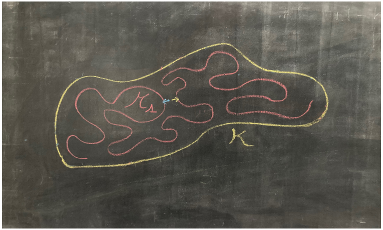

DeVore-Howard-Michelli (1989) : impose continuous selection by defining

$$\delta_n(\mathcal{K})_V := \inf_{D,E} \max_{u\in\mathcal{K}} \|u-D(E(u))\|_V,$$

where infimum is taken on all continuous pairs $E: V \to \mathbb{R}^n$ and $D: \mathbb{R}^n \to V$.

Estimating nonlinear width

Both library and manifold widths match known rates of nonlinear approximation (DeVore-Popov, 1980-1990's) by wavelets or adaptive finite elements : if $V = L^p(D)$ and $\mathcal{K} = \mathcal{U}(B^s_{q,q}(D))$ for $\frac{1}{q} < \frac{1}{p} + \frac{s}{d}$, one has

$$d_{n,N}(\mathcal{K})_V \sim \delta_n(\mathcal{K})_V \sim n^{-s/d}.$$

Upper bounds obtained by these classical nonlinear approximation results.

Library widths satisfy Carl's inequality (for the regimes $N = A^n$ or $N = n^{an}$).

$$(n+1)^s \varepsilon_n(\mathcal{K})_V \leq C_s \sup_{m=0,\dots,n} (m+1)^s d_{m,N}(\mathcal{K})_V, \quad n \geq 0.$$

Manifold widths do not satisty Carl's inequality but are bounded by below by Bernstein widths (by the Borsuk-Ulam argument).

$$\delta_n(\mathcal{K})_V \geq b_n(\mathcal{K})_V.$$

For example, if $V = L^\infty(I)$ and $\mathcal{K} = \mathcal{U}(\mathrm{Lip}(I))$, one has $\delta_n(\mathcal{K})_V \sim n^{-1}$.

Yarotzki, Shen-Yang-Zhang (2020) : Neural networks approximation of functions in $\mathrm{Lip}(I)$ converge in L^∞ with rate n^{-2} ! Parameter selection cannot be stable.

Stable nonlinear widths

Cohen-DeVore-Petrova-Wojtaszczyk (2020) : for some fixed $L > 1$ define

$$\delta_{n,L}(\mathcal{K})_V := \inf_{D,E} \max_{u \in \mathcal{K}} \|u - D(E(u))\|_V,$$

where the infimum is taken on all pairs $E : V \to \mathbb{R}^n$ and $D : \mathbb{R}^n \to V$, that satisfy

$$\|D(x) - D(y)\|_V \leq L\|x - y\|_n \quad \text{and} \quad \|E(u) - E(v)\|_n \leq L\|u - v\|_V, \quad x, y \in \mathbb{R}^n, \; u, v \in V.$$

Here $\|\cdot\|_n$ is an arbitrary norm on $\mathbb{R}^n$.

This notion of stable width now satisfies Carl's inequality : for any $L > 1$,

$$(n+1)^s \varepsilon_n(\mathcal{K})_V \leq C_s \sup_{m=0,\dots,n} (m+1)^s \delta_{m,L}(\mathcal{K})_V, \quad n \geq 0.$$

in addition to the lower bound by Gelfand width $\delta_{n,L}(\mathcal{K})_V \geq b_n(\mathcal{K})_V$

Open problem : with $V = L^p(D)$ and $\mathcal{K} = \mathcal{U}(B^s_{q,q}(D))$ for $\frac{1}{q} < \frac{1}{p} + \frac{s}{d}$, do we have $\delta_{n,L}(\mathcal{K})_V \sim n^{-s/d}$? Positive answer known only when $p = 2$.

When V is a Hilbert space, stable widths are strongly tied to entropy numbers.

Theorem : Let V be a Hilbert space, then for any $L > 1$, there exists a constant $c = c(L)$ such that, for any compact set $\mathcal{K}$,

$$\delta_{cn,L}(\mathcal{K})_V \leq 3\varepsilon_n(\mathcal{K})_V.$$

With $L = 2$ one can take $c = 26$.

Together with Carl's inequality, this means that

$$\sup_{n\geq 0} n^s \delta_{n,L}(\mathcal{K})_V < \infty \iff \sup_{n\geq 0} n^s \varepsilon_n(\mathcal{K})_V < \infty,$$

for all $s > 0$.

We do not know if this result holds for Banach spaces. Proof for Hilbert spaces :

1. Consider $\mathcal{N}$ an ε_n-net of $\mathcal{K}$ with $\#(\mathcal{N}) = 2^n$.
2. Johnson-Lindenstrauss projection as encoder : $E = P_W$ where $\dim(W) \leq cn$

$$L^{-1}\|u_i - u_j\|_V \leq \|P_W(u_i - u_j)\|_V \leq \|u_i - u_j\|_V, \quad u_i, u_j \in \mathcal{N}.$$

3. This gives an exact decoding map that is L-Lipschitz from $P_W\mathcal{N}$ to $\mathcal{N}$.
4. Extend this map from $W \sim \mathbb{R}^{cn}$ to V with same Lipschitz constant (Kirszbraun).

Stable width of solution manifolds

For the linear transport equation manifold $\mathcal{H} = \Big\{\chi_{[a,a+1]} \,:\, a \in [a_{\min}, a_{\max}]\Big\}$ it is easily established that entropy numbers in L^p spaces have exponential decay

$$\varepsilon_n(\mathcal{H})_{L^p} \leq C \exp(-cn), \quad n \geq 0.$$

This implies in particular that $\delta_{n,L}(\mathcal{H})_{L^2} \leq \tilde{C} \exp(-\tilde{c}n)$ while $d_n(\mathcal{H})_{L^2} \sim n^{-1/2}$.

Similar results hold for manifolds resulting from more general hyperbolic equations.

A general result : if $F : V_1 \to V_2$ is a L-Lipschitz mapping between Banach spaces, then an ε-net of $\mathcal{K}_1 \subset V_1$ is mapped into an $L\varepsilon$-net of $\mathcal{K}_2 := F(\mathcal{K}_1)$ and therefore

$$\varepsilon_n(\mathcal{K}_2)_{V_2} \leq L\varepsilon_n(\mathcal{K}_1)_{V_1}, \quad n \geq 0.$$

This implies in particular that when V_2 is a Hilbert space

$$\sup_{n\geq 0} n^s \delta_{n,L}(\mathcal{K}_1)_{V_1} < \infty \implies \sup_{n\geq 0} n^s \delta_{n,L}(\mathcal{K}_2)_{V_2} < \infty.$$

Benchmark : develop concrete stable numerical methods that meet these rates.

Approximation of high-dimensional parametric/stochastic PDEs

We are interested in PDE's of the general form

$$\mathcal{P}(u, y) = 0,$$

where $\mathcal{P}$ is a partial differential operator, u is the unknown and $y = (y_j)_{j=1,\dots,d}$ is a parameter vector of dimension $d >> 1$ or $d = \infty$ ranging in some domain Y.

We assume well-posedness of the solution in some Banach space V for every $y \in Y$,

$$y \mapsto u(y)$$

is the solution map from Y to V.

Solution manifold $\mathcal{K} := \{u(y) \; : \; y \in Y\} \subset V$.

The parameters may be deterministic (control, optimization, inverse problems) or random (uncertainty modeling and quantification, risk assessment). In the second case the solution $u(y)$ is a V-valued random variable.

Objective : numerical approximation to the parameter to solution map $y \mapsto u(y)$.

Related objectives : numerical approximation of scalar quantities of interest $y \mapsto Q(y) = Q(u(y))$, or of averaged quantities $\overline{u} = \mathbb{E}(u(y))$ or $\overline{Q} = \mathbb{E}(Q(y))$.

Guiding example : elliptic PDEs

We consider the steady state diffusion equation

$$-\mathrm{div}(a\nabla u) = f \quad \text{on} \quad D \subset \mathbf{R}^{\mathrm{m}} \quad \text{and} \quad \mathrm{u}_{|\partial \mathrm{D}} = 0,$$

set on a domain $D \subset \mathbb{R}^m$, where $f = f(x) \in L^2(D)$ and $a \in L^\infty(D)$

Lax-Milgram lemma : assuming $a_{\min} := \min_{x\in D} a(x) > 0$, unique solution $u \in V = H_0^1(D)$ with

$$\|u\|_V := \|\nabla u\|_{L^2(D)} \leq \frac{1}{a_{\min}} \|f\|_{V'}.$$

Proof of the estimate : multiply equation by u and integrate

$$a_{\min}\|u\|_V^2 \leq \int_D a\nabla u \cdot \nabla u = -\int_D u\,\mathrm{div}(a\nabla u) = \int_D uf \leq \|u\|_V \|f\|_{V'}.$$

We may extend this theory to the solution of the weak (or variational) formulation

$$\int_D a\nabla u \cdot \nabla v = \langle f, v\rangle, \quad v \in V = H_0^1(D),$$

if $f \in V' = H^{-1}(D)$. If $f \in L^2(D)$ one has $\|f\|_{V'} \leq C_P \|f\|_{L^2}$ by Cauchy-Schwarz and Poincaré inequalities.

Parametrization

Assume diffusion coefficients in the form of an expansion

$$a = a(y) = \bar{a} + \sum_{j\geq 1} y_j \psi_j, \quad y = (y_j)_{j\geq 1} \in U,$$

with $d >> 1$ or $d = \infty$ terms, where $\bar{a}$ and $(\psi_j)_{j\geq 1}$ are functions from L^∞,

Note that $a(y)$ is a function for each given y. We may also write

$$a = a(x,y) = \bar{a}(x) + \sum_{j\geq 1} y_j \psi_j(x), \quad x \in D, y \in Y,$$

where x and y are the spatial and parametric variable, respectively. Likewise, the corresponding solution $u(y)$ is a function $x \mapsto u(y,x)$ for each given y. We often ommit the reference to the spatial variable.

Up to a change of variable, we assume that all y_j range in $[-1,1]$, therefore

$$y \in Y = [-1,1]^d \text{ or } [-1,1]^{\mathbb{N}}.$$

Uniform ellipticity assumption :

$$(UEA) \qquad 0 < r \leq a(x,y) \leq R, \;\; x \in D, \; y \in Y \qquad \text{(equivalently } \left\| \frac{\sum_{j\geq 1} |\psi_j|}{\bar{a}} \right\|_{L^\infty} < 1).$$

Then the solution map is bounded from Y to $V := H_0^1(D)$, that is, $u \in L^\infty(Y,V)$:

$$\|u(y)\|_V \leq M := \frac{\|f\|_{V'}}{r}, \;\; y \in Y,$$

Example of parametrization : piecewise constant coefficients

Assume that a is piecewise constant over a partition $\{D_1, \dots, D_d\}$ of D, and such that on each D_j the value of a varies on $[c - c_j, c + c_j]$ for some $c > 0$ and $0 < c_j < c$.

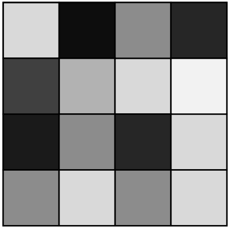

Then a natural parametrization is

$$a(y) = \bar{a} + \sum_{j=1}^{d} y_j \psi_j, \quad \bar{a} = c, \quad \psi_j = c_j \chi_{D_j},$$

with $y = (y_j)_{j=1,\dots,d} \in Y = [-1, 1]^d$.

How to defeat the curse of dimensionality ?

The map $y \mapsto u(y)$ is high dimensional, or even infinite dimensional $y = (y_j)_{j\geq 1}$.

We are thus facing the curse of dimensionality when trying to approximate it with conventional discretization tools in the y variable (Fourier series, finite elements).

A general function of d variable with m bounded derivatives cannot be approximated in L^∞ with rate better than $n^{-m/d}$ where n is the number of degrees of freedom.

A possible way out : exploit anisotropic features in the function $y \mapsto u(y)$.

The PDE is parametrized by a function a (diffusion coefficient, velocity, domain boundary) and y_j are the coordinates of a in a certain basis representation $a = \overline{a} + \sum_{j\geq 1} y_j \psi_j$.

If the ψ_j decays as $j \to +\infty$ (for instance if a has some smoothness) then the variable y_j are less active for large j.

We shall see that in certain relevant instances, this mechanism allows to break the curse of dimensionality by using suitable expansions : we obtain approximation rates $\mathcal{O}(n^{-s})$ that are independent of d in the sense that they hold when $d = \infty$.

One key tool for obtaining such result is the concept of sparse approximation.

Let $\mathcal{F}$ be a countable set and $\mathbf{u} = (u_\nu)_{\nu\in\mathcal{F}}$ be sequence. The best n-sparse approximation of $\mathbf{u}$ in ℓ^q norm is the sequence $\mathbf{u}_n$ obtained by keeping the n largest u_ν and setting to 0 all others entries.

$$\|\mathbf{u}-\mathbf{u}_n\|_{\ell^q} = \Big(\sum_{k>n}(u_k^*)^q\Big)^{1/q}$$

where $(u_k^*)_{k\geq 1}$ is the decreasing rearrangement of the sequence $(|u_\nu|)$

Theorem (Stechkin's lemma) : if $0 < p < q$ and $\mathbf{u} = (u_\lambda) \in \ell^p(\mathcal{F})$, one has

$$\|\mathbf{u}-\mathbf{u}_n\|_{\ell^q} \leq Cn^{-s}, \qquad s = \frac{1}{p}-\frac{1}{q},$$

and $C = \|\mathbf{u}\|_{\ell^p}$.

Proof : we combine

$$\|\mathbf{u}-\mathbf{u}_n\|_{\ell^q} = (\sum_{k>n}|u_k^*|^q)^{\frac{1}{q}} = (\sum_{k>n}|u_k^*|^{q-p}|u_k^*|^p)^{\frac{1}{q}} \leq |u_{n+1}^*|^{1-\frac{p}{q}}\|\mathbf{u}\|_{\ell^p}^{\frac{p}{q}}$$

and

$$(n+1)|u_{n+1}^*|^p \leq \sum_{k=1}^{n+1}|u_k^*|^p \leq \|\mathbf{u}\|_{\ell^p}^p.$$

We consider the expansion of $u(y) = \sum_{\nu \in \mathcal{F}} u_\nu y^\nu$, where

$$y^\nu := \prod_{j \geq 1} y_j^{\nu_j} \quad \text{and} \quad u_\nu := \frac{1}{\nu!} \partial^\nu u_{|y=0} \in V \quad \text{with} \quad \nu! := \prod_{j \geq 1} \nu_j! \quad \text{and} \quad 0! := 1.$$

where $\mathcal{F}$ is the set of all finitely supported sequences of integers (finitely many $\nu_j \neq 0$). The sequence $(t_\nu)_{\nu \in \mathcal{F}}$ is indexed by countably many integers.

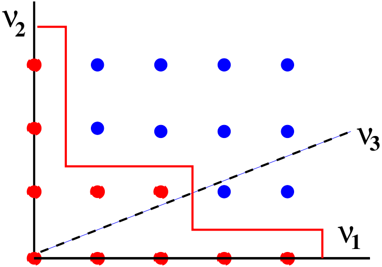

Objective : identify a set $\Lambda \subset \mathcal{F}$ with $\#(\Lambda) = n$ such that u is well approximated by the polynomial partial expansion

$$u_\Lambda(y) := \sum_{\nu \in \Lambda} t_\nu y^\nu.$$

Best n-term approximation

A-priori choices for Λ have been proposed, e.g. (anisotropic) sparse grid defined by restrictions of the type $\sum_j \alpha_j \nu_j \leq A(n)$ or $\prod_j (1+\beta_j \nu_j) \leq B(n)$.

Instead we want to choose Λ optimally adapted to u. By triangle inequality we have

$$\|u - u_\Lambda\|_{L^\infty(Y,V)} = \sup_{y\in Y} \|u(y) - u_\Lambda(y)\|_V \leq \sup_{y\in Y} \sum_{\nu\notin\Lambda} \|u_\nu y^\nu\|_V = \sum_{\nu\notin\Lambda} \|u_\nu\|_V$$

Best n-term approximation in $\ell^1(\mathcal{F})$ norm : use $\Lambda = \Lambda_n$ index set of n largest $\|u_\nu\|_V$.

Stechkin lemma : if $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$ for some $p < 1$, then for this Λ_n,

$$\sum_{\nu\notin\Lambda_n} \|u_\nu\|_V \leq Cn^{-s}, \quad s := \frac{1}{p} - 1, \quad C := \|(\|u_\nu\|_V)\|_{\ell^p}.$$

Question : do we have $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$ for some $p < 1$?

Theorem (Cohen-DeVore-Schwab, 2011) : under the uniform ellipticity assumption (UEA), then for any $p < 1$,

$$(\|\psi_j\|_{L^\infty})_{j>0} \in \ell^p(\mathbb{N}) \implies (\|u_\nu\|_V)_{\nu \in \mathcal{F}} \in \ell^p(\mathcal{F}).$$

(i) The Taylor expansion of $u(y)$ inherits the sparsity properties of the expansion of $a(y)$ into the ψ_j.

(ii) We approximate $u(y)$ in $L^\infty(U, V)$ with algebraic rate $\mathcal{O}(n^{-s})$ despite the curse of (infinite) dimensionality, due to the fact that y_j is less influencial as j gets large.

(iii) The solution manifold $\mathcal{M} := \{u(y)\ ;\ y \in U\}$ is uniformly well approximated by the n-dimensional space $V_n := \mathrm{span}\{t_\nu\ :\ \nu \in \Lambda_n\}$. Its n-width satisfies the bound

$$d_n(\mathcal{M})_V \leq \max_{y \in U} \mathrm{dist}(u(y), V_n)_V \leq \max_{y \in U} \|u(y) - u_{\Lambda_n}(y)\|_V \leq Cn^{-s}.$$

Such approximation rates cannot be proved for the usual a-priori choices of Λ.

Same result for more general linear equations $Au = f$ with affine operator dependance $A = \overline{A} + \sum_{j \geq 1} y_j A_j$ uniformly invertible over $y \in U$, and $(\|A_j\|_{V \to W})_{j \geq 1} \in \ell^p(\mathbb{N})$.

Similar results for other models : parabolic evolution, saddle-point problems, some nonlinear problems, but not hyperbolic problems.

Estimates on $\|u_\nu\|_V$ by complex analysis : extend $u(y)$ to $u(z)$ with $z = (z_j) \in \mathbb{C}^{\mathbb{N}}$.

Uniform ellipticity $\sum_{j\geq 1} |\psi_j| \leq \overline{a} - r$ implies that with $a(z) = \overline{a} + \sum_{j\geq 1} z_j \psi_j$,

$$0 < r \leq \Re(a(x,z)) \leq |a(x,z)| \leq 2R, \quad x \in D,$$

for all $z \in \mathcal{U} := \{|z| \leq 1\}^{\mathbb{N}} = \otimes_{j\geq 1}\{|z_j| \leq 1\}$.

Lax-Milgram theory applies : $\|u(z)\|_V \leq M = \frac{\|f\|_{V^*}}{r}$ for all $z \in \mathcal{U}$.

The function $z \mapsto u(z)$ is holomorphic in each variable z_j at any $z \in \mathcal{U}$: its first derivative $\partial_{z_j} u(z)$ is the unique solution to

$$\int_D a(z)\nabla \partial_{z_j} u(z) \cdot \nabla v = -\int_D \psi_j \nabla u(z) \cdot \nabla v, \quad v \in V.$$

Note that ∇ is with respect to spatial variable $x \in D$.

Extended domains of holomorphy : if $\rho = (\rho_j)_{j\geq 0}$ is any positive sequence such that for some $\delta > 0$

$$\sum_{j\geq 1} \rho_j |\psi_j(x)| \leq \overline{a}(x) - \delta, \quad x \in D,$$

then u is holomorphic with uniform bound $\|u(z)\| \leq C_\delta = \frac{\|f\|_{V^*}}{\delta}$ in the polydisc

$$\mathcal{U}_\rho := \otimes_{j\geq 1}\{|z_j| \leq \rho_j\},$$

If $\delta < r$, we can take $\rho_j > 1$.

Use Cauchy formula. In 1 complex variable if $z \mapsto u(z)$ is holomorphic and bounded in a neighbourhood of disc $\{|z| \leq b\}$, then for all z in this disc

$$u(z) = \frac{1}{2i\pi} \int_{|z'|=b} \frac{u(z')}{z - z'} dz',$$

which leads by n differentiation at $z = 0$ to $|u^{(n)}(0)| \leq n! b^{-n} \max_{|z| \leq b} |u(z)|$.

This yields exponential convergence rate $b^{-n} = \exp(-cn)$ of Taylor series for 1-d holomorphic functions. Curse of dimensionality : in d dimension, this yields sub-exponential rate $\exp(-cn^{1/d})$ where n is the number of retained terms.

Recursive application of this to all variables z_j such that $\nu_j \neq 0$, with $b = \rho_j$ gives

$$\|\partial^\nu u_{|z=0}\|_V \leq 2M\nu! \prod_{j \geq 1} \rho_j^{-\nu_j},$$

and thus

$$\|u_\nu\|_V \leq C_\delta \prod_{j \geq 1} \rho_j^{-\nu_j} = 2M\rho^{-\nu},$$

for any sequence $\rho = (\rho_j)_{j \geq 1}$ such that

$$\sum_{j \geq 1} \rho_j |\psi_j(x)| \leq \bar{a}(x) - \frac{r}{2}.$$

Optimization

Since ρ is not fixed we have

$$\|u_\nu\|_V \le 2M \inf\{\rho^{-\nu} \ : \ \rho \ \text{s.t.} \ \sum_{j\ge 1} \rho_j |\psi_j(x)| \le \bar{a}(x) - \frac{r}{2}, \ \ x \in D\}.$$

We do not know the general solution to this problem, except in particular case, for example when the ψ_j have disjoint supports.

Instead design a particular choice $\rho = \rho(\nu)$ satisfying the constraint, for which we prove that

$$(\|\psi_j\|_{L^\infty})_{j\ge 1} \in \ell^p(\mathbb{N}) \implies (\rho(\nu)^{-\nu})_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F}),$$

therefore proving the main theorem.

Assume that the ψ_j have disjoint supports. Then we maximize separately the ρ_j so that

$$\sum_{j\geq 1} \rho_j |\psi_j(x)| \leq \overline{a}(x) - \frac{r}{2}, \quad x \in D,$$

which leads to

$$\rho_j := \min_{x\in D} \frac{\overline{a}(x) - \frac{r}{2}}{|\psi_j(x)|}.$$

We have,

$$\|u_\nu\|_V \leq 2M\rho^{-\nu} = 2Mb^\nu,$$

where $b = (b_j)$ and

$$b_j := \rho_j^{-1} = \max_{x\in D} \frac{|\psi_j(x)|}{\overline{a}(x) - \frac{r}{2}} \leq \frac{\|\psi_j\|_{L^\infty}}{R - \frac{r}{2}}.$$

Therefore $b \in \ell^p(\mathbb{N})$. From (UEA), we have $|\psi_j(x)| \leq \overline{a}(x) - r$ and thus $\|b\|_{\ell^\infty} < 1$.

We finally observe that

$$b \in \ell^p(\mathbb{N}) \text{ and } \|b\|_{\ell^\infty} < 1 \iff (b^\nu)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F}).$$

Proof : factorize

$$\sum_{\nu\in\mathcal{F}} b^{p\nu} = \prod_{j\geq 1} \sum_{n\geq 0} b_j^{pn} = \prod_{j\geq 1} \frac{1}{1 - b_j^p}.$$

Recursive computation of Taylor coefficients

By differentiating with respect to the y_j in the variational formulation, we find a recursive formula for the Taylor coefficients : with $e_j = (0, \dots, 0, 1, 0, \dots)$ the Kroeneker sequence of index j, the coefficient u_ν is solution to

$$\int_D \bar{a} \nabla u_\nu \nabla v = - \sum_{j:\ \nu_j \neq 0} \int_D \psi_j \nabla u_{\nu - e_j} \nabla v, \quad v \in V.$$

This will lead to improved estimates. We introduce the quantities

$$d_\nu := \int_D \bar{a} |\nabla u_\nu|^2 \quad \text{and} \quad d_{\nu,j} := \int_D |\psi_j| \, |\nabla u_\nu|^2.$$

Recall that (UEA) implies that $\left\| \frac{\sum_{j \geq 1} |\psi_j|}{\bar{a}} \right\|_{L^\infty(D)} \leq \theta < 1$. In particular

$$\sum_{j \geq 1} d_{\nu,j} \leq \theta d_\nu.$$

We use here the equivalent norm $\|v\|_V^2 := \int_D \bar{a} |\nabla v|^2$.

Lemma : under (UEA), one has $\sum_{\nu \in \mathcal{F}} d_\nu = \sum_{\nu \in \mathcal{F}} \|u_\nu\|_V^2 < \infty$.

Proof

Taking $v = u_\nu$ in the recursion gives

$$d_\nu = \int_D \bar{a}|\nabla u_\nu|^2 = -\sum_{j:\ \nu_j \neq 0} \int_D \psi_j \nabla u_{\nu - e_j} \nabla u_\nu .$$

Apply Young's inequality on the right side gives

$$d_\nu \leq \sum_{j:\ \nu_j \neq 0} \left(\frac{1}{2} \int_D |\psi_j|\, |\nabla u_\nu|^2 + \frac{1}{2} \int_D |\psi_j|\, |\nabla u_{\nu - e_j}|^2 \right) = \frac{1}{2} \sum_{j:\ \nu_j \neq 0} d_{\nu,j} + \frac{1}{2} \sum_{j:\ \nu_j \neq 0} d_{\nu - e_j, j} .$$

The first sum is bounded by θd_ν, therefore

$$\left(1 - \frac{\theta}{2}\right) d_\nu \leq \frac{1}{2} \sum_{j:\ \nu_j \neq 0} d_{\nu - e_j, j} .$$

Now summing over all $|\nu| = k$ gives

$$\left(1 - \frac{\theta}{2}\right) \sum_{|\nu|=k} d_\nu \leq \frac{1}{2} \sum_{|\nu|=k} \sum_{j:\ \nu_j \neq 0} d_{\nu - e_j, j} = \frac{1}{2} \sum_{|\nu|=k-1} \sum_{j \geq 1} d_{\nu,j} \leq \frac{\theta}{2} \sum_{|\nu|=k-1} d_\nu .$$

Therefore $\sum_{|\nu|=k} d_\nu \leq \kappa \sum_{|\nu|=k-1} d_\nu$ with $\kappa := \frac{\theta}{2-\theta} < 1$, and thus $\sum_{\nu \in \mathcal{F}} d_\nu < \infty$.

Rescaling

Now let $\rho = (\rho_j)_{j\geq 1}$ be any sequence with $\rho_j > 1$ such that $\sum_{j\geq 1} \rho_j|\psi_j| \leq \bar{a} - \delta$ for some $\delta > 0$, or equivalently such that $\left\| \frac{\sum_{j\geq 1} \rho_j|\psi_j|}{\bar{a}} \right\|_{L^\infty(D)} \leq \theta < 1$.

Consider the rescaled solution map $\tilde{u}(y) = u(\rho y)$ where $\rho y := (\rho_j y_j)_{j\geq 1}$ which is the solution of the same problem as u with ψ_j replaced by $\rho_j \psi_j$.

Since (UEA) holds for for these rescaled functions, the previous lemma shows that

$$\sum_{\nu\in\mathcal{F}} \|\tilde{u}_\nu\|_V^2 < \infty,$$

where

$$\tilde{u}_\nu := \frac{1}{\nu!}\partial^\nu \tilde{u}(0) = \frac{1}{\nu!}\rho^\nu \partial^\nu u(0) = \rho^\nu u_\nu.$$

This therefore gives the weighted ℓ^2 estimate

$$\sum_{\nu\in\mathcal{F}} (\rho^\nu \|u_\nu\|_V)^2 \leq C < \infty.$$

In particular, we retrieve the estimate $\|u_\nu\|_V \leq C\rho^{-\nu}$ that can be obtained more directly by the complex variable approach, using Cauchy formula, however the above estimate is stronger.

Applying Hölder's inequality gives

$$\sum_{\nu\in\mathcal{F}} \|u_\nu\|_V^p \leq \Big(\sum_{\nu\in\mathcal{F}} (\rho^\nu \|u_\nu\|_V)^2\Big)^{p/2} \Big(\sum_{\nu\in\mathcal{F}} \rho^{-q\nu}\Big)^{1-p/2},$$

with $q = \frac{2p}{2-p} > p$, or equivalently $\frac{1}{q} = \frac{1}{p} - \frac{1}{2}$.

The sum in second factor is finite provided that $(\rho_j^{-1})_{j\geq 1} \in \ell^q$. Therefore, the following result holds.

Theorem : Let p and q be such that $\frac{1}{q} = \frac{1}{p} - \frac{1}{2}$. Assume that there exists a sequence $\rho = (\rho_j)_{j\geq 1}$ of numbers larger than 1 such that

$$\sum_{j\geq 1} \rho_j |\psi_j| \leq \overline{a} - \delta,$$

for some $\delta > 0$ and

$$(\rho_j^{-1})_{j\geq 1} \in \ell^q.$$

Then $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$.

Disjoint supports

Assume that the ψ_j have disjoint supports.

Then with $\delta = \frac{r}{2}$, we choose

$$\rho_j := \min_{x \in D} \frac{\overline{a}(x) - \frac{r}{2}}{|\psi_j(x)|} > 1.$$

so that $\sum_{j \geq 1} \rho_j |\psi_j| \leq \overline{a} - \delta$ holds.

We have

$$b_j := \rho_j^{-1} = \frac{|\psi_j(x)|}{\overline{a}(x) - \frac{r}{2}} \leq \frac{\|\psi_j\|_{L^\infty}}{R - \frac{r}{2}}.$$

Thus in this case, the new result gives for any $0 < q < \infty$,

$$(\|\psi_j\|_{L^\infty})_{j \geq 1} \in \ell^q(\mathbb{N}) \implies (\|u_\nu\|_V)_{\nu \in \mathcal{F}} \in \ell^p(\mathcal{F}),$$

with $\frac{1}{q} = \frac{1}{p} - \frac{1}{2}$.

Instead of Taylor, we may consider the tensorized Legendre expansion

$$u(y) = \sum_{\nu\in\mathcal{F}} v_\nu L_\nu(y),$$

where $L_\nu(y) := \prod_{j\geq 1} L_{\nu_j}(y_j)$ and $(L_k)_{k\geq 0}$ are the Legendre polynomials normalized in $L^2\left([-1,1], \frac{dt}{2}\right)$. Thus $(L_\nu)_{\nu\in\mathcal{F}}$ is an orthonormal basis for $L^2(Y,\mu)$ with $\mu := \otimes_{j\geq 1}\frac{dy_j}{2}$ the uniform probability measure and we have $v_\nu = \int_Y u(y)L_\nu(y)d\mu(y)$

Theorem : Let p and q be such that $\frac{1}{q} = \frac{1}{p} - \frac{1}{2}$. Assume that there exists a sequence $\rho = (\rho_j)_{j\geq 1}$ of numbers larger than 1 such that

$$\sum_{j\geq 1} \rho_j|\psi_j| \leq \bar{a} - \delta,$$

for some $\delta > 0$ and

$$(\rho_j^{-1})_{j\geq 1} \in \ell^q.$$

Then $(\|v_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$.

Implies best n-term approximation error $\|u - u_n\|_{L^2(Y,V,\mu)} \leq Cn^{-s}$ with $s = \frac{1}{p} - \frac{1}{2}$.

If y_i are i.i.d. uniform, this implies $\kappa_n(u)_V^2 = \min_{\dim(V_n)=n} \mathbb{E}(\|u - P_{V_n}u\|_V^2) \leq Cn^{-2s}$.

References for day 1

A. Pinkus, *n*-width in Approximation Theory, Springer 2012.

R. DeVore, R. Howard, C. Micchelli, Optimal nonlinear approximation, Manuscripta Mathematica, 1989.

V. Temlyakov, Nonlinear Kolmogorov widths, Mathematical Notes, 1998.

A. Cohen, R. DeVore, G. Petrova, and P. Wojtaszczyk, Optimal stable nonlinear approximation, Foundation of Computational Mathematics, 2021.

Zuowei Shen, Haizhao Yang, and Shijun Zhang, Deep network approximation characterized by number of neurons, Communications in Computational Physics, 2020.

A. Cohen and R. DeVore, Approximation of high-dimensional PDEs, Acta Numerica, 2015.

M. Bachmayr, A. Cohen, and G. Migliorati, "Sparse polynomial approximation of parametric elliptic PDEs - Part I : affine coefficients", M2AN, 2017.

An ubiquitous numerical problem

Reconstruct an unknown multivariate function

$$u : x \mapsto u(x), \quad x = (x_1, \dots, x_d) \in D \subset \mathbb{R}^d,$$

from (possibly noisy) observations $y^i \approx \ell_i(u) \in \mathbb{R}$ for $i = 1, \dots, m$.

Here the ℓ_i are linear forms.

An important case : evaluation $y^i \approx u(x^i)$ at sample points $x^i \in D$ for $i = 1, \dots, m$.

Distinction between two data acquisition settings :

Passive setting : we do not choose the x^i (or the ℓ_i).

Active setting : we choose the x^i (or the ℓ_i).

Optimal recovery

Let V be a general Banach space of functions defined on D, and let $\mathcal{K} \subset V$ a class that describes the prior information on u (for example smoothness).

We define the deterministic optimal recovery numbers

$$r_m^{\det}(\mathcal{K})_V := \inf_{\mathbf{x},\Phi_{\mathbf{x}}} \max_{u\in\mathcal{K}} \|u - \Phi_{\mathbf{x}}(u(x^1),\dots,u(x^m))\|_V,$$

where infimum is taken on all $\mathbf{x} = (x^1,\dots,x^m) \in D^m$ and maps $\Phi_{\mathbf{x}} : \mathbb{R}^m \to V$.

Randomized setting (random sampling) :

$$r_m^{\mathrm{rand}}(\mathcal{K})_V^2 := \inf_{\mathbf{x},\Phi_{\mathbf{x}}} \max_{u\in\mathcal{K}} \mathbb{E}_{\mathbf{x}}(\|u - \Phi_{\mathbf{x}}(u(x^1),\dots,u(x^m))\|_V^2),$$

where infimum is taken on all random variable $\mathbf{x} \in D^m$ and linear $\Phi_{\mathbf{x}} : \mathbb{R}^m \to V$.

Linear recovery : define $\rho_m^{\det}(\mathcal{K})_V$ and $\rho_m^{\mathrm{rand}}(\mathcal{K})_V$ similarly but with $\Phi_{\mathbf{x}}$ linear.

Obviously : $r_m^{\det}(\mathcal{K})_V \le \rho_m^{\det}(\mathcal{K})_V$ and $r_m^{\mathrm{rand}}(\mathcal{K})_V \le \rho_m^{\det}(\mathcal{K})_V$.

Also : $r_m^{\mathrm{rand}}(\mathcal{K})_V \le r_m^{\det}(\mathcal{K})_V$ and $\rho_m^{\mathrm{rand}}(\mathcal{K})_V \le \rho_m^{\det}(\mathcal{K})_V$.

Approximation

Error measure : $\|u - \tilde{u}\|_V$, where $V := L^2(D, \mu)$, or other Banach space of interest.

Most often, the reconstruction $\tilde{u}$ takes place within a family $V_n \subset V$ that can be parametrized by $n \leq m$ numbers.

So it is relevant to compare $\|u - \tilde{u}\|_V$ with

$$e_n(u)_V = \min_{v \in V_n} \|u - v\|_V.$$

We restrict our attention to linear families : V_n is a linear space with $n = \dim(V_n)$.

If V is a Hilbert space, $e_n(u) = \|u - P_{V_n} u\|_V$ with P_{V_n} the V-orthogonal projection.

Classical choices : algebraic polynomials, spline spaces, trigonometric polynomials, piecewise constant functions on a given partition of D.

Optimized choices : if our prior information is that $u \in \mathcal{K}$ where $\mathcal{K} \subset V$ is some compact class we are interested in spaces V_n that perform close to the Kolmogorov n-width, that is defined for a general Banach space V by

$$d_n(\mathcal{K})_V := \inf_{\dim(V_n) = n} \max_{u \in \mathcal{K}} e_n(u)_V.$$

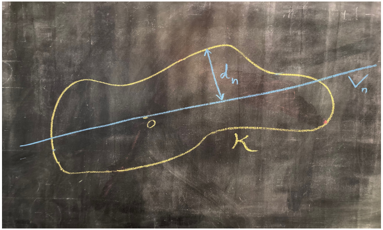

An optimal space achieving the infimum is not easy to construct.

It can be emulated by reduced basis spaces $V_n = \mathrm{span}\{u^1, \dots, u^n\}$, with $u^i \in \mathcal{K}$.

Greedy selection : given V_{k-1} pick next u^k such that

$$\|u^k - P_{V_{k-1}} u^k\| = \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V,$$

or in practice $\|u^k - P_{V_{k-1}} u^k\| \geq \gamma \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V$ for fixed $\gamma \in]0, 1[$.

Such algorithms have been proposed in the particular context of reduced order modeling, where the class $\mathcal{K}$ consists of solutions u to a PDE as we vary certain physical parameters (solution manifold). The reduced basis spaces are proved to perform as good as the optimal n-width spaces in terms of convergence rate.

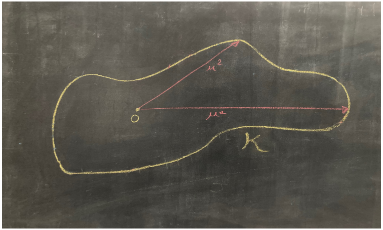

An optimal space achieving the infimum is not easy to construct.

It can be emulated by reduced basis spaces $V_n = \mathrm{span}\{u^1, \dots, u^n\}$, with $u^i \in \mathcal{K}$.

Greedy selection : given V_{k-1} pick next u^k such that

$$\|u^k - P_{V_{k-1}} u^k\| = \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V,$$

or in practice $\|u^k - P_{V_{k-1}} u^k\| \geq \gamma \max_{u \in \mathcal{K}} \|u - P_{V_{k-1}} u\|_V$ for fixed $\gamma \in]0, 1[$.

Such algorithms have been proposed in the particular context of reduced order modeling, where the class $\mathcal{K}$ consists of solutions u to a PDE as we vary certain physical parameters (solution manifold). The reduced basis spaces are proved to perform as good as the optimal n-width spaces in terms of convergence rate.

General objectives

Ideally we would like to combine

Instance optimality : achieve $\|u - \tilde{u}\|_V \leq Ce_n(u)_V$ for any u, for some fixed C.

Budget optimality : use $m \sim n$ samples (up to log factors).

Progressivity : when using $V_1 \subset V_2 \subset \dots V_n$ cumulated budget stays $m \sim n$.

In recent years, significant progresses have been made on randomized sampling and least-squares reconstruction strategies from various angles, allowing to reach the above (and other related) objectives.

Information based complexity : Wozniakowski, Wasilkowski, Kuo, Krieg, M. Ullrich, Kämmerer, Volkmer, Potts, T. Ullrich, Oettershagen, ...

Uncertainty quantification and model reduction : Doostan, Hampton, Narayan, Jakeman, Zhou, Nobile, Tempone, Chkifa, Webster, Harberstisch, Nouy, Perrin...

Approximation theory : Cohen, Davenport, Leviatan, Migliorati, Bachmayr, Arras, Adcock, Huybrechs, Temlyakov...

These results lead to natural comparison between sampling numbers and n-widths.

A simple example : interpolation by univariate polynomials

Consider $D = [-1, 1]$ and $V = \mathcal{C}(D)$ equipped with the max norm $\| \cdot \|_V = \| \cdot \|_{L^\infty}$.

Take $V_n = \mathbb{P}_{n-1}$ univariate polynomials of degree $n - 1$.

With $(x^1, \dots, x^n) \in [-1, 1]$ pairwise distincts, reconstruct by the interpolation operator

$$\tilde{u} = I_n u \in \mathbb{P}_{n-1}, \quad s.t. \quad I_n u(x^i) = u(x^i), \quad i = 1, \dots, n.$$

Budget is optimal : $m = n$ points have been used.

Instance optimality : governed by Lebesgue constant $C_n = \max_{u \neq 0} \frac{\|I_n u\|_{L^\infty}}{\|u\|_{L^\infty}}$, since

$$\|u - I_n u\|_{L^\infty} \leq \|u - v\|_{L^\infty} + \|I_n v - I_n u\|_{L^\infty} \leq (1 + C_n)\|u - v\|_{L^\infty}, \quad v \in V_n,$$

thus bounded by $(1 + C_n) e_n(u)_{L^\infty}$.

Equispaced points are known to yield $C_n \sim 2^n$.

Chebychev points $\left\{ \cos\left(\frac{2k\pi}{2n+1} \right) \; : \; k = 1, \dots, n \right\}$ yield optimal value $C_n \sim \ln(n)$.

Limitations

Multivariate case : no general theory for optimal points on a general domain $D \subset \mathbb{R}^d$.

What about other types of spaces V_n ?

Fekete points : if V_n is a linear space with basis $(\phi_1, \dots, \phi_n)$, then the points

$$(x^1, \dots, x^n) = \operatorname{argmax}\Big\{ \det(\phi_i(z_j))_{i,j=1,\dots,n} \; : \; (z_1, \dots, z_n) \in D^n \Big\},$$

yields $C_n \leq n$ but are not simply computable : non-convex optimization in $\mathbb{R}^{dn}$.

For univariate polynomials these points maximizes $\prod_{j \neq i} |x^i - x^j|$.

Progessivity : the Chebychev and Fekete points are not nested as $n \to n+1$!

The Clenshaw-Curtis points $G_n = \Big\{ \cos\Big(\frac{k\pi}{n-1}\Big) \; : \; k = 0, \dots, n-1 \Big\}$ are partially nested :

$$G_3 \subset G_5 \subset G_9 \subset \cdots \subset G_{2^j+1} \subset G_{2^{j+1}+1} \subset \cdots$$

How to fill-in by intermediate points and preserve a well-behaved Lebesgue constant ?

Lebesgue constant for nested sets

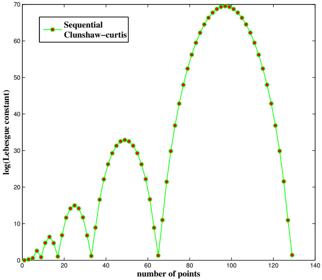

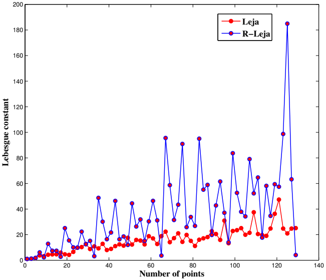

Left : fill-in by increasing order.

Right (blue) : fill-in by Van der Corput enumeration $C_n \leq n^2$ (Chkifa, 2013).

Right (red) : greedy Fekete (Leja) $\max \prod_{j=1}^{k-1} |x - x^j| \to x^k$. Open problem : $C_n \sim n$?

The behaviour $C_n \sim \ln(n)$ does not seem achievable with nested sets.

From now on, $V = L^2(D,\mu)$. Notation : $\|v\| = \|v\|_{L^2(D,\mu)}$, and $e_n(u) = \|u - P_{V_n}u\|$.

The $L^2(D,\mu)$-projection

$$P_{V_n}u := \operatorname{argmin}\Big\{ \int_D |u(x) - v(x)|^2 d\mu \,:\, v \in V_n \Big\},$$

is out of reach $\implies$ replace the integrals by a discrete sum

$$\int_D v(x)d\mu \approx \frac{1}{m}\sum_{i=1}^m w(x^i)v(x^i).$$

where w is a weight function. This is the (weighted) least-squares method

$$u_n := \operatorname{argmin}\Big\{ \frac{1}{m}\sum_{i=1}^m w(x^i)|y^i - v(x^i)|^2 \,:\, v \in V_n \Big\}.$$

In the noiseless case $y^i = u(x^i)$, the solution is the orthogonal projection of u onto V_n for the discrete (semi)-norm

$$\|v\|_m^2 := \frac{1}{m}\sum_{i=1}^m w(x^i)|v(x^i)|^2,$$

that should in some sense be close to $\|v\|^2$.

Randomized sampling

Draw $(x^1, \dots, x^m)$ i.i.d. according to a sampling probability measure σ.

Use a weight w such that

$$w(x)d\sigma(x) = d\mu(x).$$

The random norm $\|v\|_m^2 = \frac{1}{m}\sum_{i=1}^m w(x^i)|v(x^i)|^2$ then satisfies, for any function v,

$$\mathbb{E}\Big(\|v\|_m^2\Big) = \mathbb{E}_\sigma(w(x)|v(x)|^2) = \int_D w(x)|v(x)|^2 d\sigma = \int_D |v(x)|^2 d\mu = \|v\|^2.$$

Unweighted choice : $w = \mu$ and $d\sigma = d\mu$ may lead to suboptimal results

Optimality results will be achieved by appropriate choices of w and σ.

The weighted least-squares approximation u_n is now a random object. Its accuracy should be studied in some probabilistic sense, for instance $\mathbb{E}(\|u - u_n\|^2)$.

Accuracy analysis

General strategy : study the probabilistic event E_δ of the equivalence

$$(1-\delta)\|v\|^2 \leq \|v\|_m^2 \leq (1+\delta)\|v\|^2, \quad v \in V_n,$$

for some $0 < \delta < 1$, for example $\delta = \frac{1}{2}$.

This is an instance ($p = 2$ and $w_i = m^{-1}w(x^i)$) of a Marcinkiewicz-Zygmund inequality :

$$(1-\delta)\int_D |v(x)|^p d\mu \leq \sum_{i=1}^m w_i |v(x^i)|^p \leq (1-\delta)\int_D |v(x)|^p d\mu, \quad v \in V_n.$$

Let $(L_1, \dots, L_n)$ be an $L^2(D,\mu)$-orthonormal basis of V_n and consider the random Gramian matrix

$$\mathbf{G} = (G_{k,j})_{k,j=1,\dots,n}, \quad G_{k,j} := \frac{1}{m}\sum_{i=1}^m w(x^i)L_k(x^i)L_j(x^i) = \langle L_k, L_j\rangle_m.$$

Then

$$E_\delta \iff (1-\delta)\mathbf{I} \leq \mathbf{G} \leq (1+\delta)\mathbf{I} \iff \|\mathbf{G}-\mathbf{I}\|_2 \leq \delta.$$

Note that $\mathbf{G} = \frac{1}{m}\sum_{j=1}^m \mathbf{X}^i$, where $\mathbf{X}^i$ are i.i.d. realizations of

$$\mathbf{X} = (w(x)L_k(x)L_j(x))_{k,j}, \qquad x \sim \sigma, \quad \text{so} \quad \mathbb{E}(\mathbf{G}) = \mathbf{I}$$

A first accuracy bound

Under the event $E_{1/2}$, one has $\frac{1}{2}\|v\|^2 \leq \|v\|_m^2 \leq \frac{3}{2}\|v\|^2$ for all $v \in V_n$, and so

$$\|u - u_n\|^2 = e_n(u)^2 + \|P_n u - u_n\|^2 \leq e_n(u)^2 + 2\|P_n u - u_n\|_m^2.$$

In addition $\|P_n u - u\|_m^2 = \|u - u_n\|_m^2 + \|P_n u - u_n\|_m^2$, and so

$$\|u - u_n\|^2 \leq e_n(u)^2 + 2\|u - P_n u\|_m^2.$$

Since $\mathbb{E}(\|u - P_n u\|_m^2) = e_n(u)^2$, we reach

$$\mathbb{E}(\|u - u_n\|^2 \chi_{E_{1/2}}) \leq 3e_n(u)^2.$$

We can test the validity of $E_{1/2}$ by checking if $\|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}$.

First choice : define $\tilde{u} = u_n$ if $E_{1/2}$ holds and $\tilde{u} = 0$ gives the estimate

$$\mathbb{E}(\|u - \tilde{u}\|^2) \leq 3e_n(u)^2 + \delta\|u\|^2, \quad \delta := \Pr(E_{1/2}^c).$$

Is δ small with $m \sim n$?

Key tools : Christoffel functions and matrix concentration.

Haberstisch-Nouy-Perrin (2019) : redraw $\{x^1,\dots,x^m\}$ until $E_{1/2}$ holds and take $\tilde u = u_n$

If $\delta = \Pr(E^c_{1/2})$ then the number of needed redraws k^* follows a Poisson law : one has $k^* > k$ with probability δ^k and $\mathbb{E}(k^*) = \frac{1}{1-\delta}$.

The resulting sample $x^1,\dots,x^m$ follows the law $\otimes^m\sigma$ conditionned to $E_{1/2}$ and therefore, by Bayes rule

$$\mathbb{E}(\|u-\tilde u\|^2) = \mathbb{E}(\|u-u_n\|^2 \,|\, E_{1/2}) = \Pr(E_{1/2})^{-1}\mathbb{E}(\|u-u_n\|^2\chi_{E_{1/2}}),$$

which gives for all $u \in V$ (non uniform result : first fix u, then draw sample),

$$\mathbb{E}(\|u-\tilde u\|^2) \le C e_n(u)^2, \quad C := \frac{3}{1-\delta}.$$

Assume V_n contains constants and that $M := \mu(D) = \int |1|^2 d\mu < \infty$. Then under $E_{1/2}$, we have $\frac{1}{m}\sum_{i=1}^m w(x^i) = \|1\|_m^2 \le \frac{3M}{2}$, so both $\|\cdot\|$ and $\|\cdot\|_m$ dominated by $\|\cdot\|_{L^\infty}$.

Therefore, for the boosted sample $x^1,\dots,x^m$, we are ensured that for all $u \in \mathcal{C}(D)$,

$$\|u-u_n\| \le \|u-v\|+\|v-u_n\|_m \le \|u-v\|+\|u-v\|_m \le C\|u-v\|_{L^\infty}, \quad C := \sqrt{M}(1+\sqrt{3/2}),$$

and therefore (uniform result : first fix a deterministic sample, then pick any u)

$$\|u-\tilde u\| \le C e_n(u)_{L^\infty}.$$

Christoffel functions

With $L_1, \dots, L_n$ an $L^2(D,\mu)$-orthonormal basis of V_n, define

$$k_n(x) := \sum_{j=1}^n |L_j(x)|^2,$$

the inverse of the Christoffel function, also defined as

$$k_n(x) = \max_{v \in V_n} \frac{|v(x)|^2}{\|v\|^2}.$$

We use the notation

$$K_n := \|k_n\|_{L^\infty} := \sup_{x \in D} \sum_{j=1}^n |L_j(x)|^2 = \max_{v \in V_n} \frac{\|v\|_{L^\infty}^2}{\|v\|^2}.$$

These quantities only depends on V_n and μ.

For the given weight w, we introduce

$$k_{n,w}(x) := w(x) k_n(x),$$

and $K_{n,w} := \|k_{n,w}\|_{L^\infty}$, which only depends on (V_n, μ, w).

Since $\int_D k_{n,w} d\sigma = \sum_{j=1}^n \int_D |L_j|^2 d\rho = n$, one has

$$K_{n,w} \geq n.$$

Matrix Chernoff bound (Ahlswede-Winter 2000, Tropp 2011) : let $\mathbf{G} = \frac{1}{m}\sum_{i=1}^m \mathbf{X}^i$ where $\mathbf{X}^i$ are i.i.d. copies of an $n \times n$ symmetric matrix $\mathbf{X}$ such that $\mathbb{E}(\mathbf{X}) = \mathbf{I}$ and $\|\mathbf{X}\| \leq K$ a.s. Then

$$\Pr\left\{\|\mathbf{G} - \mathbf{I}\| \geq \delta\right\} \leq 2n \exp\left(-\frac{mc_\delta}{K}\right),$$

where $c_\delta := (1+\delta)\ln(1+\delta) - \delta > 0$.

In our case of interest,

$$\mathbf{X} = w(x)(L_k(x)L_j(x))_{j,k=1,\dots,n} = \mathbf{x}\mathbf{x}^T, \quad \mathbf{x} = (w(x)^{1/2}L_k(x))_{k=1,\dots,n},$$

with x distributed according to σ, which has expectation $\mathbb{E}(\mathbf{X}) = \mathbf{I}$, and

$$K = \sup \|\mathbf{X}\| = \sup |\mathbf{x}|^2 = \sup_{x \in D} w(x) \sum_{j=1}^n |L_j(x)|^2 = K_{n,w}.$$

This gives the sampling budget condition

$$m \geq cK_{n,w} \ln(2n/\varepsilon) \implies \Pr(E_{1/2}^c) = \Pr\left\{\|G - I\| \geq \frac{1}{2}\right\} \leq \varepsilon,$$

with $c = c_{1/2}^{-1} \leq 10$. For the boosted sample, take $\varepsilon = \frac{1}{2}$, and so $m \geq 10K_{n,w}\ln(4n)$.

Optimal estimation and sampling budget

Using the boosted sample, we achieve near optimal non-uniform estimate

$$\mathbb{E}(\|u-\tilde{u}\|^2) \leq Ce_n(u)^2$$

as well as uniform estimate (assuming $\mu(D) < \infty$ and $\frac{1}{m}\sum_{i=1}^m w(x^i) < \infty$)

$$\|u-\tilde{u}\| \leq Ce_n(u)_{L^\infty}$$

under a sampling budget $m \sim K_{n,w} \geq n$ up to multiplicative logarithmic factor.

In the presence of noise of variance $\kappa(x)^2$, the estimation bound has an additional term

$$e_n(u)^2 + \frac{n}{m}\kappa^2, \qquad \kappa^2 = \int_D |\kappa(x)|^2 d\mu.$$

Unweighted least-squares : $w = 1$ and $\sigma = \mu$ requires $m \sim K_n = \max_{x\in D} \sum_{j=1}^n |L_j(x)|^2$

Sometimes $K_n >> n$. leading to an excessive sampling budget.

Illustration on univariate polynomials $V_n = \mathbb{P}_{n-1}$

Regime of stability : probability that $\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2}$, white if 1, black if 0.

Unweighted case requires at least $m \sim K_n$.

Left : $D = [-1,1]$ with $d\mu = \frac{dx}{\pi\sqrt{1-x^2}}$ (Chebychev polynomials $K_n = 2n+1 \sim n$).

Center : $D = [-1,1]$ with $d\mu = \frac{dx}{2}$ (Legendre polynomials $K_n = n^2$)

Right : $D = \mathbb{R}$ with $d\mu = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx$ (Hermite polynomials $K_n = \infty$).

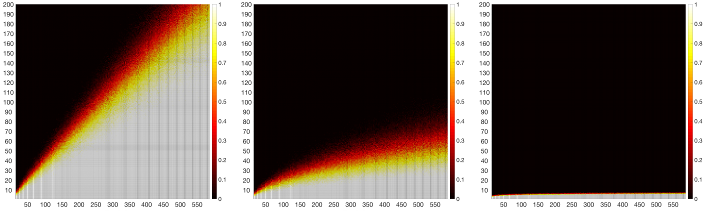

For the gaussian case, a more ad-hoc analysis shows that stability holds if $m \gtrsim \exp(cn)$

Parametric PDE's and multivariate polynomials

Prototype example : elliptic PDE's on some domain $D \subset \mathbb{R}^2$ or $\mathbb{R}^3$ with affine parametrization of the diffusion function by $x = (x_1, \dots, x_d) \in X = [-1,1]^d$

$$-\mathrm{div}(a\nabla u) = f, \quad a = \bar{a} + \sum_{j=1}^{d} x_j \psi_j,$$

with ellipticity assumption $0 < r < a < R$ for all $x \in X$, so $x \mapsto u(x) \in V = H_0^1(D)$.

With $\Lambda \subset \mathbb{N}^d$, approximation by multivariate polynomial space

$$V_\Lambda := \left\{ \sum_{\nu\in\Lambda} v_\nu x^\nu, \;\; v_\nu \in V \right\} = V \otimes \mathbb{P}_\Lambda,$$

where $x^\nu = x_1^{\nu_1} \cdots x_d^{\nu_d}$.

We only consider downward closed index sets : $\nu \in \Lambda$ and $\mu \leq \nu \Rightarrow \mu \in \Lambda$.

Basis of $\mathbb{P}_\Lambda$: tensorized orthogonal polynomials $L_\nu(x) = \prod_{j=1}^d L_{\nu_j}(x_j)$ for $\nu \in \Lambda$.

Downward closed multivariate polynomials

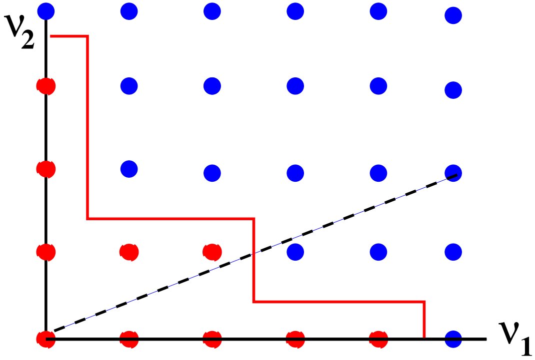

Cohen-DeVore-Schwab (2011) + Bachmayr-Migliorati (2016) : approximation results.

Under suitable summability conditions on $(|\psi_j|)_{j\geq 1}$, there exists a sequence of downward closed sets $\Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_n \dots$, with $n := \#(\Lambda_n)$ such that

$$\inf_{v\in V_n} \|u - v\|_{L^2(X,V,\mu)} \leq Cn^{-s},$$

with $V_n := V_{\Lambda_n} = \mathbb{P}_{\Lambda_n} \otimes V$, where μ is any tensorized Jacobi measure. The exponent $s > 0$ is robust with respect to the dimension d.

Chkifa-Cohen-Nobile-Tempone (M2AN, 2014) : estimate K_n for $\mathbb{P}_{\Lambda_n}$.

With $\rho = \otimes^d(\frac{dx}{2})$ the uniform distribution over X, one has $K_n \leq n^2$ for all downward closed sets Λ_n such that $\#(\Lambda_n) = n$. Up to log factor, the stability regime is $m \gtrsim n^2$.

With the tensor-product Chebychev measure, improvement $K_n \leq n^\alpha$ with $\alpha := \frac{\log 3}{\log 2}$.

The theory and least-square method is not capable of handling lognormal diffusions :

$$a = \exp(b), \quad b = \sum_{i=1}^{d} x_j \psi_j, \quad x_i \sim \mathcal{N}(0,1) \text{ i.i.d.}$$

which corresponds to the tensor product Gaussian measure over $X = \mathbb{R}^d$.

Optimal sampling measure

Narayan-Jakeman (2015), Doostan-Hampton (2015), Cohen-Migliorati (2017) : use sampling measure

$$d\sigma := \frac{k_n}{n} d\mu = \frac{1}{n}\Big(\sum_{j=1}^n |L_j|^2\Big) d\mu \implies w(x) = \frac{n}{k_n(x)}.$$

σ is a probability measure and we have $k_{n,w}(x) = w(x)k_n(x) = n$, thus $K_{n,w} = n$.

With this sampling strategy, optimal error bounds can be achieved with near optimal sampling budget $m \sim n$ up to logarithmic factors.

Observation by T. Ullrich (2020) : if μ has finite mass $\mu(D) = M < \infty$, one can also use $d\tilde{\sigma} := (\frac{1}{2M} + \frac{k_n}{2n})d\mu$ ensuring both $K_{n,w} \leq 2n$ and $\frac{1}{m}\sum_{i=1}^m w(x_i) \leq 2M$.

Optimal sampling measure

Narayan-Jakeman (2015), Doostan-Hampton (2015), Cohen-Migliorati (2017) : use sampling measure

$$d\sigma := \frac{k_n}{n} d\mu = \frac{1}{n}\Big(\sum_{j=1}^{n} |L_j|^2\Big) d\mu \implies w(x) = \frac{n}{k_n(x)}.$$

σ is a probability measure and we have $k_{n,w}(x) = w(x)k_n(x) = n$, thus $K_{n,w} = n$.

With this sampling strategy, optimal error bounds can be achieved with near optimal sampling budget $m \sim n$ up to logarithmic factors.

Observation by T. Ullrich (2020) : if μ has finite mass $\mu(D) = M < \infty$, one can also use $d\tilde{\sigma} := (\frac{1}{2M} + \frac{k_n}{2n})d\mu$ ensuring both $K_{n,w} \leq 2n$ and $\frac{1}{m}\sum_{i=1}^{m} w(x_i) \leq 2M$.

Narayan-Jakeman (2015), Doostan-Hampton (2015), Cohen-Migliorati (2017) : use sampling measure

$$d\sigma := \frac{k_n}{n} d\mu = \frac{1}{n}\Big(\sum_{j=1}^n |L_j|^2\Big) d\mu \implies w(x) = \frac{n}{k_n(x)}.$$

σ is a probability measure and we have $k_{n,w}(x) = w(x)k_n(x) = n$, thus $K_{n,w} = n$.

With this sampling strategy, optimal error bounds can be achieved with near optimal sampling budget $m \sim n$ up to logarithmic factors.

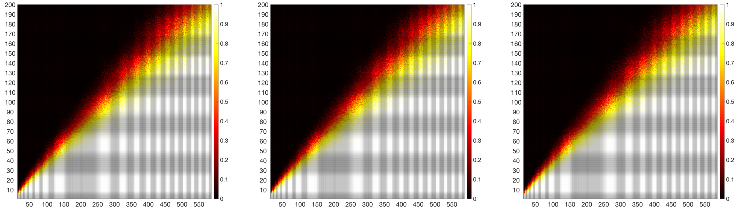

Stability regime for univariate polynomials with μ Chebychev, uniform, and Gaussian.

Observation by T. Ullrich (2020) : if μ has finite mass $\mu(D) = M < \infty$, one can also use $d\tilde{\sigma} := (\frac{1}{2M} + \frac{k_n}{2n})d\mu$ ensuring both $K_{n,w} \le 2n$ and $\frac{1}{m}\sum_{i=1}^m w(x_i) \le 2M$.

Optimal sampling measure

Narayan-Jakeman (2015), Doostan-Hampton (2015), Cohen-Migliorati (2017) : use sampling measure

$$d\sigma := \frac{k_n}{n} d\mu = \frac{1}{n}\Big(\sum_{j=1}^{n} |L_j|^2\Big) d\mu \implies w(x) = \frac{n}{k_n(x)}.$$

σ is a probability measure and we have $k_{n,w}(x) = w(x)k_n(x) = n$, thus $K_{n,w} = n$.

With this sampling strategy, optimal error bounds can be achieved with near optimal sampling budget $m \sim n$ up to logarithmic factors.

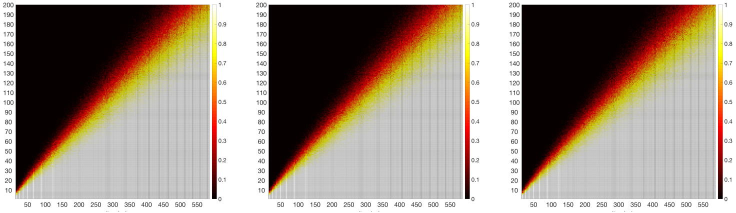

Stability regime for univariate polynomials with μ Chebychev, uniform, and Gaussian.

Observation by T. Ullrich (2020) : if μ has finite mass $\mu(D) = M < \infty$, one can also use $d\tilde{\sigma} := (\frac{1}{2M} + \frac{k_n}{2n})d\mu$ ensuring both $K_{n,w} \leq 2n$ and $\frac{1}{m}\sum_{i=1}^{m} w(x_i) \leq 2M$.

A first comparison between sampling number and n-width

By optimizing the choice of V_n in the estimate

$$\mathbb{E}(\|u - \tilde{u}\|^2) \leq C e_n(u)^2,$$

and using the optimal sampling measure, one find that

$$\rho^{\mathrm{rand}}_{cn\ln(n)}(\mathcal{K})_{L^2} \leq C d_n(\mathcal{K})_{L^2},$$

for any compact set $\mathcal{K}$ of $L^2(D, \mu)$.

Likewise, optimizing the choice of V_n in the estimate

$$\|u - \tilde{u}\| \leq C e_n(u)_{L^\infty},$$

one finds that

$$\rho^{\mathrm{det}}_{cn\ln(n)}(\mathcal{K})_{L^2} \leq C d_n(\mathcal{K})_{L^\infty},$$

for any compact set $\mathcal{K}$ of $\mathcal{C}(D)$.

Questions : remove logarithmic surplus ? deterministic sampling numbers vs. $d_n(\mathcal{K})_{L^2}$?

The optimal density is not fixed

When using a sequence $(V_n)_{n\geq 1}$ of approximation spaces

$$d\sigma = d\sigma_n := \frac{k_n}{n} d\mu.$$

Illustration : sampling densities σ_n for $n = 5, 10, 20$.

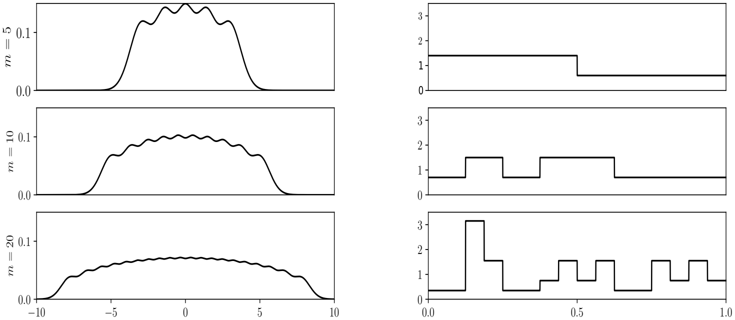

Left : Polynomials of degrees $0, \ldots, m-1$ and μ Gaussian.

Right : Piecewise constant functions on locally refined partitions and μ uniform.

Dependence on the domain geometry

Consider the space $V_n = \mathbb{P}_k$ of polynomials of total degree k on a multivariate domain $D \subset \mathbb{R}^d$, so that

$$n = \binom{k+d}{d}$$

and use the uniform probability measure $d\mu = |D|^{-1} dx$.

The local behaviour of k_n and thus of σ_n depends on closeness to the boundary of D and on the smoothness of this boundary.

Cohen-Dolbeault (2020) : For smooth domains $k_n(x) = \mathcal{O}(n^{\frac{d+1}{d}})$ on boundary, for Lipschitz domains $k_n(x) = \mathcal{O}(n^2)$ on exiting corners, for domains with cusps $k_n(x) = \mathcal{O}(n^r)$ at exiting cusps where r depends on the order of cuspitality.

Consider the space $V_n = \mathbb{P}_k$ of polynomials of total degree k on a multivariate domain $D \subset \mathbb{R}^d$, so that

$$n = \binom{k+d}{d}$$

and use the uniform probability measure $d\mu = |D|^{-1}dx$.

The local behaviour of k_n and thus of σ_n depends on closeness to the boundary of D and on the smoothness of this boundary.

Cohen-Dolbeault (2020) : For smooth domains $k_n(x) = \mathcal{O}(n^{\frac{d+1}{d}})$ on boundary, for Lipschitz domains $k_n(x) = \mathcal{O}(n^2)$ on exiting corners, for domains with cusps $k_n(x) = \mathcal{O}(n^r)$ at exiting cusps where r depends on the order of cuspitality.

Consider the space $V_n = \mathbb{P}_k$ of polynomials of total degree k on a multivariate domain $D \subset \mathbb{R}^d$, so that

$$n = \binom{k+d}{d}$$

and use the uniform probability measure $d\mu = |D|^{-1}dx$.

The local behaviour of k_n and thus of σ_n depends on closeness to the boundary of D and on the smoothness of this boundary.

Cohen-Dolbeault (2020) : For smooth domains $k_n(x) = \mathcal{O}(n^{\frac{d+1}{d}})$ on boundary, for Lipschitz domains $k_n(x) = \mathcal{O}(n^2)$ on exiting corners, for domains with cusps $k_n(x) = \mathcal{O}(n^r)$ at exiting cusps where r depends on the order of cuspitality.

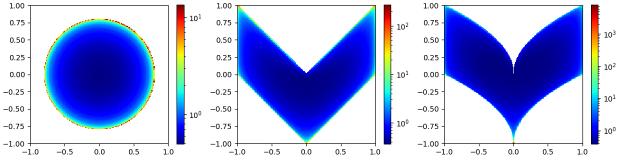

Inverse Christoffel function $k_n(x)$ for $n = 231$ (total degree $k = 20$)

Examples of draw according to optimal sample distribution

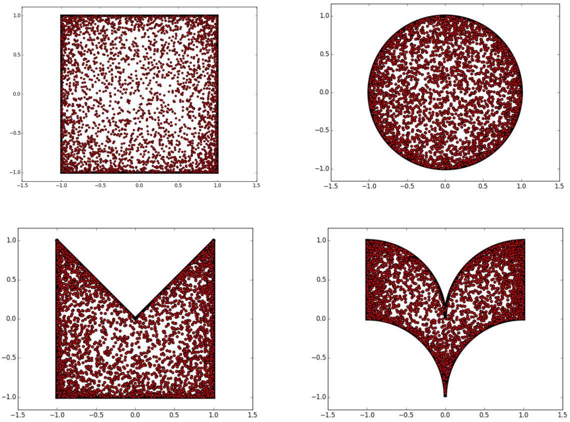

Sampling the optimal density

Problem : generate efficiently i.i.d. samples according to the optimal sampling measure

$$d\sigma = d\sigma_n = \frac{k_n}{n} d\mu = \frac{1}{n}\Big(\sum_{j=1}^n |L_j|^2\Big) d\mu.$$

This problem might be non-trivial in a multivariate setting $D \subset \mathbb{R}^d$.

In many relevant instances μ is a product measure (such as uniform, gaussian) and thus easy to sample, but $d\sigma_n$ is not. Sampling strategies :

(i) Rejection sampling : draw x^i according to μ and a uniform random variable z^i in $[0, M]$ where $M \geq \frac{\|k_n\|_{L^\infty}}{n}$. Reject x^i if $z^i > \frac{k_n(x^i)}{n}$.

(ii) Conditional sampling : obtains first component by sampling the marginal $d\sigma_1(y_1)$, then the second component by sampling the conditional marginal probability $d\sigma_{y_1}(y_2)$ for this choice of the first component, etc...

Strategy (ii) is more efficient in cases where the L_j have tensor product structure.

(iii) Mixture sampling : draw uniform variable $j \in \{1, \dots, n\}$, then sample with probability $|L_j|^2 d\mu$.

Migliorati (2018) : one can also split the sample into n batches of size $\mathcal{O}(\ln(n))$ each of them sampled according to $d\nu_j = |L_j|^2 d\mu$, with same final estimation bounds.

Sampling on general domains

Optimal sampling may become unfeasible when $D \subset \mathbb{R}^d$ is a domain with a general geometry : the $L_1, \dots, L_n$ have no simple expression and cannot be computed exactly.

General assumptions : χ_D is easily computable $\Rightarrow$ sampling according to the uniform measure μ is easy (sample uniformly on a bounding box, reject if $x \notin D$).

Migliorati, Adcock-Cardenas (2019), Cohen-Dolbeault (2020) : two-step strategies

1. With $M \sim K_n \ln(n)$ sample $z^1, \dots, z^M$ according to the uniform measure, and define

$$\tilde{\mu} := \frac{1}{M} \sum_{i=1}^{M} \delta_{z^i}.$$

Construct an orthonormal basis $\tilde{L}_1, \dots, \tilde{L}_n$ of V_n for the $L^2(X, \tilde{\mu})$ inner product and define $\tilde{k}_n = \sum_{j=1}^{n} |\tilde{L}_j|^2$.

2. With $m \sim n \ln(n)$ sample $x^1, \dots, x^m$ according to

$$d\tilde{\sigma} = \frac{\tilde{k}_n}{n} d\tilde{\mu},$$

that is, select z^i with probability $p_i = \frac{\tilde{k}_n(z^i)}{Mn}$.

Adaptivity

Update adaptively the polynomial space $\Lambda_{n-1} \to \Lambda_n$, while increasing the amount of sample necessary for stability $m = m(n) \sim n \log n$.

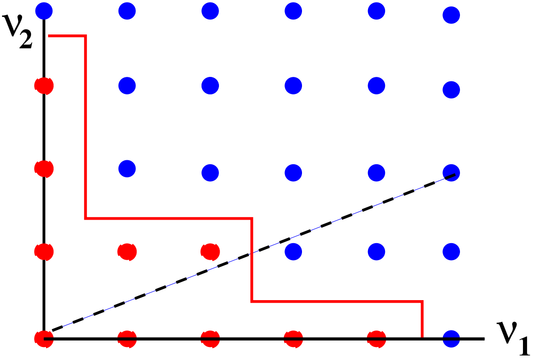

Problem : the optimal measure $\mu = \mu_m$ changes as we vary m. How should we recycle the previous samples ?

Adaptivity

Update adaptively the polynomial space $\Lambda_{n-1} \to \Lambda_n$, while increasing the amount of sample necessary for stability $m = m(n) \sim n \log n$.

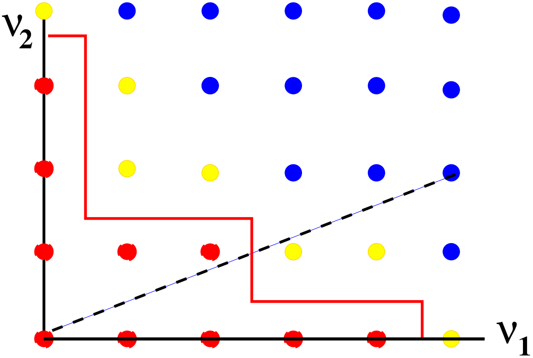

Problem : the optimal measure $\mu = \mu_m$ changes as we vary m. How should we recycle the previous samples ?

Adaptivity

Update adaptively the polynomial space $\Lambda_{n-1} \to \Lambda_n$, while increasing the amount of sample necessary for stability $m = m(n) \sim n \log n$.

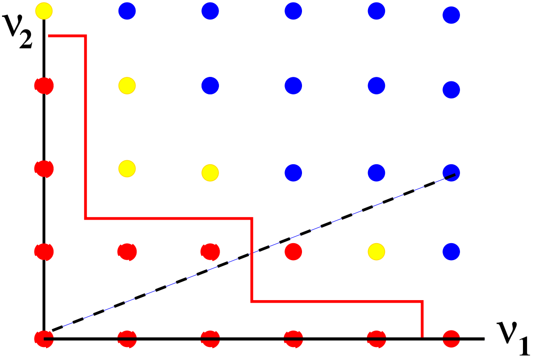

Problem : the optimal measure $\mu = \mu_m$ changes as we vary m. How should we recycle the previous samples ?

Update adaptively the polynomial space $\Lambda_{n-1} \to \Lambda_n$, while increasing the amount of sample necessary for stability $m = m(n) \sim n \log n$.

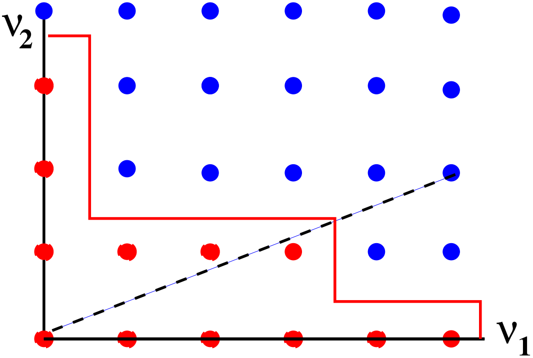

Problem : the optimal measure $\mu = \mu_m$ changes as we vary m. How should we recycle the previous samples ?

Adaptivity

Update adaptively the polynomial space $\Lambda_{n-1} \to \Lambda_n$, while increasing the amount of sample necessary for stability $m = m(n) \sim n \log n$.

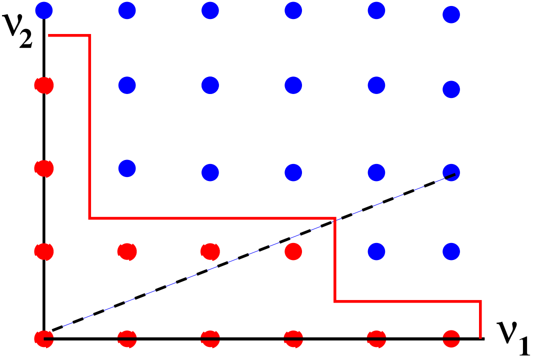

Problem : the optimal measure $\mu = \mu_m$ changes as we vary m. How should we recycle the previous samples ?

Sequencial sampling

For a given hierarchy $V_1 \subset V_2 \subset \cdots \subset V_n$, note that

$$d\sigma_n = \frac{1}{n}\Big(\sum_{j=1}^{n} |L_j|^2\Big) d\mu = \Big(1 - \frac{1}{n}\Big) d\sigma_{n-1} + \frac{1}{n} d\nu_n \quad \text{where } d\nu_n = |L_n|^2 d\mu.$$

We use this mixture property to generate the sample in an incremental manner.

Assume that the sample $S_{n-1} = \{x^1, \dots, x^m\}$ has been generated by independent draw according to the distribution $d\sigma_{n-1}$ with $m = m(n-1)$ sampling budget

Then we generate a new sample $S_n = \{x^1, \dots, x^{m(n)}\}$ as follows :

For each $i = 1, \dots, m(n)$, pick Bernoulli variable $b_i \in \{0, 1\}$ with probability $\{\frac{1}{n}, 1 - \frac{1}{n}\}$.

If $b_i = 0$, generate new x^i according to $d\nu_n$.

If $b_i = 1$, recycle x^i incrementally from S_{n-1}.

Arras-Bachmayr-Cohen (2018) : the cumulated number of sample C_n used at stage n satisfies $C_n \sim n$ up to logarithmic factors with high probability for all values of n.

With high probability, the matrix $\mathbf{G}$ satisfies $\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2}$ for all values of n.

Adaptive selection strategies ?

Sparsification

Reducing further sampling budget to $\mathcal{O}(n)$: logarithmic factors removable ?

Batson-Spielman-Srivastava (2014) : let $\mathbf{x}_1, \dots, \mathbf{x}_m$ be $m \geq n$ be vectors of $\mathbb{R}^n$ such that

$$(1-\delta)\mathbf{I} \leq \sum_{i=1}^{m} \mathbf{x}_i \mathbf{x}_i^T \leq (1+\delta)\mathbf{I}.$$

For any $c > 1$ there exists $S \subset \{1, \dots, m\}$ with $\#(S) \leq cn$ and weights s_i such that

$$\left(1 - \frac{1}{\sqrt{c}}\right)^2 (1-\delta)\mathbf{I} \leq \sum_{i \in S} s_i \mathbf{x}_i \mathbf{x}_i^T \leq (1+\delta)\left(1 + \frac{1}{\sqrt{c}}\right)^2 \mathbf{I}$$

Apply this to $\mathbf{x}_i = \left(\sqrt{\frac{w(x^i)}{m}} L_j(x^i)\right)_{j=1,\dots m}$ with $\{x^1, \dots, x^m\}$ a boosted sample.

Leads to a sample $(x^1, \dots, x^{2n})$ and weights $w_i = s_i \frac{w(x^i)}{m}$ such that

$$\alpha \|v\|^2 \leq \|v\|_{2n}^2 \leq \beta \|v\|^2, \qquad v \in V_n,$$

where $\|v\|_{2n}^2 = \sum_{i=1}^{2n} w_i |v(x^i)|^2$ and $\alpha = \frac{1}{2}\left(1 - \frac{1}{\sqrt{2}}\right)^2$, $\beta = \frac{3}{2}\left(1 + \frac{1}{\sqrt{2}}\right)^2$.

Sparsified weighted least-squares

Based on these new samples and weights, we define a weighted least-squares estimate

$$\tilde{u} := \operatorname{argmin}\Big\{ \frac{1}{2n} \sum_{i=1}^{2n} w_i |u(x^i) - v(x^i)|^2 \Big\}.$$

for which we have for all $u \in \mathcal{C}(D)$

$$\|u - \tilde{u}\| \leq C e_n(u)_{L^\infty},$$

assuming that μ is a finite measure.

The sparsification strategy of Batson-Spielman-Srivastava is performed by a deterministic greedy algorithm of total complexity $\mathcal{O}(mn^3)$: additional offline cost.

Temlyakov (2019) : comparison between deterministic linear optimal recovery numbers in L^2 and Kolmogorov n-width in L^∞ for any compact class $\mathcal{K}$ of $\mathcal{C}(D)$.

By optimizing the choice of V_n, one obtains

$$\rho_{2n}^{\det}(\mathcal{K})_{L^2} \leq C d_n(\mathcal{K})_{L^\infty}.$$

We cannot prove $\mathbb{E}(\|u-\tilde{u}\|^2) \leq Ce_n(u)^2$ with the above strategy.

We miss the averaging property $\mathbb{E}(\|v\|_{2n}^2) = \|v\|^2$ for any $v \in V$.

Use a variant result : solution to the Weaver and Kadison-Singer conjectures.

Marcus-Spielman-Srivastava (2015) : if $\mathbf{x}_1, \dots, \mathbf{x}_m$ are m vectors from $\mathbb{R}^n$ of norm $|\mathbf{x}_i|^2 \leq \delta$ and such that

$$\alpha \mathbf{I} \leq \sum_{i=1}^{m} \mathbf{x}_i \mathbf{x}_i^T \leq \beta \mathbf{I}$$

then there exists a partition $S_1 \cup S_2 = \{1, \dots, m\}$ such that

$$\frac{1-5\sqrt{\delta/\alpha}}{2} \alpha \mathbf{I} \leq \sum_{i \in S_j} \mathbf{x}_i \mathbf{x}_i^T \leq \frac{1+5\sqrt{\delta/\alpha}}{2} \beta \mathbf{I}, \quad j = 1, 2.$$

Nitzan-Olevskii-Ulanovskii (2016) apply this process recursively in order to identify a $J \subset \{1, \dots, m\}$ such that $|J| \leq cn$ and

$$C^{-1} \alpha \mathbf{I} \leq \sum_{i \in J} \mathbf{x}_i \mathbf{x}_i^T \leq C \beta \mathbf{I}.$$

for some universal constant $C > 1$.

Cohen-Dolbeault (2021) : if the $\mathbf{x}_i$ have equal norms $|\mathbf{x}_i|^2 = \frac{n}{m}$, then iterative splitting delivers for some $L = \mathcal{O}(\ln(m/n))$ a partition $J_1 \cup J_2 \cup \cdots \cup J_{2^L} = \{1, \dots, m\}$ such that

$$c_0 \mathbf{I} \leq \sum_{i \in J_k} \mathbf{x}_i \mathbf{x}_i^T \leq C_0 \mathbf{I}, \quad k = 1, \dots, 2^L,$$

with (c_0, C_0) universal constants and $|J_k| \leq C_0 n$ for all k.

Apply to $\mathbf{x}_i = \left(\sqrt{\frac{w(x^i)}{m}} L_j(x^i)\right)_{j=1,\dots m}$ with $Y = \{x^1, \dots, x^m\}$ the random boosted sample with $m \geq 10n\ln(4n)$.

Let κ be the random variable taking value $k \in \{1, \dots, 2^L\}$ with probability $p_k = \frac{|J_k|}{m}$. Define weighted least-square estimate $\tilde{u}$ with random sample $X = \{x^i \in Y \,:\, i \in J_\kappa\}$.

$$\mathbb{E}_X\left(\frac{1}{\#(X)} \sum_{x^i \in X} w(x^i)|v(x^i)|^2\right) = \mathbb{E}_Y\left(\frac{1}{m} \sum_{i=1}^m w(x^i)|v(x^i)|^2\right) \leq 2\|v\|^2, \quad v \in V.$$

This allows us to prove $\mathbb{E}(\|u - \tilde{u}\|^2) \leq C e_n(u)^2$, with sample size $|X| \leq C_0 n$.

Consequence : for any compact $\mathcal{K} \subset L^2$,

$$\rho^{\mathrm{rand}}_{C_0 n}(\mathcal{K})_{L^2} \leq C d_n(\mathcal{K})_{L^2}.$$

The RKHS setting (Krieg-M.Ullrich, Nagel-Schäffer-T.Ullrich)

Favorable comparison between deterministic sampling numbers and $d_n(\mathcal{K})_{L^2}$ can be achieved under additional assumption : $\mathcal{K}$ is a unit ball of a Reproducing Kernel Hilbert Space $\mathcal{H} \subset L^2(D,\mu)$ (that is, point evaluation is continuous on $\mathcal{H}$, as opposed to L^2).

Assuming compact embedding of $\mathcal{H}$ into L^2, the n-widths $d_n = d_n(\mathcal{K})_{L^2}$ coincide with the decreasing eigenvalues of this embedding, associated to an L^2 orthonormal basis of eigenvectors $(L_j)_{j\geq 1}$. Optimal n-width spaces are $V_n := \mathrm{span}\{L_1,\dots,L_n\}$.

Assume that $\sum_{n\geq 1} d_n^2 < \infty$ and introduce the modified optimal measure

$$d\sigma_n := \frac{1}{2}\Big(\frac{1}{n}\sum_{j=1}^n |L_j|^2 + \frac{\sum_{j>n} d_j^2 |L_j|^2}{\sum_{j>n} d_j^2}\Big)d\mu.$$

Draw $x^1,\dots,x^m$ according to σ_n and reconstruct by weighted least-squares on V_n,

If $m \sim n\ln(n)$, then with high probability, one has the estimate

$$\|u-\tilde u\|_{L^2}^2 \leq C\frac{\ln(n)}{n}\sum_{j\geq n} d_j^2, \quad u\in\mathcal{K}.$$

Dolbeault, Krieg, M. Ullrich (2022) : with additional Kadison-Singer sparisfication improved bound $n^{-1}\sum_{j\geq n} d_j^2$ and budget $m = cn$. In turn

$$\rho^{\det}_{C_0 n}(\mathcal{K})_{L^2} \leq C\Big(n^{-1}\sum_{k\geq n} d_k(\mathcal{K})_{L^2}^2\Big)^{-1/2}.$$

Summary

We can improve sparsity of the sample up to near-optimality $m \sim n$.

This comes at the prize of computational feasability of the offline sample generation.

sampling complexity	sample cardinality m	offline complexity	$\mathbb{E}(\|u-\tilde{u}\|^2) \leq Ce_n(u)^2$	$\|u-\tilde{u}\|^2 \leq Ce_n(u)_\infty^2$
conditionned $\rho^{\otimes m} \mid E$	$10n\ln(4n)$	$\mathcal{O}(n^3\ln(n))$	✓	✓
+ deterministic sparsification	$(1+\varepsilon)n$	$\mathcal{O}(n^4\ln(n))$	✗	✓
+ randomized sparsification	$C_0 n$	$\mathcal{O}(n^{cn}) \to \mathcal{O}(n^r)$?	✓	✓

Conflict between reducing sampling budget and limiting offline computational cost.

Haberstisch-Nouy-Perrin : cheap greedy sparsification but no theoretical guarantee.

Sparsification strategies do not seem to combine well with hierarchical sampling.

More general measurement models

Can we develop a similar sampling theory for other types of measurements

$$y^i = \ell_i(u), \qquad i = 1, \dots, m,$$

where ℓ_i are linear forms of some particular type ? Examples :

- Local averages $\ell_i(u) = \int_{\mathbb{R}^d} u(x)\varphi(x - x^i)$,

- Fourier samples $\ell_i(u) = \int_{\mathbb{R}^d} u(x)\exp(-i\omega^i \cdot x)$

- Radon samples $\ell_i(u) = \int_{L^i} u(s)ds$ where L^i are lines in $\mathbb{R}^2$,...

In all these examples, the linear forms are picked in a certain dictionnary where we want to make an optimal selection.

This may be viewed as applying point evaluation after a certain transformation.

$$y^i = \ell_i(u) = Ru(x^i), \qquad x^1, \dots, x^m \in D,$$

where D is now the transformed domain. For example $D = [0, \pi[\times\mathbb{R}$ for the Radon transform on $\mathbb{R}^2$.

Optimal measurement selection in transformed space

We assume $u \mapsto Ru$ to be a "stable" representation of u for a Hilbert space V of interest, in the sense that for a certain measure μ

$$\|u\|_V^2 = \int_D |Ru(x)|^2 d\mu = \|Ru\|_{L^2(D,\mu)}^2.$$

This is the case in all above examples.

For picking the approximation $u_n \in V_n \subset V$, we now solve

$$\min_{v \in V_n} \sum_{i=1}^{m} w(x^i)|y^i - Rv(x^i)|^2.$$

The optimal sampling measure on the transformed domain is again defined by

$$d\sigma = \frac{k_n}{n} d\mu, \qquad k_n(x) = \sum_{j=1}^{n} |L_j(x)|^2,$$

however with $\{L_1, \dots, L_n\}$ now an orthonormal basis of $W_n := R(V_n)$.

With $\{x^1, \dots, x^m\}$ picked according to this sampling measure and $m \sim n$, we retrieve

$$\mathbb{E}(\|u - u_n\|_V^2) \le C e_n(u)_V^2, \qquad e_n(u)_V = \min_{v \in V_n} \|u - v_n\|_V.$$

Choosing the error norm

Several possible choices of (V,μ) lead to different sampling strategies.

For the Fourier transform : $V = H^s(\mathbb{R}^d) \iff d\mu(\omega) = (1+|\omega|^{2s})d\omega$.

For the Radon transform : taking $d\mu$ the Lebesgue measure,

$$\int_D |Ru(x)|^2 d\mu = \int_R \int_0^\pi |Ru(t,\theta)|^2 dtd\theta = \int_0^\pi \int_{\mathbb{R}} |\hat{u}(te_\theta)|^2 dsd\theta \sim \int_{\mathbb{R}^2} |\omega|^{-1} |\hat{u}(\omega)|^2 d\omega.$$

This leads to a very weak error norm $V = H^{-1/2}(\mathbb{R}^2)$.

If we want to control the error in $V = L^2(\mathbb{R}^2)$, we have

$$\|u\|_V^2 \sim \int_0^\pi |R(\theta,\cdot)|^2_{H^{1/2}(\mathbb{R})} d\theta.$$

Sobolev semi-norms may be viewed as weighted L^2 norms after applying the finite difference operator : for $0 < s < 1$

$$|v|^2_{H^s(\mathbb{R})} = \int_{\mathbb{R}\times\mathbb{R}} \frac{|v(t)-v(t')|^2}{|t-t'|^{1+2s}} dtdt' = \int_{\mathbb{R}^2} |V|^2 d\mu, \quad V(t,t') = v(t) - v(t').$$

Similar definitions for $s \geq 1$ using higher-order finite differences.

General measurements : the PBDW method

Developed by Maday and Patera in the context of parametric PDEs $\mathcal{P}(u, y) = 0$.

We observe linear measurements of the solution $u(y)$ with unknown parameter y. Here, we work in a Hilbert space V and assume that the ℓ_i are given continuous linear functionals, that is $\ell_i \in V'$.

State estimation : recover $u = u(y)$ from measurements $\ell_i(u)$.

We may write

$$\ell_i(u) = \langle u, \omega_i \rangle, \quad i = 1, \dots, m,$$

and define the measurement space

$$W := \mathrm{span}\{\omega_1, \dots, \omega_m\}.$$

The measurement data determine

$$w = P_W u \in W.$$

So we retain a low dimensional information on the complex manifold $\mathcal{M}$.

A solution algorithm is a computable map $A : w \mapsto A(w)$.

Optimal recovery : choose A to make the error $\|u - A(P_W u)\|$ uniformly small over $\mathcal{M}$.

Our prior information is that u lies in the solution manifold $\mathcal{M} := \{u(y) \ : \ y \in Y\}$.

The solution manifold is complex and its exact description is numerically out of reach.

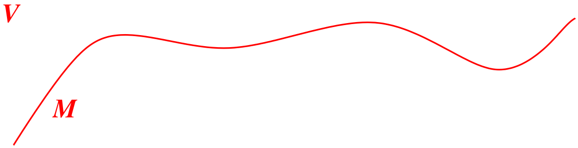

Reduced modeling methods allow us to construct (nested) linear finite dimensional spaces V_n which approximate $\mathcal{M}$ up to certified tolerances ε_n.

$$V_0 \subset V_1 \subset \cdots \subset V_n \subset \cdots, \quad \dim(V_n) = n, \quad \varepsilon_0 \geq \varepsilon_1 \geq \cdots \geq \varepsilon_n \geq \cdots \geq 0,$$

such that

$$\mathrm{dist}(u, V_n) := \min_{w \in V_n} \|u - w\| \leq \varepsilon_n, \quad n \geq 1, \quad u \in \mathcal{M},$$

Sparse polynomial methods : $u(y) \approx \sum_{\nu \in \Lambda} u_\nu y^\nu \in V_n := \mathrm{span}\{u_\nu \ : \ \nu \in \Lambda\}$.

Reduced basis methods : $V_n := \mathrm{span}\{u_i \ : \ i = 1, \dots, n\}$ with $u_i = u(a_i)$ snapshots.

Reduced model prior

Binev-Cohen-Dahmen-DeVore-Petrova-Wojtaszczyk (SIAM UQ, 2017) : replace the assumption $u \in \mathcal{M}$ by these simpler assumptions on approximability.

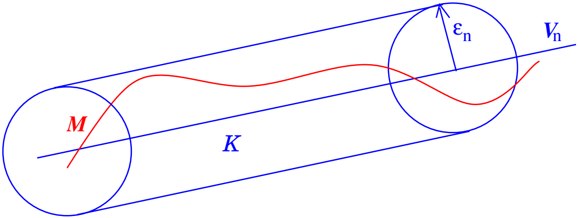

One space model : $u \in \mathcal{K} := \mathcal{K}(V_n, \varepsilon_n) := \{u \in V \ : \ \mathrm{dist}(u, V_n) \leq \varepsilon_n\}$.

Remark : affine spaces $\overline{u} + V_n$ may often be more relevant, and our discussion also applies to this case.

Our knowledge about the function u is thus that it belongs to

$$\mathcal{K}_w := \{u \in V \ : \ P_W u = w\} = \mathcal{K} \cap V_w, \quad V_w := \{u \ : \ P_W u = w\} = w + W^\perp.$$

which is an ellipsoid : intersection of the cylinder $\mathcal{K}$ with affine space V_w.

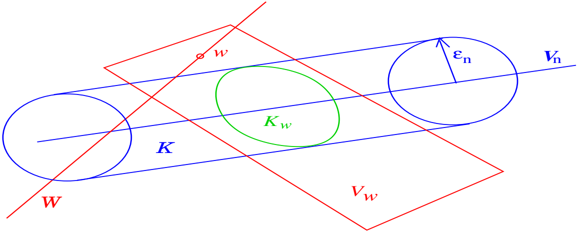

Ambiguity : all elements $u \in \mathcal{K}_w$ are assigned the same approximation $A(w)$.

Optimal recovery algorithm over $\mathcal{K}$: take $A(w)$ to be the center of the ellipsoid $\mathcal{K}_w$.

This is the "one space method", a.k.a. PBDW (Maday-Patera-Penn-Yano, 2015)

$$A(w) = \mathrm{Argmin}\{\mathrm{dist}(u, V_n) \ : \ u \in V_w\}.$$

Can be computed by a linear system, does not require the knowledge of ε_n.

Based on the inf-sup constant

$$\beta_n := \inf_{v \in V_n} \sup_{w \in W} \frac{\langle v, w \rangle}{\|v\| \, \|w\|} = \inf_{v \in V_n} \frac{\|P_W v\|}{\|v\|}.$$

or its inverse

$$\mu_n := \sup_{v \in V_n} \frac{\|v\|}{\|P_W v\|}.$$

also introduced by Adcock-Hansen (JFAA, 2012).

The error of the optimal recovery algorithm for any $u \in V$ is

$$\|u - A(P_W u)\| \leq \mu_n e_n(u)_V, \qquad e_n(u)_V = \|u - P_{V_n} u\|,$$

therefore instance optimal if we can control μ_n. On the set $\mathcal{K}$,

$$\sup_{u \in \mathcal{K}} \|u - A(P_W u)\| = \mu_n \varepsilon_n.$$

Remarks :

$\mu_n \geq 1$ and $\beta_n \leq 1$, with equality if and only if $V_n \subset W$.

$\mu_n = +\infty$ means that $V_n \cap W^\perp \neq \{0\}$. This happens if $n > m = \dim(W)$.

μ_n grows as n grows, while ε_n decreases.

Example : elliptic equation with piecewise constant diffusion field

$$-\mathrm{div}(a\nabla u) = 1 \quad \text{on} \quad [0,1]^2, \quad a = a(y) = 1 + 0.9\sum_{j=1}^{16} y_j \chi_{D_j}, \quad y = (y_j) \in [-1,1]^{16}.$$

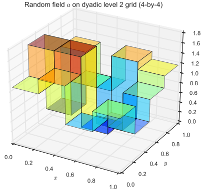

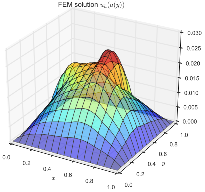

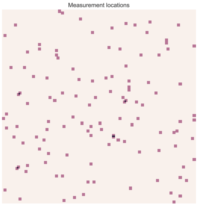

Solution with unknown parameter measured (local averages) at 120 random locations

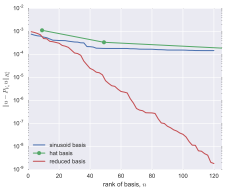

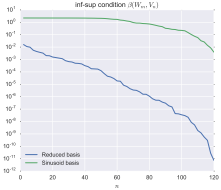

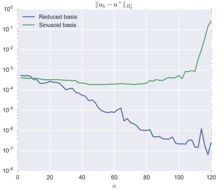

Projection error (left), Inf-Sup constant (center), Recontruction error (right)

Active data aquisition setting : optimized measurements ?

Given a reduced model space V_n we want to select the measurement functions ω_i out of a dictionnary $\mathcal{D}$ (a set of norm 1 functions, complete in V).

Objective : control $\mu(V_n, W)$, that is, guarantee a lower bound $\beta(V_n, W) > \gamma > 0$ (for example $\gamma = \frac{1}{2}$) for $W = \mathrm{span}\{\omega_1, \dots, \omega_m\}$, with $m \geq n$ as small as possible.

Benchmark : $m^*(\gamma)$ the smallest value of $m \geq n$ such that such a selection exists.

Evaluation of $\beta(V_n, W)$ requires SVD of an $n \times m$ matrix and its maximization over all possible choices of $\{\omega_1, \dots, \omega_m\}$ is computationally intensive.

Recall that

$$\beta(V_n, W) := \inf_{v \in V_n, \|v\|=1} \|P_W v\|.$$

Therefore $\beta(V_n, W) \geq \gamma > 0$ if and only if

$$\sup_{v \in V_n, \|v\|=1} \|v - P_W v\| \leq \delta := \sqrt{1-\gamma^2} < 1,$$

that is, all element of V_n should have a fixed portion of their energy captured in W.

Greedy selection : orthonormal matching pursuit

Binev-Cohen-Mula-Nichols 2017 : introduce OMP-type algorithms for selecting dictionnary elements ω_i for the collective approximation of the elements of V_n.

Collective OMP : let $\Phi = (\phi_1, \dots, \phi_n)$ be an orthonormal basis of V_n.

Having selected $\{\omega_1, \dots, \omega_k\}$ and with $W_k = \mathrm{span}\{\omega_1, \dots, \omega_k\}$, we define

$$\omega_{k+1} := \mathrm{argmax}\Big\{ \sum_{j=1}^{n} |\langle \phi_j - P_{W_k}\phi_j, \omega\rangle|^2 \ : \ \omega \in \mathcal{D} \Big\}.$$

Convergence results : if $\mathcal{D}$ is complete we always have

$$\lim_{k\to\infty} \sup_{v \in V_n, \|v\|=1} \|v - P_{W_k} v\| \to 0$$

Convergence rate $k^{-1/2}$ holds if $\Phi = \sum_{\omega\in\mathcal{D}} c_\omega \omega$ with $\sum_{\omega\in\mathcal{D}} \|c_\omega\|_2 < \infty$.

Analysis uses similar ideas as for standard OMP (DeVore-Temlyakov 1998, Barron-Cohen-Dahmen-DeVore 2007).

Example

With $V = H^1_0(]0,1[)$, consider the dictionnary $\mathcal{D}$ of (the Riesz representers of) point evaluation functionals $\ell_x(u) = u(x)$ for $x \in]0,1[$.

Take $V_n := \mathrm{span}\{s_1, \dots, s_n\}$ with $s_k(x) = \sin(\pi k x)$. In this case we can prove that a uniform sampling gives $\beta(V_n, W_m) > \gamma > 0$ with $m \sim n$.

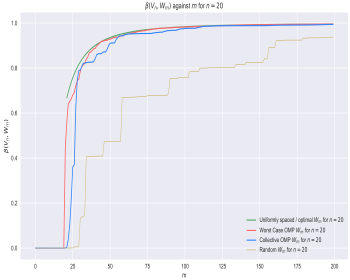

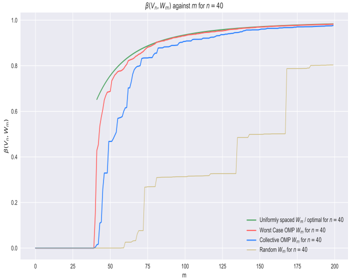

Open problem : can a greedy algorithm achieve some fixed lower bound γ with a number of measurements $m(\gamma)$ of comparable size as $m^*(\gamma)$? When can we ensure budget optimality $m(\gamma) \sim \mathcal{O}(n)$?

Limitations of the PBDW / one space method

The one space method replaces $\mathcal{M}$ by the simpler containement set $\mathcal{K}$ for which optimal recovery can be performed by simple algorithms.

1. It is a linear (or affine method). Its performance is therefore limited by below by the Kolmogorov width

$$d_m(\mathcal{M}) = d_m(\mathcal{M})_V := \min_{\dim(E)=m} \max_{u\in\mathcal{M}} \mathrm{dist}(u, E)_V.$$

2. The containement set $\mathcal{K}$ is convex, and therefore fails to capture the subtle geometry of $\mathcal{M}$.

Objective : break this limitation by nonlinear algorithms associated to non-convex reduced models.

Benchmark for optimal recovery of $\mathcal{M}$

We define the diameter of $\mathcal{M}$ from W by

$$\sigma_0 = \sigma_0(\mathcal{M}, W) = \max\{\|u - v\| \ : \ u, v \in \mathcal{M}, \ u - v \in W^\perp\}.$$

Any reconstruction algorithm A cannot achieve performance better than $\frac{1}{2}\sigma_0$.

This benchmark is not achievable by practical algorithms.

For algorithms based on a model that approximates $\mathcal{M}$ with accuracy ε, a more reasonable benchmark is

$$\sigma_\varepsilon = \sigma_\varepsilon(\mathcal{M}, W) := \max\{\|u - v\| \ : \ u, v \in \mathcal{M}_\varepsilon, \ u - v \in W^\perp\},$$

where $\mathcal{M}_\varepsilon := \mathcal{M} + B(0, \varepsilon)$ is the ε-offset of $\mathcal{M}$.

These quantities could be much smaller than the Kolmogorov width $d_m(\mathcal{M})$.

By splitting the parameter domain

$$\mathcal{A} = \mathcal{A}_1 \cup \cdots \cup \mathcal{A}_K$$

we may construct a family of reduced models $V^1, \cdots, V^K$ each of them of dimension

$$n_k = n(V^k) \leq m,$$

such that each of them approximates the corresponding portion $\mathcal{M}_k$ of the manifold with accuracy

$$\max_{u \in \mathcal{M}_k} \mathrm{dist}(u, V^k) \leq \varepsilon_k.$$

and has bounded inverse inf-sup constant

$$\mu(V^k, W) \leq \mu_k < \infty.$$

For any prescribed $\varepsilon > 0$ and $\mu > 1$, by taking K large enough, we may impose that

$$\max_{k=1,\dots,K} \varepsilon_k \leq \varepsilon$$

and

$$\max_{k=1,\dots,K} \mu_k \leq \mu$$

An oracle estimate

To each V^k corresponds a one space algorithm A_k.

From the given data $w = P_W u$, we need to select between the reconstructions

$$u_k = A_k(w), \quad k = 0, \dots, K.$$

Note that since $u \in \mathcal{M}$ there exist $k = k(u)$ such that $u \in \mathcal{M}_k$. Therefore, for this particular k,

$$\|u - A_k u\| \le \mu_k \varepsilon_k \le \mu \varepsilon.$$

This is an oracle estimate, not feasible, since it uses the knowledge of $k(u)$.

Instead, we only know the data w and want to use it for selecting a $\overline{k} = k(w)$.

Ideally we would like to select the reconstruction that is closest to the solution manifold

$$k^* = k(w) = \mathrm{argmin}_{k=1,\dots,K} \mathrm{dist}(A_k(w), \mathcal{M}),$$

but

$$\mathrm{dist}(A_k(w), \mathcal{M}) := \min_{a \in \mathcal{A}} \|u(a) - A_k(w)\|$$

is not exactly computable.

Instead, we use the minimized residual of the parametrized PDE in the dual norm

$$\delta(A_k(w), \mathcal{M}) := \min_{a \in \mathcal{A}} \|\mathcal{P}(a, A_k(w))\|_{V'},$$

which is an equivalent quantity to $\mathrm{dist}(A_k(w), \mathcal{M})$ for uniformly coercive problems. Then for $k^* = k(w)$ minimizing $\delta(A_k(w), \mathcal{M})$, one has

$$\mathrm{dist}(A_{k^*}(w), \mathcal{M}) \le C \min_{1,\dots,K} \mathrm{dist}(A_k(w), \mathcal{M}).$$

for some fixed $C > 1$.

Theorem (Cohen, Dahmen, DeVore, Mula, Nichols, 2019) : for the above selection $k^* = k(w)$, one has the estimate

$$\|u - A_{k^*}(u)\| \le \sigma_{C\mu\varepsilon}$$

References for day 2

A. Cohen and R. DeVore, *High dimensional approximation of parametric PDEs*, Acta Numerica, 2015.

G. W. Wasilkowski and H. Wozniakowski, *The power of standard information for multivariate approximation in the randomized setting*, Math. of Comp, 2006.

A. Doostan and M. Hadigol, *Least squares polynomial chaos expansion : A review of sampling strategies*, Computer Methods in Applied Mechanics and Engineering 2018.

A. Cohen and G. Migliorati, *Optimal weighted least-squares methods*, SMAI J. of Comp. Math. 2017.

C. Haberstich, A. Nouy, and G. Perrin, *Boosted optimal weighted least-squares*, Math. Comp. 2021.

G. Migliorati, *Adaptive approximation by optimal weighted least-squares methods*, SINUM, 2019.

A. Marcus, D. Spielman and N. Srivastava, *Interlacing families II : Mixed characteristic polynomials and the Kadison-Singer problem*, Annals of Maths., 2015.

S. Nitzan, A. Olevskii and A. Ulanovskii, *Exponential frames on unbounded sets*, Proc. of the AMS, 2016.

V. N. Temlyakov, *On optimal recovery in* L^2, 2020.

N. Nagel, M. Schäfer and T. Ullrich, *A new upper bound for sampling numbers*, 2020.

A. Cohen and M. Dolbeault, *Optimal pointwise sampling for* L^2 *approximation*, 2021.

M. Dolbeault, D. Krieg and M. Ullrich, *A sharp upper bound for sampling numbers in* L_2, 2022.

P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova and P. Wojtaszczyk, *Data assimilation in reduced modeling*, SIAM JUQ, 2017.

P. Binev, A. Cohen, J. Nichols and O. Mula, *Greedy algorithms for optimal measurements selection in state estimation using reduced models* , SIAM JUQ, 2018.

A. Cohen, W. Dahmen, J. Nichols and O. Mula, *Nonlinear reduced models for state and parameter estimation* , to appear in SIAM JUQ, 2021.