Sparse methods and high dimensional parametric PDE's

Albert Cohen

Laboratoire Jacques-Louis Lions
Université Pierre et Marie Curie
Paris

Beijing 18-11-2014

What is sparsity ?

Small dimesional phenomenon in high dimensional context

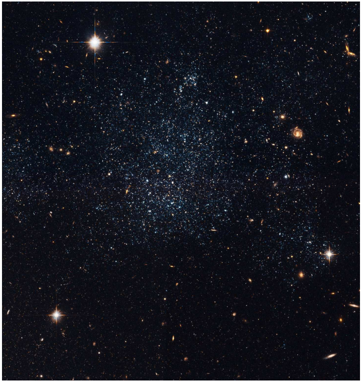

Simple example : vector $x = (x_1, \cdots, x_N) \in \mathbf{R}^{\mathrm{N}}$ representing a signal, image or function, discretized with $N >> 1$.

The vector x is sparse if only few of its coordinates are non-zero.

How to quantify this ?

The set ot k-sparse vectors

$$\Sigma_k := \{x \in \mathbb{R}^N \; ; \; \#\{i \; ; \; x_i \neq 0\} \leq k\}$$

As k gets smaller, $x \in \Sigma_k$ gets sparser.

More realistic : a vector is quasi-sparse if only a few numerically significant coordinates concentrate most of the information. How to measure this notion of concentration ?

Remarks :

A vector in Σ_k is characterized by k non-zero values and their k positions.

Intrinsically nonlinear concepts : $x, y \in \Sigma_k$ does not imply $x + y \in \Sigma_k$.

Sparsity is often hidden, and revealed through an appropriate representation (change of basis).

Importance of the concept of representation : David Marr ("Vision", Freeman, 1982).

"A representation is a formal system for making explicit certain entities or types of information, together with a specification of how the system does this... For example, the Arabic, Roman and binary numerical systems are all formal systems for representing numbers. The Arabic representation consists in a string of symbols drawn from the set 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 and the rule for constructing the description of a particular integer n is that one decomposes n into a sum of multiple of powers of 10...the alphabet allows the construction of a written representation of words... A representation, therefore is not a foreign idea at all, we all use representations all the time. However, the notion that one can capture some aspects of reality by making a description of it using a symbol and that to do so can be useful seems to me a fascinating and powerful idea...
...This issue is important, because how information is presented can greatly affect how easy it is to do different things with it. This is evident even from our number example : it is easy to add, to substract and even to multiply if the Arabic or binary representation are used, but it is not at all easy to do these things - especially multiplication - with Roman numerals. This is a key reason why the Roman culture failed to develop mathematics in the way the Arabic culture had."

The choice of an appropriate representation of a function can be fundamental to solve a specific task.

Agenda

1. Sparsity and wavelet representations (90-00)
2. Sparsity in PDE's and Images, compressed sensing (00-10)
3. High dimensional parametric PDE's (10-)

References

R. DeVore, “Nonlinear approximation”, Acta Numerica, 1998.

A. Cohen, “Numerical analysis of wavelet methods”, Studies in Mathematics and its Applications, Elsevier, 2003.

A. Cohen, W. Dahmen, I. Daubechies and R. DeVore, “Harmonic Analysis of the space BV”, Revista Matematica Iberoamericana, 2003.

A. Cohen, W. Dahmen and R. DeVore, “Compressed sensing and best k-term approximation”, Journal of the AMS, 2009.

A. Cohen, R. DeVore and C. Schwab, “Analytic regularity and polynomial approximation of parametric and stochastic PDEs”, Analysis and Application, 2011.

Fourier representations

- Analysis : $\hat{f}(\omega) = \int_{-\infty}^{+\infty} f(t) e^{-i\omega t} dt$.
- Synthesis : $f(t) = (2\pi)^{-1} \int_{-\infty}^{+\infty} \hat{f}(\omega) e^{i\omega t} d\omega$.

Representation of f in terms of the pure waves $e_\omega(t) = e^{i\omega t}$, $\omega \in \mathbb{R}$.

For 1-periodic functions :
- Analysis : $c_n(f) = \int_0^1 f(t) e^{-i2\pi nt} dt$.
- Synthesis : $f(t) = \sum_{n\in\mathbb{Z}} c_n(f) e^{i2\pi nt}$.

Discrete Fourier transform : $(x[k])_{k=0,\cdots,N-1}$ and $(\hat{x}[k])_{k=0,\cdots,N-1}$ connected by

$$\hat{x}[k] = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x[n] e^{-i2\pi nk/N} \text{ and } x[k] = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \hat{x}[n] e^{i2\pi nk/N}.$$

Implemented in $\mathcal{O}(N \log N)$ operations by FFT.

Fourier representations and computation

Approximation of a (1-periodic) function by its partial sum $S_N f(t) = \sum_{n=-N}^{N} c_n(f) e^{i2\pi nt}$.

Problem : fast convergence ?

If $f, f', \cdots, f^{(m)}$ are continuous over $\mathbb{R}$, we can apply n times the integration by part to obtain

$$\begin{aligned} |c_n(f)| &= |(i2\pi n)^{-1} c_n(f')| \\ &= \cdots |(i2\pi n)^{-m} c_n(f^{(m)})| \\ &\leq |i2\pi n|^{-m} \int_0^1 |f^{(m)}| \leq C_m n^{-m}. \end{aligned}$$

$\Rightarrow$ Fast decay if f is smooth.

However, if f is smooth everywhere except at some discontinuity point $x \in [0, 1]$, we cannot hope better than $|c_n(f)| \leq Cn^{-1}$ (also Gibbs phenomenon for $S_N f$ near the singularity).

Better representations are needed for such functions.

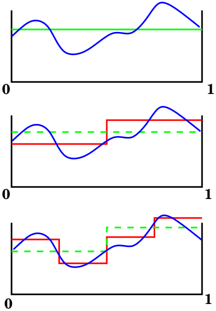

$$f = \langle f, e_0 \rangle e_0$$

$$+ \langle f, e_1 \rangle e_1$$

$$+ \langle f, e_2 \rangle e_2 + \langle f, e_3 \rangle e_3$$

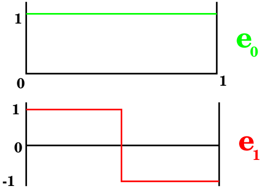

$$\cdots = \sum_\lambda f_\lambda \psi_\lambda \qquad f_\lambda := \langle f, \psi_\lambda \rangle$$

$$\psi_\lambda(x) := 2^{j/2}\psi(2^j x - k), \quad \lambda = (j,k),\ j \geq 0,\ k \in \mathbb{Z},\quad |\lambda| = \mathrm{j} = \mathrm{j}(\lambda).$$

More general wavelets are constructed from similar multiscale approximation processes, using smoother functions such as splines, finite elements...

In d dimension $\psi_\lambda(x) := 2^{dj/2}\psi(2^j x - k)$, $k \in \mathbb{Z}^{\mathrm{d}}$.

Discrete signals : fast decomposition/reconstruction algorithms

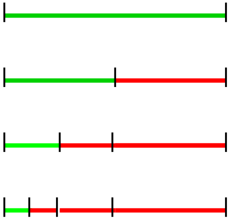

1D array $(f_0, \cdots, f_N)$
$\Rightarrow$ Two half array : averages $\frac{f_{2k}+f_{2k+1}}{2}$
and differences $\frac{f_{2k}-f_{2k+1}}{2}$

$\Rightarrow$ Iterate on the half array of averages...

Multiscale processing of 2D data : separable algorithm

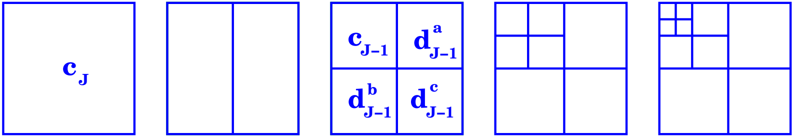

Image $f(k, l) \Rightarrow$ process lines $\Rightarrow$ process columns $\Rightarrow$ Iterate ...

Digital Image 512x512

Multiscale Decomposition

Multiscale decompositions of natural images are sparse : a few numerically significant coefficients concentrate most of the energy and information.

Application to Image Compression

Basic idea : encode with more precision
the few numerically significant coefficients
⇒Resolution is locally adapted
Example : 1 % largest coefficients encoded

Compression standard JPEG 2000 :
- Same basic principles
- Based on smoother wavelets
- Good quality with compression 1/40

Application to image denoising

Noisy digital image

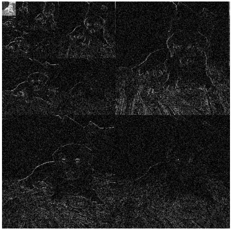

Multiscale decomposition

Natural strategy : thresholding
i.e. put to zero the coefficients which are smaller than the noise level.

Two other applications

Statistical learning : given a set of data (x_i, y_i), $i = 1, 2, \cdots, m$, drawn independently according to a probability law, build a function f such that $|f(x) - y|$ is small in the average ($E(|f(x) - y|^2)$ as small as possible).

Difficulty : build the adaptive grid from uncertain data, update it as more and more samples are received.

Adaptive numerical simulation of PDE's : Computing on a non-uniform grid is justified for solutions which displays isolated singularities (shocks).
Difficulty : the solution f is unknown. Build the grid or set of wavelet coefficients which is best adapted to the solution. Use a-posteriori information, gained throughout the numerical computation.

Measuring sparsity in a representation $f = \sum f_\lambda \psi_\lambda$

Intuition : the number of coefficients above a threshold η should not grow too fast as $\eta \to 0$.

Weak spaces : $(f_\lambda) \in w\ell^p$ if and only if

$$\mathrm{Card}\{\lambda \text{ s.t. } |f_\lambda| > \eta\} \le C\eta^{-p},$$

or equivalently, the decreasing rearrangement $(f_n)_{n>0}$ of $(|f_\lambda|)$ satisfies

$$f_n \le Cn^{-1/p}.$$

The $w\ell^p$ quasi-norm can be defined by

$$\|(f_\lambda)\|_{w\ell^p} := \sup_{n>0} n^{1/p} f_n.$$

Obviously $\ell^p \subset w\ell^p$. The representation is sparser as $p \to 0$.

If $p < 2$ and (ψ_λ) is (any) orthonormal basis in a Hilbert space H, an equivalent statement is in terms of best N-term approximation : with $f_N = \sum_{N \text{ largest } |f_\lambda|} f_\lambda \psi_\lambda$,

$$\|f - f_N\|_H = \Big(\sum_{n \ge N} |f_n|^2\Big)^{1/2} \le \|(f_\lambda)\|_{w\ell^p} \Big(\sum_{n \ge N} n^{-2/p}\Big)^{1/2} \le C\|(f_\lambda)\|_{w\ell^p} N^{-s}, \quad s = \frac{1}{p} - \frac{1}{2}.$$

Older observation by Stechkin

For the strong ℓ^p space one has

$$(f_\lambda)_{\lambda\in\Lambda} \in \ell^p \Rightarrow \|f - f_N\|_H \leq \|(f_\lambda)\|_{\ell^p}(N+1)^{-s}, \quad s = \frac{1}{p} - \frac{1}{2}.$$

Proof : using the decreasing rearrangement, we combine

$$\|f - f_N\|_H = (\sum_{n>N} f_n^2)^2 = (\sum_{n>N} f_n^{2-p} f_n^p)^{1/2} \leq f_{N+1}^{1-p/2} \|(f_\lambda)\|_{\ell^p}^{p/2}$$

and

$$(N+1) f_{N+1}^p \leq \sum_{n=1}^{N+1} f_n^p \leq \|(f_\lambda)\|_{\ell^p}^p.$$

Note that a large value of s corresponds to a value $p < 1$ (non-convex spaces).

For concrete choices of bases (such as wavelets) a relevant question is therefore : what smoothness properties of f ensure that the coefficient sequence (f_λ) belongs to ℓ^p or $w\ell^p$ for small values of p ?

Central problems in approximation theory

- X normed space.

- $(\Sigma_N)_{N \geq 0} \subset X$ approximation subspaces : $g \in \Sigma_N$ described by N (or $\mathcal{O}(N)$) parameters.

- Best approximation error $\sigma_N(f) := \inf_{g \in \Sigma_N} \|f - g\|_X$.

Problem 1 : characterise those functions in $f \in X$ having a certain rate of approximation

$$f \in X^r \Leftrightarrow \sigma_N(f) \lesssim N^{-r}$$

Here $A \lesssim B$ means that $A \leq CB$, where the constant C is independent of the parameters defining A and B.

Examples

Linear approximation : Σ_N linear space of dimension N (or $\mathcal{O}(N)$).

- $\Sigma_N := \Pi_N$ polynomials of degree N in dimension 1

- $\Sigma_N := \{f \in C^r([0,1])\ ;\ f_{|[\frac{k}{N},\frac{k+1}{N}]} \in \Pi_m,\ \ k = 0, \cdots, N-1\}$ with $0 \leq r \leq m$ fixed, splines with uniform knots.

- $\Sigma_N := \mathrm{Vect}(e_1, \cdots, e_N)$ with $(e_k)_{k>0}$ a functional basis.

Nonlinear approximation : $\Sigma_N + \Sigma_N \neq \Sigma_N$.

- $\Sigma_N := \{\frac{p}{q},\ p, q \in \Pi_N\}$ rational fractions

- $\Sigma_N := \{f \in C^r([0,1])\ ;\ f_{|[x_k,x_{k+1}]} \in \Pi_m,\ \ 0 = x_0 < \cdots < x_N = 1\}$ with $0 \leq r \leq m$ fixed, free knots splines.

- $\Sigma_N := \{\sum_{\lambda \in E} d_\lambda \psi_\lambda\ ;\ \#(E) \leq N\}$ set of all N-terms combination of a basis (ψ_λ).

Central problem in computational approximation

Problem 2 : practical realization of $f \mapsto f_N \in \Sigma_N$ such that

$$\|f - f_N\|_X \lesssim \sigma_N(f).$$

If Σ_N are linear spaces and $P_N : X \to \Sigma_N$ are uniformly bounded projectors $\|P_N\|_{X \to X} \leq C$, then $f_N := P_N f$ is a good choice, since for all $g \in \Sigma_N$,

$$\begin{aligned} \|f - f_N\|_X &\leq \|f - g\|_X + \|g - f_N\|_X \\ &= \|f - g\|_X + \|P_N(g - f))\|_X \\ &\leq (1 + C)\|g - f\|_X, \end{aligned}$$

and therefore $\|f - f_N\|_X \leq (1 + C)\sigma_N(f)$.

What about nonlinear spaces ?

A basic example

Approximation of $f \in C([0,1])$ by piecewise constant functions on a partition $I_1, \cdots, I_N$, defining

$$f_N(x) = a_k := |I_k|^{-1} \int_{I_k} f, \ \text{if} \ x \in I_k.$$

Local error : $\|f - a_k\|_{L^\infty(I_k)} \le \max_{x,y \in I_k} |f(x) - f(y)|$

Linear case : $I_k = [\frac{k}{N}, \frac{k+1}{N}]$ uniform partition.

$$f' \in L^\infty \Leftrightarrow \|f - f_N\|_{L^\infty} \le CN^{-1} \quad (C = \sup |f'|).$$

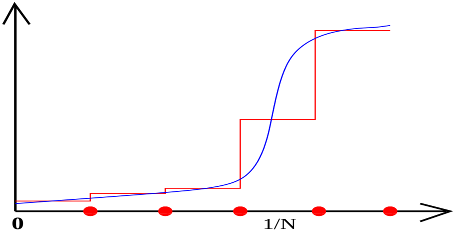

Nonlinear case : I_k free partition. If $f' \in L^1$, choose the partition such that equilibrates the total variation $\int_{I_k} |f'| = N^{-1} \int_0^1 |f'|$.

$$f' \in L^1 \Leftrightarrow \|f - f_N\|_{L^\infty} \leq CN^{-1} \quad (C = \int_0^1 |f'|).$$

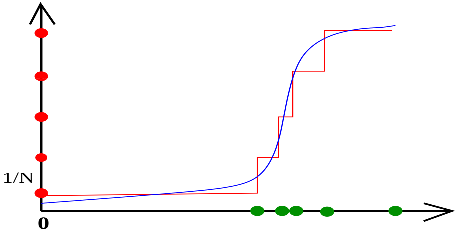

Approximation rate governed by differents smoothness spaces !

Example : $f(t) = t^\alpha$ with $0 < \alpha < 1$, then $f'(t) = \alpha t^{\alpha-1}$ is in L^1, not in L^∞.
Nonlinear approximation rate N^{-1} outperforms linear approximation rate $N^{-\alpha}$.

Adaptive greedy splitting

Split intervals I into two equal parts as long as $\|f - a_I\|_{L^\infty(I)} > \varepsilon$, the final adaptive partition is built when $\|f - a_I\|_{L^\infty(I)} \leq \varepsilon$ holds for all intervals (leaves of the decision tree).

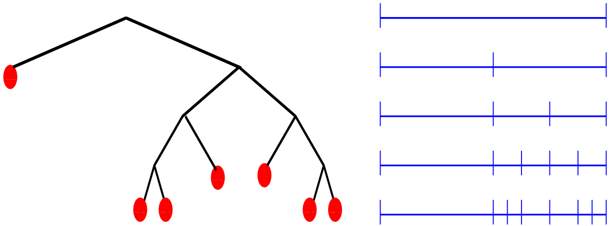

Limitation to dyadic intervals. In turn $f' \in L^1$ is not sufficient to ensure that $\|f - f_N\|_{L^\infty} \lesssim N^{-1}$, but it can be shown that the slightly stronger condition on the Hardy-Littlewood maximal function $\mathcal{M}(f') \in L^1$ suffices (holds if $f' \in L^p$ for some $p > 1$)

Approximating functions by wavelet bases

- Linear (uniform) approximation at resolution level j by taking the truncated sum $f \mapsto P_j f := \sum_{|\lambda|<j} f_\lambda \psi_\lambda$.

- Nonlinear (adaptive) approximation obtained by thresholding

$$f \mapsto \mathcal{T}_\Lambda f := \sum_{\lambda \in \Lambda} f_\lambda \psi_\lambda, \quad \Lambda = \Lambda(\eta) = \{\lambda \ \text{ s.t. } \ |f_\lambda| \geq \eta\}.$$

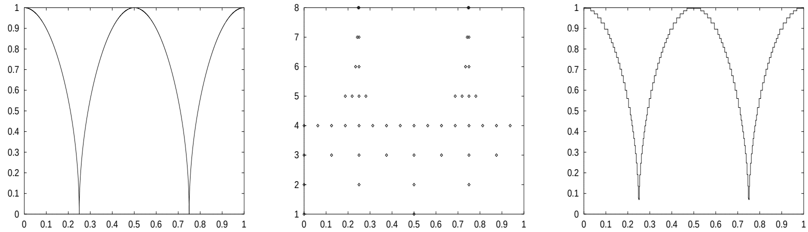

Wavelet analysis of local smoothness

- If f is bounded on $S_\lambda := \mathrm{Supp}(\psi_\lambda)$, an obvious estimate is

$$|f_\lambda| = |\langle f, \psi_\lambda\rangle| \le \sup_{t\in S_\lambda} |f(t)| \int |\psi_\lambda| = 2^{-|\lambda|/2} \sup_{t\in S_\lambda} |f(t)|.$$

- If f is C^1 on S_λ, a finer estimate is

$$\begin{aligned} |f_\lambda| &= \inf_{c\in\mathbb{R}} |\langle f - c, \psi_\lambda\rangle| \\ &\le \inf_{c\in\mathbb{R}} \|f - c\|_{L^\infty(S_\lambda)} \|\psi_\lambda\|_{L^1} \\ &\le 2^{-3|\lambda|/2} \sup_{t\in S_\lambda} |f'(t)|. \end{aligned}$$

- If f is Hölder continuous of exponent s on S_λ, i.e. $|f(x) - f(y)| \le C|x-y|^s$, for some $s \in (0,1]$, we have the intermediate estimate $|f_\lambda| \le C 2^{-|\lambda|(s+1/2)}$.

Decay of wavelet coefficients influenced by local smoothness (as opposed to that of Fourier coefficients).

A general framework

Mallat and Meyer (1986) : a multiresolution approximation (MRA) is a sequence of nested spaces $V_j \subset V_{j+1} \subset \cdots$ of $L^2(\mathbb{R}^d)$, such that :

- $\overline{\cup V_j} = L^2$, i.e. $\lim_{j\to+\infty} \|f - P_j f\|_{L^2} = 0$ for all $f \in L^2$ where P_j is the L^2-orthogonal projector.

- There exists a scaling function $\varphi \in V_0$ such that

$$\varphi_{j,k}(t) = 2^{j/2}\varphi(2^j t - k), \quad k \in \mathbb{Z}^d,$$

constitute a Riesz basis of V_j.

Riesz basis in Hilbert spaces : basis (e_n) such that $\|(x_n)\|_{\ell^2} \sim \|\sum x_n e_n\|_H$.

For piecewise constant functions we had $\varphi = \chi_{[0,1]}$. In this case

$$\|f - P_j f\|_{L^p} \leq 2^{-j}\|f'\|_{L^p},$$

but no better rate such as $2^{-mj}\|f^{(m)}\|_p$ (first order accuracy).

Raising the accuracy : V_j should contain higher order polynomials. Example : B-spline of degree N

$$\varphi(x) = \chi_{[0,1]} * \cdots * \chi_{[0,1]} = (*)^{N+1}\chi_{[0,1]},$$

Remark : except for $N = 0$, the functions $\varphi_{j,k}$ are not orthogonal. In turn the orthogonal projector P_j is not local. New difficulties :

- Define numerically simple projectors P_j onto V_j.

- Construct wavelet bases (ψ_λ) which characterize the difference between two successive levels of projection so that

$$f = P_0 f + \sum_{j\geq 0} Q_j f, \quad Q_j f := P_{j+1} f - P_j f = \sum_{|\lambda|=j} f_\lambda \psi_\lambda$$

Recall that $\psi_\lambda(x) = 2^{dj/2}\psi(2^j x - k)$ and $|\lambda| := j$ when $\lambda = (j, k)$.

Several approaches : orthogonal wavelets, biorthogonal wavelets, finite element wavelets...

Can be adapted to a bounded domain $\Omega \subset \mathbb{R}^d$. Then $\dim(V_j) \sim 2^{jd}$.

Wavelet characterizations of functions spaces

Let $f = \sum f_\lambda \psi_\lambda$, $f_\lambda = \langle f, \tilde{\psi}_\lambda \rangle$.

- L^2 characterized by $\|f\|_{L^2}^2 \sim \|P_0 f\|_{L^2}^2 + \sum_{j\geq 0} \|Q_j f\|_{L^2}^2 \sim \sum |f_\lambda|^2$.

- Sobolev space $H^s = W^{s,2}$ characterized by

$$\|f\|_{H^s}^2 \sim \|P_0 f\|_{L^2}^2 + \sum_{j\geq 0} 2^{2sj} \|Q_j f\|_{L^2}^2 \sim \sum 2^{2s|\lambda|} |f_\lambda|^2 \sim \sum \|f_\lambda \psi_\lambda\|_{H^s}^2.$$

Hint : $\|f\|_{H^s}^2 \sim \int (1 + |\omega|^{2s}) |\hat{f}(\omega)|^2 \sim \|S_0 f\|_{L^2}^2 + \sum_{j\geq 0} 2^{2sj} \|\Delta_j f\|_{L^2}^2$ with $\widehat{S_j f}(\omega) \sim \hat{f}(\omega) \chi_{|\omega| \leq 2^j}$ and $\Delta_j f = S_{j+1} f - S_j f$ (Littlewood-Paley analysis).

- Besov space $B^s_{p,p}$ characterized by

$$\begin{aligned} \|f\|_{B^s_{p,p}}^p &\sim \|P_0 f\|_{L^p}^p + \sum_{j\geq 0} 2^{psj} \|Q_j f\|_{L^p}^p \sim \sum 2^{ps|\lambda|} \|f_\lambda \psi_\lambda\|_{L^p}^p \\ &\sim \sum 2^{ps|\lambda|} 2^{pd(\frac{1}{2} - \frac{1}{p})|\lambda|} |f_\lambda|^p \sim \sum \|f_\lambda \psi_\lambda\|_{B^s_{p,p}}^p. \end{aligned}$$

Remark : $B^s_{p,p} = W^{s,p}$ if $s \notin \mathbb{N}$ or $p = 2$ and $B^s_{\infty,\infty} = C^s$ if $s \notin \mathbb{N}$.

All this holds provided that ψ_λ has enough smoothness

Linear multiscale approximation

From the characterization of H^s, we get $\|Q_j f\|_{L^2} \lesssim 2^{-js}\|f\|_{H^s}$ and therefore

$$f \in H^s = B^s_{2,2} \Rightarrow \|f - P_j f\|_{L^2} \leq \sum_{l \geq j} \|Q_l f\|_{L^2} \lesssim 2^{-tj}.$$

and in a similar manner

$$f \in B^s_{p,p} \Rightarrow \|f - P_j f\|_{L^p} \lesssim 2^{-sj}.$$

We actually have a finer result

$$f \in B^s_{p,q} \Leftrightarrow (2^{sj}\|f - P_j f\|_{L^p})_{j \geq 0} \in \ell^q.$$

Besov spaces are thus characterized from the rate of linear multiscale approximation.

These results are very similar to finite element approximation on uniform meshes ($V_j \sim V_h$ with $h \sim 2^{-j}$).

On a bounded domain, they roughly say that s order of smoothness in L^p corresponds to a linear approximation rate $\mathcal{O}(N^{-s/d})$ in $\Sigma_N = V_j$ where $N = \dim(V_j) \sim 2^{dj}$.

Nonlinear wavelet approximation in L^2

Recall that $B^s_{p,p}$ is characterized by

$$\|f\|^p_{B^s_{p,p}} \sim \sum 2^{ps|\lambda|} 2^{pd(\frac{1}{2}-\frac{1}{p})|\lambda|} |f_\lambda|^p$$

Assume that $f \in B^s_{p,p}$ with $\frac{1}{p} = \frac{1}{2} + \frac{s}{d}$. In this case

$$\|f\|_{B^s_{p,p}} \sim \|(f_\lambda)\|_{\ell^p},$$

and therefore $(f_\lambda) \in \ell^p \subset w\ell^p$. If $f_N := \sum_{N \text{ largest } |f_\lambda|} f_\lambda \psi_\lambda$, we have

$$\|f - f_N\|_{L^2} \lesssim N^{-s/d}.$$

For linear approximation, the same rate is achieved under the stronger condition $f \in H^s$.

Note that the relation $\frac{1}{p} = \frac{1}{2} + \frac{s}{d}$ corresponds to the critical (non-compact) embedding $B^s_{p,p} \subset L^2$, expressed in the wavelet representation by the elementary inclusion $\ell^p \subset \ell^2$.

N-terms approximations : $\Sigma_N := \{\sum_{\lambda\in\Lambda} d_\lambda\psi_\lambda \; ; \; \#(\Lambda) \leq N\}$.

- Rate of decay governed by weaker smoothness conditions (DeVore) : with $\frac{1}{q} = \frac{1}{p} + \frac{s}{d}$

$$f \in B^s_{q,q} \Rightarrow \inf_{g\in\Sigma_N} \|f - g\|_{L^p} \leq CN^{-s/d}.$$

- Similar results when approximation is measured in smoother norms ($W^{s,p}$, $B^s_{p,q}$..)

- Similar theory for adaptive finite element on N simplices with isotropy constraints (minimal angle condition).

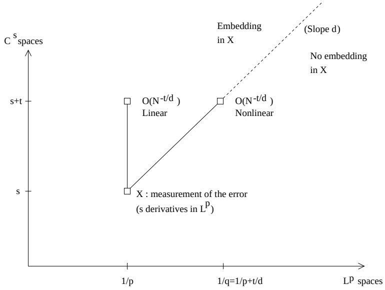

Greedy bases

Let (ψ_λ) be a basis in a Banach space X with $\|\psi_\lambda\|_X = 1$ for all λ.

The basis is greedy if and only if for all $f \in X$ and $N > 0$,

$$\|f - \sum_{N \text{ largest } |f_\lambda|} f_\lambda \psi_\lambda\|_X \leq C \inf_{g \in \Sigma_N} \|f - g\|_X.$$

The basis is unconditional if and only there exists $C > 0$ such that

$$|x_\lambda| \leq |y_\lambda| \text{ for all } \lambda \Rightarrow \|\sum x_\lambda \psi_\lambda\|_X \leq C \|\sum y_\lambda \psi_\lambda\|_X.$$

The basis is democratic if and only if there exists $C > 0$ such that

$$\#(E) = \#(F) \Rightarrow \|\sum_{\lambda \in E} \psi_\lambda\|_X \leq C \|\sum_{\lambda \in F} \psi_\lambda\|_X.$$

Two results due to Temlyakov (2003) :

1. Greedy $\Leftrightarrow$ unconditional and democratic.

2. Conveniently normalized wavelet are greedy in $X = L^p$ or $X = W^{m,p}$ when $1 < p < +\infty$, and in all Besov spaces $X = B^s_{p,q}$.

General program for PDE's

- Theoretical : revisit regularity theory for PDE's. Solutions of certain PDE's might have substantially higher regularity in the scale governing nonlinear approximation than in the scale governing linear approximation. Examples : hyperbolic conservation laws (DeVore and Lucier 1987), elliptic problems on corner domains (Dahlke and DeVore, 1997).

- Numerical : develop for the unknown u of the PDE $\mathcal{F}(u) = 0$ appropriate adaptive resolution strategies which perform essentially as well as thresholding : produce $\tilde{u}_N$ with N terms such that $\|u - \tilde{u}_N\|$ has the same rate of decay N^{-s} as $\|u - u_N\|$ in some prescribed norm, if possible in $\mathcal{O}(N)$ computation.

Remark : similar goals can be formulated for adaptive finite elements with N being the number of elements.

Revisiting regularity theory for PDE's

Solutions of certain PDE's might have substantially higher regularity in the scale governing nonlinear approximation than in the scale governing linear approximation.

Example : 1D nonlinear conservation law

$$\partial_t u + \partial_x F(u) = 0, \quad u(x,0) = u_0(x), .$$

with F smooth and strictly convex (e.g. Burger $F(u) = u^2/2$).

- Smoothness for linear approximation in L^1 : for large t, $u(\cdot, t) \in BV$ but not smoother.

- Smoothness for nonlinear approximation (DeVore & Lucier, 1987) : for all $s > 0$ and $\frac{1}{p} = 1 + s$, if $u_0 \in B^s_{p,p}$ then $u(\cdot, t) \in B^s_{p,p}$ for all $t > 0$.

Similar results are available for elliptic PDE's on non-smooth domains (DeVore & Dahlke)

Pictorial interpretation

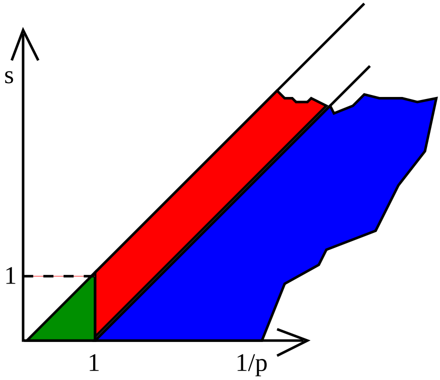

Classical theory : $s < 1/p$ for $s \leq 1$

DeVore-Lucier : $s < 1/p - 1$ for all $s > 0$

Interpolation : $s < 1/p$ for all $s > 0$

Principle of proof : approximation by adaptive piecewise polynomials

Simplest case : Burgers' equation ($F(u) = u^2/2$) and piecewise affine ($s < 2$).

$u_0 \in B^s_{p,p} \Rightarrow \|u_0 - u_0^N\|_{L^1} \lesssim N^{-s}$, with u_0^N piecewise affine on N intervals.

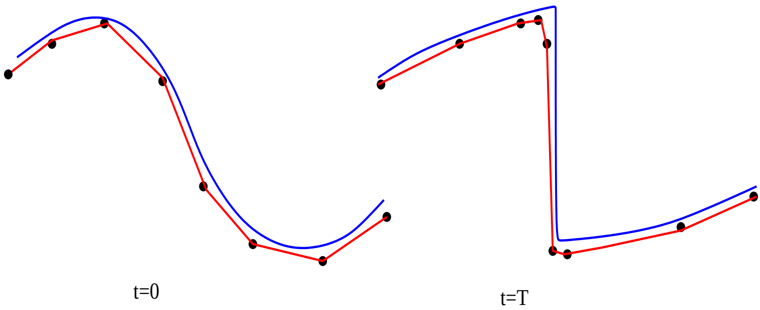

Evolution to time $T > 0$ is L^1 contractive. $\Rightarrow \|u_T - u_T^N\|_{L^1} \lesssim N^{-s} \Rightarrow u_T \in B^s_{p,p}$.

$f \in BV$ if and only if $f \in L^1$ and ∇I is a finite measure.

$\|f\|_{BV} = \|f\|_{L^1} + |f|_{BV}$ with $|f|_{BV} = \sup_{\|g\|_{L^\infty} \le 1} \int f \operatorname{div} g$.

If $f \in W^{1,1}$, i.e. $\nabla f \in L^1$, then $|f|_{BV} = \int |\nabla f|$.

Prototype : χ_Ω where $\partial\Omega$ has finite length.

In d dimensions $BV \subset L^{d^*}$ with $d^* = \frac{d}{d-1}$. In 2d, this space is often used as a model to describe real images. Intuition : Images are "piecewise smooth" and their singularities (edges) have finite total length.

Co-area formula : $|f|_{BV} = \int_{-\infty}^{+\infty} |\chi_{\Omega_t}|_{BV} dt$ ($= \int_{-\infty}^{+\infty} \mathcal{H}^1(\partial\Omega_t) dt$ for smooth functions), with $\Omega_t := \{x \, ; \, f(x) > t\}$.

This is an instance of an atomic decomposition in a Banach space B : dense set of functions $(\varphi_\gamma)_{\gamma \in \Gamma}$ such that

$$\|f\|_B \sim \inf \Big\{ \sum_{\gamma \in \Gamma} \|c_\gamma \varphi_\gamma\|_B \; : \; \sum_{\gamma \in \Gamma} c_\gamma \varphi_\gamma = f \Big\}.$$

Here, the sum is replaced by an integral with the atoms being the characteristic functions χ_Ω since we may write for $f(x) = \lim_{A \to -\infty} \Big(A + \int_A^{+\infty} \chi_{\Omega_t}(x) \Big)$.

Wavelet analysis of BV

Theorem (DeVore, Petrushev, Xu, Dahmen, Daubechies, AC)

$$f \in BV([0,1]^2) \Rightarrow (f_\lambda) \in w\ell^1$$

where (f_λ) are its wavelet coefficients, or equivalently

$$\|f - f_N\|_{L^2} \lesssim N^{-1/2}.$$

BV is almost characterized by wavelets since $(f_\lambda) \in \ell^1 \Rightarrow f \in BV([0,1]^2)$ (no simple exact characterization : BV has no unconditional basis).

Optimal estimate for wavelets : if $f = \chi_\Omega$ then at scale 2^{-j} there are $\mathcal{O}(2^j)$ nonzero coefficients (edges) estimated by $\mathcal{O}(2^{-j})$.

Optimal estimate among all bases

The case of Fourier coefficients (Lebeau) :

$$f \in BV([0,1]^2) \Rightarrow \sum_{n \in \mathbb{Z}^2} \frac{|c_n(f)|}{1+|n|} < \infty.$$

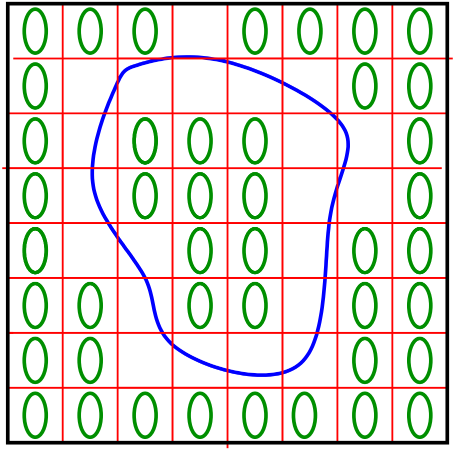

Proof by co-area formula ?

For the expansion of a single atom $\chi_\Omega = \sum d_\lambda(\Omega)\psi_\lambda$, one has

$$\|(d_\lambda(\Omega))\|_{w\ell^1} \leq C|\chi_\Omega|_{BV}.$$

Now we use the representation $f(x) = \lim_{A\to-\infty}\Big(A + \int_A^{+\infty}\chi_{\Omega_t}(x)\Big)$ and write

$$f_\lambda = \int_{-\infty}^{+\infty} d_\lambda(\Omega_t))dt.$$

Then use the co-area formula to write

$$\|(f_\lambda)\|_{w\ell^1} \leq \int_{-\infty}^{+\infty} \|(d_\lambda(\Omega_t))\|_{w\ell^1} dt \leq C\int_{-\infty}^{+\infty} |\chi_{\Omega_t}|_{BV} dt = C|f|_{BV}.$$

Unfortunately, not so simple...

In dimension $d = 2$, one has the continuous embedding $BV \subset L^2$ with

$$\|f\|_{L^2} \leq C|f|_{BV}.$$

Not sharp for oscillatory functions : if $f_\omega(x) = e^{i\omega \dot{x}}\varphi(x)$ with $\varphi \in \mathcal{D}(\mathbb{R}^2)$, one has $\|f_\omega\|_{L^2} = \|\varphi\|_{L^2}$ and $|f_\omega|_{BV} \sim |\omega|$.

We introduce the Besov space $B^{-1}_{\infty,\infty}$ which is defined by Littlewood-Paley theory or by the wavelet characterization

$$\|f\|_{B^{-1}_{\infty,\infty}} \sim \|(f_\lambda)\|_{\ell^\infty}$$

Then the Gagliardo-Nirenberg type inequality

$$\|f\|_{L^2} \leq C|f|_{BV}\| \, \||f\|_{B^{-1}_{\infty,\infty}},$$

follows from $\|(f_\lambda)\|_{\ell^2} \leq C\|(f_\lambda)\|_{w\ell^1}\|(f_\lambda)\|_{\ell^\infty}$. Sharper : $\|f_\omega\|_{B^{-1}_{\infty,\infty}} \sim |\omega|^{-1}$

One also has the real interpolation result

$$L^2 = [B^{-1}_{\infty,\infty}, BV]_{\frac{1}{2},2}.$$

following from the fact that $\ell^2 = [\ell^1, \ell^\infty]_{\frac{1}{2},2} = [w\ell^1, \ell^\infty]_{\frac{1}{2},2}$.

Practical implications in image processing

Optimal performances of wavelet adaptive denoising and compression methods when the images are modeled by BV functions.

Yves Meyer : "In a world where images are BV functions and the eye measures the error in L^2, wavelets are the best tool".

Toward better models : Image = geometry + texture.

Geometry (objects) : should take into account the smoothness of edges (ignored in BV modeling).

Texture (or noise) : should involve statistical modeling, and a different error measure than for geometry.

Image : $f = \chi_\Omega$, with $\partial\Omega$ smooth.

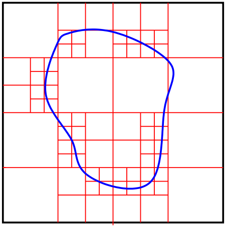

f_N = approximation by N largest wavelet coefficients

$\Rightarrow \|f - f_N\|_{L^2} \sim N^{-1/2}$

Problem : imposes isotropic refinement

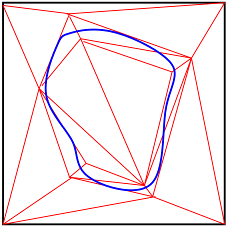

f_N = piecewise linear interpolation on N optimaly selected triangles

$\Rightarrow \|f - f_N\|_{L^2} \sim N^{-1}$

Problem : non-supervised algorithm ?

Greedy algorithms for adaptive triangulations

Optimal triangulation : NP hard problem.

Adaptive refinement algorithms : from an initial coarse triangulation $\mathcal{T}_0$, add points iteratively, e.g. at the location where the interpolation error is the largest (A.C., Dyn, Hecht, Mirebeau).

Adaptive coarsening algorithms : from a very fine triangulation $\mathcal{T}_0$, remove points iteratively. Criterion for point removal : minimize the anticipated approximation error when retriangulating (using e.g. Delaunay triangulation, Dyn-Floater, Iske).

Algorithms stop when reaching the minimal number of triangles N for which a prescribed L^2 error D is ensured.

Open problem : do greedy algorithms allow to obtain the rate $D \leq CN^{-1}$ for piecewise smooth functions such as χ_Ω ?

Sparse geometric representations

- Donoho and Candes : sparse representations based on curvelets frames allow us to recover $\|f - f_N\|_{L^2} \sim N^{-1}[\log N]^{3/2}$ with a thresholding algorithm for piecewise C^2 functions with C^2 edges (curvelet coefficients are - roughly - in $w\ell^{2/3}$). Closely related : contourlets (Do and Vetterli), shearlets (Kuttyniok).

- Other approaches : bandlets (Mallat and Le Pennec), nonlinear multiscale representations (Arandiga, Donat, A.C.), wedgelets (Donoho), wedgeprints (Baraniuk, Romberg, Wakin), nonlinear lifting scheme (Baraniuk, Claypoole, Davis and Sweldens).

This area of research lack solid functional analytic foundations : are there simple function spaces describing piecewise smooth functions with geometrically smooth edges ?

Exploiting sparsity in a different way

Assume that f is a sparse signal or image (in some basis).

Classical way to encode f : retain its k largest coordinates in the basis and encode them. This requires to compute all coordinates before discarding the small one.

Compressed sensing (Donoho, Candes-Tao-Romberg) : use m linear measurements of f prescribed in advance, and exploit that f is sparse in order to reconstruct it accurately from these measurements.

In other word, we observe $y = \Phi(f) \in \mathbb{R}^{\mathrm{m}}$ with Φ a fixed measurement matrix and we want to build $g = \Delta(y)$ close to f.

Key ingredient : Δ should be nonlinear.

An instructive example : 2D tomography (Candes-Romberg-Tao)

The Radon transform captures partial Fourier information.

 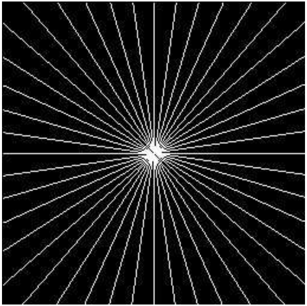

Left : the Logan-Shep phantom test image

Right : position of the observed Fourier coefficients (white)

Two different reconstructions

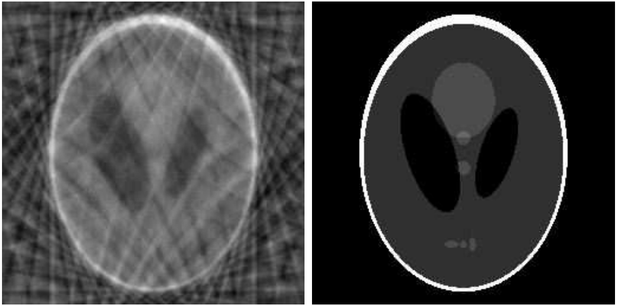

Left : put the unknown coefficient to zero (minimum ℓ^2 norm) and reconstruct the partial Fourier serie $\Rightarrow$ oscillation artifacts.

Right : adjust the unknown coefficients so to minimize the total variation of the image $|f|_{TV} = \int |\nabla f|$ $\Rightarrow$ exact reconstruction !

Questions

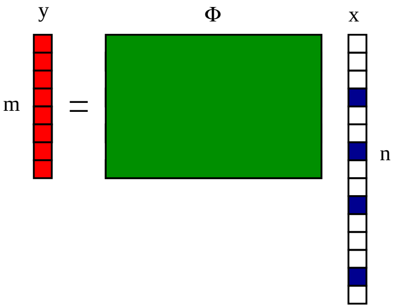

Minimal number m of measures which is sufficient to characterize any $x \in \Sigma_k$.

With which matrices Φ ? Which decodes Δ ?

Robustestness ? In practice, $y = \Phi x + e$ with $\|e\|_{\ell^2} \leq \varepsilon$ and $x \in \mathbb{R}^n$ close to Σ_k.

Available results

With $m = 2k$ measures and generic choice of Φ, one can reconstruct exactly any $x \in \Sigma_k$, but...

(i) Complex decoder : $\Delta(y) := \text{Argmin}\{\|y - \Phi z\| \; ; \; z \in \Sigma_k\}$, and therefore $\mathcal{O}(N^k)$ least square systems to solve. Alternative : $\Delta(y) := \text{Argmin}\{\|z\|_0 \; ; \; \Phi z = y\}$, with $\|z\|_0 = \#\{i \; ; \; z_i \neq 0\}$, same complexity.

(ii) No robustness to noise and deviation from Σ_k.

With $m \sim ck \log(N/k)$ measures and specific choices of Φ, one can reconstruct exactly any $x \in \Sigma_k$, with (Candes-Tao)

(i) Simple decoder : $\Delta(y) := \text{Argmin}\{\|z\|_1 \; ; \; \Phi z = y\}$ with $\|z\|_1 := |z_1| + \cdots + |z_n|$, convex optimization, linear programming.

(ii) Robustness : $\|x - \Delta(\Phi x)\|$ controlled by noise and deviation of x from Σ_k.

but... Φ obtained by probabilistic techniques. Example : $\Phi = (\Phi_{i,j})$ with $\Phi_{i,j}$ independant random draws of Bernoulli ± 1 or gaussians $\mathcal{N}(0, 1)$.

Deterministic constructions ?

Consider a continuous function $y \mapsto u(y)$ with $y \in [0,1]$.
Sample at equispaced points.
Reconstruct, for example by piecewise linear interpolation.

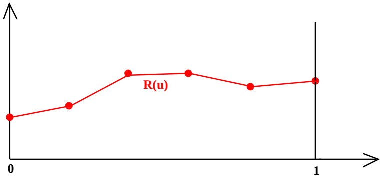

Error in terms of point spacing $h > 0$: if u has C^2 smoothness

$$\|u - R(u)\|_{L^\infty} \le C\|u''\|_{L^\infty} h^2.$$

Using piecewise polynomials of higher order, if u has C^m smoothness

$$\|u - R(u)\|_{L^\infty} \le C\|u^{(m)}\|_{L^\infty} h^m.$$

In terms of the number of samples $N \sim h^{-1}$, the error is estimated by N^{-m}.

In d dimensions : $u(y) = u(y_1, \cdots, y_d)$ with $y \in [0,1]^d$. With a uniform sampling, we still have

$$\|u - R(u)\|_{L^\infty} \le C\|d^m u\|_{L^\infty} h^m,$$

but the number of samples is now $N \sim h^{-d}$, and the error estimate is in $N^{-m/d}$.

Other sampling/reconstruction methods cannot do better

Can be explained by N-width

Let X be a normed space and $K \subset X$ a compact set.

Linear N-width (Kolmogorov) :

$$d_N(K)_X := \inf_{\dim(E)=N} \max_{u \in K} \min_{v \in E} \|u - v\|_X.$$

Benchmark for linear approximation methods applied to the elements from K.

If $X = L^\infty([0,1]^d)$ and K is the unit ball of $C^m([0,1]^d)$ it is known that

$$cN^{-m/d} \leq d_N(K)_X \leq CN^{-m/d}.$$

Upper bound : approximation by a specific method.

Lower bound : diversity in K.

Exponential growth in d of the needed complexity to reach a given accuracy.

Non-linear methods cannot do better

Use a notion of nonlinear N-width (Alexandrov, DeVore-Howard-Micchelli).

Consider maps $E : K \mapsto \mathbb{R}^N$ (encoding) and $R : \mathbb{R}^N \mapsto X$ (reconstruction).

Introducing the distorsion of the pair (E, R) over $\mathcal{K}$

$$\max_{u \in K} \|u - R(E(u))\|_X,$$

we define the nonlinear N-width of $\mathcal{K}$ as

$$\delta_N(K)_X := \inf_{E,R} \max_{u \in K} \|u - R(E(u))\|_X,$$

where the infimum is taken over all continuous maps (E, R). Comparison with the Kolmorgorov N-width : $\delta_N \leq d_N$ and sometimes substantially smaller.

If $X = L^\infty([0,1]^d)$ and K is the unit ball of $C^m([0,1]^d)$ it is known that

$$cN^{-m/d} \leq \delta_N(K)_X \leq CN^{-m/d}.$$

Many other variants of N-widths exist (book by A. Pinkus).

Infinitely smooth functions

Nowak and Wozniakowski : if $X = L^\infty([0,1]^d)$ and

$$K := \{u \in C^\infty([0,1]^d) \ : \ \|\partial^\nu u\|_{L^\infty} \leq 1 \ \text{for all} \ \nu\}.$$

then, for the linear width,

$$\min\{N \ : \ d_N(K)_X \leq 1/2\} \geq c2^{d/2}.$$

High dimensional problems occur frequently :

PDE's with solutions $u(x, v, t)$ defined in phase space : $d = 7$.

Post-processing of numerical codes : u solver with imput parameters $(y_1, \cdots, y_d)$.

Learning theory : u regression function of imput parameters $(y_1, \cdots, y_d)$

In these applications d may be of the order up to 10^3.

Approximation of stochastic-parametric PDEs : $d = +\infty$.

Smoothness properties of functions should be revisited by other means than C^m classes, and appropriate approximation tools should be used.

Key tools : (i) Sparsity, (ii) Variable reduction, (iii) Anisotropy

Parametric and stochastic PDE's

We are interested in PDE's of the general form

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

where u is the unkown and a is a parameter which is either finite or infinite dimensional. Typically

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

where V, X, W are Banach spaces and a ranges in some compact set $K \subset X$.

Model 1 : steady state linear diffusion equation.

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

where $f \in L^2(D)$ is fixed.

In this example $\mathcal{P}(u, a) = f + \mathrm{div}(a\nabla u)$ and the spaces are

$$X = L^\infty(D), \quad V = H^1_0(D), \quad W = V' = H^{-1}(D).$$

The solution map

We also assume well-posedness of the problem in the Banach space V for every $a \in K$. This allows us to define

$$a \mapsto u(a)$$

which is the solution map from K to V.

For Model 1, this is done by assuming that

$$0 < r \leq a \leq R, \qquad a \in K \subset X = L^\infty(D).$$

Then Lax-Milgram theory ensures existence in $V = H_0^1(D)$.

A priori bound : the solution map is bounded from K to V. :

$$\|u(a)\|_V \leq C_r := \frac{\|f\|_{V'}}{r}, \quad a \in K, \quad \text{where} \quad \|v\|_V := \|\nabla v\|_{L^2}.$$

The parameter may be deterministic (control, optimization, inverse problems) or random (uncertainty modeling and propagation, risk assessment).

These applications often requires many queries of $u(a)$, and therefore in principle running many times a numerical solver. We want to avoid this.

We expand a according to

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

where $\overline{a} \in X$, and $(\psi_j)_{j\geq 1}$ is a given basis of functions from X, and assume that $y_j \in [-1,1]$ for all $a \in K$, so that $y \in U := [-1,1]^{\mathbb{N}}$ (always possible up to renormalizing the ψ_j).

Therefore, we have

$$K \subset Q = \{a(y) \ : \ y \in U\} = \Big\{\overline{a} + \sum_{j\geq 1} y_j \psi_j, \quad (y_j)_{j\geq 1} \in U\Big\}$$

In what we shall simply assume that K is exactly of the form Q (big geometrical simplification, often used as a model though).

For Model 1 well posedness over Q is ensured by the uniform ellipticity assumption for

$$(UEA) \qquad 0 < r \leq a(x,y) \leq R, \ \ x \in D, \, y \in U.,$$

where

$$a(x,y) = a(y)(x) = \overline{a}(x) + \sum_{j\geq 1} y_j \psi_j(x).$$

or equivalently $\bar{a} \in L^\infty(D)$ and $\sum_j |\psi_j(x)| \leq \bar{a}(x) - r, \ \ x \in D$.

We may thus define a new solution map

$$y \mapsto u(y) := u(a(y)),$$

from U to V, which we would like to approximate numerically.

Infinite number of variables : we face the curse of dimensionality.

Anisotropy : the variables y_j are less active when ψ_j is small.

Approximation by truncated expansions

$$u \approx \sum_{j=1}^{N} u_j \varphi_j,$$

with $\varphi_j : U \to \mathbb{R}$ and $u_j \in V$. Therefore, separable format :

$$u(x,y) = u(y)(x) \approx \sum_{j=1}^{N} u_j(x)\varphi_j(y),$$

Optimal expansion ? In $L^2(D \times U)$ provided by singular value decomposition (best low rank approximation). However, generally not accessible and the functions φ_j and u_j could be numerically complicated.

Instead, we use functions φ_j which are sparsely picked from a predefined simple dictionnary.

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

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

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.

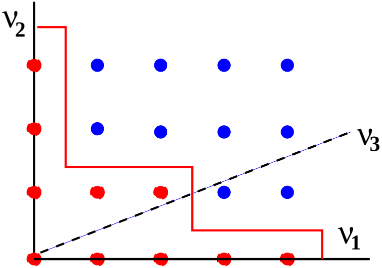

Objective : identify a set $\Lambda \subset \mathcal{F}$ with $\#(\Lambda) = N$ such that u is well approximated by the partial expansion

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

Best N-term approximation

By triangle inequality we have

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

Best N-term approximation in the $\ell^1(\mathcal{F})$ norm : use for Λ the N largest $\|t_\nu\|_V$.

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

$$\sum_{\nu\notin\Lambda} \|t_\nu\|_V \leq CN^{-s}, \quad s := \frac{1}{p}-1, \quad C := \|(\|t_\nu\|_V)\|_p.$$

Proof : with $(t_n)_{n>0}$ the decreasing rearrangement, we combine

$$\sum_{\nu\notin\Lambda} \|t_\nu\|_V = \sum_{n>N} t_n = \sum_{n>N} t_n^{1-p} t_n^p \leq t_N^{1-p} C^p \quad \text{and} \quad N t_N^p \leq \sum_{n=1}^{N} t_n^p \leq C^p.$$

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

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

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

Interpretations :

(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.

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 $(\|A_j\|_{V\to W})_{j\geq 1} \in \ell^p(\mathbb{N})$.

Key idea in the proof : holomorphic extension $z \mapsto u(z)$ with $z = (z_j) \in \mathbb{C}^{\mathbb{N}}$. Domains of holomorphy : if $\rho = (\rho_j)_{j\geq 0}$ is any positive sequence such that

$$\sum_{j>0} \rho_j |\psi_j(x)| \leq \overline{a}(x) - \delta, \quad x \in D,$$

for some $\delta > 0$, then u is holomorphic, with bound $\|u(z)\|_V \leq C_\delta$, in the polydisc

$$\mathcal{U}_\rho := \otimes\{|z_j| \leq \rho_j\},$$

Hint : Lax-Milgram theorem applies with $\Re(a(x,z)) \geq \delta > 0$.

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 $|u^{(n)}(0)| \leq n! b^{-n} \max_{|z|\leq b} |u(z)|$.

Recursive application of this to all variables z_j such that $\nu_j \neq 0$, with $b = \rho_j$, for a δ-admissible sequence ρ gives

$$\|\partial^\nu u_{|z=0}\|_V \leq C_\delta \nu! \prod_{j>0} \rho_j^{-\nu_j}.$$

and therefore

$$\|t_\nu\|_V \leq C_\delta \prod_{j>0} \rho_j^{-\nu_j} = C_\delta \rho^{-\nu}.$$

Since ρ is not fixed we have

$$\|t_\nu\|_V \leq C_\delta \inf\Big\{\rho^{-\nu} \;\; ; \;\; \rho \;\; \text{s.t.} \;\; \sum_{j\geq 1} \rho_j |\psi_j(x)| \leq \bar{a}(x) - \delta, \;\; x \in D\Big\}.$$

We do not know the general solution to this problem, except when the ψ_j have disjoint supports. Instead design a particular choice $\rho = \rho(\nu)$ satisfying the constraint with $\delta = r/2$, for which we prove that

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

A simple case

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

$$\sum_{j>0} \rho_j |\psi_j(x)| \leq \overline{a}(x) - \frac{r}{2}, \quad x \in D,$$

which leads to

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

We have

$$\|t_\nu\|_V \leq 2C_0 \rho^{-\nu} = 2C_0 b^\nu,$$

where $b = (b_j)$ and

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

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

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

Proof : factorize

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

Other models

Model 2 : same PDE but no affine dependence, e.g. $a(x,y) = \overline{a}(x) + (\sum_{j\geq 0} y_j\psi_j(x))^2$. Assuming that $\overline{a}(x) \geq r > 0$ guarantees ellipticity uniformly over $y \in U$.

Model 3 : similar problems + non-linearities, e.g.

$$g(u) - \mathrm{div}(a\nabla u) = f \quad \text{on} \quad D = D(y) \quad u_{|\partial D} = 0,$$

with same assumptions on a and f. Well-posedness in $V = H_0^1(D)$ for all $f \in V^*$ is ensured for certain nonlinearities, e.g. $g(u) = u^3$ of u^5 in dimension $m = 3$ ($V \subset L^6$).

Model 4 : PDE's on domains with parametrized boundaries, e.g.

$$-\Delta v = f \quad \text{on} \quad D = D_y \quad u_{|\partial D} = 0.$$

where the boundary of D_y is parametrized by y, e.g.

$$D_y := \{(x_1, x_2) \in \mathbb{R}^2 \;\; : \;\; 0 < x_1 < 1 \;\; \text{and} \;\; 0 < x_2 < b(x_1, y)\},$$

where $b = b(x,y) = \overline{b}(x) + \sum_j y_j\psi_j(x)$ satisfies $0 < r < b(x,y) < R$. We transport this problem on the reference domain $[0,1]^2$ and study

$$u(y) := v(y) \circ \phi_y, \quad \phi_y : [0,1]^2 \to D_y, \quad \phi_y(x_1, x_2) := (x_1, x_2 b(x_1, y)).$$

which satisfies a diffusion equation with coefficient $a = a(x,y)$ non-affine in y.

Polynomial approximation results for these models

In contrast to model 1, bounded holomorphic extension is generally not feasible in a complex domain containing the polydisc $\mathcal{U} = \otimes\{|z_j| \leq 1\}$. For this reason,Taylor series are not expected to converge. Instead we consider the tensorized Legendre expansion

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

where $L_\nu(y) := \prod_{j \geq 1} L_{\nu_j}(y_j)$ and $(L_k)_{k \geq 0}$ are L^∞ normalized Legendre polynomials.

Theorem (Chkifa-Cohen-Schwab, 2013) : For models 2, 3 and 4 and for any $p < 1$,

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

with $X = L^\infty$ for models 2, 3, and $X = W^{1,\infty}$ for model 4.

Therefore, there exists polynomial approximations with uniform rate $\mathcal{O}(N^{-s})$ where $s = \frac{1}{p} - 1$ and mean square rate $\mathcal{O}(N^{-r})$ where $r = \frac{1}{p} - \frac{1}{2}$.

Key ingredient in the proof : estimates of Legendre coefficients for holomorphic functions in a "small" complex neighbourhood of U.

In one variable :

- If u is holomorphic in an open neighbourhood of the disc $\{|z| \le b\}$ and bounded by M on this disc, then the n-th Taylor coefficient of u is bounded by

$$|t_n| := \left| \frac{u^{(n)}(0)}{n!} \right| \le M b^{-n}$$

- If u is holomorphic in an open neighbourhood of the domain $\mathcal{E}_b$ limited by the ellipse of semi axes of length $(b + b^{-1})/2$ and $(b - b^{-1})/2$, for some $b > 1$, and bounded by M on this domain, then the n-th Legendre coefficent of u is bounded by

$$|u_n| := |\langle u, L_n \rangle| \le M b^{-n} \phi(b), \qquad \phi(b) := \frac{\pi b}{b-1}$$

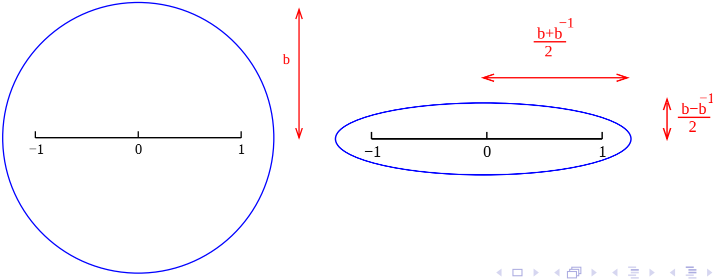

We say that the solution to a parametric PDE $\mathcal{D}(u, y) = 0$ satisfies the (p, ε)-holomorphy property if and only if there exist a sequence $(c_j)_{j\geq 1} \in \ell^p(\mathbb{N})$, a constant $\varepsilon > 0$ and $C_0 > 0$, such that : for any sequence $\rho = (\rho_j)_{j\geq 1}$ such that $\rho_j > 1$ and

$$\sum_{j\geq 1} (\rho_j - 1) c_j \leq \varepsilon,$$

the solution map has a complex extension

$$z \mapsto u(z),$$

of the solution map that is holomorphic with respect to each variable on a domain of the form $\mathcal{O}_\rho = \otimes_{j\geq 1} \mathcal{O}_{\rho_j}$ where $\mathcal{O}_{\rho_j}$ is an open neigbourhood of the elliptical domain $\mathcal{E}_{\rho_j}$, with bound

$$\sup_{z\in\mathcal{E}_\rho} \|u(z)\|_V \leq C_0,$$

where $\mathcal{E}_\rho = \otimes_{j\geq 1} \mathcal{E}_{\rho_j}$.

Under such an assumption, one has (up to additional harmless factors) an estimate of the form

$$\|u_\nu\|_V \leq C_0 \inf\{\rho^{-\nu} \ ; \ \rho \ \text{s.t.} \ \sum_{j\geq 1} (\rho_j - 1) c_j \leq \varepsilon\},$$

allowing us to prove that $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$.

A general framework for establishing the (p,ε)-holomorphy assumption

Assume a general problem of the form

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

with $a=a(y)=\overline{a}+\sum_{j\geq 1} y_j\psi_j$, where

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

with V,X,W a triplet of complex Banach spaces, and $\overline{a}$ and ψ_j are functions in X.

Theorem (Chkifa-Cohen-Schwab, 2013) : assume that

(i) The problem is well posed for all $a\in Q=a(U)$ with solution $u(y)=u(a(y))\in V$.

(ii) The map $\mathcal{P}$ is differentiable (holomorphic) from $X\times V$ to W.

(iii) For any $a\in Q$, the differential $\partial_u\mathcal{P}(u(a),a)$ is an isomorphism from V to W

(iv) One has $(\|\psi_j\|_X)_{j\geq 1}$ in $\ell^p(\mathbb{N})$ for some $0<p<1$,

Then, for $\varepsilon>0$ small enough, the (p,ε)-holomorphy property holds.

Idea of proof

Based on the holomorphic Banach valued version of the implicit function theorem (see e.g. Dieudonné).

1. For any $a \in Q = \{a(y) \; : \; y \in U\}$ we can find a $\varepsilon_a > 0$ such that the map $a \longrightarrow u(a)$ has an holomorphic extension on the ball $B(a, \varepsilon_a) := \{\tilde{a} \in X \; : \; \|\tilde{a} - a\|_X < \varepsilon_a\}$.

2. Using the decay properties of the $\|\psi_j\|_X$, we find that Q is compact in X. It can be covered by a finite union of balls $B(a_i, \varepsilon_{a_i})$, for $i = 1, \dots, M$.

3. Thus $a \longrightarrow u(a)$ has an holomorphic extension on a complex neighbourhood $\mathcal{N}$ of Q of the form

$$\mathcal{N} = \cup_{i=1}^{M} B(a_i, \varepsilon_{a_i}).$$

4. For ε small enough, one proves that if $\sum_{j \geq 1} (\rho_j - 1) c_j \leq \varepsilon$ with $c_j := \|\psi_j\|_L$ then with $\mathcal{O}_\rho = \otimes_{j \geq 1} \mathcal{O}_{\rho_j}$ where $\mathcal{O}_b := \{z \in \mathbb{C} \; : \; \mathrm{dist}(z, [-1, 1])_{\mathbb{C}} \leq b - 1\}$ is a neighborhood of $\mathcal{E}_b$, one has

$$z \in \mathcal{O}_\rho \Rightarrow a(z) \in \mathcal{N}.$$

This gives holomorphy of $z \mapsto u(z) = u(a(z))$ in each variable for $z \in \mathcal{O}_\rho$.

Strategies to build the set Λ :

(i) Non-adaptive, based on the available a-priori estimates for the $\|t_\nu\|_V$.

(ii) Adaptive, based on a-posteriori information gained in the computation $\Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_N$.

Objective : develop adaptive strategies that converge with optimal rate (similar to adaptive wavelet methods for elliptic PDE's : Cohen-Dahmen-DeVore, Stevenson).

Recursive computation of the Taylor coefficients for Model 1 : with e_j the Kroenecker sequence

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

We compute the t_ν on downward closed (or "lower") sets Λ : $\nu \in \Lambda$ and $\mu \leq \nu \Rightarrow \mu \in \Lambda$.

Given such a Λ_k and the $(t_\nu)_{\nu\in\Lambda_k}$ we compute the t_ν for ν in the margin

$$\mathcal{M}_k := \{\nu \notin \Lambda_k \,;\, \nu - e_j \in \Lambda_k \ \text{ for some } j\},$$

and build the new set by bulk search : $\Lambda_{k+1} = \Lambda_k \cup \mathcal{S}_k$, with $\mathcal{S}_k \subset \mathcal{M}_k$ smallest such that $\sum_{\nu\in\mathcal{S}_k} \|t_\nu\|_V^2 \geq \theta \sum_{\nu\in\mathcal{M}_k} \|t_\nu\|_V^2$, with $\theta \in (0,1)$.

Such a strategy can be proved to converge with optimal convergence rate $\#(\Lambda_k)^{-s}$.

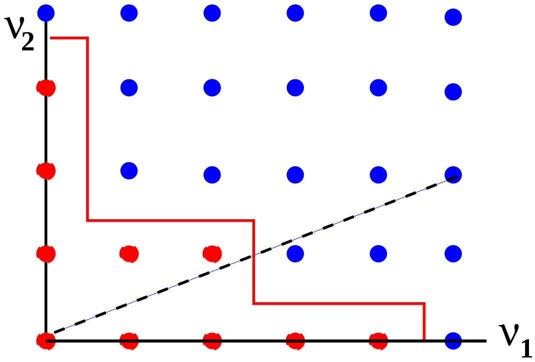

ν_2
ν_1

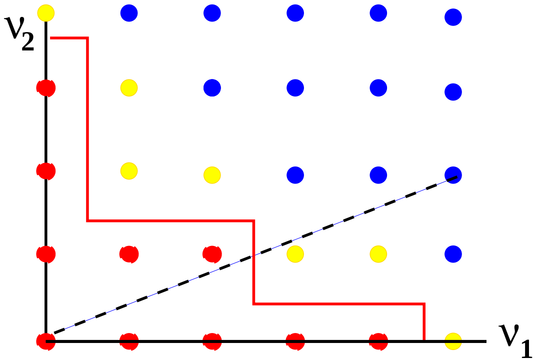

v_2
v_1

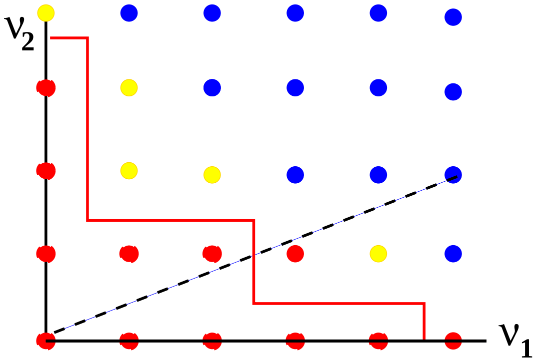

ν_2
ν_1

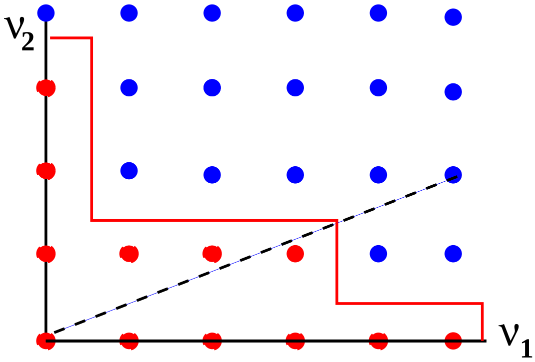

ν_2
ν_1

Test case in high dimension $d = 255$

Physical domain $D = [0,1]^2$.

Diffusion coefficients $a(x,y) = 1 + \sum_{j=1}^{d} y_j \psi_j$ where ψ_j are weighted Haar wavelets of the form

$$\psi = \beta_l h_{k,l}, \quad h_{k,l} = h(2^l \cdot - k), \quad \beta_l = c 2^{-\gamma l}$$

with $\gamma > 0$ (correlation in the diffusion field) and $c > 0$ such that (UEA) holds.

Adaptive search of Λ implemented in C++, spatial discretization by FreeFem++.

Comparison between the Λ_k generated by adaptive algorithms (red, green) and non-adaptive choices $\{\sup \nu_j \leq k\}$ (black) or $\{\sum \nu_j \leq k\}$ (purple) or k largest a-priori bounds on the $\|t_\nu\|_V$ (blue). Left $\gamma = 0.5$, right $\gamma = 3$.

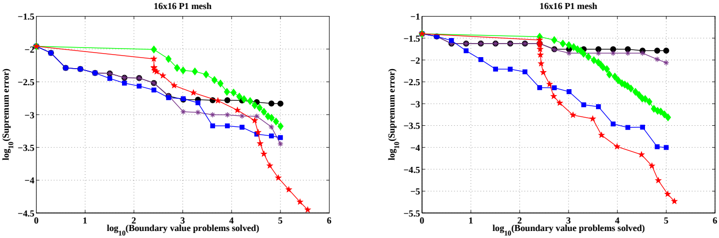

Highest polynomial degree with $\#(\Lambda) = 1000$ coefficients : 1, 2, 16 and 13.

Numerical methods : strategies to build the polynomial approximation

(i) Intrusive : exact computation of the Taylor coefficients $\|t_\nu\|_V$ for the linear-affine model (Chkifa-Cohen-DeVore-Schwab) or Galerkin approximation of the Legendre coefficients (Gittelson-Schwab). Adaptive algorithms with optimal theoretical guarantees.

(ii) Non-intrusive : based on snapshots $u_i := u(y^i)$ for $i = 1, \dots, m$. Interpolation (Chkifa-Cohen-Schwab) or Least Squares (Chkifa-Cohen-Migliorati-Nobile-Tempone). Adaptive algorithms seem to work well, however with no theoretical guarantees.

Additional prescriptions for non-intrusive methods :

(i) Progressive : enrichment $\Lambda_N \to \Lambda_{N+1}$ requires only one or a few new snapshots.

(ii) Stable : moderate growth with N of the Lebesgue constant relative to the interpolation operator.

Sparse interpolation

Let $\{t_0, t_1, t_2 \dots\}$, be an infinite sequence of pairwise distinct points in $[-1, 1]$ and let I_k be the univariate interpolation operator on $\mathbb{P}_k$ associated to the section $\{t_0, \dots, t_k\}$.

Hierarchical (Newton) form : $I_k = \sum_{l=0}^{k} \Delta_l$, with $\Delta_l := I_l - I_{l-1}$ and $I_{-1} := 0$.

Tensorization and sparsification : for $\nu \in \mathcal{F}$, we define the point

$$z_\nu := (t_{\nu_1}, t_{\nu_2}, \dots) \in U.$$

Theorem (Kuntzmann 1959) : if Λ is downward closed, the set

$$\Gamma_\Lambda := \{z_\nu \ : \ \nu \in \Lambda\},$$

is unisolvent for $\mathbb{P}_\Lambda = \mathrm{Span}\{y \mapsto y^\nu \ : \ \nu \in \Lambda\}$ and the interpolant is

$$I_\Lambda := \sum_{\nu \in \Lambda} \Delta_\nu, \quad \Delta_\nu := \otimes_{j \geq 1} \Delta_{\nu_j}.$$

Theorem (Chkifa-Cohen-Schwab, 2012) : if $\mathbb{L}_k = \|I_k\|_{L^\infty \to L^\infty} \leq (1+k)^a$, then $\mathbb{L}_\Lambda = \|I_\Lambda\|_{L^\infty \to L^\infty} \leq \#(\Lambda)^{1+a}$. Moderate growth of $\mathbb{L}_k$ for Leja points ($a = 1$).

A straightforward adaptive algorithm : given Λ_N, define $\Lambda_{N+1} := \Lambda_N \cup \{\nu^*\}$ with $\nu^* \notin \Lambda_N$ such that Λ_{N+1} is downward closed and maximizing $\|\Delta_\nu u\|_{L^\infty}$.

Remark : the same principles apply to the tensorization of other systems, such as hierarchical piecewise linear finite elements.

We apply the adaptive interpolation algorithm to

$$u(y) := (1 + \sum_{j=1}^{d} \gamma_j y_j)^{-1}, \quad \gamma_j = \frac{3}{5j^3},$$

for different numbers d of variables.

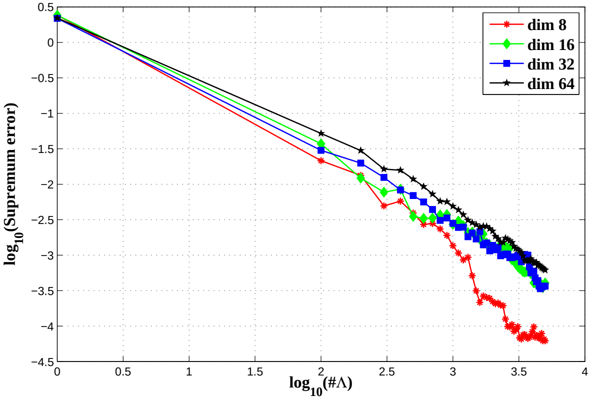

Robustness to noise

Same function u in dimension $d = 16$, with noisy samples (noise level $= 10^{-2}$). using adaptive interpolation based on different univariate sequences.

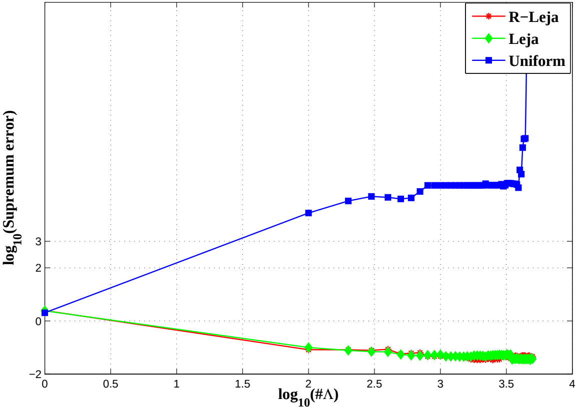

The Lebesgue constant for the Clemshaw-Curtis point with sequencial intermediate filling.

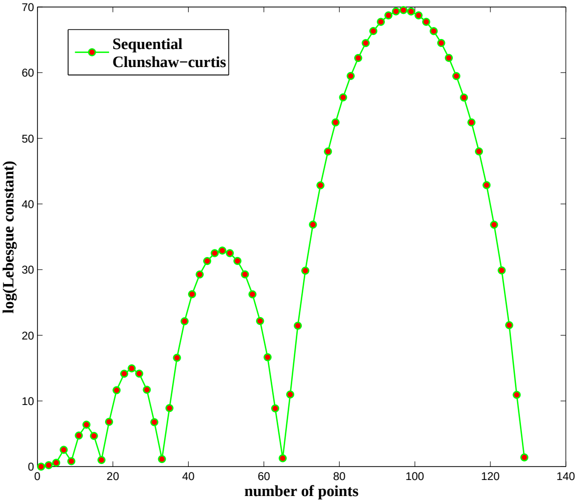

Stability

The Lebesgue constant for
- the Leja points on $[-1, 1]$.
- the R-Leja points (Clemshaw-Curtis points with intermediate Van der Corput filling).

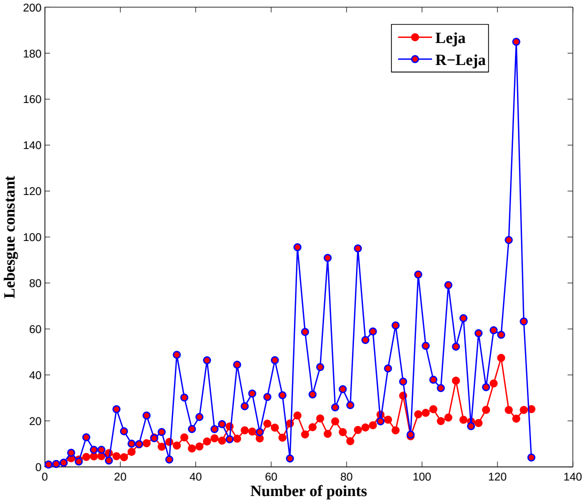

Comparison with Kriging interpolation algorithms

Test case : $y = (y_1, y_2, y_3, y_4)$ shape parameters in the design of an airfoil and $u(y)$ is the lift to drag ratio (scalar quantity of interest) obtained by ONERA numerical solver.

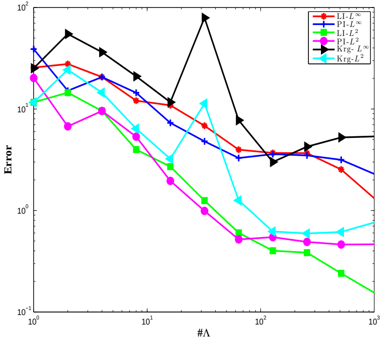

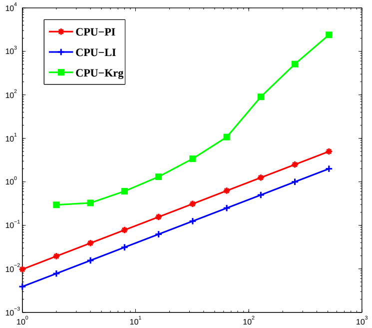

Error curves in terms of number of points are comparable.

The CPU cost for sparse interpolation scales linearly with the number of points.

This contrasts with Kriging methods which require solving ill-conditionned linear systems of growing size + optimization of the parameters of a Gaussian kernel.

Approximation of the solution map and reduced order modeling

For a parametric PDE $\mathcal{P}(u,a)=0$ with a ranging in $K\subset X$, we define the solution manifold

$$M := u(K) = \{u(a)\ :\ a\in K\}\subset V.$$

Reduced modeling : find low dimension spaces that simultaneously approximate well all solutions to the parametric PDE.

Benchmark : Kolmogorov N-width

$$d_N(M)_V = \inf_{\dim(E)=N}\max_{v\in M}\min_{w\in E}\|v-w\|_V.$$

If K is of the form $K=Q=\{a(y)=\overline{a}+\sum_{j\geq 1} y_j\psi_j\ :\ y\in U\}$, we have

$$d_N(M)_V = \inf_{\dim(E)=N}\max_{y\in U}\min_{w\in E}\|u(y)-w\|_V.$$

Uniform approximation estimates of the solution map $y\mapsto u(y)$ by truncated separable expansions of the form

$$\max_{y\in U}\|u(y)-\sum_{i=1}^N u_i\varphi_i(y)\|\leq \varepsilon_N\sim N^{-s},$$

with $u_i\in V$ and $\varphi_i : U\to\mathbb{R}$, imply similar estimates on the Kolmogorov width of the solution manifold :

$$d_N(M)_V\leq \max_{v\in M}\min_{w\in E_N}\|v-w\|_V\leq\varepsilon_N,\qquad E_N:=\mathrm{span}\{u_1,\dots,u_N\}.$$

Reduced bases (Maday, Patera)

Define a reduced modeling space $E_N = \mathrm{span}\{u_1, \dots, u_N\}$, where the u_i are particular instances (snapshots) from the solution manifold

$$u_i = u(a_i)$$

for some $a_1, \dots, a_N \in K$.

Greedy selection : having selected $u_1, \dots, u_{N-1} \in M$, choose the next instance by

$$u_N = \mathrm{argmax}\{\|v - P_{E_{N-1}} v\|_V \ : \ v \in M\},$$

where P_E is the orthogonal projector onto E, or equivalently $u_N = u(a_N)$, with

$$a_N = \mathrm{argmax}\{\|u(a) - P_{E_{N-1}} u(a)\|_V \ : \ a \in K\}.$$

This algorithm is not realistic : $\|u(a) - P_{E_{N-1}} u(a)\|_V$ is unknown, however can be estimate at moderate cost by a-posteriori error analysis. Therefore, one rather apply a weak-greedy algorithm : u_N such that

$$\|u_N - P_{E_{N-1}} u_N\|_V \geq \gamma \max\{\|v - P_{E_{N-1}} v\|_V \ : \ v \in M\},$$

for some fixed $0 < \gamma < 1$.

Comparison with N-width

Performance of reduced bases : $\sigma_N(M)_V := \max\{\|v - P_{E_N} v\|_V \;:\; v \in M\}$

Comparison with N-width : $\sigma_N(M)_V$ can be much larger than $d_N(M)_V$ for an individual N and M.

There exists M and N such that $\sigma_N(M)_V \geq 2^N d_N(M)_V$.

However, a more favorable comparison is possible in terms of convergence rates :

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

$$\sup_{N\geq 1} N^s d_N(M)_V < \infty \Rightarrow \sup_{N\geq 1} N^s \sigma_N(M)_X < \infty,$$

and for any $a > 0$ there exists $b > 0$ such that

$$\sup_{N\geq 1} e^{aN^s} d_N(M)_V < \infty \Rightarrow \sup_{N\geq 1} e^{bN^s} \sigma_N(M)_X < \infty.$$

For a compact set $K \subset X$ and a continuous mapping $u : K \to V$, we would like to control the decay $d_N(u(K))_V$ from $d_N(K)_X$.

Note that if u was a linear mapping, we would simply have

$$d_N(u(K))_V \leq C d_N(K)_X, \qquad C := \|u\|_{\mathcal{L}(X,V)}.$$

The following result shows that nonlinear holomorphic maps behave almost like linear maps with respect to the asymptotic decay of N-widths.

Theorem (Cohen-DeVore, 2014) : Let X, V be complex Banach spaces and let

$$K \subset O \subset X,$$

with K compact and O open sets. Assume that

$$u : O \to V$$

is uniformly bounded and holomorphic (Frechet differentiable in the sense of complex Banach space). Then, for all $t > 0$,

$$\sup_{N\geq 1} N^t d_N(K)_X < \infty \Rightarrow \sup_{N\geq 1} N^s d_N(u(K))_V < \infty, \qquad s < t - 1.$$

Proof uses scalar parametrizations of K and polynomial approximations.

Conclusions

The curse of dimensionality can be "defeated" by exploiting both smoothness and anisotropy in the different variables.

For certain models, this can be achieved by sparse polynomial approximations.

Adaptive algorithms with optimal theoretical guarantees are still to be developed, in particular for non-intrusive approaches (interpolation, collocation, least-squares).

The choice of parametrization and representation of the solution are critical in this analysis since it affects the properties of the map $y \mapsto u(y)$.

Other approaches to evaluate Kolmogorov width of solution manifold ? Lower bounds ?