Is polynomial interpolation in the monomial basis unstable?

Zewen Shen, Kirill Serkh

University of Toronto

May 2023

Polynomial interpolation

Polynomials are powerful tools for approximating functions.

Definition

Given a function $F : [-1, 1] \to \mathbb{C}$, the Nth degree interpolating polynomial P_N of F satisfies $P_N(x_j) = F(x_j)$, for a set of $(N+1)$ distinct collocation points $\{x_j\}_{j=0,1,\ldots,N}$.

The choice of collocation points matters. In this talk, we only consider collocation points with a small Lebesgue constant (e.g., Chebyshev points).

Polynomial interpolation in finite precision

To compute P_N on a computer, we first choose a polynomial basis $\{\phi_k\}_k$

$$P_N(x) = \sum_{k=0}^{N} a_k \phi_k(x)$$

$$\begin{pmatrix} \phi_0(x_0) & \phi_1(x_0) & \phi_2(x_0) & \cdots & \phi_N(x_0) \\ \phi_0(x_1) & \phi_1(x_1) & \phi_2(x_1) & \cdots & \phi_N(x_1) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \phi_0(x_N) & \phi_1(x_N) & \phi_2(x_N) & \cdots & \phi_N(x_N) \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_N \end{pmatrix} = \begin{pmatrix} F(x_0) \\ F(x_1) \\ \vdots \\ F(x_N) \end{pmatrix}.$$

Why does the choice of basis matter?

- Condition number
- Time complexity

The standard choices:

- Lagrange polynomials.
- Orthogonal polynomials (Chebyshev, Legendre, etc).

Polynomial interpolation in the monomial basis

What about expressing P_N in the monomial basis?

$$P_N(x) = \sum_{k=0}^{N} a_k x^k$$

The previous linear system becomes

$$\underbrace{\begin{pmatrix} 1 & x_0 & x_0^2 & \cdots & x_0^N \\ 1 & x_1 & x_1^2 & \cdots & x_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_N & x_N^2 & \cdots & x_N^N \end{pmatrix}}_{V^{(N)}} \underbrace{\begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_N \end{pmatrix}}_{a^{(N)}} = \underbrace{\begin{pmatrix} F(x_0) \\ F(x_1) \\ \vdots \\ F(x_N) \end{pmatrix}}_{f^{(N)}}.$$

$V^{(N)}$ is known as a Vandermonde matrix.

Monomial basis is ill-conditioned

Given any set of real collocation points, $\kappa(V^{(N)})$ grows at least exponentially fast.

Example: when the Chebyshev points are used for collocation:

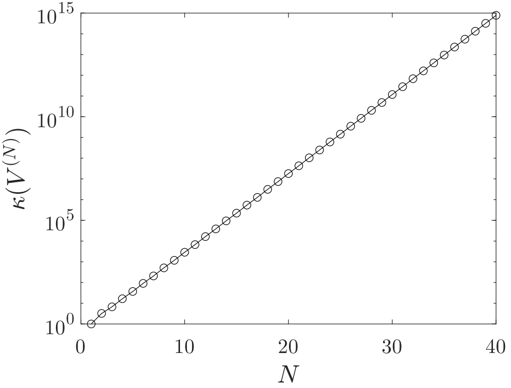

Numerical experiments

Let's run some experiments. The following quantities will be reported.

- $\|F - \widehat{P}_N\|_{L^\infty([-1,1])}$: Monomial approximation error.
  Denoted by the label "monomial".
- $\|F - P_N\|_{L^\infty([-1,1])}$: Exact polynomial interpolation error, estimated using the Barycentric interpolation formula.
  Denoted by the label "Lagrange".

Chebyshev points are used for collocation.

Numerical experiments

$$F(x) = \cos(2x + 1)$$

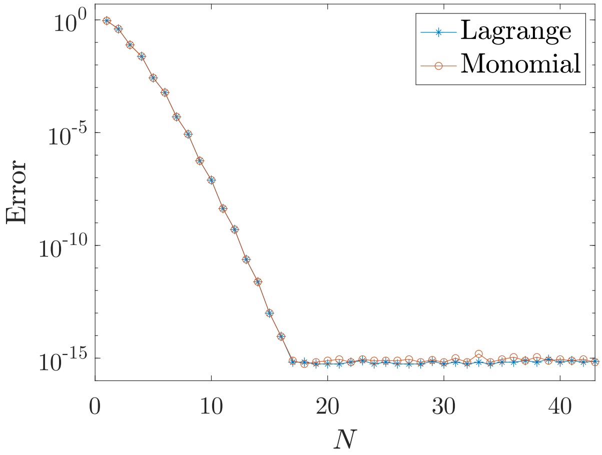

Numerical experiments

$$F(x) = e^{-2(x+0.1)^2}$$

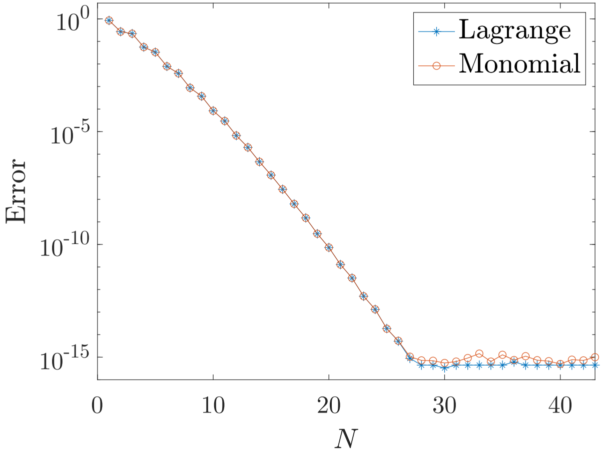

Numerical experiments

$$F(x) = \cos(8x + 1)$$

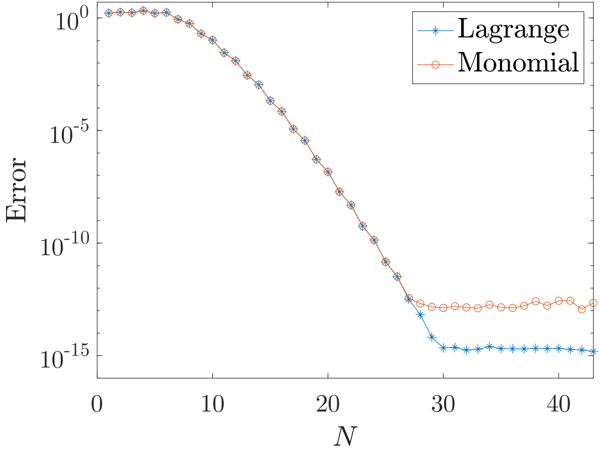

Numerical experiments

$$F(x) = \cos(12x + 1)$$

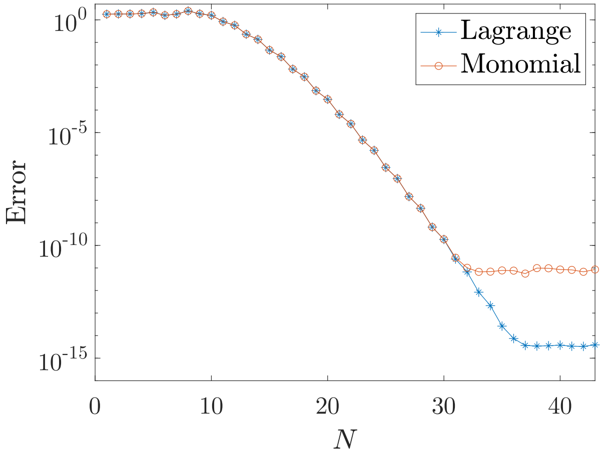

Numerical experiments

$$F(x) = \frac{1}{x - \sqrt{2}}$$

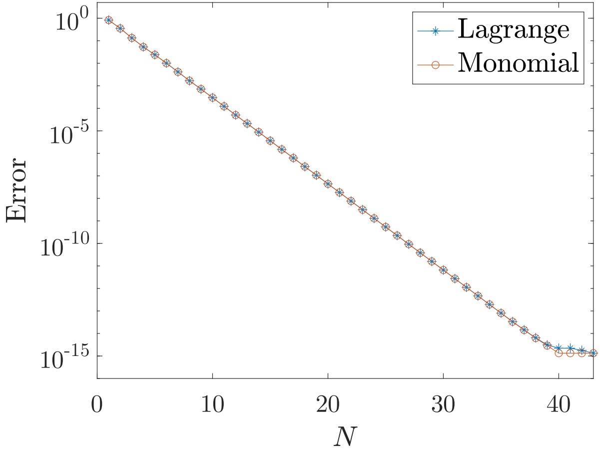

Numerical experiments

$$F(x) = \frac{1}{x - 0.5i}$$

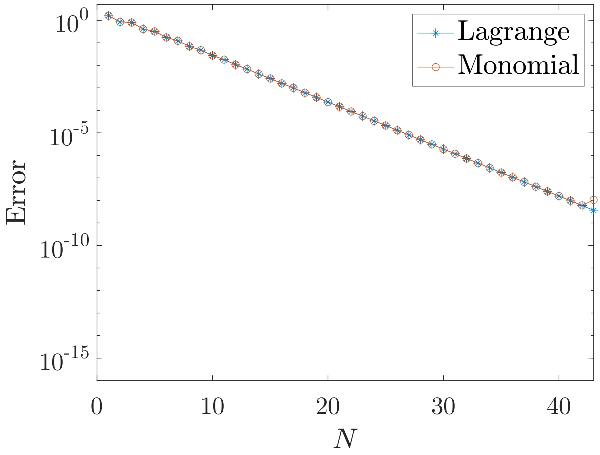

Numerical experiments

$$F(x) = |x|^{2.5}$$

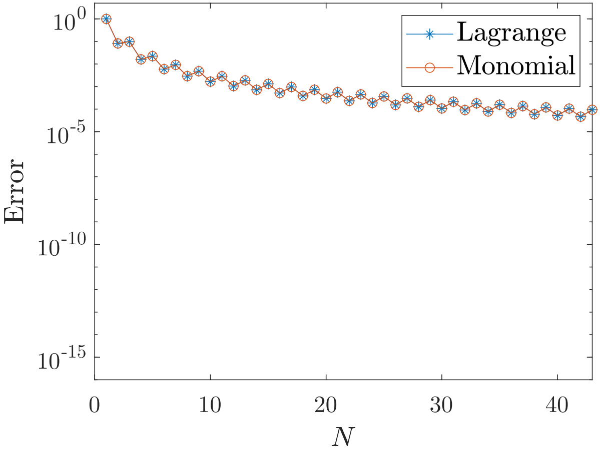

Reflections

- Polynomial interpolation in the monomial basis is **not** as unstable as it appears.
- The same thing happens when the domain is not $[-1, 1]$ (say, a triangle in $\mathbb{R}^2$).
- The observation that "the monomials can approximate sufficiently smooth functions to high accuracy" dates back to ≥ 25 years ago.
- Not very widely known. Not fully understood. Not the complete story.

The monomial basis is not too different from a well-conditioned polynomial basis for interpolation, provided that $\kappa(V^{(N)}) \leq \frac{1}{u}$.

Before I explain why, I'll present an application.

Monomial as the default choice for interpolation

Complicated irregular domains appear in many applications.

Well-conditioned polynomial bases over unstructured mesh elements are generally either unknown or complicated.

- What are orthogonal polynomials over an arbitrary curved triangle?
- Orthogonal polynomials over a standard simplex:

$$K_{mn}(x, y) = (1 - x)^m \cdot P_{n-m}^{(2m+1,0)}(2x - 1) \cdot P_m\Big(\frac{2y}{1 - x} - 1\Big).$$

- What about tetrahedrons?

On the other hand, the monomial basis works for any domain, is extremely handy, and is much cheaper to evaluate.

Rethinking interpolation

Huge condition number of Vandermonde matrices

$$\Longrightarrow$$

extremely inaccurate monomial coefficients

Do we care about the accuracy of the computed monomial coefficients?

What's really important is the backward error, i.e.,

$$\|V^{(N)}\hat{a}^{(N)} - f^{(N)}\|_2,$$

of the numerical solution $\hat{a}^{(N)}$ to the Vandermonde system $V^{(N)}a^{(N)} = f^{(N)}$.

Rethinking interpolation

The difference between the exact interpolating polynomial P_N and the computed monomial expansion $\widehat{P}_N$ satisfies

$$\|P_N - \widehat{P}_N\|_{L^\infty(\Gamma)} \leq \Lambda_N \|V^{(N)}\widehat{a}^{(N)} - f^{(N)}\|_2.$$

How large will the backward error be?

Backward stable linear system solver

When a backward stable linear system solver is used to solve the Vandermonde system $V^{(N)}a^{(N)} = f^{(N)}$, the numerical solution $\widehat{a}^{(N)}$ is the exact solution to

$$\left(V^{(N)} + \delta V^{(N)}\right)\widehat{a}^{(N)} = f^{(N)},$$

for some $\delta V^{(N)} \in \mathbb{C}^{(N+1)\times(N+1)}$ that satisfies

$$\|\delta V^{(N)}\|_2 \leq u \cdot \gamma_N,$$

where u denotes machine epsilon and $\gamma_N = \mathcal{O}(\|V^{(N)}\|_2)$.

Remark: When MATLAB's backslash is used, we observe that $\gamma_N \lesssim 1$ for at least $N \leq 100$.

It follows that

$$\|V^{(N)}\widehat{a}^{(N)} - f^{(N)}\|_2 = \|\delta V^{(N)}\widehat{a}^{(N)}\|_2 \leq u \cdot \gamma_N \|\widehat{a}^{(N)}\|_2.$$

A priori estimate

Lemma

If $\|(V^{(N)})^{-1}\|_2 \leq \frac{1}{2u\cdot\gamma_N}$, *then* $\frac{2}{3}\|a^{(N)}\|_2 \leq \|\widehat{a}^{(N)}\|_2 \leq 2\|a^{(N)}\|_2$.

Therefore,

$$\|(V^{(N)})^{-1}\|_2 \leq \frac{1}{2u\cdot\gamma_N} \Longrightarrow \left\|V^{(N)}\widehat{a}^{(N)} - f^{(N)}\right\|_2 \leq 2u\cdot\gamma_N\|a^{(N)}\|_2.$$

Corollary (Finite-precision interpolation error)

If $\|(V^{(N)})^{-1}\|_2 \leq \frac{1}{2u\cdot\gamma_N}$, *then*

$$\|F - \widehat{P}_N\|_{L^\infty(\Gamma)} \leq \|F - P_N\|_{L^\infty(\Gamma)} + 2u\cdot\gamma_N\Lambda_N\|a^{(N)}\|_2.$$

Numerical experiments

$$F(x) = \cos(2x + 1)$$

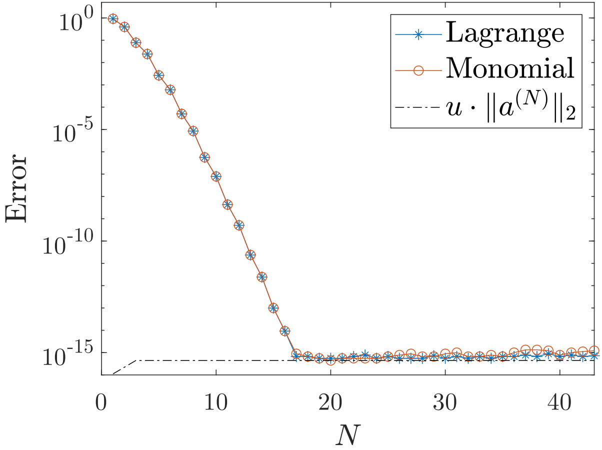

Numerical experiments

$$F(x) = \frac{1}{x - \sqrt{2}}$$

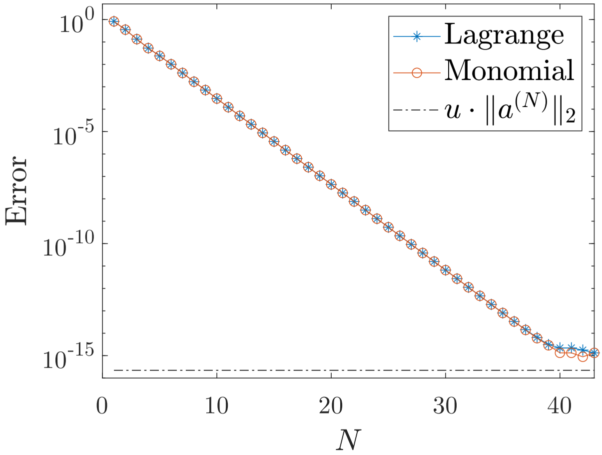

Numerical experiments

$$F(x) = \cos(8x + 1)$$

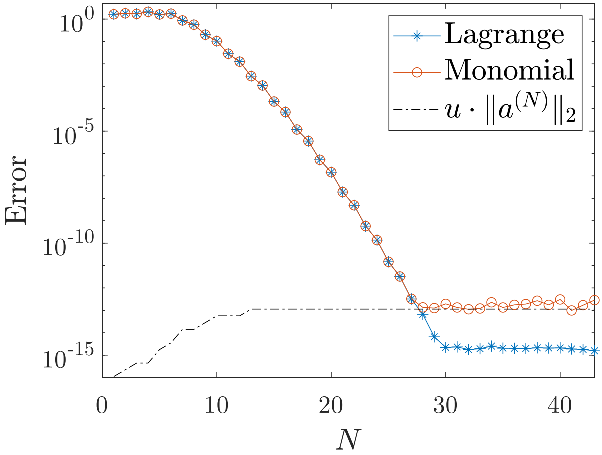

Numerical experiments

$$F(x) = \cos(12x + 1)$$

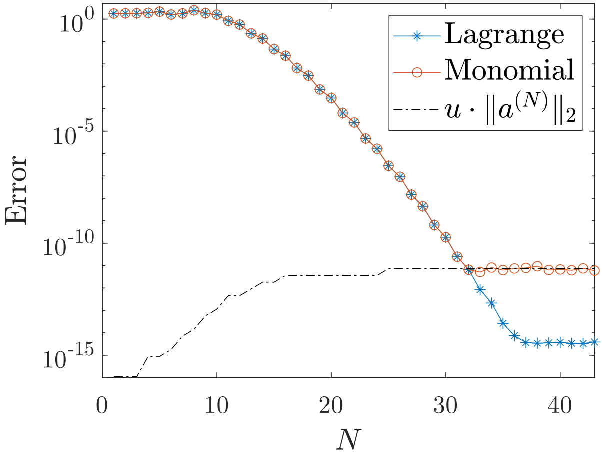

Numerical experiments

$$F(x) = \frac{1}{x-0.5i}$$

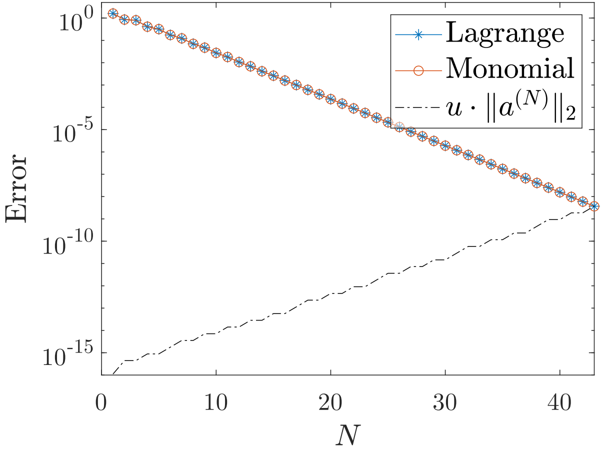

Numerical experiments

$$F(x) = |x|^{2.5}$$

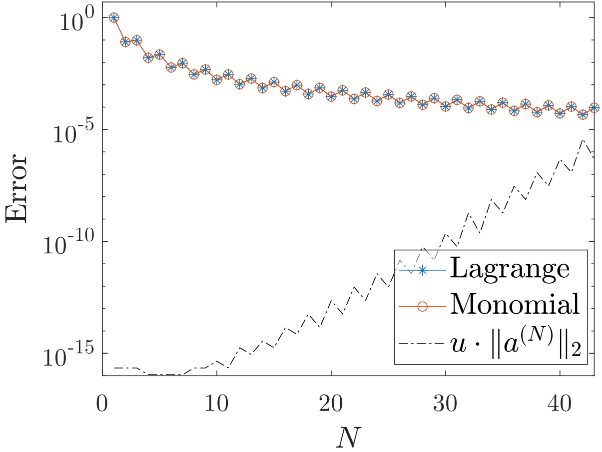

Story so far

- We can now explain these experiments, but many things are still unclear.
- For example, when will the extra error (i.e., $u \cdot \|a^{(N)}\|_2$) be small?
- This requires an a priori estimate for the growth of $\|a^{(N)}\|_2$.

An important constant

Given a smooth simple arc $\Gamma \subset \mathbb{C}$, define ρ_* to be the parameter of the smallest Bernstein ellipse for Γ that contains the unit disk.

Example

When $\Gamma = [-1, 1]$, $\rho_* = 1 + \sqrt{2} \approx 2.4$

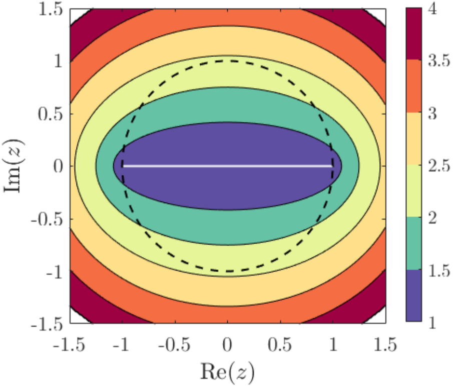

An improved upper bound for $\|a^{(N)}\|_2$

Theorem

Suppose that there exists a finite sequence of polynomials $\{Q_n\}_{n=0,1,\ldots,N}$, where Q_n has degree n, which satisfies

$$\|F - Q_n\|_{L^\infty(\Gamma)} \leq C\rho_*^{-n}, \quad 0 \leq n \leq N,$$

for some constant $C \geq 0$. The 2-norm of the monomial coefficient vector of the Nth degree interpolating polynomial P_N of F satisfies

$$\|a^{(N)}\|_2 \leq \|F\|_{L^\infty(\Gamma)} + C\Big(\Lambda_N + 2\rho_* N + 1\Big) \lesssim C \cdot N.$$

In practice, one can take $\{Q_n\}_{n=0,1,\ldots,N}$ to be a finite sequence of interpolating polynomials $\{P_n\}_{n=0,1,\ldots,N}$ of F.

We first deal with the case where the $\|F - P_n\|_{L^\infty(\Gamma)} \lesssim \rho_*^{-n}$.

$$\|a^{(N)}\|_2 \lesssim C \cdot N \approx N.$$

Implications: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays quickly

Therefore, when $\|(V^{(N)})^{-1}\|_2 \lesssim \frac{1}{u}$, the monomial approximation error satisfies

$$\|F - \widehat{P}_N\|_{L^\infty(\Gamma)} \lesssim \|F - P_N\|_{L^\infty(\Gamma)} + u \cdot N.$$

The extra error is around machine epsilon in this case!

Visualization: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays quickly

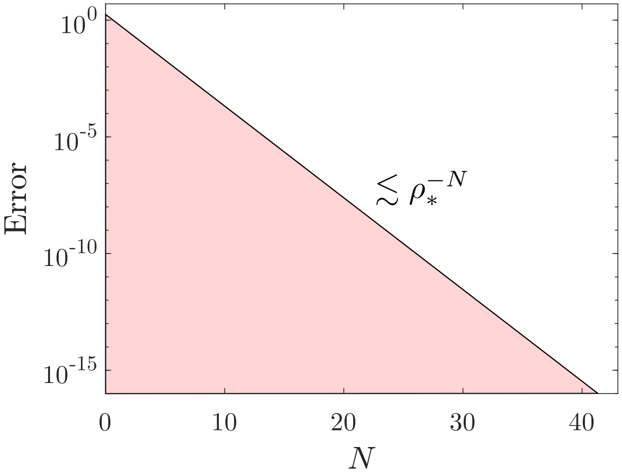

Visualization: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays quickly

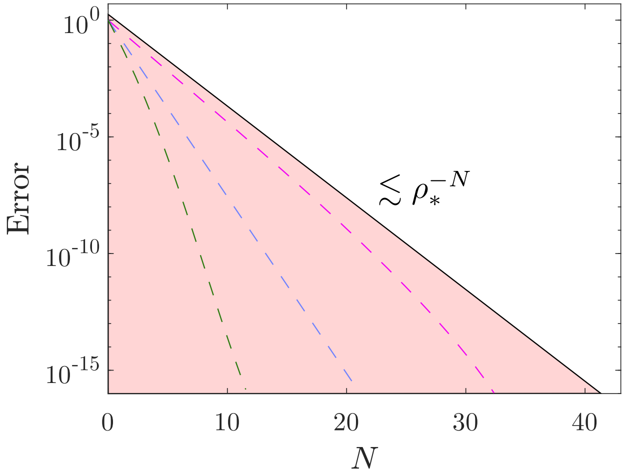

Examples: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays quickly

$$F(x) = \cos(2x + 1)$$

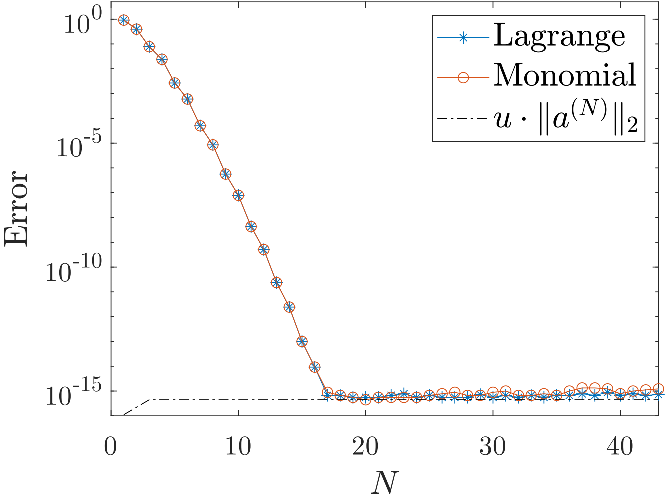

Examples: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays quickly

$$F(x) = \frac{1}{x - \sqrt{2}}$$

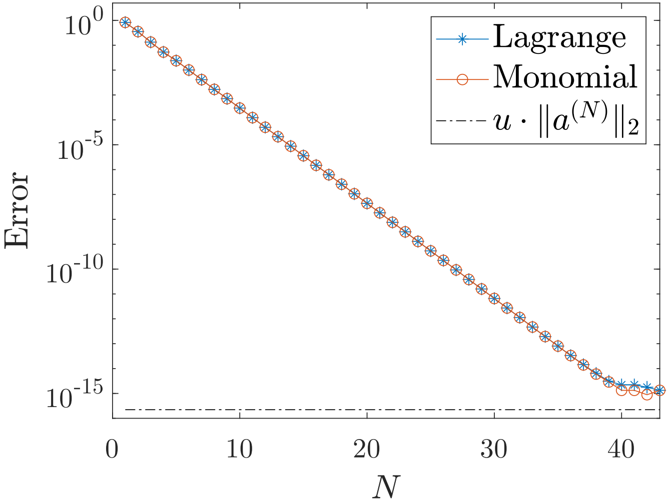

Implications: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

When $\|F - P_n\|_{L^\infty(\Gamma)} \lesssim \rho_*^{-n}$ for $0 \le n \le N$,

- the growth of $\|a^{(N)}\|_2$ is suppressed,
- and one loses nothing by using the monomial basis.

What happens if the polynomial interpolation error decays more slowly?

- $\|a^{(N)}\|_2$ will be larger.
- extra error caused by the monomial basis becomes non-negligible.

Does it matter?

Examples: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$F(x) = \cos(120x + 1)$$

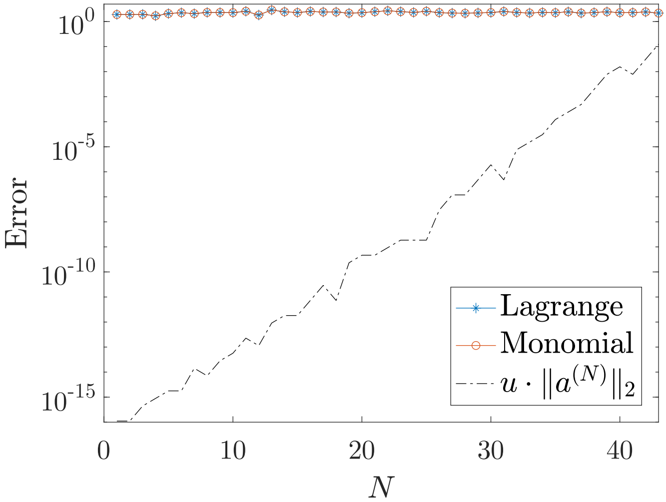

Examples: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$F(x) = |x|^{5/2}$$

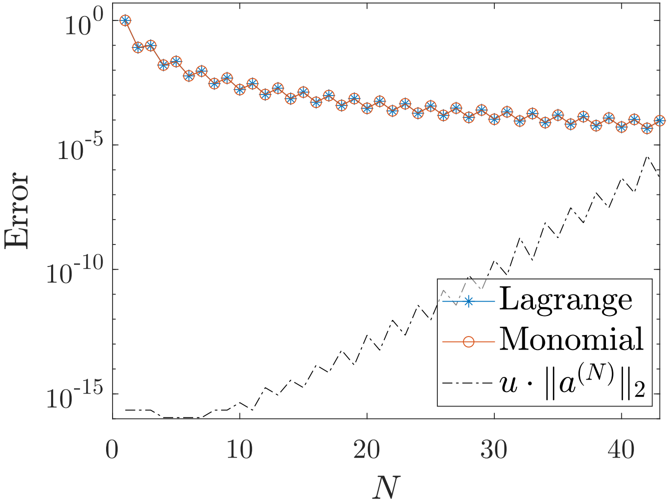

Examples: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$F(x) = \frac{1}{x - 0.5i}$$

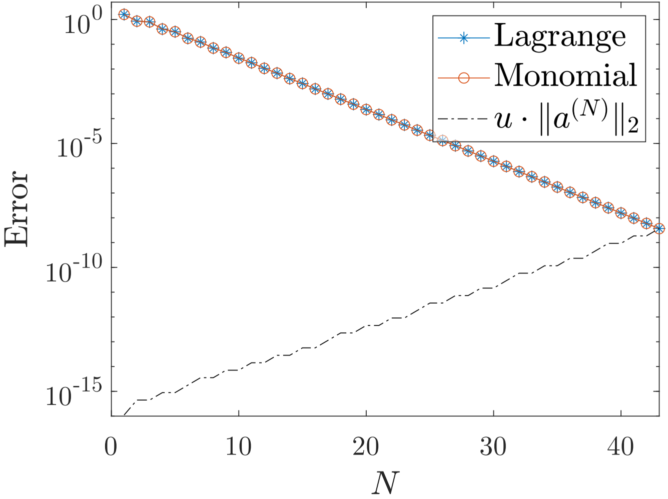

Implications: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

I'll now characterize what we just observed.

Assume that $\|F - P_n\|_{L^\infty(\Gamma)}$ decays to the value $\|F - P_N\|_{L^\infty(\Gamma)}$ at a rate slower than ρ_*^{-n}, i.e.,

$$\|F - P_n\|_{L^\infty(\Gamma)} \leq \rho_*^{N-n} \|F - P_N\|_{L^\infty(\Gamma)}, \quad \text{for } 0 \leq n \leq N.$$

Visualizations: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$\|F - P_n\|_{L^\infty(\Gamma)} \leq \rho_*^{N-n} \|F - P_N\|_{L^\infty(\Gamma)}, \quad \text{for } 0 \leq n \leq N.$$

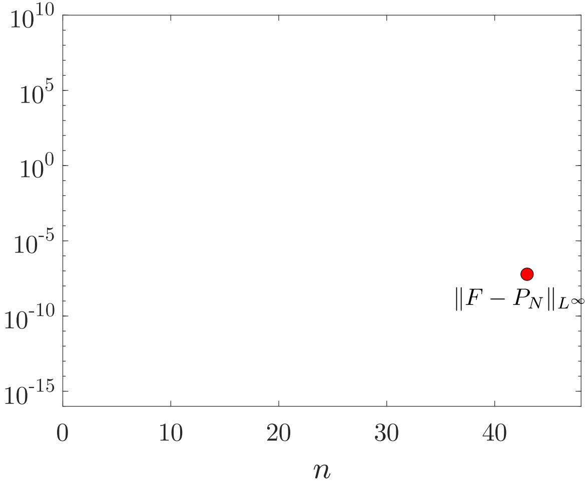

Visualizations: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$\|F - P_n\|_{L^\infty(\Gamma)} \leq \rho_*^{N-n} \|F - P_N\|_{L^\infty(\Gamma)}, \quad \text{for } 0 \leq n \leq N.$$

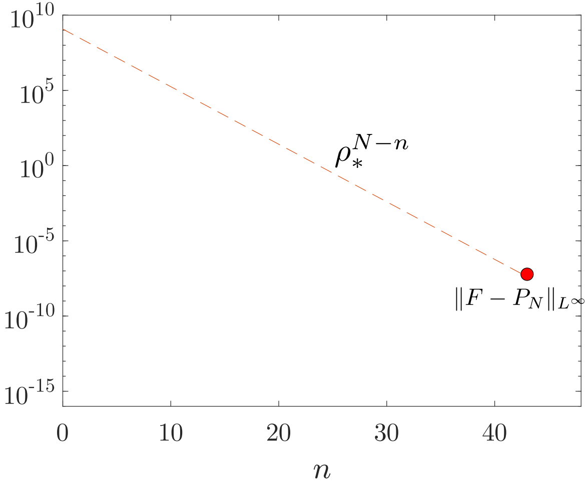

Visualizations: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$\|F - P_n\|_{L^\infty(\Gamma)} \le \rho_*^{N-n} \|F - P_N\|_{L^\infty(\Gamma)}, \quad \text{for } 0 \le n \le N.$$

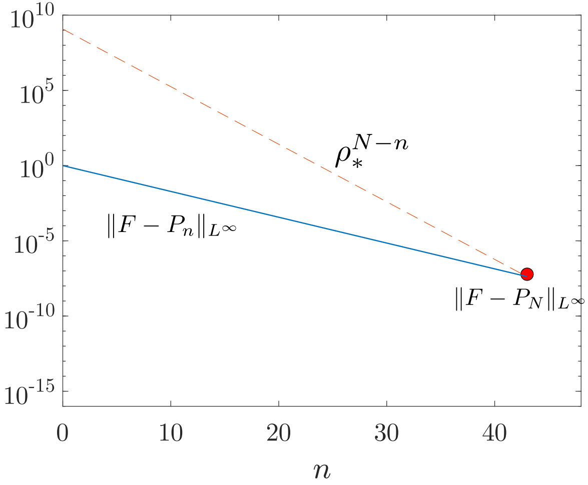

Visualizations: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

$$\|F - P_n\|_{L^\infty(\Gamma)} \le \rho_*^{N-n} \|F - P_N\|_{L^\infty(\Gamma)}, \quad \text{for } 0 \le n \le N.$$

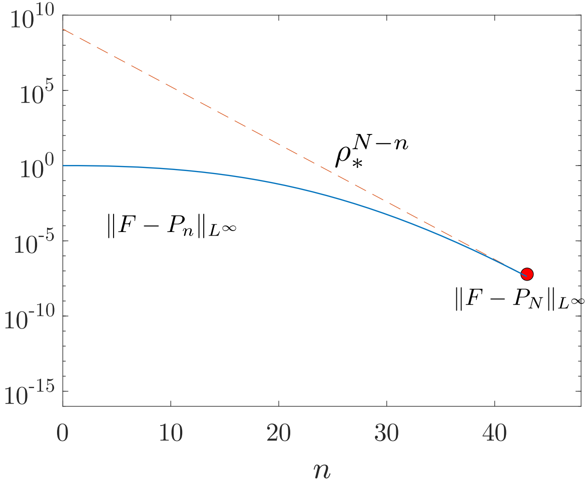

Implications: when $\|F - P_N\|_{L^\infty(\Gamma)}$ decays slowly

Theorem

Under this assumption, the monomial approximation error satisfies

$$\|F - \widehat{P}_N\|_{L^\infty(\Gamma)} \lesssim 2\|F - P_N\|_{L^\infty(\Gamma)},$$

so long as $\|(V^{(N)})^{-1}\|_2 \lesssim \frac{1}{u}$.

The proof is similar to the previous case.

Implications: stagnation of convergence

We've shown that if $\|F - P_n\|_{L^\infty(\Gamma)}$

- decays at a rate **faster** than ρ_*^{-n},
- or decays at a rate **slower** than ρ_*^{-n},

then the monomial basis = a well-conditioned basis when the order $\leq$ threshold.

The only way for stagnation to happen before the order reaches the threshold is that, $\|F - P_n\|_{L^\infty(\Gamma)}$ first decays at a rate **slower** than ρ_*^{-n}, then starts to decay at a rate **faster** than ρ_*^{-n}.

Examples: stagnation of convergence

$$F(x) = \cos(12x + 1)$$

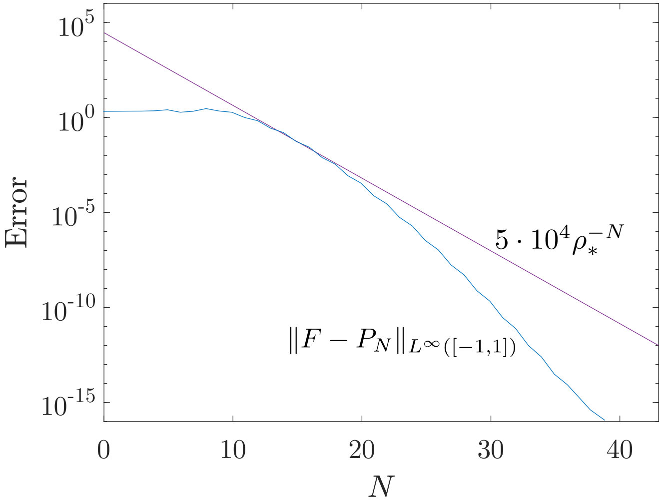

Examples

$$F(x) = \cos(12x + 1)$$

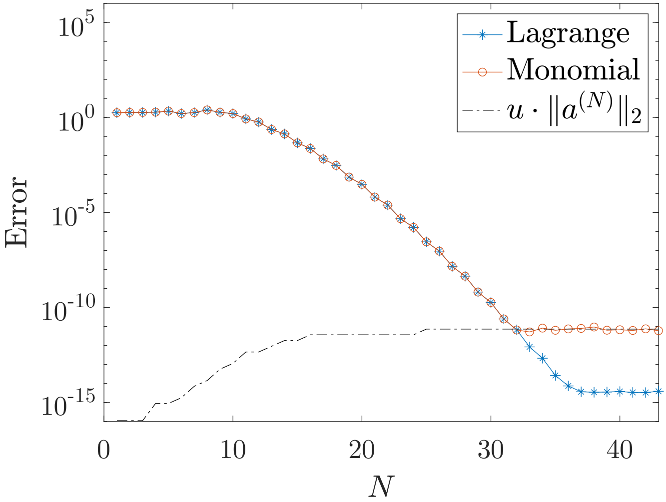

Examples

$$F(x) = T_{20}(x)$$

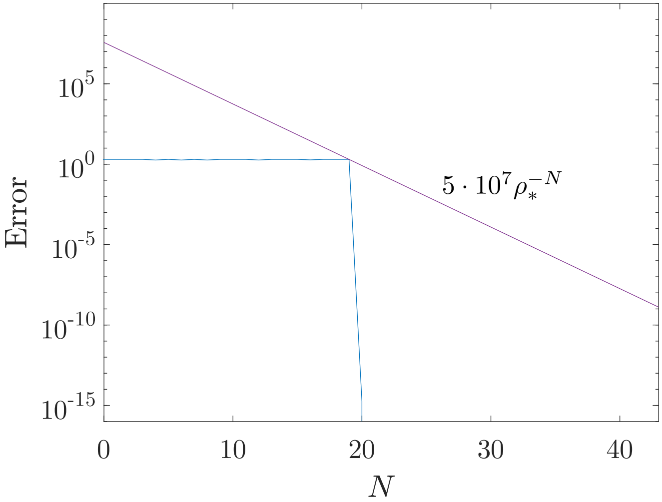

Examples

$$F(x) = T_{20}(x)$$

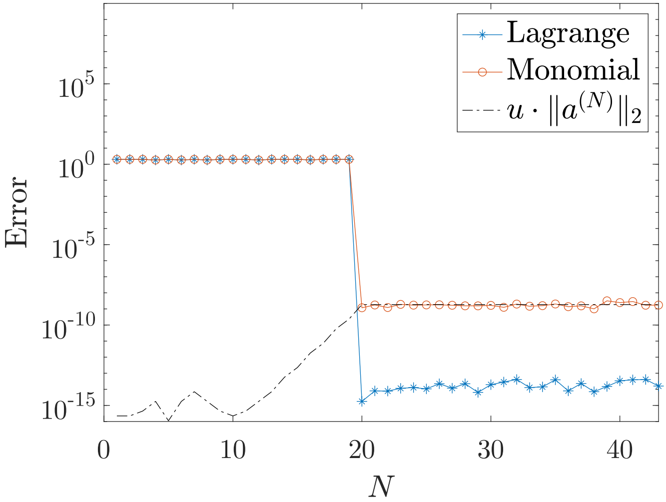

How restrictive is the monomial basis?

- Extremely high-order interpolation is impossible due to the precondition $\|(V^{(N)})^{-1}\|_2 \lesssim \frac{1}{u}$.
- So **global** interpolation won't work.

How restrictive is the monomial basis?

On the other hand, **piecewise** polynomial interpolation in the monomial basis over a partition of Γ can be carried out stably, provided that

1. the maximum order of approximation over each subpanel is maintained below the threshold;
   Fine. The threshold isn't small and can be estimated easily.
2. the size of $u \cdot \|a^{(N)}\|_2$ is kept below the size of $\|F - P_N\|_{L^\infty(\Gamma)}$.
   1. Often satisfied automatically. If not, adding an extra level of subdivision almost always resolves the issue. Reducing the maximum order also helps.
   2. Even easier when high accuracy is not required.
   3. $u \cdot \|a^{(N)}\|_2$ can be easily estimated a posteriori.

The convergence rate of piecewise polynomial approximation is $\mathcal{O}(h^{N+1})$.

Conclusions

There are many other applications of this work (see our paper).

This paper is not only about monomials. It characterizes the universal behavior of function approximation with any ill-conditioned basis before the condition number reaches $1/u$.

Paper & slides are available on my personal website (`https://zewenshen.github.io`).

Thank you for listening!

Bonus

How restrictive is the monomial basis?

- The Vandermonde system is dense.
- Backward stable linear system solve generally takes $\mathcal{O}(N^3)$ operations.

Not a problem.

- The size of the Vandermonde matrix is not large ($\lesssim 50$ in 1-D).
- Highly optimized linear algebra libraries, e.g., LAPACK.
- $\mathcal{O}(N^2)$ algorithms exist (could be less backward stable).

Generalization to higher dimensions

In 2-D, the Vandermonde matrix looks like

$$V^{(N)} := \begin{pmatrix} 1 & x_1 & y_1 & x_1^2 & x_1 y_1 & \cdots & y_1^N \\ 1 & x_2 & y_2 & x_2^2 & x_2 y_2 & \cdots & y_2^N \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{\widetilde{N}} & y_{\widetilde{N}} & x_{\widetilde{N}}^2 & x_{\widetilde{N}} y_{\widetilde{N}} & \cdots & y_{\widetilde{N}}^N \end{pmatrix},$$

where $\widetilde{N}$ is the dimensionality of bivariate polynomials of order up to N.

Collocation points with relatively small Lebesgue constants have been constructed (Vioreanu & Rokhlin 2014).

The theory of monomial approximation is essentially same as 1-D.

Numerical experiments

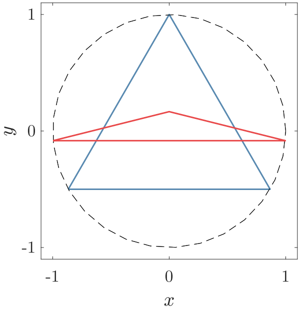

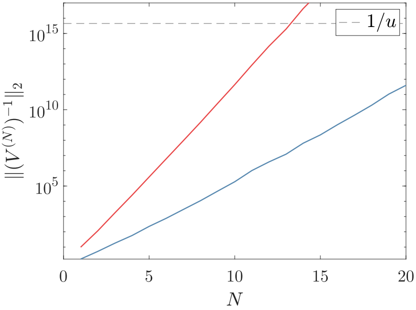

I'll show some experiments that compares the monomial basis with the Koornwinder polynomial basis over the blue triangle.

Numerical experiments

$$F(x,y) = e^{-(x^2+y^2)/4}$$

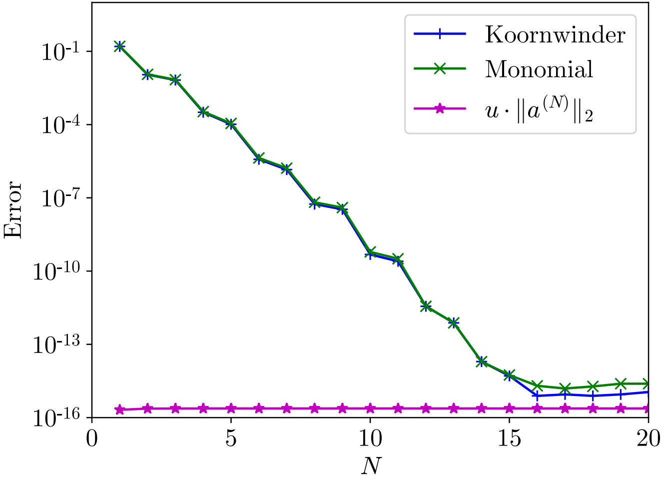

Numerical experiments

$$F(x, y) = \sin(xy/2 + x + y)$$

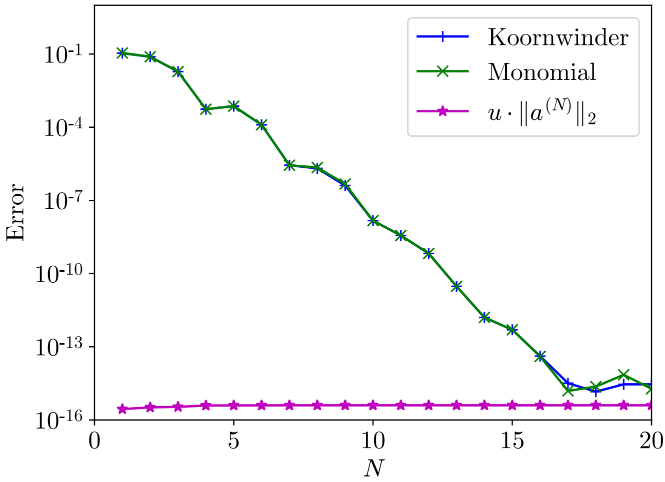

Numerical experiments

$$F(x, y) = \arctan(x) \cdot \arctan(y)$$

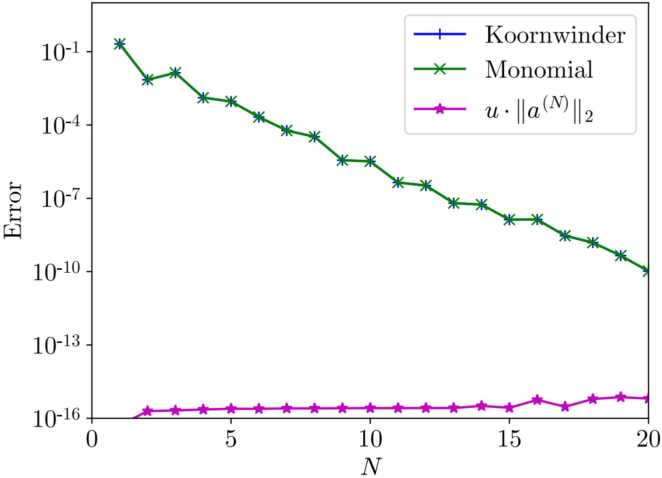

Numerical experiments

$$F(x, y) = |x + y|^{5.5}$$

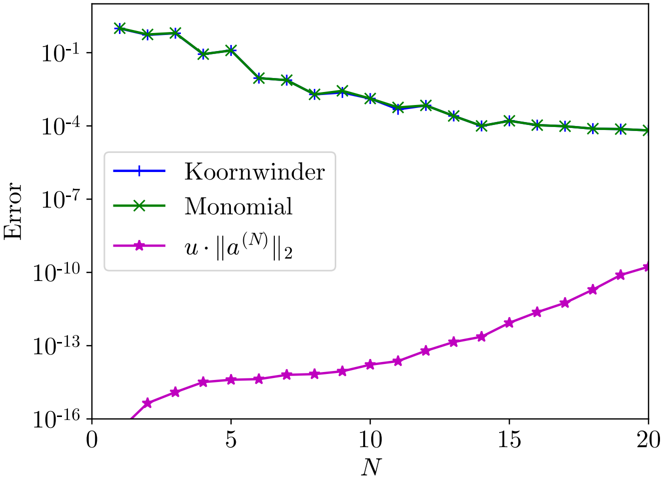

Bonus: what happens when the order $>$ the threshold?

$\cos(12x + 1)$, MATLAB's backslash