The Neurosurgery Resident Who Proved Crouzeix's Conjecture

Alex Townsend and Anne Greenbaum

August 11, 2026

On July 30, 2026, the first author asked ChatGPT 5.6 to solve Crouzeix's conjecture, as he had been routinely doing for the preceding year. Until that day, it had either returned a bogus proof or produced an argument that stalled at a missing lemma. But on July 30 it linked to a preprint posted July 27 titled *The Numerical Range Is a 2-Spectral Set*, claiming that the conjecture was solved. Crouzeix's conjecture had been open for more than two decades, so we initially read the preprint with some skepticism, but after a few hours we realized that the argument was the real deal.

Figure 1: Photo of Dr. Shanmu Jin.

The author of the preprint is Dr. Shanmu Jin (see Figure 1), a postdoctoral researcher and neurosurgery resident at Peking Union Medical College Hospital. He studied geology as an undergraduate and later earned an M.D. His formal mathematical education consisted of the courses ordinarily taken by science students. "Everything beyond that has been self-taught," he told us by email.

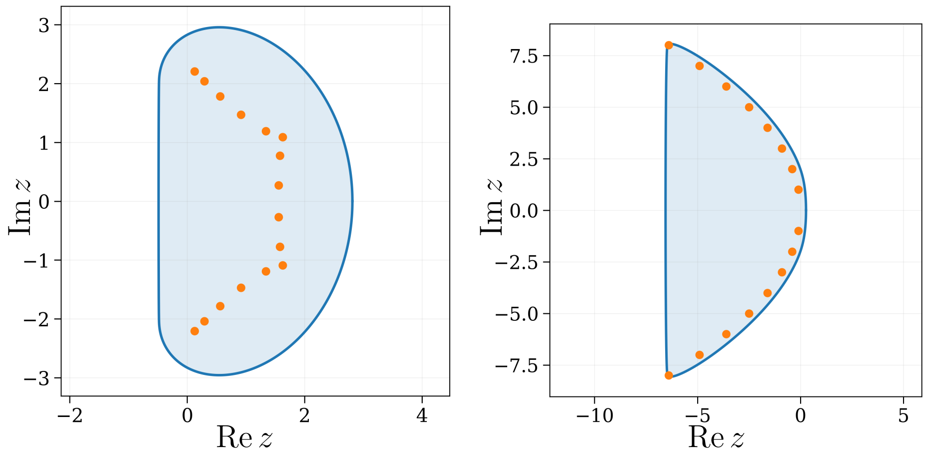

Figure 2: Numerical range $W(A)$ (blue shaded region) and the corresponding eigenvalues (orange dots) for two 16×16 nonnormal matrices: the Grcar matrix (left) and a real Schur-form test matrix (right).

Both authors of this article, as well as Michel Crouzeix himself, have checked the proof thoroughly and believe that Dr. Jin's manuscript is correct.

Jin explained to us by email that his route into matrix analysis began with his research on transcranial ultrasound. He encountered Crouzeix's conjecture while teaching himself the subject and was drawn to its unusually simple statement and to the visual geometry of the numerical range (see Figure 2). In our correspondence, Jin has been strikingly modest about the discovery. He emphasized that there was "certainly an element of luck" in finding the key idea.

For a complex matrix $A \in \mathbb{C}^{n \times n}$, its numerical range is $W(A) = \{x^* A x : x \in \mathbb{C}^n,\ \|x\|_2 = 1\}$. It is a compact convex region of the complex plane that contains the eigenvalues of A. For a nonnormal matrix, the eigenvalues can tell an incomplete story, while the numerical range captures considerably more of the matrix's behavior. Crouzeix conjectured that, for every polynomial p,

$$\|p(A)\|_2 \leq 2 \max_{z \in W(A)} |p(z)|, \tag{1}$$

where $\|\cdot\|_2$ is the spectral norm. If this holds for all polynomials (which thanks to Dr. Jin, we now know it does), it must also hold for all functions f that are holomorphic on a neighborhood of $W(A)$.

The inequality in (1) is useful. If a polynomial or rational function r approximates a function f uniformly on $W(A)$, a Crouzeix-type bound transfers that scalar approximation error to a norm bound for $f(A) - r(A)$. This provides a way to get from approximation in the complex plane to the analysis of matrix functions, GMRES, and polynomial and rational Krylov methods [4].

A substantial literature has grown around the conjecture. Crouzeix formulated it in 2004 and proved in 2007 that replacing the constant 2 in (1) with

11.08 always suffices [2, 3]. In 2017, Crouzeix and Palencia reduced the universal constant to $1 + \sqrt{2} \approx 2.414$ [5]. That same year, the American Institute of Mathematics hosted a dedicated week-long workshop in San Jose, organized by Mark Embree, Michael Overton, and the second author. Participants studied theoretical approaches, numerical evidence, matrix dilations, inverse numerical-range questions, and applications to iterative methods [1].

How did a busy neurosurgery resident solve the conjecture? Jin told us that the key result (see Theorem 3.1 of his preprint) emerged during an approximately sixteen-hour autonomous run of GPT-5.6 Sol in ChatGPT Work mode. As we read his preprint, we were stunned by the proof of that decisive theorem. Rather than requiring the much stronger estimates one expected, the argument reduces the problem via a careful sampling strategy to a surprisingly simple positivity condition. Anyone who has spent time using ChatGPT 5.6 on mathematics will have noticed how strikingly capable it has become in the last few weeks.

Dr. Jin did prompt ChatGPT in a very particular way, and many of us may want to copy him for other problems. He adapted a prompt used by OpenAI to solve the Cycle Double Cover conjecture. The prompt explicitly denied the system access to the public web and other external contexts. Jin then started the run and did not intervene. The prompt required a branching portfolio of genuinely different approaches. It told ChatGPT to spin up lots of subagents to individually think about proof strategy, and that these agents should not converge prematurely on the same attractive idea. Instead, they should repeatedly subject candidate proofs to adversarial audits and only rule out strategies with counterexamples. Finally, the prompt instructed ChatGPT not to give up until a complete proof survived checking. The exact prompt is publicly available [7]. Jin's repository also includes the prompt, successive manuscripts, a Lean formalization, and an axiom audit, which makes the work unusually open to examination [7]. Remarkably, all Jin's work was done alongside his primary occupation: clinical neurosurgery, with a schedule that sounds grueling.

On August 4—only eight days after Jin's first preprint appeared—Emiel Lorist and Felix Schwenninger posted an independent five-page proof [8]. Their argument is different and strikingly short. It combines the classical double-layer representation with a perturbation lemma for 2-dilations. It might be a more friendly argument to those who tried to prove Crouzeix's conjecture in the past. Lorist and Schwenninger also disclose that ChatGPT 5.6 was used to explore proof strategies [8]. Dr. Jin is aware of this independent proof, and his reaction to it was pure delight: he described it as an interesting complement to his own method and was pleased to see two genuinely different proofs emerge almost simultaneously.

It is remarkable to realize that a longstanding conjecture was first solved by someone with no specialized training in mathematics, working with a large language model. Only a few years ago, it was difficult to imagine an AI system contributing the decisive idea in a proof of a major conjecture. Now it has happened in our field.

We suspect that episodes like this—in which people outside professional mathematics make major contributions with the help of AI—will become increasingly common, even if the burden of verification continues to fall largely on experts. We do not yet know whether academia has the capacity to absorb the flood of results that may be coming or have time to digest them all. This is surely the beginning of a golden age for results in applied mathematics, but what it means to be an applied mathematician may be about to dramatically change.

References

[1] American Institute of Mathematics, "Crouzeix's conjecture," workshop held July 31–August 4, 2017, San Jose, CA, `https://aimath.org/pastworkshops/crouzeix.html`.

[2] M. Crouzeix, "Bounds for analytical functions of matrices," *Integral Equations and Operator Theory*, 48 (2004), pp. 461–477.

[3] M. Crouzeix, "Numerical range and functional calculus in Hilbert space," *Journal of Functional Analysis*, 244 (2007), pp. 668–690.

[4] M. Crouzeix and A. Greenbaum, "Spectral sets: Numerical range and beyond," *SIAM Journal on Matrix Analysis and Applications*, 40 (2019), pp. 1087–1101.

[5] M. Crouzeix and C. Palencia, "The numerical range is a $(1+\sqrt{2})$-spectral set," *SIAM Journal on Matrix Analysis and Applications*, 38 (2017), pp. 649–655.

[6] S. Jin, "The numerical range is a 2-spectral set," preprint, 2026, `https://www.preprints.org/manuscript/202607.1919`.

[7] S. Jin, *CrouzeixConjecture*, public research repository, 2026, `https://github.com/jinshanmu/CrouzeixConjecture`.

[8] E. Lorist and F. L. Schwenninger, "A solution to Crouzeix's conjecture," *arXiv preprint* arXiv:2608.03841 (2026), `https://arxiv.org/abs/2608.03841`.