A Close Encounter of the First Kind* with ChatGPT

James Ferry
Metron, Inc.
Reston, Virginia, USA
ferry@metsci.com

How can Large Language Models (LLMs) like ChatGPT assist us with our technical work? True believers in Artificial Intelligence (AI) insist that those who leverage LLMs will displace those who don't during a chaotic transition period before AI makes human intelligence obsolete. Skeptics are unconvinced. They see ChatGPT as a souped-up search engine, rather than agent of societal upheaval. Whether harbinger of the AI singularity or mere shiny new toy, I didn't want to miss out. So, in early 2023, as LLMs were beginning to make headlines, I signed up for a paid subscription to OpenAI and looked for excuses to use it. Like many people, I found it helpful for generating snippets of code with the correct syntax. But could it do "real" math?

Shortly thereafter, I received the following late-night email from my colleague Roy Streit. "Jim, I encountered this problem today and thought of you. An orthogonal matrix Q is used to rotate an axis-aligned, origin-centered box in $R\hat{}n$. What is the smallest axis-aligned box that contains the rotated box? It seems like this ought to be easy, but it's late and I don't see a way to do it that doesn't require $O(2\hat{}n)$ calculations."

Hmm. Me neither. Perhaps the day has arrived when we handle such questions as offhandedly as using a calculator or performing a Google search. I pasted the question into ChatGPT 4, but left off the final sentence about efficiency. This rendered the question rather trivial. Would it be able to give a sensible mathematical answer to this easy version of the question?

Yes. ChatGPT noted that the box's "corners can be represented by points whose coordinates are all possible combinations of $\pm a_i$" where "a_i represents the half-length of the box along axis i." Apply Q to each of these corner points, find the maximum absolute value in each coordinate i, and then double the results to convert to side lengths. Very sensible.

But now came the real test. I responded, "Your algorithm is $O(2\hat{}n)$. Can you do this in polynomial time in n?"

"Yes," ChatGPT responded without hesitation. Its response was long and thorough, so I'll summarize. It noted that the values of the i^{th} coordinate of the transformed box are the 2^n values $\pm Q_{i1}a_1 \pm Q_{i2}a_2 \pm \cdots \pm Q_{in}a_n$, where the Q_{ij} comprise the i^{th} row of Q. The maximum over these 2^n values is achieved when the signs are chosen to make each coefficient $\pm Q_{ij}$ non-negative. Therefore if we let $\vec{a}$ and $\vec{b}$ be the respective half-lengths (or, if you prefer, lengths) of (a) the original box and (b) the bounding box of its transformed version, then $\vec{b} = M\vec{a}$, where $M = \text{abs}(Q)$ is the matrix of entry-wise absolute values of Q.

Roy and I were delighted by the result. Neither of us had encountered, before now, the entry-wise absolute values of a matrix forming a useful operator. We were also curious about whether ChatGPT was truly reasoning, or leveraging a pre-existing solution from its training data, or something in-between. ChatGPT does not offer insights into how it generates its results. Microsoft Copilot does: it answered the original question, which did not ask for efficiency, with the efficient $O(n^2)$ solution, and when asked how it did this, it asserted that it reasoned it out for itself. I almost expected it to add, "and I'm offended that you would even ask."

LLMs are strange tools. They can pass for human in a sufficiently casual Turing test but are prone to outputting all-too-plausible nonsense. So far, I remain a casual user—I haven't tried to hone my prompt engineering skills, for example, or played with power tools like Auto-GPT. But I look forward enhancing my productivity by staying plugged into the technology as it evolves, whether it shepherds us into some brave new world or just the latest iteration of our current one.

* "Visual sightings of an unidentified flying object, seemingly less than 500 feet (150 m) away, that show an appreciable angular extension and considerable detail." https://en.wikipedia.org/wiki/Close_encounter

James Ferry, Ph.D., is a Principal Research Scientist at Metron, Inc. He specializes in applying Bayesian methods to solve Defense and Intelligence problems in a variety of domains, such as test and evaluation, tracking and data association, machine learning, and network science. He has 30+ publications in these and other areas. Prior to joining Metron in 2004, Dr. Ferry developed supercomputer simulations of multiphase fluid flows as a Research Scientist at the Center for Simulation of Advanced Rockets at the University of Illinois. He holds a Ph.D. in applied mathematics from Brown University and an S.B. in mathematics from MIT.