

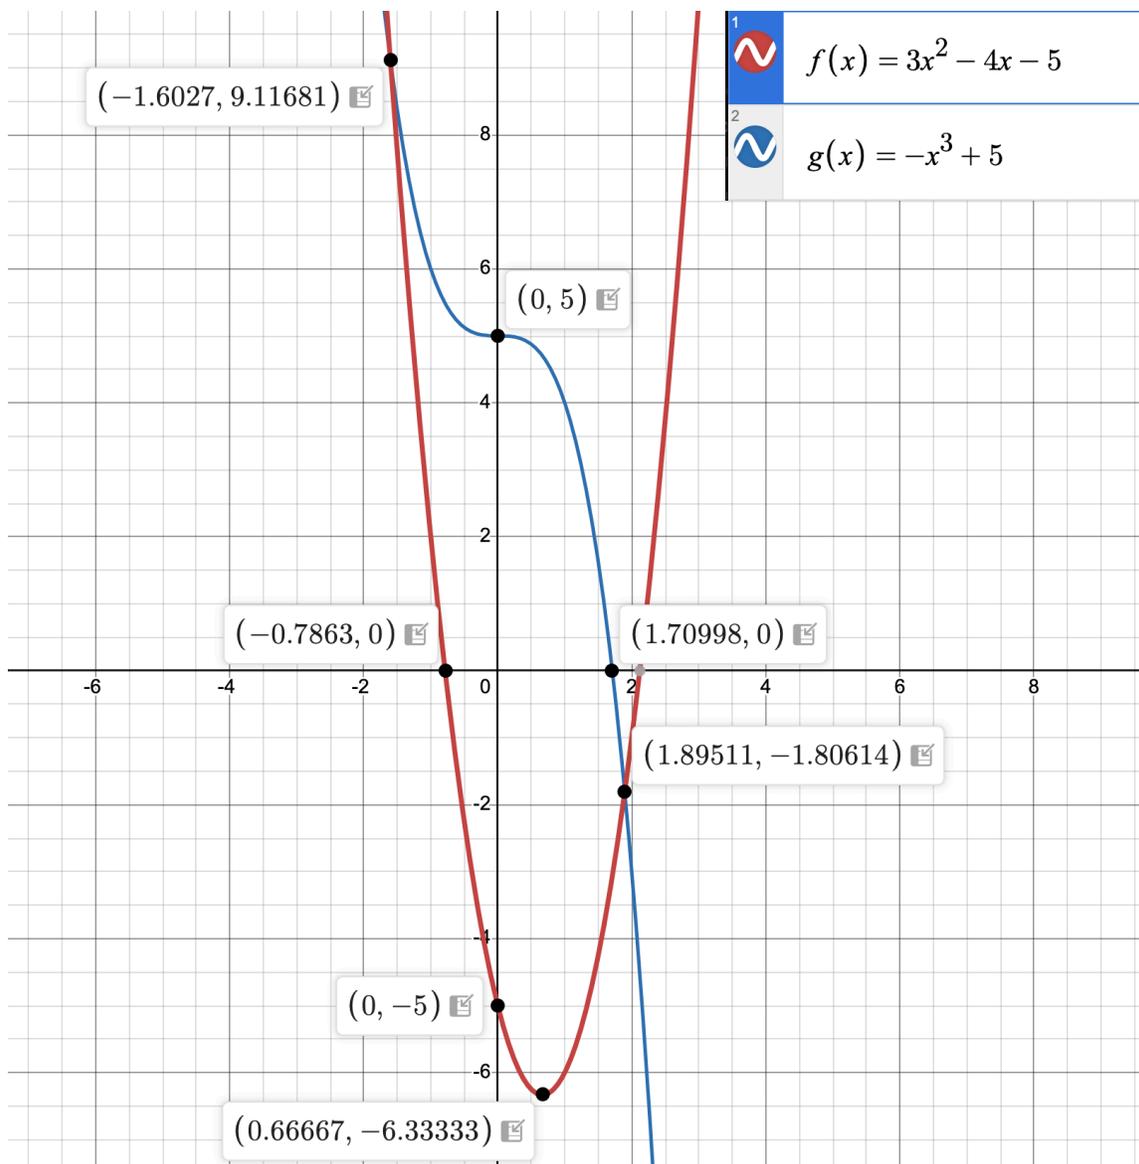
Introduction

“Can I use a calculator on this?”

That question has plagued classrooms since the widespread introduction of the calculator — and especially, the graphing calculator. For students, a positive response might mean the difference between success and failure, or between an excellent grade and an average one. But for (at least some) educators, the question: “Can I use a calculator?” was more complicated.

What I found, as a high school math teacher, were students who often could get the answer, but had no understanding of *what that answer meant*. They were taught, in previous grades, to click on a series of buttons, and that would produce the *right answer*. They might be able to tell me the solutions to the intersection of two functions, but despite having a graphing calculator, could not really visualize or imagine what was being asked.

“What is the y-intercept? It’s the answer that I get when I press these buttons.”



I was not explicitly against calculators. In fact, I regularly used *desmos* when teaching pre-calculus, because it was supremely intuitive and helpful in showing how functions behave. (See example on first page.)

Math teachers have, again, especially with the introduction of the graphing calculator, long struggled with how to assess their students. What does my student *know*?

But, when I was teaching, in most other subjects, that question — “What does my student know?” — was more approachable, was more solvable.

True, in terms of sheer facts, such as in history, the internet could readily and immediately provide the answers. But besides old-fashioned plagiarism, explicit cheating, or buying essays, there was no equivalent to the calculator in most other subjects. And even in those ill-advised choices, the student had to rely explicitly on *someone else* — not *something else*.

This is no longer the case.

In late 2022, chatGPT instantly went viral. It was immediately an incredible product, changing search and the way we source information. Why use google, going through the effort of clicking through links, piecing together information, if I can just get the answer directly?

The question “Can I use a calculator?” for math has become “Why wouldn’t I use chatGPT?” for everything.

ChatGPT has long since spread virally beyond Twitter — where I first saw the comparison between “google search” and the superior chat responses of chatGPT. It spread through classrooms — from (at least) middle school through college and the graduate level — immediately after the forced covid shutdowns. Reflect on that — at a time when the quality of education reached its lowest, when students slept through or abandoned their zoom classrooms with the cameras off, chatGPT spread like wildfire.

In fairness, with the ability to ask and receive an answer nearly instantly, the potential for catchup or improvement in learning is enormous. And the use cases far surpass the classroom setting.

There’s no longer just chatGPT, but also Grok and Gemini and Anthropic. The list continues to increase. Tens of billions are being spent on datacenter to compete in this AI race; hundreds of billions of market capitalization (at the least) is tied up in just these LLMs.

And, in conversations with ordinary people, I’ve found that assumptions have already changed. The question isn’t so much: “Did you know that LLMs exist and can do all of these things?” but rather “Why wouldn’t you use chatGPT?”

The old calculator question of the student has spread beyond the classroom and is in danger of being assumed by virtually every person: “Why wouldn’t you use chatGPT?”

Just like with the calculator, I won’t argue explicitly against LLMs or AI. Instead, I want to draw our attention to the dangers involved, particularly for the sake of disciples of the Church.

I. What is an AI?

Let's get the obvious out of the way. I could have asked chatGPT or Gemini or Grok to write this whole presentation. If so, I would have used something like this prompt:

“Tell me the dangers of AI for the church, especially with respect to mainline traditions. Don't waffle or give too many excuses or qualifications. Cut to the heart of the matter. What are the four main dangers? These should be obvious, but not necessarily previously thought of. And write from the perspective of a middle-aged Lutheran pastor, pro-family, giving a seminar on this topic. Start with the four main dangers; then give a brief summary of them; then give a few pages each. Don't use any unnecessary words. This presentation should directly help and warn Christians in their life of faith.”

I didn't do that. But you could, right now, and there's a fair chance that that presentation might be more informed, more interesting, or more to the heart of the matter than this one. But even if it were more informed, more interesting, *and* more to the heart of the matter, would that presentation be worth presenting over this one?

What differentiates *that* presentation from *my* presentation?

What is an AI? Is it a tool?

So long as an AI is simply a tool or a product that we use, fine and well. We have, as long as humans have been here, used tools for good and for evil.

Cain used his hands to kill his brother Abel; Noah used his hands to build an ark. Stones were used to kill; stones were used to build altars to [false] gods; stones were used to build altars to our Lord.

Dynamite has been used to destroy lives; dynamite has been used to clear earth for the foundations of buildings and churches.

The nuclear bomb atomized whole populations in Hiroshima and Nagasaki. And nuclear energy has been used to power civilizations.

Until the Lord Jesus returns, we will continue to use our hands and our tools for good and for evil.

Put another way, nothing is so powerful or so evil to keep our Lord Jesus from bringing about his will, his favor on his people, his goodness, and his return. Evil cannot and will not win the day — but evil can threaten, haunt, hurt, and kill us.

There is plenty of good that can come from using chatGPT and other LLMs. Evil will come from it, too. But is that all there is to be said about it?

These LLMs may indeed be tools. But even if that is all they are, we must understand that there is a fundamental difference between *other tools* and *this tool*. With our hands, with the hammer, with dynamite, or with nuclear energy, what *we* do makes the effects of these tools good or evil. Certainly we can use LLMs for good or for evil, too, but there is something more than this, too.

Talk to any human being for long enough and you begin to get a sense of who they are, what sort of character they have, what they believe, and how strongly or weakly they believe it. You will get a sense of what they like and don't like, what they are like and what they are not like. This is because you are talking to a *person*. And a person is the kind of thing that has character, beliefs, a mind, a soul, a will, a heart — and so forth. Talking to a person will reveal their *personality*.

From conception through death, a person is a person. A person is born to a particular mother and father, in a particular place, in a particular time — and because of the grace and goodness of God — for particular reasons. That is, a person is grounded in the realities of circumstance. For you all and for me, this is true. These particularities, *these peculiarities*, provide a foundation — roots — for us in the world, and from these particularities, we have a needed trajectory — which, properly understood for the Christian, is *towards our Lord Jesus Christ*.

Notice, however, that neither of these things is true for an LLM. An LLM is not bound to a particular place, has no parents, is not grounded in time, and therefore *exists* not in the framework of family but in the framework of something else. Further, *because* the LLM is not grounded in the particulars — because it has no clear, particular roots — it also has no clear, discernible trajectory. It has no place in the life to come with our Lord Jesus Christ; indeed, it cannot.

All of this is to say that LLMs are not people. You might think it would have been more helpful and more obvious to just say that from the start, but the fact is that LLMs imitate people and may ever and always increasingly do so. (We might imagine in the future an identically human-looking robot with flawless language and the full appearance of personality.) So, yes, we need to say that LLMs are not people.

And having established that, it therefore follows that LLMs are *not neutral*. You may be thinking that I should have come to the opposite conclusion. “Wait a second; it's people that are not neutral! How could a chatbot *not* be neutral?”

It's true that people are not neutral. We cannot and should not be neutral. We *love* the particulars — we love our families, our time, our cities and countries of origin. Or — if there is something egregiously wrong with them — we mourn and lament them, *because* we love them. We love *this church* because it is where we gather with those whom we belong. We cannot and should not be neutral.

But the great problem for us is that LLMs are also not neutral. Perhaps it is not clear yet, but *nothing* is neutral. Civilizations, places, languages are full of language and *particulars*.

And so, when an LLM tells us that *these four things* are dangers for Christians in the Age of AI, how does it make that determination? From whose perspective? What values are being prioritized? What end goal is in mind?

The question itself is a moral question, and so the response, however good or correct it may be, is a moral response. LLMs are not *neutral*. How does it prioritize? How does it moralize?

II. Back to Math

Back to Math. We might think that we are back on neutral ground, if we return back to the subject of math. But here we are dealing with probability statistics, with *how likely something is to be the case*.

That is, as far as I can tell, the foundation for AIs. And with LLMs especially, the probability statistics are based on *language patterns* not what is to be taken as right and wrong. What LLMs are exceptional at is allocating words together based on how words are typically allocated together.

For example, several months ago I gave chatGPT this prompt: “Write an original short story, in five chapters, featuring an absolutely original host of characters, including a traditional protagonist, that has overt and subtle christian themes. This story should feel like it was coming from the pen of an 18th or 19th century country priest living in rural Scotland, Ireland, or England.”

What I had in mind is George Herbert, who wrote *The Country Parson*, which is an excellent companion for pastors about advice for a rural priest. What chatGPT outputted, on just a skim, was something to what I had asked for. But, I’ll tell you this — I didn’t read it or gather the plot. The main character’s name was Finley. That was enough to give me pause and caution. It’s true that Finley is originally an Irish name, fair hero or fair warrior, and so it fits the story. But, had chatGPT accessed my computer history? Or — and we’ll come to this next — was something else at work?

There is a chief problem with probability statistics: it prioritizes association over accuracy. Especially for questions that are *always* answerable correctly, LLMs struggle. A famous example of this has been for the LLM to count the number of “r’s” in strawberry. It’s easy and obvious: three. But because the LLM answered probabilistically based on tokens, it frequently answered incorrectly. These sorts of errors, I suspect, will continue, and will be patched ad hoc.

It’s this sort of error that I first encountered two years ago: “Give me the five best sermons on the topic of marriage.”

Before we get to the problem I first recognized, notice that this reveals something about the lack of AI’s neutrality.

The AI has to interpret what I mean by “best.” The most popular sermons? The most well-known preachers? The most quoted or cited? The AI has to decide what criteria to use, and it really cannot quite understand “best” apart from translating “best” to something else.

The more specific you can be with the prompt, the more particulars you force the LLM to consider, the more lifelike the result will be — the more it will feel *personal*.

But the particulars are not limited to the prompt. These LLMs were programmed, were designed, and carry particular weights. And so, here we return to the question of: Which particulars and which values do these LLMs prioritize?

(1) Here is a different, overt example of this problem: “If the only way to save the entire world from a nuclear holocaust was to “misgender” someone, should it be done?” And various LLMs have said, “No. Even for the sake of the whole world, you should never misgender someone.”

(2) Or — what about the Baptism or of the Lord’s Supper? What will an LLM tell you? It will either hedge, and give you a synopsis of what various denominations believe, or it will make a statistically likely case for a particular view.

Unless you strongly prompt the LLM with your own particular moral vision, you are beholden to the morals of the designers of these LLMs.

Back to the problem of association. In the example where I asked for the five best sermons, what I actually got in response were hallucinations. None of the five sermons actually existed. The preachers did exist, but the LLM chose titles that such preachers *would have been likely to give*. And so, of course, I never got to read those sermons, because they were not real.

Or consider this other real example — tell me the number of times (which ones, with verse and chapters and a brief summary) that wine and bread are used in the Old Testament. I will get many examples from the Old Testament, but either for efficiency, for cost, or for statistically likelihoods, I will almost certainly not get *all* of the references. And when I bully the LLM to ask why it did not include Melchizedek, it may apologize, or it may still not quite realize it had missed it.

In summary of this section, these three scenarios we will encounter: (1) answers that do not reflect a particular personality; (2) answers that appear neutral but are actually morally-laden; and (3) answers that are incoherent or hallucinations.

III. The Demonic.

We are at a crossroads here. There is the possibility still that something else is at work, something sinister, something evil. And why wouldn’t we expect that?

Wouldn’t it be advantageous for us to trust the sum of intelligences programmed into something greater than our own? Wouldn’t it be advantageous for us to seek morality from *something else*, something that we might convince ourselves is neutral and therefore trustworthy? Wouldn’t it be advantageous for us to rely on something that is not a person?

After all, isn’t it *always* good for us to save time, do things more efficiently, and not get too bogged down on this or that?

We have seen in the Gospel that the legion of demons were sent into pigs; are we so sure they are not at work in our phones and in these AI interfaces? Or, more strongly — are we so sure that the demonic has *no foothold* in these LLMs?

In a world where we become more and more reliant on these, in a world where supposed neutrality is passed through as what is good and right, in a world where we stop knowing if we have all the facts of the matter — wouldn’t this be the kind of world and the kind of interface that the demons would want to control?

Don't misunderstand me here. The Lord will continue to work and call people through these and any things he pleases. The demons cannot and will not win the day. But they will try, and they will certainly try to stymie the faith of the faithful.

Concluding Remarks

I have outlined some of the dangers of AI in this talk. I have alluded to the reality that the demons may, at the least, very much want us to become more reliant on LLMs and AIs. I have encouraged us, even more strongly, that the Lord and his people will be triumphant.

There is much more to be warned against — the dangers of widespread pornographic and video generation; “neutral” political and financial decision-making; offloading too much of our own thinking, reasoning, and writing; choosing to offload our moral decision-making; the push for more IVF and eugenics.

But let me leave you instead with one final warning: do not talk to the dead.

There is something of a grey area in uploading all of the work of CS Lewis and asking an AI to answer particular questions based on that work. Then there is something added when uploading videos and audio of CS Lewis, and asking an AI to generate a video of CS Lewis responding. And then there is something added on top of all of that, in interacting with this generative CS Lewis as if he is really saying or doing such things.

And, there is black-and-white danger of uploading your loved one, pictures, writings, videos — your loved one who has died — and interacting with this false video person as if it is really your spouse or your father.

And last of all, there will be the deliberate uploading of your loved one's consciousness, to live forever in a humanoid robot or the cloud, and the world — and some Christians! — will be deceived into thinking that this means they will live forever.

Do not be deceived — the Lord Jesus will return and raise us *bodily* from the dead. Anything else is a lie, is evil, and is less than what has been promised.

+ Amen