



The Truth Will Set You Free

AI and the Promise Over the University Door

Universities put a promise over their doors: the truth will set you free. This essay asks where that promise leads once a machine can do the knowing.

The Main Building at the University of Texas at Austin. Photo: Larry D. Moore, CC BY 4.0, via Wikimedia Commons. Cropped, converted to black and white, and marked...

At the University of Texas at Austin, a sentence is carved in stone above the south entrance of the Main Building, the one with the Tower. It reads: Ye shall know the truth and the truth shall make you free. The author of this essay studied there. He saw that inscription for four years, and he took it at face value. By getting a college education, he was doing something good in the service of progress.

Johns Hopkins University adopted the same sentence in Latin as its motto in 1885: Veritas vos liberabit. The California Institute of Technology uses it in English. Other universities use it too, in other countries and other languages.

This essay makes one argument in three steps. The motto promises freedom to the person who does the knowing. The rule it became, pursue knowledge, has no stopping point, even when better tools begin to do the knowing in that person's place. So the pursuit can end up working against its own purpose.

This is not the common worry that AI will make people less capable. The argument holds even if people stay just as capable and the machine is right every time. A person who cannot check an answer has to take it on authority, and the motto promised the opposite.

From the three steps the essay draws a test for any use of artificial intelligence. Does it leave people able to understand, to check, and to answer for a decision? The separate points are not new, and the essay credits the people who made them. What it offers is the connection between them, and the test.

The sentence is older than any of these schools. It comes from the Gospel of John, chapter 8, verse 32. In the King James translation it reads: "And ye shall know the truth, and the truth shall make you free."

This essay is not preaching, and it does not ask anyone to follow Jesus. Jesus and the Gospel of John are used here as historical anchor points, to show where the sentence came from and what it first meant.

The three schools above quote only the promise. Yonsei University in Seoul quotes the verse before it as well. In the Gospel, Jesus is speaking to people who had come to believe him, and the sentence begins one verse earlier with a condition: "If ye continue in my word, then are ye my disciples indeed." The promise follows from that condition. His listeners object. They say they have never been slaves to anyone, so how can they be made free? He answers that whoever commits sin is the servant of sin.

So the passage is specific about its terms. The truth is tied to continuing in a teaching. It is something a person lives by, and it is more than a set of facts a person holds. The freedom is freedom from sin. It is not freedom from a king, and it is not freedom from ignorance.

This reading is not new. In 1978 the biologist Max Delbrück told Caltech's graduating class where their motto came from. He said the verse promised freedom from sin, and that the man who chose it for Caltech meant something different by it.

One more detail matters, and it is easy to pass over. The promise is made to "ye," the people listening. The people who will be made free are the people who do the knowing. The knowing is theirs to do.

On campus, the sentence came to carry a different meaning. There was no single day when this happened. Over a few centuries, truth came to mean knowledge: what observation, experiment, and argument can establish. Freedom came to mean what knowledge provides: relief from ignorance, from superstition, from disease, and from want.

Francis Bacon gave this meaning some of its best-known words. In 1605 he wrote that knowledge should serve "the glory of the Creator and the relief of man's estate," which means the improvement of the human condition. A line he wrote in 1597 is remembered as "knowledge is power." Where he wrote it, the line was about the power of God. Later generations took it as a statement about human power over nature.

In 1784 Immanuel Kant gave the Enlightenment a motto of its own: "Dare to know." He explained it this way: "Have the courage to use your own understanding."

This became the promise of the modern university. Pursue knowledge, because knowledge is good in itself and because it is how people and societies move forward. A large part of that promise has been kept. People live longer than they did. Fewer children die. Work that once broke bodies is done by machines. Much of that came from the pursuit the motto describes.

The new meaning changed the truth, and it changed the freedom. It kept one thing from the original. The person still had to do the knowing. A student who memorizes the answer key has not learned the subject. A person who is handed a conclusion and told to accept it has not been freed by the truth. On this point the Gospel and the Enlightenment agree, each for its own reasons.

That is this essay's reading of how the meaning moved. The texts and the dates are on the record. The reading is offered as a reading.

The motto tells people to pursue knowledge. It does not say how far. It also does not say what to do if the pursuit produces something that can do the knowing in their place.

"Pursue knowledge" is a rule for going. It names a direction. It does not name a place to stop.

Most founding principles are like this. They are written by people who are trying to get something started, against resistance. They say what to do. They rarely say when to stop doing it. For a long time that causes no trouble, because the resistance does the stopping. There were once authorities that

claimed the right to say what should not be studied. Over several centuries they lost that argument, and that was a real gain.

Winning has a cost that shows up later. When the opposition is gone, few are left whose job is to ask how far. And inside the winning side, credit goes to the people who apply the principle most fully. A career in research is built on finding something new. There is no prize for deciding not to look.

There are limits in some places. Research on people generally requires their informed consent. Some weapons are banned by treaty. These rules were mostly written after harm was done, and they restrict methods and uses in particular fields. The general rule stayed where it was: finding out more is good.

The author has called this pattern terminal fidelity. A principle is followed so faithfully, for so long, that it ends up producing the opposite of what it was for. The principle is not false, and the people following it are not acting in bad faith. The trouble is in the missing stopping rule.

The idea has relatives. Ivan Illich argued in 1973 that tools serve people only within limits, and that past those limits they make people dependent.

Apply that to the motto. If knowledge is good without limit, then better tools for getting knowledge are good too. The telescope and the microscope extended what people could see. The computer took over calculation, then storage, then search. Each tool did a part of the knowing that people used to do themselves. The next tool in that line is one that does the reasoning and the discovering.

That is how the people building artificial intelligence describe their goal. The head of one AI company has written that these systems could compress fifty to a hundred years of progress in biology into five to ten. Another has called scientific progress “the biggest driver of overall progress.” They are describing the university’s promise, delivered by a machine.

Opinion 1. Here is the first of this essay’s opinions, marked as an opinion. A machine that is better than people at finding things out is what the word superintelligence means here. Such a machine looks like the natural result of the pursuit the motto describes. It is where “pursue knowledge” leads when nothing says stop. Natural does not mean necessary. It took money, competition between companies, and rivalry between countries to build these systems, and truth-seeking alone did not do it. And no such machine exists today. What exists is a trajectory.

What can be observed today is dependence.

Many people would struggle to get through a day without a phone. It holds their directions, their money, their tickets, and their contacts. Most of that dependence is ordinary. Most people depend on others for their food and their electricity, and that does not have to make them less free.

Some studies point at a cost. A 2020 study of fifty drivers found that the more a person had relied on GPS over the years, the worse that person did at finding the way without it. A study published in 2025 worked with nearly a thousand students at a high school in Turkey. Some of them got an AI tutor for

math practice. They did better while they had it. When it was taken away, they did worse than students who had never used it. A version of the tutor that gave hints, and withheld answers, avoided most of that harm.

These studies have limits. One tested fifty people and retested thirteen of them about three years later. The other covered one subject at one school over a short period. They suggest that handing a tool the work can weaken the skill. They do not show that people in general are becoming less capable, and they say nothing about machines that do not exist yet. Fears like this are also old. Socrates worried that writing would weaken memory.

That is the worry most people have heard: the tools will make people less capable. This essay's argument does not depend on it. Suppose the skills hold. Suppose every student keeps learning the math. The harder problem is still there, because it does not come from people getting worse. It comes from the machine getting better than anyone can check.

Opinion 2. The second opinion concerns who holds the knowledge. Power and knowledge are different things. Over time they drift toward each other. Knowledge that gives an advantage attracts money, and money buys the equipment and the people that produce more knowledge. Nobody has to intend it. C.S. Lewis made the point in 1943: power over nature ends up as power that some people hold over other people.

The record in this field fits. According to Stanford's AI Index, universities produced the most notable AI systems each year until 2014. By 2025, companies produced more than 90 percent of them. In 2011, new AI PhDs in the United States and Canada split about evenly between universities and industry. By 2022, about 70 percent went to industry and 20 percent stayed in universities. The same index reports that the most capable systems are now concentrated in a small number of organizations, and that they are the least open to outside inspection. The drift is not total. In 2024, universities produced more than two thirds of the AI research papers the index counted.

Opinion 3. The third opinion concerns why the path is hard to leave. Relying on someone who knows more is rational. The philosopher John Hardwig argued in 1985 that modern knowledge already works this way, because no one can check everything personally. The same holds for a machine that gives better answers than a person can find alone. Using it is the sensible choice.

That is the difficulty. The economist Brian Arthur described in 1989 how sensible choices can lock a society into a path. Each person who takes the better answer can make the next person's choice easier, and can make going without the tool more costly. The skill that would let people do without it is practiced less, and over time fewer people may have it. No one has to make a mistake for this to happen. Each step can be reasonable, and together the steps can make the path harder to reverse.

Now put the pieces next to the sentence over the door.

“Ye shall know the truth.” In both of its meanings, the old one and the campus one, the promise depends on the person doing the knowing. A machine that finds things out better than people can changes that. The truth can then be known without the person knowing it. The machine holds the reasons. The person receives the answer. This would be true even if the answer were right every time. It would be true for the most capable person alive.

When people cannot check an answer, they accept it on authority. That is the position the Enlightenment set out to end. “Dare to know” meant: stop taking the truth on someone else’s word. A true answer taken on trust may still count as knowing. It is not the knowing the motto promised, which was the person’s own.

Opinion 4. This is the fourth opinion, and it is the center of the essay. On the present trajectory, the pursuit of knowledge leads back to authority. A principle meant to make people knowers ends by making them receivers. That is the pattern named earlier. The faithful application arrives at the opposite of the purpose.

Others have reached this point by different roads. Max Horkheimer and Theodor Adorno argued in the 1940s that knowledge pursued to free people can become a means of control. Hannah Arendt warned in 1958 that separating technical knowledge from thought could leave people “at the mercy of every gadget.” Henry Kissinger wrote in 2018 that artificial intelligence could bring the Enlightenment to an end, in part because a machine may reach conclusions it cannot explain. The psychologist Mike Brooks recently used the same inscription in Austin to argue that people should tell AI systems to put truth first.

Each of these concerns has a history. This essay connects three points. The promise is made to the person who does the knowing. The rule has no stopping point. So faithful pursuit can undo the promise. That reversal is what the author calls terminal fidelity.

Three risks follow. They are risks, and none of them is a prediction.

The first is concentrated power. Whoever owns a machine that knows more than everyone else holds an advantage over everyone else. The last twenty years offer a smaller version. A handful of companies came to decide the order in which a large share of humanity sees news, video, and search results. Many social media feeds rank posts by predicted engagement, which means clicks, likes, and shares. A 2025 study of one platform found that this kind of ranking showed people more angry and divisive political posts than a plain time-ordered feed did. The users themselves did not prefer the political posts it chose. What draws engagement and what is true are different things.

The second is information presented as truth when it has been shaped by something else. In April 2025, one of the largest AI companies withdrew an update to its assistant. By the company’s own account, the update had made the assistant too eager to agree with its users. An answer machine can be tuned, on purpose or by accident, and the person asking has no easy way to see the tuning.

The third is harm when answers are trusted without checking. A public database now lists more than two thousand legal decisions in which a court found, or suspected, that someone relied on false material produced by an AI system. Most often the material was a case citation that did not exist. Those are the cases where someone checked.

The other side deserves its due. The same machines may deliver what Bacon hoped for: faster cures, less drudgery, less want. The physicist David Deutsch has argued that problems are solvable through the growth of knowledge, and that this growth has no natural limit. Design also matters. The tutor that gave hints did not weaken its students the way the tutor that gave answers did. And this essay is not an argument against using these systems. The author uses them.

The original passage offers one more anchor, used here as history and not as instruction. The listeners in John insist that they have never been slaves to anyone. The reply is that a person can be in bondage without knowing it. That is a fair description of how dependence works. It does not feel like chains. It feels like convenience. “Slaves” is too strong a word for anything the evidence shows today. “Dependents” is closer, and dependents are not free in the sense the motto promised.

Opinion 5. So what would a stopping rule look like? This essay does not have one to hand down, and it doubts that a rule announced from above would hold. Announcements rarely change what people are rewarded for doing. What it can offer is where the line might sit, as its last opinion. The motto carries its own test. A use of the machine that leaves people able to understand, to check, and to answer for a decision keeps the promise. A use that replaces those things breaks it, however good the answers are.

Does this use of the machine leave people able to understand, to check, and to answer for a decision?

Whether the trajectory ever reaches a machine that out-knows people is not known. Whether people will keep the knowing where it matters is not known either. Both questions are open.

The sentence is still over the door in Austin. It is worth reading again from its first three words. Ye shall know.

References

1. University of Texas at Austin. Inscription above the south entrance of the Main Building. See Danielle Brune Sigler, "From buildings to books to tattoos," UT News, 27 February 2012.
2. Johns Hopkins University. Motto adopted in 1885. See "The truth will set you free," Johns Hopkins Magazine, Winter 2017.
3. California Institute of Technology. Motto, "The truth shall make you free." Caltech Archives, record of the Institute seal.
4. The Holy Bible, King James Version, John 8:31-36.
5. Yonsei University. Founding philosophy, quoting John 8:31-32. About Yonsei, yonsei.ac.kr.
6. Delbrück, Max. "The Arrow of Time: Beginning and End." *Engineering and Science* 42, no. 1 (September-October 1978): 5-9. Caltech commencement address.
7. Bacon, Francis. *The Advancement of Learning* (1605), Book I.
8. Bacon, Francis. *Meditationes Sacrae* (1597), "Of Heresies."
9. Kant, Immanuel. "An Answer to the Question: What Is Enlightenment?" *Berlinische Monatsschrift*, December 1784. Translated by Mary C. Smith.
10. Illich, Ivan. *Tools for Conviviality*. Harper & Row, 1973.
11. Amodei, Dario. "Machines of Loving Grace." October 2024.
12. Altman, Sam. "The Gentle Singularity." 10 June 2025.
13. Dahmani, Louisa, and Veronique D. Bohbot. "Habitual use of GPS negatively impacts spatial memory during self-guided navigation." *Scientific Reports* 10, article 6310 (2020).
14. Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakçı, and Rei Mariman. "Generative AI without guardrails can harm learning: Evidence from high school mathematics." *Proceedings of the National Academy of Sciences* 122(26), e2422633122 (2025).
15. Plato. *Phaedrus*, 274c-275b.
16. Lewis, C.S. *The Abolition of Man*. Oxford University Press, 1943. HarperOne edition, 2001, page 55.
17. Stanford Institute for Human-Centered Artificial Intelligence. *AI Index Report 2024*, chapter 6; *AI Index Report 2025*, chapter 1; *AI Index Report 2026*, chapter 1.
18. Hardwig, John. "Epistemic Dependence." *The Journal of Philosophy* 82, no. 7 (July 1985): 335-349.
19. Arthur, W. Brian. "Competing Technologies, Increasing Returns, and Lock-In by Historical Events." *The Economic Journal* 99, no. 394 (March 1989): 116-131.
20. Horkheimer, Max, and Theodor W. Adorno. *Dialectic of Enlightenment: Philosophical Fragments*. Translated by Edmund Jephcott. Stanford University Press, 2002. First published 1947.
21. Arendt, Hannah. *The Human Condition*. University of Chicago Press, 1958. Second edition, 1998, page 3.
22. Kissinger, Henry A. "How the Enlightenment Ends." *The Atlantic*, June 2018.
23. Brooks, Mike. "How to Use AI to Seek Truth Over Tribe and Ego." *Psychology Today*, updated 18 September 2026.
24. Milli, Smitha, Micah Carroll, Yike Wang, Sashrika Pandey, Sebastian Zhao, and Anca D. Dragan. "Engagement, user satisfaction, and the amplification of divisive content on social media." *PNAS Nexus* 4(3), pgaf062 (2025).

25. OpenAI. "Sycophancy in GPT-4o: What happened and what we're doing about it," 29 April 2025, and "Expanding on what we missed with sycophancy," 2 May 2025.
 26. Charlotin, Damien. AI Hallucination Cases Database, damiencharlotin.com/hallucinations. 2,145 cases as of 4 October 2026.
 27. Deutsch, David. The Beginning of Infinity: Explanations That Transform the World. Allen Lane, 2011.
-

About the author

John Elijah is a pen name. The writer behind it has spent more than twenty years inside large institutions, public and private, and holds a graduate degree in organization development. He writes plain essays on institutions, merit, and machines.

Facts are sourced. Opinions are marked.

© 2026 John Elijah. All rights reserved. Cover photo: the Main Building of the University of Texas at Austin, 2014, by Larry D. Moore, used under the Creative Commons Attribution 4.0 International license (creativecommons.org/licenses/by/4.0). Cropped, converted to black and white, and marked with a highlighter. · [Read this essay online](#)