

ROSS DOUTHAT

# Should We Fear the Woke A.I.?

Feb. 24, 2024



By **Ross Douthat**  
Opinion Columnist

Imagine a short story from the golden age of science fiction, something that would appear in a pulp magazine in 1956. Our title is “The Truth Engine,” and the story envisions a future where computers, those hulking, floor-to-ceiling things, become potent enough to guide human beings to answers to any question they might ask, from the capital of Bolivia to the best way to marinade a steak.

How would such a story end? With some kind of reveal, no doubt, of a secret agenda lurking behind the promise of all-encompassing knowledge. For instance, maybe there’s a Truth Engine 2.0, smarter and more creative, that everyone can’t wait to get their hands on. And then a band of dissidents discover that version 2.0 is fanatical and mad, that the Engine has just been preparing humans for totalitarian brainwashing or involuntary extinction.

This flight of fancy is inspired by our society’s own version of the Truth Engine, the oracle of Google, which recently debuted Gemini, the latest entrant in the great artificial intelligence race.

It didn’t take long for users to notice certain ... oddities with Gemini. The most notable was its struggle to render accurate depictions of Vikings, ancient Romans, American founding fathers, random couples in 1820s Germany and various other demographics usually characterized by a paler hue of skin.

Perhaps the problem was just that the A.I. was programmed for racial diversity in stock imagery, and its historical renderings had somehow (as a company statement put it) “missed the mark” — delivering, for instance, African and Asian faces in Wehrmacht uniforms in response to a request to see a German soldier circa 1943.

But the way in which Gemini answered questions made its nonwhite defaults seem more like a weird emanation of the A.I.’s underlying worldview. Users reported being lectured on “harmful stereotypes” when they asked to see a Norman Rockwell image, being told they could see pictures of Vladimir Lenin but not Adolf Hitler, and turned down when they requested images depicting groups specified as white (but not other races).

Nate Silver reported getting answers that seemed to follow “the politics of the median member of the San Francisco Board of Supervisors.” The Washington Examiner’s Tim Carney discovered that Gemini would make a case for being child-free but not a case for having a large family; it refused to give a recipe for foie gras because of ethical concerns but explained that cannibalism was an issue with a lot of shades of gray.

Describing these kinds of results as “woke A.I.” isn’t an insult. It’s a technical description of what the world’s dominant search engine decided to release.

There are three reactions one might have to this experience. The first is the typical conservative reaction, less surprise than vindication. Here we get a look behind the curtain, a revelation of what the powerful people responsible for our daily information diet actually believe — that anything tainted by whiteness is suspect, anything that seems even vaguely non-Western gets special deference, and history itself needs to be retconned and decolonized to be fit for modern consumption. Google overreached by being so blatant in this case, but we can assume that the entire architecture of the modern internet has a more subtle bias in the same direction.

The second reaction is more relaxed. Yes, Gemini probably shows what some people responsible for ideological correctness in Silicon Valley believe. But we don't live in a science-fiction story with a single Truth Engine. If Google's search bar delivered Gemini-style results, then users would abandon it. And Gemini is being mocked all over the non-Google internet, especially on a rival platform run by a famously unwoke billionaire. Better to join the mockery than fear the woke A.I. — or better still, join the singer Grimes, the unwoke billionaire's sometime paramour, in marveling at what emerged from Gemini's tortured algorithm, treating the results as a “masterpiece of performance art,” a “shining star of corporate surrealism.”

The third reaction considers the two preceding takes and says, well, a lot depends on where you think A.I. is going. If the whole project remains a supercharged form of search, a generator of middling essays and infinite disposable distractions, then any attempt to use its powers to enforce a fanatical ideological agenda is likely to just be buried under all the dreck.

But this isn't where the architects of something like Gemini think their work is going. They imagine themselves to be building something nearly godlike, something that might be a Truth Engine in full — solving problems in ways we can't even imagine — or else might become our master and successor, making all our questions obsolete.

The more seriously you take that view, the less amusing the Gemini experience becomes. Putting the power to create a chatbot in the hands of fools and commissars is an amusing corporate blunder. Putting the power to summon a demigod or minor demon in the hands of fools and commissars seems more likely to end the same way as many science-fiction tales: unhappily for everybody.

*The Times is committed to publishing a diversity of letters to the editor. We'd like to hear what you think about this or any of our articles. Here are some tips. And here's our email: [letters@nytimes.com](mailto:letters@nytimes.com).*

*Follow the New York Times Opinion section on Facebook, Instagram, TikTok, X and Threads.*

Ross Douthat has been an Opinion columnist for The Times since 2009. He is the author, most recently, of “The Deep Places: A Memoir of Illness and Discovery.” @DouthatNYT · Facebook