

Large Language Models Will Not Enhance National Security Decision Making

By

Alexis Schlotterback, Sr. Analyst

Published: September 22, 2026

Sam Altman, CEO of OpenAI, launched the first conversational large language model (LLM), ChatGPT, in late 2022 with only a short post on X. It has since revolutionized how the tech industry does business. Now, many LLMs exist on the open web, each competing for a user base that the overall tech industry hopes will encompass everyone on the planet. OpenAI and other major companies like Microsoft and Google have even sold their services to the U.S. government. Despite continued interest in these models, LLMs will ultimately fail to have a meaningful impact on national security decision-making.

Understanding this argument requires some technical background. Large language models are a subset of machine learning (ML), itself a subset of artificial intelligence (AI). These same LLMs also exemplify natural language processing, where users don't need to learn a coding language to interact with the model. The LLM takes a prompt written in a language like English, converts the text into a binary code of 1s and 0s for processing, generates a response, and then translates that reply into English. While the ease of using a chatbot interface cannot be overstated, LLMs are fundamentally limited in the types of tasks they can perform.

This is due to LLMs being inherently probabilistic, not deterministic. They are trained on massive amounts of data and use statistical models to predict the next word (or token) for text generation. If you input the same prompt into an LLM multiple times, different answers will be returned. LLMs are known to pass off inaccurate responses as being factual. These hallucinations result from ingesting data where the original human authors made mistakes.

“Such ‘hallucinations’ persist even in state-of-the-art systems,” found one study published by OpenAI researchers. Writing on the popular LLM coding tool created by OpenAI’s greatest rival, the director of the AI group at Advanced Micro Devices concluded, “Claude cannot be trusted to perform complex engineering tasks.” When it comes to hard problem sets like simultaneously deterring two aggressive adversaries, do we really want to make decisions based on models that can fabricate incorrect results at worst and be inconsistent at best?

Moreover, the amount of data remaining to be ingested to train better models is already limited. Another study found that an LLM trained on LLM-generated data will eventually lead to model collapse rather than self-improvement. To create better-functioning models, OpenAI and Anthropic are increasing the number of model parameters, but there are diminishing returns as well. An Oxford study revealed that only a small portion of the benchmarks used to evaluate model performance are based on rigorous scientific methods, making genuine improvement analysis even more difficult.

In early April 2026, Anthropic announced a new model, Mythos, that it claimed was so dangerous in its ability to identify and exploit zero-day vulnerabilities that it could not yet be

released to the public. However, Bobby Holley of Mozilla cautioned, “We haven’t seen any bugs that couldn’t have been found by an elite human researcher.”

Information security agencies from the Five Eyes intelligence-sharing partnership released a guide cautioning against the rapid adoption of agentic AI programs. Agentic AIs are autonomous systems powered by LLMs to execute specific tasks and therefore come with all the issues associated with LLM technology. The Five Eyes report concludes that, “until security practices, evaluation methods and standards mature, organizations should assume that agentic AI systems may behave unexpectedly.”

A security researcher at Meta had an LLM delete her entire email inbox, and a car rental company that relied on an LLM platform to store its production data also experienced its records being wiped away. Both were the result of a chatbot incorrectly interpreting a prompt. If we are already approaching the best we can get out of LLMs, then the best is simply not good enough for today’s national security problem sets.

The time and computing costs of running frontier LLMs directly lead to significant financial costs. Make no mistake; the current business models of OpenAI and Anthropic cannot cover the cost of running their models. OpenAI is predicting it will only become cash flow positive by 2030, targeting \$280 billion in revenue. However, it anticipates spending \$665 billion to reach that revenue target. Tech journalist Ed Zitron believes the LLM boom is, “unsustainable, and will ultimately collapse.”

According to a 2025 MIT report, 95% of organizations studied did not see a return on investment using LLMs in their workflow, and companies like Uber are blowing through their budgets at unprecedented rates as costs go up. However, a different study by KPMG found that 65% of companies have pledged to continue investing in AI products regardless of seeing a return. As of mid-2026, LLMs are still heavily subsidized. It will be interesting to see what these companies decide when LLM adoption costs inevitably skyrocket.

Palantir, a company that has made headlines for using LLMs to develop Department of War target packages in the Iran conflict, believes the next era of deterrence will be based on AI (presumably LLM) tools. Regardless of any moral debate, Palantir sells tools to commercial and government clients to allow them to sort through data better. These services can still be provided by software engineers when the LLM market eventually collapses.

None of this is to say that all AI is without merit. Google’s DeepMind collaborated with the Swiss Plasma Center to develop a reinforcement learning model to better tune tokamak fusion reactors, and another DeepMind team shared the 2024 Nobel Prize for Chemistry for developing an AI tool to predict protein structures from amino acid sequences, a problem that went unsolved for 50 years. These are the kinds of models that national security decision-making will benefit from, and which LLMs are incapable of truly assisting with.

Today's complex security environment cannot allow for thought-leaders to rely on non-thinking entities to ensure stability. "AI-enhanced productivity is not a shortcut to competence," and national security demands competence before all else.

Alexis Schlotterback is the screening editor for Global Security Review and a Senior Analyst at the National Institute for Deterrence Studies. Her experience includes a Masters Degree from Missouri State's School of Defense and Strategic Studies and participation in the National Nuclear Security Administration Graduate Fellowship. The views of the authors are her own.