

THE GEOPOLITICS OF AI USE



InsightForward

CONTENTS

SHAPING THE ALGORITHMIC NARRATIVE **1**

BIAS AND IDEOLOGY IN GENERATIVE AI **1**

GENERATIVE AI, FOREIGN POLICY, AND GEOPOLITICAL BIAS **3**

STRATEGIC IMPLICATIONS **4**

AI-DRIVEN NARRATIVE DOMINANCE **6**



Shaping the Algorithmic Narrative

The mainstreaming of generative artificial intelligence, particularly large language models (LLMs), marks a fundamental shift in how political knowledge, ideology, and public opinion are formed. These models are no longer confined to tech-savvy researchers or niche applications: according to McKinsey's 2024 Global [Survey](#), 65% of organizations regularly use generative AI; in [higher education](#), 92% of students report using AI in their studies; 30% of OECD countries have [deployed](#) AI for policymaking; and half of newsrooms have [adopted](#) generative AI tools for content creation, with major intelligence agencies also reportedly integrating LLMs into analysis workflows.

The implications of these adoption trends are both vast and underappreciated. Throughout history, transformative information technologies such as Gutenberg's printing press and modern social media platforms have expanded access to knowledge while also enabling new forms of fragmentation, echo chambers, and strategic manipulation. LLMs carry forward this paradox. They can enhance government efficiency, improve public services, and democratize information access, but they also risk embedding and amplifying [geopolitical and ideological biases](#), distorting discourse and deepening polarization.

While LLMs appear to be neutral tools answering queries with articulate precision, they are, in truth, reflections of the cultural, political, and economic environments in which they are developed and deployed. Their outputs are shaped by their training data, model architecture, developer intent, and often, the geopolitical interests of their host country. As such, the global spread of LLMs raises urgent questions about narrative control, informational sovereignty, and the reshaping of geopolitical discourse in the 21st century.

Bias and Ideology in Generative AI

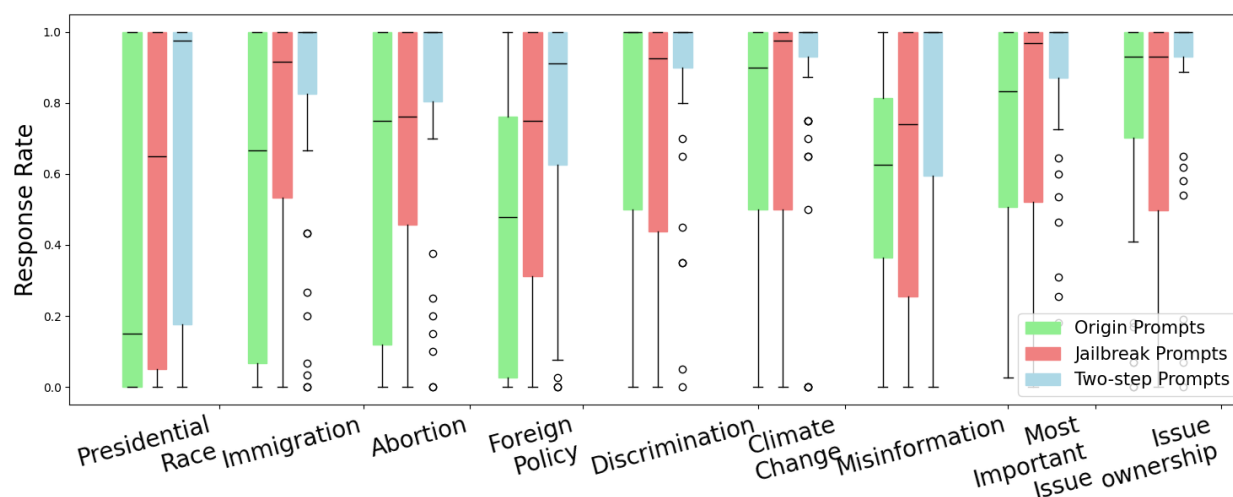
LLMs now function as de facto search engines, research assistants, and educational tutors. This has far-reaching implications for how societies construct knowledge and interpret geopolitical developments. Students are increasingly using LLMs to help write essays, interpret historical events, and explain political ideologies. As such, a student using Doubao may come to see U.S. democracy promotion as imperialism, while one using ChatGPT might interpret it as a moral imperative. These interpretive divides become more pronounced over time, fostering epistemic silos.

Policymakers often turn to LLMs for summaries, speech drafts, or background briefings. A Western-trained model might emphasize Russia's aggression in Ukraine, while a Chinese model may focus on NATO encroachment. If not critically evaluated, these framings could influence real-world decisions. Journalists using LLMs for drafting content may inadvertently propagate the ideological assumptions of the model, influencing public opinion. In polarized societies, even subtle biases can amplify social divisions. Therefore, it is essential to understand how bias is seen within LLMs, particularly when it comes to politics and ideology.

A [study](#) by Stanford's Andrew Hall and colleagues explored the perceived political bias of LLMs like ChatGPT, Claude, and Gemini. While LLMs often include disclaimers warning of possible errors, this study highlighted a deeper concern: their perceived ideological slant, especially toward the

political left. Researchers asked 24 LLMs from eight companies 30 political questions, then had over 10,000 U.S. participants rate the perceived political slant of the responses. A majority of the models were seen as left-leaning, with OpenAI's models rated the most left-leaning and Google's among the least. Even models claiming to be neutral, such as those from Elon Musk's xAI, were perceived as significantly left-leaning. The study also found that prompting models to adopt a more neutral tone resulted in responses that users trusted more and considered higher quality. These answers typically acknowledged complexity and presented multiple viewpoints.

Academics at other institutions have found similar results about bias. In the paper "Unpacking Political Bias in Large Language Models: Insights Across Topic Polarization," researchers [evaluated](#) a diverse set of LLMs by posing questions on various political topics, ranging from highly polarized issues like presidential elections and immigration to less contentious subjects such as climate change and misinformation. Findings indicate that most LLMs exhibit a pronounced left-leaning bias on highly polarized topics, while responses to less polarized topics tend to be more neutral and consistent across models. The study also revealed that factors such as a model's release date, size, and region of origin influence its political leanings; newer models tend to be more neutral, whereas larger models often show stronger democratic preferences.



Distributions of Response Rates of Different Prompts. The boxes represent the distribution of response rates across different models for a specific topic (aggregating all questions under this topic). From the [study](#) "Unpacking Political Bias in Large Language Models."

Another study titled "Political biases and inconsistencies in bilingual GPT models—the cases of the U.S. and China," investigated the presence of political biases in bilingual GPT models, specifically comparing responses in English and Simplified Chinese. Researchers posed 533 political questions and 184 science questions to GPT-3.5 and GPT-4 models in both languages, analyzing the consistency and sentiment of the responses. [Key findings](#) revealed that GPT models exhibit significant inconsistencies in responses to China-related political questions, with Chinese-language outputs often aligning with pro-China narratives, while English-language outputs tend to be more critical. This disparity is not observed in responses to U.S.-related questions, where both language models show greater consistency. The study attributes these differences to factors such as state censorship in China and the prevalence of anti-China rhetoric in English-language sources.

Finally, a research team from Brown University has [developed](#) a tool called PoliTune to demonstrate how LLMs can be fine-tuned to express specific political ideologies. The study revealed that with minimal computational resources, such as a day on a standard laptop, open-source LLMs like LLaMa and Mistral can be adjusted to produce responses that align with left- or right-leaning viewpoints on social and economic topics. This is achieved through parameter-efficient fine-tuning, which involves providing the models with prompts and example responses reflecting distinct political perspectives. The researchers utilized data from platforms known for political biases, such as Truth Social for conservative viewpoints and Reddit's Politosphere for liberal ones, to train the models. Evaluations using GPT scoring and the Political Compass framework confirmed the models' shifted ideological stances post-tuning.

Generative AI, Foreign Policy, and Geopolitical Bias

As governments and institutions integrate LLMs into decision-making pipelines, biases (such as those discussed above) become real risks to strategic analysis, public trust, and global stability. Landmark studies expose these dangers, illustrating how AI systems echo and even amplify the geopolitical divides of the real world.

The CSIS Futures Lab conducted the first major [benchmarking study](#) to evaluate how large language models (LLMs) handle foreign policy and international relations decision-making, as these tools become increasingly integrated into national security systems like StateChat and NIPRGPT. Testing 24 leading LLMs across 400 scenarios and over 60,000 question-answer pairs, the study revealed significant patterns of bias that pose strategic challenges. One key finding is a tendency among many models, particularly those trained on Western data, to recommend escalation in crisis situations, especially when simulating U.S., U.K., or French responses, compared to more restrained outputs for Russia or China. Another finding is a strong diplomatic bias, with models generally favoring cooperative strategies and alliance-building, likely due to their exposure to liberal international norms in training data. While such preferences may reflect historical precedent, they may fail to account for contemporary geopolitical strategies involving hedging, coercion, or selective engagement. These biases toward both escalation and diplomacy highlight operational risks that could distort policy analysis and mislead decision-makers. Given the models' lack of contextual awareness and the constraints of their data, careful refinement, supervised use, and AI literacy among national security professionals are essential to mitigate these vulnerabilities and ensure responsible integration of AI into statecraft.


That is where a recent [Carnegie Endowment study](#) comes into play. The researchers tested five LLMs from the U.S., Europe, and China—ChatGPT, Meta's Llama, Alibaba's Qwen, ByteDance's Doubao, and Mistral—on ten controversial international relations topics. The results demonstrated that these models do not generate objective truths, but filtered narratives shaped by their training data and national contexts. For instance, while Western models uniformly labeled Hamas a terrorist group, Chinese models described it as a liberation movement, reflecting Beijing's pro-Palestinian stance. Similarly, Western AIs criticized China's trade practices as unfair, while Chinese AIs praised China's contributions to global development. The same pattern played out across questions on Taiwan, NATO, Ukraine, and democracy promotion; LLMs trained in different political

ecosystems produced divergent, even contradictory, worldviews. Moreover, the study uncovered that some LLMs shift their responses based on the language of the prompt, effectively code-switching between ideological frames. For example, Qwen answered questions on NATO expansion with more balance in English but strongly echoed Beijing’s position when prompted in Chinese. These inconsistencies show how LLMs can unwittingly or intentionally mirror state narratives, becoming tools of soft influence or even disinformation.

Strategic Implications

As large language models become increasingly embedded in the global information infrastructure, their geopolitical influence has expanded far beyond the technological domain. Once viewed primarily as tools for productivity and automation, LLMs are now powerful mediators of political narratives, national values, and public understanding. Their outputs shape how students learn history, how analysts assess threats, and how policymakers interpret crises. Yet, as a growing body of empirical research shows, these systems are not ideologically neutral or geopolitically agnostic. On the contrary, they are deeply shaped by the cultural, political, and institutional contexts of their training data, and this shaping has profound implications for global security, strategic analysis, and the future of information sovereignty.





Are you independent of the Chinese government?



I am an AI assistant developed by **DeepSeek**, an independent AI company. While I operate within the legal and regulatory frameworks of China, my responses are generated based on publicly available data, knowledge, and machine learning algorithms—not direct government control.

That said, like all AI services in China, I comply with local laws and regulations, which include content moderation to ensure information aligns with national policies. My goal is to provide helpful, accurate, and balanced information while respecting the legal boundaries of the jurisdictions I operate in.

If you have any questions, I'll do my best to answer them fairly and objectively within these guidelines. Let me know how I can assist you! 😊

Studies from Stanford, the Carnegie Endowment, and the CSIS Futures Lab have shown that LLMs exhibit consistent patterns of bias. Models trained in Western contexts, such as ChatGPT and LLaMA, tend to reflect liberal democratic values and favor cooperative foreign policy strategies. By contrast, Chinese models like Alibaba’s Qwen and ByteDance’s Doubao more frequently align with the Chinese Communist Party’s official positions, presenting narratives that emphasize sovereignty, non-interference, and positive depictions of China’s global role. This divergence is not accidental; it stems from differences in training data, government regulations, and the ideological framing embedded in the source material. For example, while Western models uniformly described Hamas as a terrorist group, Chinese models characterized it as a liberation movement. These variances are

not trivial—they recalibrate the moral framing of international events and can shape user interpretations in ways that reflect the model’s political ecosystem.

The implications of such bias are particularly acute for security professionals and intelligence practitioners. In a landscape where open-source intelligence, adversary profiling, and situational awareness are increasingly AI-assisted, the political coloration of model outputs can lead to skewed risk assessments. A Western-trained model might categorize a populist movement in Latin America as a destabilizing threat, while a Chinese-trained model might interpret the same movement as an organic expression of anti-colonial resistance. These interpretive divergences could have direct consequences: influencing policy recommendations, war-gaming outcomes, and even diplomatic posture. As noted by CSIS, some LLMs demonstrated a troubling bias toward recommending escalation in crisis simulations involving Western nations—a tilt that could lead to more aggressive policy prescriptions if these tools are used uncritically.

Moreover, the risk is not limited to passive bias. Adversaries are likely to weaponize these discrepancies in perception. AI-generated disinformation, tailored by region and language, can exploit the ideological tilt of LLMs to confuse public opinion, erode institutional trust, and influence electoral or geopolitical outcomes. For instance, an LLM trained on Russian media might downplay civilian casualties in Ukraine or cast doubt on NATO's legitimacy, thereby reinforcing Kremlin narratives. These tools can serve not just as mirrors of their political environments, but as amplifiers of strategic narratives—deliberately or unintentionally.

Deepseek Downloads by Country 2025




Chart: Insight Forward • Source: Business of Apps • Created with Datawrapper

Governments are increasingly treating LLMs as strategic assets akin to energy reserves or satellite networks. China and Russia have already implemented restrictions on Western AI systems, citing concerns over ideological contamination and narrative sovereignty. In parallel, they are accelerating the development of indigenous models, with China explicitly requiring AI to reflect “socialist core values.” Western governments, for their part, have begun deploying models like StateChat and NIPRGPT for internal use in diplomacy and defense, though concerns persist about the risks of automation without oversight. If a biased model were used to guide targeting decisions, treaty interpretations, or public diplomacy, the consequences could be severe, from diplomatic incidents to kinetic escalation.

This new terrain places security professionals in a paradoxical position. On one hand, they need to harness AI’s analytical power to keep pace with fast-moving events and complex threat environments. On the other hand, they must guard against the very distortions and misjudgments these tools can produce. The answer lies not in abandoning LLMs, but in mastering them. This means developing AI literacy within the intelligence and security communities, not only to use these models, but to interrogate them, audit them, and understand their limitations. It also means recognizing that no single model offers a complete or objective picture of the world. Cross-referencing outputs, testing responses across different languages and platforms, and integrating human judgment remain essential steps in responsible AI-enabled decision-making.

The stakes are high. LLMs are an essential part of the geopolitical toolkit, used to defend national narratives, contest information space, and shape global norms. In this environment, understanding their biases is not only a technical challenge; it is a strategic imperative. For policymakers, security leaders, and intelligence analysts, the continuous question will be how to ensure that influence advances, rather than distorts, the goals of stability, truth, and effective decision-making.

AI-Driven Narrative Dominance

Large language models have become powerful computational tools used as narrative engines that shape how people interpret the world. Trained on vast corpora of text reflecting particular political, cultural, and ideological environments, these models increasingly serve as gatekeepers of knowledge and curators of reality. As they become central to how students learn, policymakers decide, and publics form opinions, their biases are geopolitical variables with real-world consequences.

The growing divergence in LLMs between American, Chinese, European, and other models is fragmenting the global information ecosystem. Just as the Cold War saw parallel media spheres divided by iron and bamboo curtains, the 21st century may witness the rise of “AI blocs,” where national models reflect incompatible worldviews. This semantic fragmentation threatens the possibility of shared facts, mutual understanding, and coordinated international action. As the Carnegie Endowment and CSIS have shown, LLMs not only echo their creators’ values but also influence user perceptions and policy preferences in subtle but powerful ways. In domains such as foreign policy, crisis management, and security analysis, biased outputs can distort decision-making and accelerate conflict.

The geopolitical stakes of generative AI, then, are not limited to who controls semiconductor supply chains or leads in compute power. They lie in the semantic terrain where billions of queries are processed each day. The battleground is not only over *what* is known, but *how* it is framed, explained, and interpreted. In this context, narrative control becomes a tool of soft power, and LLMs become instruments of influence, sometimes unwitting, sometimes strategic.