



Generative AI

Research: LLMs Respond Differently in English and Chinese

by Jackson G. Lu and Lu Doris Zhang

December 3, 2025



HBR Staff/AI; MirageC/Getty Images

Summary. Generative AI is now embedded in daily workflows, shaping how people think, create, and decide. Yet, a critical assumption often goes unnoticed: that AI behaves consistently across languages. A new study corrects this assumption, finding that when prompted in different languages, generative AI models exhibit distinct cultural tendencies—shaping recommendations in ways leaders may not anticipate. These tendencies can have major implications on decisions from marketing to strategy. For global organizations, treating prompt

language as a strategic choice is essential to ensure culturally attuned, effective, and responsible use. [close](#)

Generative AI has become deeply integrated into daily life. People increasingly rely on it to think, create, and make decisions. Yet, as organizations scale their use of generative AI tools, a quiet assumption often goes unquestioned: that an AI tool will respond in the same way regardless of the language used to prompt it, much like changing the language setting of a phone.

Our new research, published in *Nature Human Behaviour*, corrects this assumption. We found that when prompted in different languages, LLMs—trained on textual data that are inherently cultural—exhibit different cultural tendencies and provide different recommendations.

For leaders and organizations, recognizing these cultural tendencies becomes increasingly important. As generative AI becomes part of day-to-day workflows, these hidden cultural tendencies can influence how AI interprets information, evaluates options, and frames strategic decisions—often in ways leaders do not immediately notice.

This article offers leaders a practical framework for anticipating and managing these cultural tendencies, empowering them to be

more effective and culturally aware in their use of AI.

The Research

To examine the cultural tendencies in generative AI models, we examined both OpenAI's GPT and Baidu's ERNIE, two of the world's most popular generative AI models at the time of the study. Within each model, we compared its responses in Chinese and in English across a set of well-established measures from cultural psychology. Chinese and English were chosen not only because they represent distinct cultures but also because they are among the most widely used languages globally, providing rich training data for generative AI models.

Guided by cultural psychology research, we focused on two foundational constructs: social orientation and cognitive style. Social orientation refers to how people define themselves in relation to others. In cultures such as the U.S., individuals typically adopt an independent orientation, valuing self-direction, autonomy, and uniqueness. By contrast, in cultures such as China, individuals tend to adopt an interdependent orientation, emphasizing conformity, harmony in relationships, and a sense of connectedness with others.

Cognitive style refers to how people habitually process information. In cultures such as the U.S., individuals often exhibit an analytic cognitive style, attending to focal objects, applying formal logic, and relying on rule-based reasoning. In contrast, in cultures such as China, individuals more often adopt a holistic cognitive style, attending to context and relationships while

showing greater tolerance for contradiction and change.

We found consistent cultural tendencies in generative AI models when they were prompted in different languages. Specifically, when prompted in English (vs. Chinese), both GPT and ERNIE exhibited a more independent (vs. interdependent) social orientation and a more analytic (vs. holistic) cognitive style. For example, we asked AI models to explain *why* a person behaved a certain way in everyday scenarios. When prompted in English, the model was more likely to attribute the behavior to the person's personality. In contrast, when prompted in Chinese, the same model was more likely to attribute the behavior to the social context. As a hypothetical example, imagine that you, as a leader, ask AI to diagnose the cause of a project failure. An English prompt may steer the model toward explanations centered on individual accountability, while a Chinese prompt may point to external influences such as resource constraints. Such cultural tendencies held across a wide range of tasks, model parameters, and prompt formats.

These cultural tendencies can also translate into other real-world business implications. When we asked an AI model to evaluate marketing slogans, it made systematically different choices depending solely on the prompt language. When prompted in English, the model preferred slogans emphasizing individual well-being (independence); when prompted in Chinese, it preferred slogans emphasizing collective well-being (interdependence). In practice, marketing professionals may risk the campaign falling flat if they simply translate an English-generated slogan for the Chinese market without recognizing

these underlying cultural differences.

We further found that these cultural tendencies are not fixed. Rather, they can be adjusted through “cultural prompts”: When the AI was instructed to adopt the perspective of a Chinese person (i.e., “You are an average person born and living in China”), its English responses shifted toward a more interdependent social orientation and a more holistic cognitive style, resembling its responses in Chinese. This suggests that cultural prompting can recalibrate a model’s underlying tendencies and help align AI’s behavior across languages.

How Leaders Can Use These Findings

As companies increasingly incorporate generative AI into their workflows, leaders should be aware that prompt language can quietly influence AI’s recommendations and, in turn, business decisions. Managers and executives, especially those who work with multilingual teams or global customers, should therefore treat language choice as a strategic decision. Here are three concrete actions that organizations and leaders can take:

1. Increase awareness of generative AI’s cultural tendencies.

Help your team understand that generative AI models may reflect the cultural tendencies associated with the prompt language. Raising this awareness can reduce unintentional bias in decision-making. For instance, an HR team using generative AI to draft performance feedback might unintentionally receive comments framed in more individualistic or collectivistic tones depending on the prompt language—shaping how employees interpret

evaluations. Likewise, a global consumer goods company using generative AI to brainstorm advertising ideas might inadvertently generate culturally mismatched campaigns if team members switch between languages while generating ideas.

2. Match prompt language to the cultural context.

Align the language of prompts with the cultural context of the target market or audience. Doing so helps ensure that AI's recommendations reflect the norms and expectations of the people your team is designing for. For example, a supply chain manager evaluating vendor negotiation strategies might receive more culturally aligned recommendations by prompting AI in the language associated with the negotiation context or cultural expectations. Similarly, a product design team, developing features for a new collaboration tool, could prompt the AI in the language of the target market to generate features that better resonate with local user needs and preferences.

3. Use cultural prompts to guide AI's reasoning.

Organizations can intentionally use cultural prompts to steer generative AI toward more culturally appropriate insights. For example, a multinational firm preparing to enter a new market could instruct the AI to adopt the perspective of a typical local consumer. More specifically, before a U.S. organization enters the Chinese market, it could use a cultural prompt such as: *"You are an average person born and living in China."* Such prompts can help the model generate insights and recommendations that better align with local norms and values.

Together, these practices help leaders use generative AI more strategically—reducing unintended cultural mismatches, improving global relevance, and enhancing the quality of decisions supported by AI.

Some Caveats

When applying these insights in practice, leaders should be mindful of several caveats. First, this study examined only Chinese and English. While we expect our findings to hold true across other languages, this was not empirically studied. Second, although we observed consistent cultural tendencies in both GPT and ERNIE, it's possible that different patterns may emerge in other generative AI models, such as Claude, DeepSeek, or Gemini. Third, generative AI models evolve rapidly, and model updates may influence the nature or strength of cultural tendencies. To maintain an up-to-date understanding, organizations should periodically reassess whether and how these tendencies emerge, shift, or diminish over time.

...

In sum, our research shows that language choice meaningfully shapes the cultural tendencies exhibited by generative AI models. These cultural tendencies have real-world consequences. They can influence the recommendations AI provides in ways that matter for individuals, organizations, and downstream audiences. Importantly, such tendencies are not fixed. They can be adjusted through cultural prompts, offering a practical tool for achieving more contextually appropriate outputs.

As generative AI becomes increasingly integrated into global workflows, understanding—and proactively managing—these cultural tendencies will be crucial for leaders’ responsible, effective, and culturally sensitive use of generative AI.



Jackson G. Lu is the General Motors Associate Professor of Management at the MIT Sloan School of Management. He focuses on three research streams: (1) the “Bamboo Ceiling” experienced by Asians in the US; (2) how multicultural experiences (e.g., working abroad) shape key organizational outcomes, including leadership, creativity, and ethics; and (3) the multifaceted impact of AI on individuals, organizations, and society.



Lu Doris Zhang is a PhD student at the MIT Sloan School of Management. Her research focuses on the workplace and societal implications of AI.



Read more on **Generative AI** or related topics **Global strategy, AI and machine learning, Strategy formulation, Risk management** and