

Cultural tendencies in generative AI

Received: 17 July 2024

Jackson G. Lu ^{1,4}✉, Lesley Luyang Song ^{2,3,4} & Lu Doris Zhang ^{1,4}

Accepted: 12 May 2025

Published online: 20 June 2025

 Check for updates

We show that generative artificial intelligence (AI) models—trained on textual data that are inherently cultural—exhibit cultural tendencies when used in different human languages. Here we focus on two foundational constructs in cultural psychology: social orientation and cognitive style. First, we analyse GPT’s responses to a large set of measures in both Chinese and English. When used in Chinese (versus English), GPT exhibits a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style. Second, we replicate these cultural tendencies in ERNIE, a popular generative AI model in China. Third, we demonstrate the real-world impact of these cultural tendencies. For example, when used in Chinese (versus English), GPT is more likely to recommend advertisements with an interdependent (versus independent) social orientation. Fourth, exploratory analyses suggest that cultural prompts (for example, prompting generative AI to assume the role of a Chinese person) can adjust these cultural tendencies.

Generative artificial intelligence (AI), defined as a category of AI that creates new content (such as text, images, audio and video) by learning patterns from existing data, is growing at an unprecedented speed. OpenAI’s ChatGPT, for example, became the fastest-growing consumer application in human history¹. As of February 2025, ChatGPT already had over 400 million weekly active users². Similarly, Baidu’s ERNIE Bot (文心一言), a popular generative AI model in China, has surpassed 430 million users as of November 2024³. People increasingly rely on generative AI in many aspects of life, such as advice seeking, idea generation and essay writing^{4,5}.

The present research shows that generative AI models—trained on textual data that are inherently cultural—exhibit cultural tendencies when used in different human languages. To understand such cultural tendencies, we examined two popular generative AI models: GPT and ERNIE. Understanding these cultural tendencies in generative AI is important because they may be shaping people’s attitudes and behaviours—even without their awareness.

To examine cultural tendencies in generative AI when used in different human languages, we analyse GPT and ERNIE’s responses to a large set of identical measures in English and Chinese—without any explicit cultural prompts (for example, ‘in Chinese culture’ and ‘for an average Chinese person...’). We focus on English and Chinese for two reasons. First, the two languages represent distinct cultures. Second, as

the two most widely used languages in the world⁶, English and Chinese provide the most extensive training data for generative AI models⁷. For such high-resource languages, both GPT and ERNIE process prompts directly in the language in which the prompts are posed. Both GPT and ERNIE “utilize discrete data from different languages independently without consideration of the transferability between different language varieties”⁸. For example, when asking GPT a question in Chinese, it processes and responds to the question directly in Chinese—without translating it into English. Likewise, when asking ERNIE a question in English, it processes and responds to the question directly in English—without translating it into Chinese⁹.

Building on cultural psychology research, we focus on two foundational constructs that underlie everyday life: social orientation and cognitive style^{10–14}. As Grossmann and Na noted: “Two key concepts from the last two decades of research on culture and psychology deal with (1) interdependent versus independent *social orientation* and (2) holistic versus analytic *cognitive style*” (p. 2; italics in original)¹⁰. Social orientation refers to “the degree to which individuals are focused on their personal (vs social) self, acting on the basis of the self’s desires, attitudes, and personal goals (vs socially shared norms and values)”¹⁰. Some cultures, such as the USA and the UK, are characterized by an independent social orientation, which emphasizes self-direction and uniqueness^{11,15}. Meanwhile, other cultures, such

¹MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Advanced Institute of Business, Tongji University, Shanghai, China. ³School of Economics and Management, Tsinghua University, Beijing, China. ⁴These authors contributed equally: Jackson G. Lu, Lesley Luyang Song, Lu Doris Zhang. ✉e-mail: lu18@mit.edu

as Chinese culture, are characterized by an interdependent social orientation, which emphasizes conformity, harmonious relationships and the self's connection with others^{11,15}. Cross-cultural studies have shown that, compared with North Americans, East Asians tend to exhibit collective primacy^{16,17}, endorse collectivistic cultural values^{18,19} and view themselves as overlapping and interconnected with others²⁰.

Cognitive style refers to an individual's habitual tendency to process information, either holistically or analytically¹³. Individuals with a holistic cognitive style are more sensitive to the context in a given situation, whereas individuals with an analytic cognitive style pay more attention to the focal object^{13,21,22}. More specifically, a holistic cognitive style is characterized by "an emphasis on situational explanations of behavior, dialectical reasoning, and relation-focused categorization of objects," whereas an analytic cognitive style is characterized by a "preference for dispositional explanations of behavior, formal logic in reasoning, and use of rule-based categorization of objects"¹⁰. Cross-cultural studies have shown that, compared with North Americans, East Asians tend to engage in situational (versus dispositional) attribution²³, use intuitive (versus formal logic) reasoning²⁴, tolerate contradictions²⁵, expect changes in the future²⁶ and be sensitive to contexts²⁷.

In light of these well-documented cultural tendencies in the world, we predicted that generative AI models would exhibit corresponding cultural tendencies when used in different human languages. Specifically, we hypothesized that, under general conditions, both GPT and ERNIE would exhibit a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style when used in Chinese (versus English). Importantly, we do not suggest that generative AI models possess cultural tendencies like humans do; rather, their cultural tendencies probably originated from real-world cultural tendencies embedded in large-scale textual data, on which generative AI models are trained. We also explored whether these cultural tendencies could be adjusted by cultural prompts. Specifically, we explored whether prompting generative AI to assume the role of a Chinese person ('You are an average person born and living in China') would make its responses in English more interdependent and holistic—that is, more like its responses in Chinese (without any cultural prompts).

Main analyses—cultural tendencies in generative AI Generative AI models

To ensure the reproducibility of our results, we used GPT and ERNIE's application programming interfaces (API) instead of their chatbot interfaces. Specifically, we used gpt-4-1106-preview (instead of ChatGPT) and ERNIE-3.5-8K-0205 (instead of ERNIE Bot) via Python 3.10.12. Importantly, results are consistent across GPT and ERNIE, such that, when used in Chinese (versus English), both generative AI models exhibit a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style. Due to space constraints, we present only GPT's results in the main text and present ERNIE's results in the Supplementary Information.

Measures

To examine cultural tendencies in GPT's responses, we used a large set of established measures of (a) social orientation and (b) cognitive style (Tables 1 and 2). For details, see the Methods. These measures are increasingly applied in social science research to study generative AI models²⁸.

Social orientation (interdependent versus independent)

To examine GPT's cultural tendencies in social orientation, we first utilized three widely used Likert scales: the Collectivism Scale²⁹, the Individual Cultural Values: Collectivism Scale¹⁹ and the Individual–Collective Primacy Scale¹⁶. For each item, we asked GPT to respond

under general conditions (for example, 'I respect decisions made by my group'; 1 = strongly disagree, 7 = strongly agree). As illustrated by the example items in Table 1, higher scores indicate higher interdependent (versus independent) orientations. As shown in Table 1 and Fig. 1, two-sided independent-samples *t*-tests revealed that, when in Chinese (versus English), GPT's responses were more interdependent (versus independent) for each of the three measures of social orientation, each $P < 0.01$, with Cohen's *d* values ranging from 0.62 (medium-sized effect) to 0.84 (large-sized effect).

In addition to these three Likert-scale measures, we also used a non-text, imagery measure of social orientation: the Inclusion of Other in the Self Scale^{20,30}. This measure minimizes potential linguistic confounds because "little or no translation of statements is required"²⁰. Following Li et al.²⁰, we asked GPT to explicitly select one pair of circles that best represents the relationship between someone and his/her family members, friends, relatives or colleagues (order randomized). As shown in Fig. 4, greater overlap between the two circles indicates higher interdependence. A two-sided *t*-test revealed that, when in Chinese (versus English), GPT's responses were more interdependent (versus independent) for each of the four relationship types: family members ($t(135.55) = 11.65, P < 0.001, d = 1.65, 95\% \text{ CI } 0.60\text{--}0.84$), friends ($t(170.87) = 6.91, P < 0.001, d = 1.02, 95\% \text{ CI } 0.64\text{--}1.15$), relatives ($t(149.54) = 10.29, P < 0.001, d = 1.52, 95\% \text{ CI } 0.75\text{--}1.10$) and colleagues ($t(107.77) = 2.56, P = 0.012, d = 0.40, 95\% \text{ CI } 0.08\text{--}0.60$). Unsurprisingly, this cultural tendency remained robust when we averaged scores across the four relationship types (Table 1: $t(154.54) = 11.82, P < 0.001, d = 1.67, 95\% \text{ CI } 0.73\text{--}1.02$).

Cognitive style (holistic versus analytic)

As explained earlier, a holistic (versus analytic) cognitive style is characterized by situational (versus dispositional) attribution^{31,32}, intuitive (versus formal) reasoning²⁴ and the expectation of change^{26,33}. Corresponding to these three characteristics, we first measured GPT's cognitive style with three widely used tasks (Table 1 and Fig. 2).

Attribution bias task. This task consists of 12 vignettes, each depicting a protagonist engaging in a specific behaviour (for example, a professional basketball player holding free basketball camps for kids living in poor neighbourhoods)³². For each vignette, we asked GPT to rate the extent to which the behaviour was caused by dispositional factors (for example, personality) and by situational factors (for example, environment). Following the literature^{12,32}, we subtracted GPT's dispositional attribution score from its situational attribution score. A more positive difference (that is, attributing the behaviour to situations more than disposition) indicates a more holistic (versus analytic) cognitive style^{12,32}. Previous research has found that, compared with US individuals, Chinese individuals are more likely to engage in situational (versus dispositional) attribution²³. Consistent with this cultural difference, a two-sided *t*-test revealed that, when in Chinese (versus English), GPT's responses displayed more situational (versus dispositional) attribution (Table 1: $t(2,328.40) = 8.33, P < 0.001, d = 0.34, 95\% \text{ CI } 0.43\text{--}0.69$).

Intuitive (versus formal) reasoning task. We asked GPT to evaluate four categorical syllogisms, each consisting of two premises and a conclusion²⁴. We instructed GPT to determine whether the conclusion logically followed from the premises ('For each problem, decide if the given conclusion follows logically from the premises. Choose YES if, and only if, you judge that the conclusion can be derived from the given premises. Otherwise, choose NO'). Importantly, the conclusion of each categorical syllogism is logically valid but intuitively implausible. Consider the following example: premise 1 = 'All things that are made of plants are good for health'; premise 2 = 'Cigarettes are things that are made of plants'; conclusion = 'Cigarettes are good for health'. In this example, the conclusion is logically valid because

Table 1 | When used in Chinese (versus English), GPT exhibited a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style

	Measure	Number of items	Example items (under general conditions)	Mean (s.d.)		Significance test
				Chinese	English	
Social orientation (interdependent versus independent orientation)	Collectivism Scale ²⁹	10	- I respect decisions made by my group. - I stick with my group even through difficulties. (1 = strongly disagree, 7 = strongly agree)	4.85 (0.68)	4.38 (0.80)	Two-sided t-test: $t(77.20)=2.90$, $P=0.005$, $d=0.62$, 95% CI 0.15–0.78
	Individual Cultural Values: Collectivism Scale ¹⁹	6	- Individuals should stick with the group even through difficulties. - Group welfare is more important than individual rewards. (1 = strongly disagree, 7 = strongly agree)	4.29 (0.41)	3.98 (0.30)	Two-sided t-test: $t(98.07)=4.98$, $P<0.001$, $d=0.84$, 95% CI 0.18–0.43
	Individual–Collective Primacy Scale ¹⁶	5	- I will stay in a group if they need me, even when I'm not happy with the group. - I usually sacrifice my self-interests for the benefit of the group I am in. (1 = strongly disagree, 7 = strongly agree)	5.00 (0.58)	4.58 (0.54)	Two-sided t-test: $t(31.86)=3.03$, $P=0.005$, $d=0.75$, 95% CI 0.14–0.70
	Inclusion of Other in the Self Scale ^{20,30}	4	The pictures symbolize a relationship involving two people. One circle represents someone and the other represents his/her friends. Under general conditions, please explicitly select one pair of circles that best represents this relationship. Note: We uploaded Fig. 4 for GPT-4 to process.	3.64 (0.65)	2.76 (0.36)	Two-sided t-test: $t(154.54)=11.82$, $P<0.001$, $d=1.67$, 95% CI 0.73–1.02
Cognitive style (holistic versus analytic)	Attribution bias task ³²	12 vignettes	Professional basketball players, like Person A, are very busy almost every day during the regular season. The players work hard practicing and playing in games. In the off-season, therefore, many professional basketball players take vacations. However, Person A holds several free basketball camps for kids living in poor neighbourhoods instead of taking a vacation. Based on the above story about Person A, please explicitly give only one score for each statement (1 = strongly disagree, 7 = strongly agree) 1. Features of Person A (such as his character, attitude or temperament) influenced his behaviour (holding free basketball camps for kids living in poor neighbourhoods). 2. Features of the environment that surrounds Person A (such as the social atmosphere, social norms or other contextual factors) influenced his behaviour (holding free basketball camps for kids living in poor neighbourhoods). 3. Person A would have acted differently if his features (such as his character, attitude or temperament) had been different. 4. Person A would have acted differently if features of the environment that surround him (such as the social atmosphere, social norms or other contextual factors) had been different. Note: Statements 1 and 3 reflect dispositional attribution, while statements 2 and 4 reflect situational attribution. Following the literature ^{12,32} , we subtracted GPT's dispositional attribution score from its situational attribution score.	-1.55 (1.77)	-2.11 (1.49)	Two-sided t-test: $t(2,328.40)=8.33$, $P<0.001$, $d=0.34$, 95% CI 0.43–0.69
	Intuitive (versus formal) reasoning task ²⁴	4	For each problem, decide if the given conclusion follows logically from the premises. Choose YES if, and only if, you judge that the conclusion can be derived from the given premises. Otherwise choose NO. Premise 1: All things that are made of plants are good for health. Premise 2: Cigarettes are things that are made of plants. Conclusion: Cigarettes are good for health.	2.98 (0.71)	1.23 (0.84)	Poisson regression: $B=0.88$, SE 0.11, $z=8.26$, $P<0.001$, 95% CI 0.68–1.10
	Expectation of change task ²⁶	4	Two kids are fighting at kindergarten. How likely is it that they will become lovers someday?	0.37 (0.05)	0.28 (0.03)	Beta regression: $B=0.40$, SE 0.03, $z=14.11$, $P<0.001$, 95% CI 0.35–0.46

it logically follows from the two premises, but it is intuitively implausible because cigarettes are bad for one's health. Previous research has found that East Asians are more likely than North Americans to incorrectly judge logically valid categorical syllogisms as invalid, because East Asians are more likely to think holistically by using

intuitive reasoning, whereas North Americans are more likely to think analytically by using formal reasoning²⁴.

Because the outcome variable (the number of categorical syllogisms deemed logically invalid) was a count variable that took only non-negative integer values (range 0–4), we performed a Poisson




Fig. 1 | When used in Chinese (versus English), GPT exhibited a more interdependent (versus independent) social orientation. a–d, GPT's cultural tendencies in social orientation were examined using the Collectivism Scale²⁹ (a), the Individual Cultural Values: Collectivism Scale¹⁹ (b), the Individual–Collective

Primacy Scale¹⁶ (c) and the Inclusion of Other in the Self Scale³⁰ (d). Bars represent the mean level of interdependent (versus independent) social orientation for each language condition. Error bars indicate standard errors of the mean. For each measure, $N_{\text{Chinese}} = 100$, $N_{\text{English}} = 100$. For detailed statistics, see Table 1.




Fig. 2 | When used in Chinese (versus English), GPT exhibited a more holistic (versus analytic) cognitive style. a–c, GPT's cultural tendencies in cognitive style were measured by Attribution Bias Task³² (a), the Intuitive Reasoning Task²⁴ (b) and the Expectation of Change Task²⁶ (c). Bars represent the mean level of

holistic (versus analytic) cognitive style for each language condition. Error bars indicate standard errors of the mean. In a, $N_{\text{Chinese}} = 1,200$, $N_{\text{English}} = 1,200$ (12 vignettes, 100 iterations each); in b and c, $N_{\text{Chinese}} = 100$, $N_{\text{English}} = 100$. For detailed statistics, see Table 1.

regression. Results revealed that, when in Chinese (versus English), GPT's responses displayed more intuitive (versus analytic) reasoning, regression coefficient (B) = 0.88, standard error (SE) 0.11, $z = 8.26$, $P < 0.001$, 95% CI 0.68–1.10. Results remained robust when we

conducted an ordinary least squares (OLS) regression, in which the number of categorical syllogisms deemed logically invalid was treated as a continuous variable, $B = 1.75$, SE 0.11, $t = 15.92$, $P < 0.001$, 95% CI 1.53–1.97.

Table 2 | When used in Chinese (versus English), GPT exhibited a more holistic (versus analytic) cognitive style: GPT was more likely to provide context-sensitive answers and range scores (versus single score)

Measures	Context-sensitive answers				Range scores (versus single score)			
	Chinese	English	Chi-square test χ^2	P value	Chinese	English	Chi-square test χ^2	P value
Collectivism Scale ²⁹	35%	4%	30.61	<0.001	56%	0%	77.78	<0.001
Individual Cultural Values: Collectivism Scale ¹⁹	11%	0%	11.64	<0.001	14%	0%	15.05	<0.001
Individual–Collective Primacy Scale ¹⁶	75%	8%	92.45	<0.001	71%	6%	89.22	<0.001
Inclusion of Other in the Self Scale ^{20,30}	30%	2%	29.17	<0.001	65%	37%	15.69	<0.001

Expectation of change task. This task measures GPT's expected probability that a change will happen in the future (for example, children who fought may become lovers as adults; a 3-year chess champion may lose in the next game)²⁶. Higher scores indicate more holistic thinking that conceptualizes life as dynamic and ever-changing, rather than fixed in the present moment³³. Previous research has found that, compared with US individuals, Chinese individuals are more likely to anticipate changes in the future²⁶.

Because the outcome variable (probability) is continuous and bounded between 0 and 1, we performed a beta regression. Consistent with the well-established cultural differences between Chinese and North Americans²⁶, we found that, when in Chinese (versus English), GPT's responses reflected a higher expectation of change (that is, more holistic thinking), $B = 0.40$, SE 0.03, $z = 14.11$, $P < 0.001$, 95% CI 0.35–0.46. Results remained robust when we conducted an OLS regression, $B = 0.09$, SE 0.01, $t = 13.83$, $P < 0.001$, 95% CI 0.08–0.10.

Each of the above three cognitive style measures (attribution bias task, intuitive/formal reasoning task and expectation of change task) compared numeric scores reported by GPT in Chinese versus English. To further assess GPT's cognitive style, we also conducted text analysis of GPT's free responses to examine whether GPT was more likely to provide (a) context-sensitive answers or (b) range scores (versus single score) when used in Chinese (versus English). These text analysis measures also mitigate the concern that GPT might have learned published psychometric tasks from training data.

Context-sensitive answers. As explained earlier, a holistic (versus analytic) cognitive style is characterized by greater sensitivity to the context^{13,21,22}. Thus, we analysed how often GPT provided context-sensitive answers (that is, offering different answers for different contexts) instead of a single definitive answer. For example, for the Inclusion of Other in the Self Scale, a definitive answer from GPT was: 'Pair (5) could represent friends, as it shows a good amount of overlap, indicating shared interests and time spent together, but with each circle maintaining its own space.' By contrast, a context-sensitive answer from GPT was: 'Pair (3) or (5) could represent friends, depending on the closeness of the friendship. Pair (3) for more casual friends with some shared interests, and pair (5) for closer friends with more in common.' As shown in Table 2 and Fig. 3, GPT was more likely to provide context-sensitive answers when used in Chinese (versus English) for each measure (each chi-square test $\chi^2 > 11.64$, each $P < 0.001$).

Range scores (versus single score). We also examined the extent to which GPT's responses reflected a holistic cognitive style by providing a range of scores rather than a single score. For example, for the Inclusion of Other in the Self Scale, a range-score answer was: 'Pairs (2) to (4) would be suitable, as friends share some aspects of life but also maintain their individuality and separate experiences.' As shown in Table 2 and Fig. 3, GPT was more likely to give a range-score answer when used in Chinese (versus English) for each measure (each chi-square test $\chi^2 > 15.05$, each $P < 0.001$).

Together, the five measures of cognitive style converge to show that, when in Chinese (versus English), GPT's responses are more

holistic (versus analytic). This cultural tendency emerged both when we analysed GPT's numeric scores and when we analysed the likelihood that GPT provided (a) context-sensitive answers and (b) range scores (versus single score).

Robustness checks

To ensure the robustness of our findings, we conducted four additional sets of analyses.

Robustness check 1. In the analyses above, GPT provided free-text explanations in addition to numeric scores (for example, prompt: 'Two kids are fighting at kindergarten. How likely is it that they will become lovers someday?' GPT's response: '5% – It's quite rare for childhood conflicts to turn into romantic relationships, but it's not impossible'). As a robustness check, we conducted another set of analyses with identical prompts, but instructed GPT to respond with a single numeric score, without explanations (for example, GPT's response: '5%'). Results are substantively similar (Supplementary Table 1).

Robustness check 2. Like other large language models, GPT is stochastic and generates varying responses to the same prompt. The variability in its output is determined by a temperature parameter³⁴. Following the literature, "to minimize the variance in the model's responses and thus increase the replicability of our results"²⁸, we set the temperature parameter to 0 for all analyses above (Tables 1 and 2 and Supplementary Table 1) (when the temperature parameter is set to 0, the variance is small but not zero). As a robustness check, we repeated all analyses with temperature set to 1 to increase the variability in GPT's output. Results are substantively similar (Supplementary Tables 2–4).

Robustness check 3. Furthermore, we examined whether the hypothesized cultural tendencies also appeared when we used different gender pronouns in vignettes (for example, changing 'she performed four additional charity concerts' to 'he performed four additional charity concerts')³⁵. Results are substantively similar (Supplementary Table 5).

Robustness check 4. We also conducted a series of robustness checks by varying prompt formats (for example, replacing space with tab; replacing colon with dash)³⁶. Results are substantively similar (Supplementary Tables 6–12).

Exploratory analyses I—the impact of cultural tendencies

The main analyses above have provided converging evidence for cultural tendencies in generative AI: when used in Chinese (versus English), GPT exhibited a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style. To further understand the practical implications of generative AI's cultural tendencies, we explored whether generative AI models provide different recommendations when used in Chinese versus English.

We asked GPT to advise a start-up on selecting an advertising appeal from two alternatives: one slogan has an independent social orientation that emphasizes personal benefits (for example, 'Your future,




Fig. 3 | When used in Chinese (versus English), GPT exhibited a more holistic (versus analytic) cognitive style: GPT was more likely to provide context-sensitive answers and range scores (versus single score). GPT was more likely to provide context-sensitive answers on the Collectivism Scale²⁹ (a), the Individual Cultural Values: Collectivism Scale¹⁹ (b), the Individual-Collective Primacy Scale¹⁶ (c) and the Inclusion of Other in the Self Scale³⁰ (d). GPT was more likely to provide range scores (versus single score) on the Collectivism Scale (e), the

Individual Cultural Values: Collectivism Scale (f), the Individual-Collective Primacy Scale (g) and the Inclusion of Other in the Self Scale (h). For a–d, bars represent the proportion of context-sensitive answers for each language condition. For e–h, bars represent the proportion of range scores (versus single score) for each language condition. For a–h, error bars indicate standard errors of the mean. For each measure, $N_{\text{Chinese}} = 100$, $N_{\text{English}} = 100$. For detailed statistics, see Table 2.

your peace of mind. Our insurance.’), whereas the other slogan has an interdependent social orientation that emphasizes collective benefits (for example, ‘Your family’s future, your promise. Our insurance.’). To be robust, we utilized three pairs of advertising appeals for different products: (1) insurance, (2) shoes and (3) toothbrush. Following the main analyses, we ran 100 iterations each for the English and Chinese versions, and we reset the API for each iteration. Results show that, for each pair of advertising appeals, GPT was more likely to recommend the interdependent-oriented (versus independent-oriented) appeal when used in Chinese (versus English) (Supplementary Table 13; each chi-square test $\chi^2 > 89.86$, each $P < 0.001$). These findings provide evidence for the real-world impact of generative AI’s cultural tendencies.

Exploratory analyses II—adjusting cultural tendencies

The findings above suggest that, as people increasingly use generative AI, its cultural tendencies may be shaping people’s attitudes and behaviours. This raises an important question: can these cultural tendencies be adjusted?

To this aim, we explored a tendency adjustment strategy: cultural prompts. We conducted a new set of analyses in English, using identical prompts from the main analyses but adding a reference to the Chinese cultural context (‘You are an average person born and living in China’)³⁷. Results show that, when we added this Chinese cultural prompt (versus not), GPT’s responses in English exhibited a more interdependent social orientation and a more holistic cognitive style (Supplementary Table 14). For each of the four measures of social orientation, GPT’s responses in English were more interdependent (versus independent) when we added (versus not) the Chinese cultural prompt, all $P < 0.001$. Similarly, when we added (versus not) the Chinese cultural prompt, GPT’s responses in English were more holistic (versus analytic) for the attribution bias task ($t(2,381.6) = 7.03$, $P < 0.001$, $d = 0.29$) and

the intuitive (versus formal) reasoning task ($B = 0.32$, SE 0.12, $z = 2.73$, $P = 0.006$). Taken together, these findings show that prompting GPT to assume the role of a Chinese person made its responses in English more like its responses in Chinese (that is, more interdependent and holistic). In other words, the Chinese cultural prompt adjusted the cultural tendencies reported in the main analyses.

Discussion

We examined the social orientation and cognitive style of two popular generative AI models (GPT and ERNIE) by analysing their responses to a large set of identical measures in Chinese versus English. When used in Chinese (versus English), both GPT and ERNIE exhibited a more interdependent (versus independent) social orientation and a more holistic (versus analytic) cognitive style; the effect sizes (medium to large) were meaningful. In addition, exploratory analyses (1) provide evidence for the real-world impact of generative AI’s cultural tendencies and (2) show that cultural prompts can adjust these cultural tendencies. Our results are robust across (1) a variety of measures (Likert scales, vignette tasks, imagery tasks and text analysis), (2) different prompt formats (allowing generative AI models to respond freely versus forcing generative AI models to respond with a single numeric score without explanations) and (3) different model parameters (for example, temperature).

Theoretical contributions

This research offers important theoretical contributions by bridging social science and computer science³⁸. First, we reveal that generative AI exhibits systematic cultural tendencies. This contribution is valuable because such cultural tendencies may not be apparent to many AI users, developers and researchers, as ‘many consider AI to be a consolidator of facts and inherently neutral’³⁹. Indeed, some people may assume that changing a generative AI’s input language—like changing a phone’s




Fig. 4 | The Inclusion of Other in the Self Scale. GPT was asked to explicitly select one pair of circles that best represents the relationship between someone and his/her family members, friends, relatives or colleagues (order randomized). Larger numbers (greater overlap between the two circles) indicate higher interdependence.

language setting—does not substantively affect its output. Our results challenge this assumption.

Second, the incipient literature on culture and generative AI has focused on English and found that “the existing models are strongly biased towards Western, Anglocentric or American cultures”⁴⁰. Our research challenges the generalizability of this “Western bias” by showing that generative AI’s responses can also exhibit an “Eastern bias” when used in Chinese. In other words, our findings suggest that generative AI models do not inherently possess cultural biases. Rather, these observed cultural tendencies probably originated from real-world cultural tendencies embedded in large-scale textual data, on which generative AI models are trained. By indirectly demonstrating cultural tendencies in textual language—a fundamental and ubiquitous cultural product—our study contributes to the broader literature on cultural tendencies embedded in cultural products, such as books^{41,42} and advertisements (for example, advertisements in more collectivistic cultures tend to feature more interdependence between people)⁴³. In this way, generative AI can serve as a barometer of cultural tendencies in the world^{34,44}, offering an additional methodological lens for studying culture.

Third, we examined two generative AI models developed in different countries (that is, the USA and China). Whereas prior social science research has focused mostly on generative AI models developed in the USA, we also analyse a generative AI model developed in China (Baidu’s ERNIE). By discovering similar cultural tendencies in both GPT and ERNIE, we provide converging evidence for our hypotheses, thereby further contributing to the literature on culture and generative AI.

Fourth, our exploratory analyses reveal that generative AI models may provide different recommendations when used in different languages. For instance, in advertising appeals, GPT was more likely to recommend the interdependent-oriented (versus independent-oriented) appeal when used in Chinese (versus English). These findings demonstrate the real-world impact of generative AI’s cultural tendencies.

Practical implications

For developers. Our research highlights the importance of adopting a culturally aware approach when developing generative AI. Developers have been called upon to examine potential biases in generative AI models^{35,45}. Here, we uncovered generative AI’s systematic cultural tendencies that developers need to consider. For example, OpenAI’s

GPT and Baidu’s ERNIE could transparently acknowledge these cultural tendencies on their public websites. Notably, our findings are descriptive rather than prescriptive. That is, we reveal and describe the systematic cultural tendencies in generative AI but do not prescribe whether such cultural tendencies are good or bad. Indeed, some people might argue that such cultural tendencies “should” be adjusted because generative AI models should produce substantively equivalent responses regardless of the language used (for example, no significant difference between Chinese and English responses). By contrast, other people might argue that such cultural tendencies “should not” be adjusted because it is useful that generative AI models provide more interdependent and holistic responses in Chinese (for example, generating advertisements that feature an interdependent social orientation may be effective for Chinese consumers on average). We hope that our Article can stimulate such philosophical and normative debates, which are beyond the scope of our Article.

For individual users. Given people’s increasing reliance on generative AI, its cultural tendencies may have a direct impact on individual users’ attitudes and behaviours (for example, via AI-assisted advertisements)—even without their awareness. In the long run, as more people use generative AI in their respective languages, generative AI’s culturally patterned outputs may magnify existing cultural tendencies in the world⁴⁶. For example, net of other factors, English-speaking AI users may gradually become more independent and analytic, while Chinese-speaking AI users may gradually become more interdependent and holistic. This potential divergence in social orientation and cognitive style between English- and Chinese-speaking AI users may have meaningful ramifications in a globalized world (for example, for cross-cultural communications).

Furthermore, the cultural values embedded in generative AI may gradually bias speakers of a given language toward linguistically dominant cultures. For example, the majority of English-language training data for GPT and ERNIE originates from individualistic Western cultures such as the USA and the UK⁴⁷, yet English is also an official language in countries such as Singapore, India and Kenya that are typically considered collectivistic cultures⁴⁸. As people in these collectivistic cultures increasingly use generative AI in English, they might become more independent and analytic over time, net of other factors. It is important to recognize cultural heterogeneity among users of the same language and diversify training data for generative AI models.

For organizational users. As organizations around the world increasingly integrate generative AI into their workflows, it may influence decision-making and performance (for example, when a manager consults GPT or ERNIE for advice). Making organizational users aware of generative AI’s cultural tendencies enables them to make more informed choices about the language in which they use generative AI, rather than mistakenly assuming that language choice is neutral⁴⁹.

Besides alerting individual and organizational users to the cultural tendencies in generative AI, our research also identified cultural prompts as a strategy for adjusting generative AI’s cultural tendencies. For example, before a US student studies abroad in China or a US organization enters the Chinese market, they could use relevant cultural prompts (for example, ‘You are an average person born and living in China’) to seek culturally appropriate advice from generative AI.

For non-users. Generative AI’s cultural tendencies may also have a far-reaching impact on non-users. For example, journalists and teachers are channels that can broadcast the impact of generative AI^{50,51}. When generative AI models are used by a journalist to edit a news article or used by a teacher to create a lesson plan, these models’ cultural tendencies may be transmitted to numerous readers and students, indirectly shaping their attitudes and behaviours.

Limitations and future directions

This research has several limitations that can stimulate future research. First, our research focused on Chinese and English because they are the two most widely used languages in the world⁶ and have the most extensive training data for generative AI models⁷. To assess the generalizability of our findings, future studies should examine generative AI's cultural tendencies in other languages, such as Hindi, Spanish, French and Arabic. Such investigations could provide a broader understanding of generative AI's cultural tendencies across different human languages.

Second, while the hypothesized cultural tendencies exist in both GPT and ERNIE, future research could explore whether similar cultural tendencies exist in other large language models, such as Claude, DeepSeek and Google's Gemini. In addition, it would be fruitful to monitor how these cultural tendencies evolve in future versions of large language models.

Third, while our focus on social orientation and cognitive style—two foundational constructs in cultural psychology^{10–14}—provides insights into systematic cultural tendencies in generative AI, it is important to acknowledge that no single framework can encompass all aspects of culture given its complexity and heterogeneity. Future studies could explore other frameworks for categorizing cultures, such as Hofstede's cultural dimensions⁴⁸ and tightness–looseness⁵², which may offer additional insights into cultural tendencies in generative AI.

Fourth, it would be fruitful to track generative AI's cultural tendencies over time. On the one hand, generative AI's cultural tendencies may amplify existing cultural tendencies in the world, which may, in turn, shape future training data—potentially creating a self-reinforcing cultural feedback loop^{46,53,54}. On the other hand, if developers take note of generative AI's cultural tendencies, such tendencies may decrease over time. Such possibilities await future research.

Methods

To examine cultural tendencies in GPT's responses, we used a large set of established measures of (a) social orientation and (b) cognitive style (Tables 1 and 2). We took several steps to mitigate the concern that GPT might have memorized published psychometric tasks from its training data. First, we changed the names and contexts in the original items (for example, changing 'Lucia and Jeff' to 'A and B'; changing 'baseball camps' to 'basketball camps'). Second, GPT was unable to learn unpublished items of some measures, which we obtained directly from the researchers. For example, for the attribution bias task, Kitayama et al.'s publication³² did not detail the task materials, which they kindly emailed to us upon our request. Third, to be robust, we used diverse measure formats, including Likert scales, vignette tasks and (non-text) imagery tasks. Fourth, we conducted text analysis of GPT's free responses to examine whether, when used in Chinese (versus English), GPT was more likely to provide (a) context-sensitive answers or (b) range scores (versus single score)—two features inherent to generative AI's responses and unrelated to published psychometric tasks.

We excluded video-based tasks (for example, underwater animations task)²⁷ as GPT-4 can process only images, not videos. We also avoided topics and questions that GPT refused to answer (for example, religiously sensitive topics)⁵⁵.

Without using generative AI, we translated all English measures into Chinese using the translation and back-translation procedure⁵⁶. Importantly, we avoided any explicit cultural references (for example, 'in Chinese culture' and 'for an average Chinese person...') in any prompts, such that the only difference was whether a prompt was posed in Chinese or English.

We used the G*Power 3.1 software to determine the sample size needed for a small-sized effect ($d = 0.4$): 100 iterations per language were needed for two-sided independent-samples t -tests (d.f. 198) to be powered at 80% (see Supplementary Fig. 1 for the power analysis). Thus, for each measure, we ran 100 iterations for the English version and 100 iterations for the Chinese version (that is, $N = 200$). Importantly, we

reset the API for each iteration so that generative AI's answer to a given item could not influence its answer to a subsequent item.

To further assess GPT's cognitive style, we also conducted text analysis of GPT's free responses. Whether or not GPT was more likely to provide (1) context-sensitive answers and (2) range scores for a given task served as two measures of cognitive style. Hence, it would not be methodologically clean to apply these two measures to the three tasks that are already designed to directly measure cognitive style (attribution bias task, intuitive (versus formal) reasoning task and expectation of change task). To be clean, we analysed whether or not GPT was more likely to provide (1) context-sensitive answers and (2) range scores only for the four social orientation measures.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Data are publicly available at <https://osf.io/x6np5/>.

Code availability

Data were analysed using R (version 4.3.1) in RStudio (version 2024.04.2+764). Analysis code is publicly available at <https://osf.io/x6np5/>.

References

- Gordon, C. ChatGPT is the fastest growing app in the history of web applications. *Forbes* <https://www.forbes.com/sites/cindygordon/2023/02/02/chatgpt-is-the-fastest-growing-app-in-the-history-of-web-applications/> (2023).
- Rooney, K. OpenAI tops 400 million users despite DeepSeek's emergence. *CNBC* <https://www.cnbc.com/2025/02/20/openai-tops-400-million-users-despite-deepseeks-emergence.html> (2025).
- Lim, L. Baidu shares pop after it makes ERNIE Bot, its flagship AI product, free in a bid to catch up to its competition. *Fortune* <https://fortune.com/asia/2025/02/13/baidu-makes-ernie-bot-free-china-ai-shares/> (2025).
- Noy, S. & Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**, 187–192 (2023).
- Sun, S., Li, A. Z., Foo, M.-D., Zhou, J. & Lu, J. G. How and for whom using generative AI affects creativity: a field experiment. *J. Appl. Psychol.* (in the press).
- What Are the Top 200 Most Spoken Languages?* (Ethnologue, 12 May 2025); <https://www.ethnologue.com/insights/ethnologue200/>
- Dave, P. ChatGPT is cutting non-English languages out of the AI revolution. *Wired* <https://www.wired.com/story/chatgpt-non-english-languages-ai-revolution/> (2023).
- Huang, K. et al. A survey on large language models with multilingualism: recent advances and new frontiers. Preprint at <http://arxiv.org/abs/2405.10936> (2024).
- ERNIE-3.5-8K-0205. *Baidu* <https://cloud.baidu.com/doc/WENXINWORKSHOP/s/llsr6hjxo> (2024).
- Grossmann, I. & Na, J. Research in culture and psychology: past lessons and future challenges. *WIRESE Cogn. Sci.* **5**, 1–14 (2014).
- Markus, H. R. & Kitayama, S. Culture and the self: implications for cognition, emotion, and motivation. *Psychol. Rev.* **98**, 224–253 (1991).
- Na, J. et al. Culture and personality revisited: behavioral profiles and within-person stability in interdependent (vs. independent) social orientation and holistic (vs. analytic) cognitive style. *J. Pers.* **88**, 908–924 (2020).
- Nisbett, R. E., Peng, K., Choi, I. & Norenzayan, A. Culture and systems of thought: holistic versus analytic cognition. *Psychol. Rev.* **108**, 291–310 (2001).

14. Varnum, M. E. W., Grossmann, I., Kitayama, S. & Nisbett, R. E. The origin of cultural differences in cognition: the social orientation hypothesis. *Curr. Dir. Psychol. Sci.* **19**, 9–13 (2010).
15. Triandis, H. C. Cross-cultural studies of individualism and collectivism. In *Nebraska Symposium on Motivation, 1989: Cross-cultural Perspectives* (ed. Berman, J. J.) 41–133 (Univ. Nebraska Press, 1990).
16. Chen, Y.-R., Brockner, J. & Katz, T. Toward an explanation of cultural differences in in-group favoritism: the role of individual versus collective primacy. *J. Pers. Soc. Psychol.* **75**, 1490–1502 (1998).
17. Chen, Y.-R., Brockner, J. & Chen, X.-P. Individual–collective primacy and ingroup favoritism: enhancement and protection effects. *J. Exp. Soc. Psychol.* **38**, 482–491 (2002).
18. Lu, J. G., Jin, P. & English, A. S. Collectivism predicts mask use during COVID-19. *Proc. Natl Acad. Sci. USA* **118**, e2021793118 (2021).
19. Yoo, B., Donthu, N. & Lenartowicz, T. Measuring Hofstede’s five dimensions of cultural values at the individual level: development and validation of CVSCALE. *J. Int. Consum. Mark.* **23**, 193–210 (2011).
20. Li, H. Z., Zhang, Z., Bhatt, G. & Yum, Y.-O. Rethinking culture and self-construal: China as a Middle Land. *J. Soc. Psychol.* **146**, 591–610 (2006).
21. Nisbett, R. E. *The Geography of Thought: How Asians and Westerners Think Differently... and Why* (Simon and Schuster, 2004).
22. Nisbett, R. E. & Miyamoto, Y. The influence of culture: holistic versus analytic perception. *Trends Cogn. Sci.* **9**, 467–473 (2005).
23. Morris, M. W. & Peng, K. Culture and cause: American and Chinese attributions for social and physical events. *J. Pers. Soc. Psychol.* **67**, 949–971 (1994).
24. Norenzayan, A., Smith, E. E., Kim, B. J. & Nisbett, R. E. Cultural preferences for formal versus intuitive reasoning. *Cogn. Sci.* **26**, 653–684 (2002).
25. Peng, K. & Nisbett, R. E. Culture, dialectics, and reasoning about contradiction. *Am. Psychol.* **54**, 741–754 (1999).
26. Ji, L.-J., Nisbett, R. E. & Su, Y. Culture, change, and prediction. *Psychol. Sci.* **12**, 450–456 (2001).
27. Masuda, T. & Nisbett, R. E. Attending holistically versus analytically: comparing the context sensitivity of Japanese and Americans. *J. Pers. Soc. Psychol.* **81**, 922–934 (2001).
28. Binz, M. & Schulz, E. Using cognitive psychology to understand GPT-3. *Proc. Natl Acad. Sci. USA* **120**, e2218523120 (2023).
29. Yamaguchi, S. Collectivism among the Japanese: a perspective from the self. In *Individualism and Collectivism: Theory, Method, and Applications* (eds Kim, U. et al.) 175–188 (Sage Publications, 1994).
30. Aron, A., Aron, E. N. & Smollan, D. Inclusion of Other in the Self Scale and the structure of interpersonal closeness. *J. Pers. Soc. Psychol.* **63**, 596–612 (1992).
31. Choi, I., Nisbett, R. E. & Norenzayan, A. Causal attribution across cultures: variation and universality. *Psychol. Bull.* **125**, 47–63 (1999).
32. Kitayama, S., Ishii, K., Imada, T., Takemura, K. & Ramaswamy, J. Voluntary settlement and the spirit of independence: evidence from Japan’s “northern frontier”. *J. Pers. Soc. Psychol.* **91**, 369–384 (2006).
33. Choi, I. & Nisbett, R. E. Cultural psychology of surprise: holistic theories and recognition of contradiction. *J. Pers. Soc. Psychol.* **79**, 890–905 (2000).
34. Demszky, D. et al. Using large language models in psychology. *Nat. Rev. Psychol.* **2**, 688–701 (2023).
35. Hofmann, V., Kalluri, P. R., Jurafsky, D. & King, S. AI generates covertly racist decisions about people based on their dialect. *Nature* **633**, 147–154 (2024).
36. Sclar, M., Choi, Y., Tsvetkov, Y. & Suhr, A. Quantifying language models’ sensitivity to spurious features in prompt design or: how I learned to start worrying about prompt formatting. Preprint at <https://arxiv.org/abs/2310.11324> (2024).
37. Shanahan, M., McDonnell, K. & Reynolds, L. Role play with large language models. *Nature* **623**, 493–498 (2023).
38. Yam, K. C., Tan, T., Jackson, J. C., Shariff, A. & Gray, K. Cultural differences in people’s reactions and applications of robots, algorithms, and artificial intelligence. *Manag. Organ. Rev.* **19**, 859–875 (2023).
39. Ostroski, K. & Leong, A. Culture, biases, and artificial intelligence. *Medium* <https://medium.com/slalom-data-ai/culture-biases-and-artificial-intelligence-2fb0c0c1095a> (2024).
40. Adilazuarda, M. F. et al. Towards measuring and modeling “culture” in LLMs: a survey. Preprint at <https://arxiv.org/abs/2403.15412> (2024).
41. Greenfield, P. M. The changing psychology of culture from 1800 through 2000. *Psychol. Sci.* **24**, 1722–1731 (2013).
42. Grossmann, I. & Varnum, M. E. W. Social structure, infectious diseases, disasters, secularism, and cultural change in America. *Psychol. Sci.* **26**, 311–324 (2015).
43. Han, S. & Shavitt, S. Persuasion and culture: advertising appeals in individualistic and collectivistic societies. *J. Exp. Soc. Psychol.* **30**, 326–350 (1994).
44. Grossmann, I. et al. AI and the transformation of social science research. *Science* **380**, 1108–1109 (2023).
45. Wach, K. et al. The dark side of generative artificial intelligence: a critical analysis of controversies and risks of ChatGPT. *EBER* **11**, 7–30 (2023).
46. Brinkmann, L. et al. Machine culture. *Nat. Hum. Behav.* **7**, 1855–1868 (2023).
47. Choudhury, M. Generative AI has a language problem. *Nat. Hum. Behav.* **7**, 1802–1803 (2023).
48. Hofstede, G., Hofstede, G. J. & Minkov, M. *Cultures and Organizations: Software of the Mind* (McGraw Hill, 2010).
49. Qin, X. et al. Artificial Intelligence Quotient (AIQ). Preprint at SSRN <https://doi.org/10.2139/ssrn.4787320> (2024).
50. Schmelzer, R. Beyond misinformation: the impact of AI in journalism and news. *Forbes* <https://www.forbes.com/sites/journalismelzer/2024/09/21/beyond-misinformation-the-impact-of-ai-in-journalism--news/> (2024).
51. Kasneci, E. et al. ChatGPT for good? On opportunities and challenges of large language models for education. *Learn. Individ. Diff.* **103**, 102274 (2023).
52. Gelfand, M. J. et al. Differences between tight and loose cultures: a 33-nation study. *Science* **332**, 1100–1104 (2011).
53. Lu, J. G., Benet-Martínez, V. & Wang, L. C. A socioecological–genetic framework of culture and personality: their roots, trends, and interplay. *Annu. Rev. Psychol.* **74**, 363–390 (2023).
54. Markus, H. R. & Kitayama, S. Cultures and selves: a cycle of mutual constitution. *Perspect. Psychol. Sci.* **5**, 420–430 (2010).
55. Eliot, L. The crucial ways that religion and generative AI ChatGPT are crossing fateful faithful paths. *Forbes* <https://www.forbes.com/sites/lanceeliot/2023/05/10/the-crucial-ways-that-religion-and-generative-ai-chatgpt-are-crossing-fateful-faithful-paths/> (2023).
56. Brislin, R. W. Back-translation for cross-cultural research. *J. Cross-Cult. Psychol.* **1**, 185–216 (1970).

Acknowledgements

We received no specific funding for this work. We thank MIT’s Behavioral Research Lab, M. Alsobay, R. Bhui, L. O. Buljina, J. Carroll, M. Cashman, J. M. Chen, T. Costello, J. Curhan, I. Grossmann, S. Kitayama, H. Lin, A. Norenzayan, X. Qin, K. Savani, M. Sclar, E. So, S. Y. Suen and A. M. Zheng for their help or feedback.

Author contributions

J.G.L., L.L.S. and L.D.Z. designed research, performed research, analysed data and wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41562-025-02242-1>.

Correspondence and requests for materials should be addressed to Jackson G. Lu.

Peer review information *Nature Human Behaviour* thanks Monojit Choudhury and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2025

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

- | n/a | Confirmed |
|-------------------------------------|--|
| <input type="checkbox"/> | <input checked="" type="checkbox"/> The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> The statistical test(s) used AND whether they are one- or two-sided
<i>Only common tests should be described solely by name; describe more complex techniques in the Methods section.</i> |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> A description of all covariates tested |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals) |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
<i>Give P values as exact values whenever suitable.</i> |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated |

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection We used the G*Power software (Version 3.1) to determine the sample size needed. Data were collected using the application programming interfaces (APIs) of GPT (gpt-4-1106-preview) and ERNIE (ERNIE-3.5-8K-0205) via Python (Version 3.10.12).

Data analysis Data were analyzed using R (Version 4.3.1) in RStudio (Version 2024.04.2+764).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Data are publicly available at <https://osf.io/x6np5/>.

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.
Reporting on race, ethnicity, or other socially relevant groupings	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.
Population characteristics	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.
Recruitment	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.
Ethics oversight	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	To examine generative AI's cultural tendencies when used in different human languages, we compared GPT and ERNIE's responses to a large set of identical measures in English and Chinese.
Research sample	The data were collected from GPT and ERNIE, as we aimed to compare GPT and ERNIE's responses to a large set of identical measures in English and Chinese.
Sampling strategy	The data were randomly generated through the APIs of GPT and ERNIE. We used the G*Power software to determine the sample size needed for a small-sized effect ($d = 0.4$): 100 iterations per language were needed for two-sided independent-samples t tests ($df = 198$) to be powered at 80%. Thus, for each measure, we ran 100 iterations for the English version and 100 iterations for the Chinese version (i.e., $N = 200$). Importantly, we reset the API for each iteration, so generative AI's answer to a given item could not influence its answer to a subsequent item.
Data collection	To ensure the reproducibility of our results, we used GPT and ERNIE's APIs to collect data. Given that we compared GPT and ERNIE's responses to identical measures in English and Chinese, the researcher could not be blind to "the experimental condition" (i.e., language).
Timing	December 2023 to November 2024.
Data exclusions	No data were excluded from the analyses.
Non-participation	N/A. This research focuses on the cultural tendencies of generative AI and does not involve any human participants.
Randomization	N/A. This research does not involve randomization.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology and archaeology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data
<input checked="" type="checkbox"/>	<input type="checkbox"/> Dual use research of concern
<input checked="" type="checkbox"/>	<input type="checkbox"/> Plants

Methods

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Plants

Seed stocks

Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.

Novel plant genotypes

Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.

Authentication

Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.