

If AI tells better jokes, have we lost last bastion of human language?

INTERDISCIPLINARY STUDIES

By LU JIA

Humor has long been regarded as a domain of human language resistant to algorithmic formalization. This view stems not only from humor's heavy reliance on context, but also from its entanglement with the recognition of intent, cultural resonance, and emotional attunement—dimensions that cannot easily be formalized.

A study conducted at Peking University found that OpenAI's GPT-4o model outperformed humans in text-based humor generation tasks, scoring particularly well in its use of humor as a coping strategy in negative situations. This counterintuitive finding compels us to ask: If humor proves empirically "computable," where does the distinctiveness of human language reside? Has humor been conquered by AI precisely because it is computable, or does its uncomputability remain the hallmark of an inalienable form of human intelligence?

Against this backdrop, we must carefully delineate the true boundaries of AI's capacity for humor and defend, from the perspective of language philosophy, the unique value of human linguistic expression.

The advent of the 'computable humor' era

Recent research is challenging conventional understandings of humor. From textual generation to visual creation, and from cultural context to character portrayal, AI has demonstrated unexpected capabilities across a wide range of humor-related tasks.

In March 2025, a "meme Turing test" conducted by researchers from LMU Munich and Swedish institutions found that AI-generated memes received higher average scores for funniness, creativity, and shareability than either entirely human-created memes or those produced through human-AI collaboration. A study published in November 2025 by a team at Mohamed bin Zayed University of Artificial Intelligence in Abu Dhabi addressed the problem of stylistic incongruity through "cross-character embedding": When cartoon and live-action characters appeared in the same frame, both retained their distinct visual styles while interacting naturally.

Although advances in text generation were perhaps foreseeable, progress in the visual domain has been more striking. Following the release of Sora 2 in October 2025, AI-generated content moved beyond the simplistic reproduction of stereotypes toward a deeper grasp of specific cultural logics. Under human

direction, AI's interpretation of cultural contexts has reached a remarkable level of precision.

These are not isolated advances. Academic research on humor has shifted from theoretical construction toward computational modeling, while frameworks for humor recognition and pun detection informed by cognitive linguistics have become increasingly mature. By learning from vast corpora, large language models can capture the statistical regularities and structural features of humor and, in certain tasks, outperform humans.

The philosophical foundations of the 'uncomputable'

Research by Hou Yubo (associate professor of psychology from Peking University) and colleagues indicates that AI has statistically outperformed humans in humor generation. Yet a deeper question follows: What does this advantage actually signify? Is the difference between AI-generated humor and human humorous expression merely one of degree, or is it a difference in kind? Answering this question requires us to approach the issue from the perspective of language philosophy.

A useful starting point is the distinction between two conceptions of language: rule-based systems and statistical prediction. American linguist Noam Chomsky distinguished between linguistic competence and linguistic performance, arguing that the core of the human capacity for language is a rule-based system capable of generating an infinite number of expressions through finite means. Large language models, by contrast, operate primarily through statistical prediction based on correlations among units of text, or tokens. By learning the probability that words will occur together across vast corpora, a model calculates the most likely next word in a given context. As technical researchers put it, large language models are systems that "guess the next word from prior text." Their "understanding" remains confined to surface-level statistical patterns and cannot reach underlying structures. This distinction is crucial to understanding AI's capacity for humor. Humor is regarded as a "privileged domain" of human intelligence precisely because it demands more than mastery of formal linguistic rules: It also requires an understanding of implied meaning, a recognition of irony, and access to accumulated cultural knowledge, none of which can be reduced entirely to statistical regularities. AI's humor is "learned," not "understood."

Research into what academics have termed "absurdology"—namely, "profound nonsense"—offers a useful test of the depth of this understanding. Experiments suggest



Humor is regarded as a "privileged domain" of human intelligence precisely because it requires an understanding of implied meaning, a recognition of irony, and access to accumulated cultural knowledge. Image generated by AI

that large language models still struggle with such texts: They often mistake them for shallow, meaningless discourse, produce incoherent explanations, or entirely overlook their implicit rhetorical functions. If AI cannot grasp these multilayered semantic structures, it may be able to identify the statistical patterns of "what is funny" without comprehending the existential implications of "why it is funny."

Research by Yao Congjun, a professor of logic from Xiangtan University, using combinatory categorial grammar (CCG) lends formal support to this distinction through the concepts of "syntactic transparency" and "semantic opacity." CCG demonstrates that syntactic analysis can be highly formalized, with syntactic structures derived by combining lexical categories according to established rules. Yet the full grasp of meaning—particularly meaning that depends on context, intention, and cultural background—can never be entirely reduced to formal computation. Syntax may be transparent, but meaning retains an irreducible opacity. This insight has important implications for our understanding of AI's linguistic capacities. AI can efficiently simulate human language at the syntactic level, but its "understanding" remains, in essence, an externalized reflection of human understanding.

From human-machine competition to human-AI symbiosis

AI may have statistically outperformed humans in humor generation, but its "understanding" remains confined to surface-level pattern matching. This prompts a further question: Once we acknowledge AI's capacity for humor, what remains uniquely valuable about human language?

Humor can be understood on three levels. The textual level encompasses formal devices such as puns and reversals. The contextual level requires sensitivity to what is appropriate in a given situation; AI is advancing rapidly at this level and is approaching human performance in tasks such as conflict mediation and emotional support. The existential level, however, concerns humor as a distinctly human way of confronting the predicaments of life. Research by teams working in this field has found that self-enhancing humor is most effective in negative situations—a capacity rooted in the resilience of human beings amid adversity, grounded in the finite nature of human life, and therefore impossible to simulate algorithmically.

Moreover, as AI becomes more proficient at generating humor, it absorbs from human corpora not only "good humor," but also discriminatory and harmful jokes. As Zeng Yi, a professor of AI from Renmin University of China, has observed, AI is a mirror of human society, laying bare our own biases.

The relationship between humans and AI in the domain of humor should not be framed as a "competition of competence" over which can tell the better joke. The value of humor has never resided in "being funny" alone, but in the human concerns it carries. The distinctiveness of human language does not derive from "doing better than AI," but from "existing in a different mode." Human language acquires its existential significance as a response by finite beings to the circumstances in which they find themselves. We speak because we exist; we laugh because existence is suffused with absurdity, suffering, and uncertainty.

This is the foundation of humor's uncomputability. It lies neither in textual structures nor in contextual parameters, but in the finitude of human life, the embodied nature of the human subject, and the historicity of human communities. The very computability of technology thus throws the uncomputability of life itself into sharper relief.

The future trajectory of human-machine relations should therefore shift from the competitive logic of "who is superior" toward the symbiotic logic of "how shall we coexist." AI's ability to generate humor with remarkable efficiency should prompt deeper human reflection: When AI imitates our humor, it helps us see more clearly what humor means to us. This is the genuine meaning of "AI excels, humanity excels further"—not a contest between humans and machines, but a dialectical relationship in which technology, by displaying its capabilities, compels humanity to turn inward and reaffirm its proper place. Defending the value of humor ultimately means defending the value of humanity itself.

"I think, therefore I am." We might add: "I speak through humor, therefore I am." In laughter, human beings do not merely contemplate the world; they transcend it. They do not merely endure fate; they reshape it. AI can generate jokes, but it cannot place its own life within them. It can simulate humor, but it cannot "exist" through humor. This is the inalienable foundation of human language—and the reason we may still smile, even in the age of artificial intelligence.

Lu Jia is an associate professor from the School of Languages and Literature at the University of South China.