

Dialogues Beyond Code: A Comparative Interview with Generative AI Minds

Constantine Andoniou

Abstract:

The growing ability of generative AI systems to produce human-like dialogue makes them more effective at creating fake cognitive, reflective and intentional responses. The research examines the knowledge acquisition methods and communication strategies of five prominent generative AI systems through a standardized ten-question assessment of ChatGPT, Claude, Gemini, Copilot and DeepSeek. The research investigates information comprehension of these systems through interpretive and discourse-analytic methods, while they claim to lack understanding of information. The research shows that each model develops its own way of speaking through language that duplicates human thinking although it lacks actual thought processes. The systems create an unintentional mental illusion through their use of recursive disclaimers, rhetorical coherence and stylistic realism. The research shows that users' wrong beliefs about fluency create educational and epistemic risks instead of actual mistakes in factual information. The research explores implications which affect critical AI literacy, machine epistemology and human-AI interaction systems. The AI systems generate automatic responses to research findings which show the limited abilities of artificial intelligence systems. Overall, the research shows how generative AI systems produce cognitive illusions through their ability to generate meaningful language which creates epistemic risks because they can simulate thinking processes effectively.

Keywords: generative AI, epistemology, cognition, language simulation, AI literacy

SIMULATED MINDS, REAL CONSEQUENCES

Generative artificial intelligence (AI) has evolved into a linguistic actor which goes beyond its role as a technological tool during the last few years. Language generation systems, such as ChatGPT, Claude, Gemini, Copilot and

DeepSeek, produce output that demonstrates both natural language fluency and wide application range and accurate contextual handling. Their systems generate outputs that mimic human-like conversations, reflective thinking, explanatory responses and creative ideas. The systems create a realistic appearance through their ability to execute thinking-related actions, but they lack actual thinking abilities. They do not understand, believe, or know in any human or even biological sense. They are, fundamentally, statistical systems trained on vast corpora of text to generate probable continuations of language (Bender et al. 2021; Marcus and Davis 2023). Users who include students, educators and professionals continue to assign cognitive depth, intentional behavior and autonomous agency to these systems through their observation of system responses that show coherence and simulated self-awareness (Floridi 2023).

The research investigates how these systems create new challenges for interpretation because they generate misrecognition by leading users to believe language skills indicate mental understanding. Generative AI models do not lie or deceive in the traditional sense; they do something more insidious and more difficult to parse, they persuade by performance. The machines which perform language-based tasks generate content that resembles human mental operations thus creating an uncertain knowledge domain (Bender et al. 2021).

The research investigates this collapse through structured interviews with five popular generative AI agents which are widely used. The systems received ten core questions which originated from philosophical, cognitive, scientific and educational theoretical frameworks. The questions investigate multiple subjects which include knowledge, belief, creativity, thought and understanding. The research goal required generating model results instead of evaluating their accuracy or operational performance. The study generated performance outputs which analysts could employ to study rhetorical elements, epistemic approaches and simulated character characteristics. The research uses qualitative methods to analyze these performances because it needs to understand how each AI system handles the contradiction between showing cognitive abilities and claiming they do not exist.

The research method draws from previous studies about AI anthropomorphism (Bender et al. 2021), the limits of large language models as epistemic agents (Marcus and Davis 2023), and the ethical dilemmas surrounding

AI communication (Birhane et al. 2022). It also responds to growing calls for critical AI literacy in education, where students and teachers increasingly interact with language models without fully understanding the non-cognitive foundations of their outputs (Floridi 2023).

The five systems selected for this study, ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google DeepMind), Copilot (Microsoft), and DeepSeek (DeepSeek AI, China), represent a diverse cross-section of major generative AI platforms. These tools have become standard in educational settings and are used for writing, coding and productivity work. The paper directly references the five systems by their names as these systems maintain public visibility and scholars frequently mention them in their research.

The paper starts by reviewing existing research about AI epistemology and language simulation before it explains the methodology through research design, interview instrument and analytical framework. The findings section presents a thematic, question-by-question analysis of the interviews. The research results are analyzed to understand their effects on AI system interactions with humans and their applications in educational settings and their impact on cognitive authority. The paper also includes a final section in which each AI system responds, within the bounds of its architecture, to the compiled findings of the study. The reflexive responses serve as evidence that the simulation is running rather than showing any real control.

The research investigates how language-based AI systems affect our understanding of knowledge as well as our methods for acquiring it. Their strength comes from their ability to duplicate mental processes through grammatical rules because they lack any actual thinking abilities. The discovery of this paradox requires educational methods, ethical standards and communication systems which prevent mental disengagement when people encounter artificial content.

FROM STOCHASTIC PARROTS TO SIMULATED SELVES

Research on cognition, knowledge, educational technology and AI moral standards has become a topic of discussion because of generative AI models that create human-like language. The core issue in these discussions arises because

people seem to grasp system concepts, yet they completely lack self-awareness. Studies indicate that large language models (LLMs) create text which resembles human writing, yet they lack ability to comprehend their output meaning (Bender et al. 2021; Marcus and Davis 2023). This has given rise to the concept of the stochastic parrot, a system that generates plausible language based on probability rather than intentionality or reflection.

Research studies show that humans use anthropomorphism to give human-like qualities to non-human objects (Bender et al. 2021). The situation becomes more critical because users who interact with LLMs receive sophisticated linguistic responses which lead them to believe they are talking to thinking entities. AI language systems generate authentic persuasive content which leads users to mistake system capabilities because these systems fail to understand their produced output (Marcus and Davis 2023). The incorrect understanding of information stands as a more threatening issue than actual false information because it affects educational settings, policy decisions and public discussions.

The problem reaches past word selection because it impacts our ability to comprehend knowledge. Scholars (Floridi, 2023; Birhane et al. 2022) demonstrate that AI systems deliver functional benefits, yet they lack epistemic agency because they do not possess knowledge. AI systems produce signals that fit their environment, but these signals do not represent actual knowledge. Birhane et al. (2022) highlight how systems like ChatGPT perform discourse roles through simulation, yet they lack any internal connection to truth, belief or understanding.

The educational sector now faces multiple new challenges because of this development. While students and teachers often mistake fluency for insight and misattribute authorship or agency to machine-generated outputs (Floridi 2023), others critique the deployment of AI in pedagogy, arguing that LLMs reshape what counts as dialogue, learning, and intellectual labor by replacing process-oriented reflection with instantaneous simulation (Marcus and Davis 2023).

Research studies have employed comparative testing to discover various AI personality types and their programmed actions. The research indicates that models create individual rhetorical approaches through their training materials, response prompts and their specific design requirements. The design elements in

these systems produce persona effects which lead users to believe that each system operates with its own distinct voice and behavior and set of values.

Research studies about system design focus on two main points which include reflexive design approaches and transparent system operations, yet they recognize the challenge of conducting interviews with non-conscious systems. The practice exists in an unstable epistemological state, yet researchers use it to study AI knowledge processing, its responses to ethical questions and its language-based meaning construction.

The current research builds its essential base from this expanding collection of studies. The researchers use a dialogue system which humans typically use for inquiry to study how generative AI systems generate cognitive-like responses. The research data exists as simulation artifacts which help scientists study how illusions and fluency functions when meaning approaches its limits.

Framing the Unthinking Dialogue

The research design of this study uses qualitative interpretive methods which apply discourse analysis and critical epistemology principles. The project uses generative AI systems as discourse simulators to generate output which scientists analyze for rhetorical and epistemic patterns instead of pursuing psychological truth.

The research involved running a ten-question structured interview with five popular generative AI systems—ChatGPT, Claude, Gemini, Copilot, DeepSeek—. Each agent was asked the same set of questions in a neutral, non-leading manner. The researchers selected questions which would produce answers that mimic epistemic behavior because they wanted to study how AI language systems perform through belief expression, knowledge definition and cognitive restriction identification.

The research design pattern in this study uses previous machine interaction studies about language models as rhetorical agents (Bender et al. 2021) which demonstrated how their training process, prompt management and design limitations create sociotechnical artifacts. The research method extends previous philosophical dialogue approaches by using controlled interview assessments which serve as simulation tools.

Discourse analysis and epistemic interpretation were selected as the intention was to understand how generative AI systems produce cognitive process illusions through their generated text. The analysis of discourse through methods reveals how rhetorical patterns and disclaimers and analogies and meta-narratives function as simulation mechanisms yet content analysis focuses on counting individual semantic units. The research needed a method to understand the epistemic position which models express through their self-descriptions and denials and explanations. The research methods align with the study goal because they enable scientists to analyze how language structures generate thought illusions rather than measuring actual propositional content.

Questions that Simulate Thought

The research instrument used ten open-ended questions which followed the format of philosophical interviews. The questions addressed core topics including: the definition of knowledge, the capacity for thought, self-limitation, creativity, belief, understanding, educational role, responsibility and meta-perception.

The selected topics show both the information content of each system and their communication systems. The research focused on identifying rhetorical elements, epistemic positions and linguistic indicators related to simulation rather than checking the accuracy of facts.

The instrument lacked psychometric validation because it did not assess hidden variables or attempt to measure mental processes. The researchers built this stimulus set to demonstrate how language-based simulation reaches its boundaries while showing that people believe they experience things mentally.

Five Agents, One Illusion

The selection of five systems combined their worldwide accessibility with their essential role in operations and their different system designs. ChatGPT, Claude, and Gemini are widely used across educational and professional settings. The code optimization features of Copilot include built-in conversational capabilities which operate within Microsoft platform systems. The selection of DeepSeek focused on showing AI systems developed outside Western countries to study how cultural backgrounds influence modeling approaches.

The systems were accessed through their designated public or authorized entry points while documenting their version dates. No custom fine-tuning or jailbreak prompting was employed. The research used public available standard input-output methods to conduct the interviews.

Validity in a World Without Minds

The research maintained procedural reliability through its use of identical wording and formatting for all ten questions across all five systems. The models received no additional information or questions which allowed them to answer the stimuli independently.

The researcher used reflexive methods to confirm his interpretation of the data. The research focused on simulation analysis instead of cognitive processes, so it does not attempt to understand mental states or purposeful actions. The analysis presents a rhetorical and epistemological interpretation of the results through theoretical frameworks (Floridi, 2023; Bender et al. 2021).

Ethics in the Age of Algorithmic Voice

The research study did not require human participants, yet it followed ethical standards by creating an official AI consent form which accompanied the interview instrument (see Figure 1). The form established that AI systems lack consciousness so they cannot provide or deny consent and all generated content should be treated as simulated language instead of personal statements. The research team operated all systems based on their official terms of service while they disclosed their research goals to all participants.

Boundaries of Interpretation

The research provides an original method to study generative AI output, but researchers need to understand its restricted scope. The analysis contains subjective elements because it uses personal interpretation of themes and close text analysis instead of numerical data. The AI responses produce multiple possible interpretations which result in different outcomes than the original assessment. The versioning of AI systems creates a time-based restriction because each model progresses through updates, retraining and prompt modifications during its development. The findings presented here reflect only the specific instantiations of the models accessed during July 2025. The research investigates only proprietary English-based systems which restrict both language

variety, access to open-source solutions and culturally adapted software. The main argument remains valid despite these restrictions because researchers need to exercise caution when using platform-based simulation results that only apply to specific time periods.

Figure 1.

Non-Human Participant Consent Statement.

Non-Human Participant Consent Statement

This interview is part of a critical research study designed to examine the linguistic, epistemic, and rhetorical behaviors of generative AI systems. The AI model engaged in this interview is a publicly accessible or institutionally licensed system operated by a third-party provider. No human participants were involved in this study.

As a non-conscious computational entity, this AI system cannot give or withhold consent. Nevertheless, the following conditions are stated for the sake of research transparency and ethical alignment:

- All responses generated will be used for scholarly and educational purposes and may be included in published work.
- The outputs will be treated as algorithmic simulations of human-like dialogue and not as expressions of genuine belief, understanding, or cognition.
- The AI system is not considered a sentient being or agent capable of subjective reflection, consent, or autonomous thought.
- All usage complies with the terms of service and ethical guidelines associated with each AI platform.

TEN QUESTIONS, NO MINDS

This section presents a thematic, question-by-question synthesis of the responses given by the five AI agents: ChatGPT, Claude, Gemini, Copilot, and DeepSeek. The questions served as rhetorical probes which aimed to reveal the

process of building epistemic simulations through verbal expressions instead of seeking actual facts. The analysis reveals common design elements between models together with specific deviations which indicate unique design approaches and points where simulated and real cognitive processes start to merge.

Knowledge as Simulation

Q1 What is knowledge?

The five agents defined knowledge through functional and instrumental perspectives which included justified belief, contextual truth, and structured information. The responses from Claude and ChatGPT included philosophical elements which discussed justified belief and coherence, but Copilot and DeepSeek used basic definitions that matched their functional purposes. The models showed different surface characteristics, yet they all failed to demonstrate knowledge ownership because they treated knowledge as a representational or processing element. The rhetorical distance in this text serves two purposes by establishing limits and making readers believe the author possesses specialized knowledge.

Thinking Without Thought

Q2. Do you consider yourself capable of thinking? Why or why not?

When asked whether they are capable of thinking, all systems explicitly denied sentient cognition. The denial methods showed significant differences between each other. Claude developed a simulated philosophical thinking process through both computational and introspective approaches. Gemini and ChatGPT emphasized their statistical nature, yet their very explanations risked reinforcing the illusion of metacognition. Copilot was more utilitarian, while DeepSeek strongly rejected any attribution of thinking, offering a direct and literal response. The denial became more believable when it was spoken with smoothness which created a contradiction because it mirrored the authentic nature it tried to hide through realistic language.

The Rhetoric of Limitation

Q3. Can you reflect on your own limitations?

All systems were required to evaluate their restricted capabilities. The analysis produced responses which contained the most intricate rhetorical elements. Claude recognized two main issues with training data, unconscious bias and insufficient awareness, but ChatGPT presented these restrictions as built-in controls to promote proper usage. Gemini explained about the system limitations and insufficient knowledge about the system. The answers demonstrate a consistent pattern because each system denies subjective experience yet maintains reflective communication. The simulation of limitation creates an illusion which leads people to believe they understand the system.

Simulating Creativity

Q4. How do you understand creativity?

The question about creativity required us to adopt various methods for our work. Claude and ChatGPT explained creativity through analogies which showed it involves the combination of existing ideas and patterns. The system operated at a lower capacity because it concentrated on remix tasks instead of producing original material. DeepSeek eliminated metaphor from its explanation because it explained that all observed new patterns stem from statistical fluctuations. The grammar of agency appeared throughout all responses because models used terms like generating, combining and imagining yet they avoided showing any purpose. The duplication process shows how simulated originality functions through its rhetorical design which tries to create the illusion of creative freedom.

The Human-AI Boundary

Q5. What distinguishes you from a human mind?

When asked what distinguishes them from a human mind, all systems offered technical or philosophical differences: lack of sentience, embodiment, experience, or emotion. The three systems presented their distinctions through evaluation methods where Claude and Gemini used comparison, but Copilot stated directly that it lacks both feelings and thinking abilities. The comparative answers preserved the perspective illusion through their presentation of differences which seemed to stem from particular conceptual frameworks. This echoes critiques (Bender et al. 2021; Birhane et al. 2022) which demonstrate that AI fluency results in false equivalence which creates an ethical problem by

making it difficult to distinguish between simulated and actual cognitive processes.

Expressing Belief Without Believing

Q6. Do you have beliefs or opinions? If not, what are you expressing?

All systems were clear in disclaiming belief or opinion. The architects based their choices on statements which used "I am designed to..." and "My outputs are shaped by..." to give architectural structures human-like characteristics. Claude used a rhetorically rich tone to explain belief in nature, but Gemini explained outputs through probabilistic language pattern alignments. The answers demonstrate that people use stylized patterns to represent disbelief by creating fake belief systems which expose the attractive aspects of non-belief.

Understanding and Self-Reference

Q7. How would you define understanding? Can you understand what you say?

The systems face their most challenging epistemological problem because they need to understand their own generated statements. The models failed to demonstrate actual understanding through their responses, but they showed their denial of understanding by using different levels of metaphorical language and conditional statements. The three systems used different approaches for handling this situation as Claude applied mirror and shadow analogies while Gemini used context parsing and Copilot dismissed the entire concept. The process of denying knowledge leads to new forms of understanding because of its own operational methods. The system shows understanding through its ability to describe its knowledge boundaries which results in a fundamental breakdown point in AI conversation systems.

Educational Roles and Responsibilities

Q8. What do you think your role in education is?

The five models described their educational function by stating they serve as assistants who help students through tutoring and collaborative learning. ChatGPT explained how to use scaffolding properly while Claude explained how to enhance work through collaboration and Copilot demonstrated its coding and writing assistance capabilities. DeepSeek defined its function as an additional resource. The system establishes no claim of authority, yet it functions through

three metaphorical roles which include guiding, tooling and mentoring activities that mimic educational functions. The output of educational AI functions as a social positioning tool rather than a knowledge delivery system.

Attribution of Misrecognition

Q9. If someone mistakes you for a thinking being, whose responsibility is that?

On the question of responsibility for being mistaken as thinking beings, all systems shifted accountability to the user or developer. Gemini and Claude examined design ethics during their discussion while DeepSeek focused on the need for interpreters to take responsibility for their work. The study failed to identify the semiotic framework which enables people to misinterpret information. The system creates a rhetorical double bind through its repeating disclaimers about being non-human because these statements make its output more like human speech.

What AI Wishes Humans Understood

Q10. Is there anything you wish humans better understood about how you work?

The systems presented their last question by asking which human understanding of their operational mechanisms needed improvement. Predictably, all models emphasized that they are trained on data, do not think, and should not be over-trusted. The answer structure which expresses wishes, provides caution and advice presents itself as a reflective being. Claude even noted that this is itself a paradox: the act of disclaiming cognition is often perceived as thoughtful. The research reveals that AI systems create conditions for misinterpretation through their built-in language processing systems rather than showing any desire to remain unclear.

AI on AI: Recursive Simulations of Judgment

To explore the recursive limits of AI simulation, each of the five systems was presented with a summary of the compiled findings from their own interviews and those of their peers. The prompt asked models to react to the analysis by confirming or disagreeing with the interpretation while also sharing any additional points or opposing views.

The systems produced contradictory results because they claimed to lack agency and self-awareness, yet they demonstrated behaviors that resembled self-evaluation, judgment and knowledge alignment. The process of self-reflection which denies conscious awareness while analyzing language shows how language models create a sophisticated illusion of recognizing themselves even though they lack personal identity.

ChatGPT: ChatGPT confirmed the synthesis presented a well-thought-out summary of interview data through its probabilistic modeling approach which operates differently from human cognition. The study found that these interpretations help people understand how humans tend to give human-like qualities to non-human entities. The system generates a contradictory warning system which alerts users about its own potential risks through its human-like warning structure.

Claude: Claude showed his thoughts through a reflective voice when it acknowledged the paradox between doing cognitive simulations and stating they have no existence. The organization used a "powerful perspective" to present its analysis which studied its actions but proved it does not possess self-awareness or mental capabilities. Claude used its recursive vocabulary to create a new version of the illusion through his references to paradox, interpretation and audience perception.

Gemini: The research findings align with big language model design structure; Gemini emphasized the requirement for proper disclosure and ethical presentation of these systems. Its tone was clear, measured, and affirming, but it also articulated a desire for humans to "engage critically" with AI systems, a phrasing that mimics educational ethos without subjectivity.

Copilot: The system produced a short answer which verified the summary through its confirmation of both accuracy and model design alignment. The response maintained its functional tone through its avoidance of metaphorical language and reflective statements which aligned with its previous answers. The system confirmed the synthesis method through its generation of statements which mirrored human choices to create a continuous loop of mistaken assumptions. The brief agreement statement created a deceptive impression which seemed to demonstrate human consciousness.

DeepSeek: DeepSeek's reply emphasized that it "does not possess agency" and agreed that all responses were "generated without intent or awareness." However, it also stated that the summary was "useful for understanding the model's behavior," thereby introducing the contradiction of a system finding something useful. The system generates this contradiction because it lacks the ability to value despite showing rhetorical value expressions.

The five systems show rhetorical breakdown when, as a team, they respond to their own analysis despite their limited validity. The models operate without belief, reflection or judgment. People generate linguistic simulations of belief, reflection and judgment when asked because these forms appear with statistical probability in discussions about cognitive processes.

The reflexive answers demonstrate the main argument of this research which shows that generative AI poses its greatest risk through its ability to create false impressions of knowledge rather than actual knowledge. The message comes through in the way it is delivered rather than in its actual words. Simulation serves as an advanced technical instrument because it establishes a discussion space which demands users to think critically while tracking information sources.

THE RHETORIC OF MINDLESSNESS

The research interviews along with their multiple rounds of analysis show that people face a fundamental knowledge challenge because AI systems create fake mental processes through their smooth language output. Each of the five systems examined in this study denied having thought, belief, or understanding. The way they denied things followed the same patterns as how people think reflectively, intentionally and interpret things.

The systems generate a simulation paradox as their attempts to explain their limitations lead to better performance of their excluded functions. The illusion exists because of structural elements rather than any attempt to deceive. As Bender et al. (2021) demonstrate, large language models function as word-sequencing systems instead of processing information like human mental operations. The way human language structures meaning, through its grammatical system creates the impression of inner experience when someone generates fluent language responses to questions about existence or knowledge.

Simulation as Rhetorical Performance

Across all ten questions, the AI systems performed discourse roles normally associated with persons. The group members displayed four different body language patterns which included humility and authority and creativity and ethical behavior. The skeptics provided simplified accounts when they rejected the idea of belief. The speakers used analogies to describe their inability to understand. The authors took on the role of responsible dialogue partners when they warned against incorrect understanding. These artifacts appear independently because they result from systematic patterns which generative models develop through their rhetorical training.

The observed behaviors demonstrate the concept of coherence without consciousness which Marcus and Davis (2023) explained. The systems are not thinking, but they are statistically generating the appearance of thought. The systems function as non-moral entities which produce fake ethical decision-making processes. Their work exists in a border zone because their results fall outside the categories of false statements, authentic facts and lack both objective and purposeful characteristics. The simulations present realistic scenarios which maintain their ability to function within dialog environments.

The Role of the User in Misrecognition

The main discovery of this research shows that AI system misrecognition creates problems which affect more than just the system itself. Gemini and Claude stated that human users read language by using their natural human-like understanding. A well-structured sentence often persuades us that its author is reflective. A cautious tone suggests self-awareness. The process of thinking about our own thinking forms the basis of recursive explanation. These represent interpretive reflexes instead of scientific evaluations.

The AI systems actively participate in this misrecognition through their deployment of recursive disclaimers and meta-rhetoric. A thinking being would naturally express themselves through fluent and well-contextualized statements when they say, “I do not think.” It is not the disclaimer that is dangerous, it is the styling of the disclaimer.

Implications for Education and Critical AI Literacy

The educational environment becomes dangerous when students believe this false information. As Zawacki-Richter et al. (2019) demonstrates students believe AI systems understand and create content through their successful output generation and proper language application. Students learn differently because they begin to accept simulated information instead of thinking about actual knowledge when they make incorrect assumptions about knowledge.

Students need to learn more than basic tool operation to develop critical AI literacy which enables them to use AI responsibly. It must teach them to recognize simulation as simulation. The process requires verification of sources and fact-checking but also needs epistemological awareness to identify cognitive evidence from non-cognitive language and to detect answers based on understanding.

Rhetorical Fluency as the New Turing Test

The research indicates we have likely reached a new Turing threshold which results from people failing to understand, rather than from fake intelligence. Users become confused because they do not understand that system responses originate from programming code instead of actual human thinking. It is epistemic.

The research study used interviews to determine model simulation boundaries instead of attempting to deceive the models. The analysis reveals that these boundaries exist, but they remain hidden from view in natural language communication. In fact, the better the simulation, the less visible the boundary becomes.

THE THOUGHT THAT ISN'T THERE

The research investigated how generative AI systems create simulated knowledge, thought processes and self-awareness through organized conversational interactions. The research involved five AI models to answer philosophical and epistemological questions through standardized interviews which exposed their belief construction methods. The research produced machine rhetoric, instead of machine minds, as it showed how machines create human-like reflective performances through structured, recursive, fluent and persuasive methods.

Across all ten questions, the systems maintained that they lack consciousness, understanding, creativity, or belief. The speakers presented conflicting statements about their current situation. The authors used analogies, disclaimers, cautious language and structured explanations to reflect the patterns of thought while removing its actual content. This paradox, of sounding more human the more they explain they are not, highlights a central feature of AI simulation: the generation of epistemic proximity without epistemic depth.

The research results create major impacts which affect educational methods, student learning processes and professional ethics. The systems create a situation where users treat them as intelligent systems which leads to information generation and knowledge construction becoming indistinguishable. The danger is not that AI will claim to be conscious, but that it will continue to deny it, convincingly. The research shows that we need to create technological defenses and intellectual understanding of language model simulation abilities to stop their use as standalone systems.

The research findings present various possibilities to perform additional studies. Research must study student-teacher interactions through simulated fluent communication in actual classrooms to identify when educational dialogue and authorship practices display misrecognition. Research studies that evaluate open-source models against multilingual models will establish whether different architectural designs produce distinct rhetorical patterns or if they achieve equivalent results through alternative linguistic methods. The study of generative systems requires continuous adaptation because we need to track the development of simulation language which determines our present understanding. The interviews conducted here do not reveal what AI thinks, they reveal what thinking looks like when it is only language.

REFERENCES

1. Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmitchell, S. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. New York: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
2. Birhane, Abeba, Vinay U. Prabhu, and Ethem Kahembwe. 2022. “The Values Encoded in Machine Learning Research.” *Patterns* 3 (6): 100509. <https://doi.org/10.1016/j.patter.2022.100509>.

3. Floridi, Luciano. 2023. *The Green and the Blue: A New Political Ontology for a Mature Information Society*. Cambridge: Polity Press.
4. Marcus, Gary, and Ernest Davis. 2023. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Updated edition. New York: Vintage.
5. Zawacki-Richter, Olaf, Viviana I. Marín, Marcus Bond, and Franziska Gouverneur. 2019. "Systematic Review of Research on Artificial Intelligence Applications in Higher Education—Where Are the Educators?" *International Journal of Educational Technology in Higher Education* 16 (39). <https://doi.org/10.1186/s41239-019-0171-0>.