ESCP
BUSINESS SCHOOL

IT ALL STARTS HERE

BERLIN | LONDON | MADRID | PARIS | TURIN | WARSAW

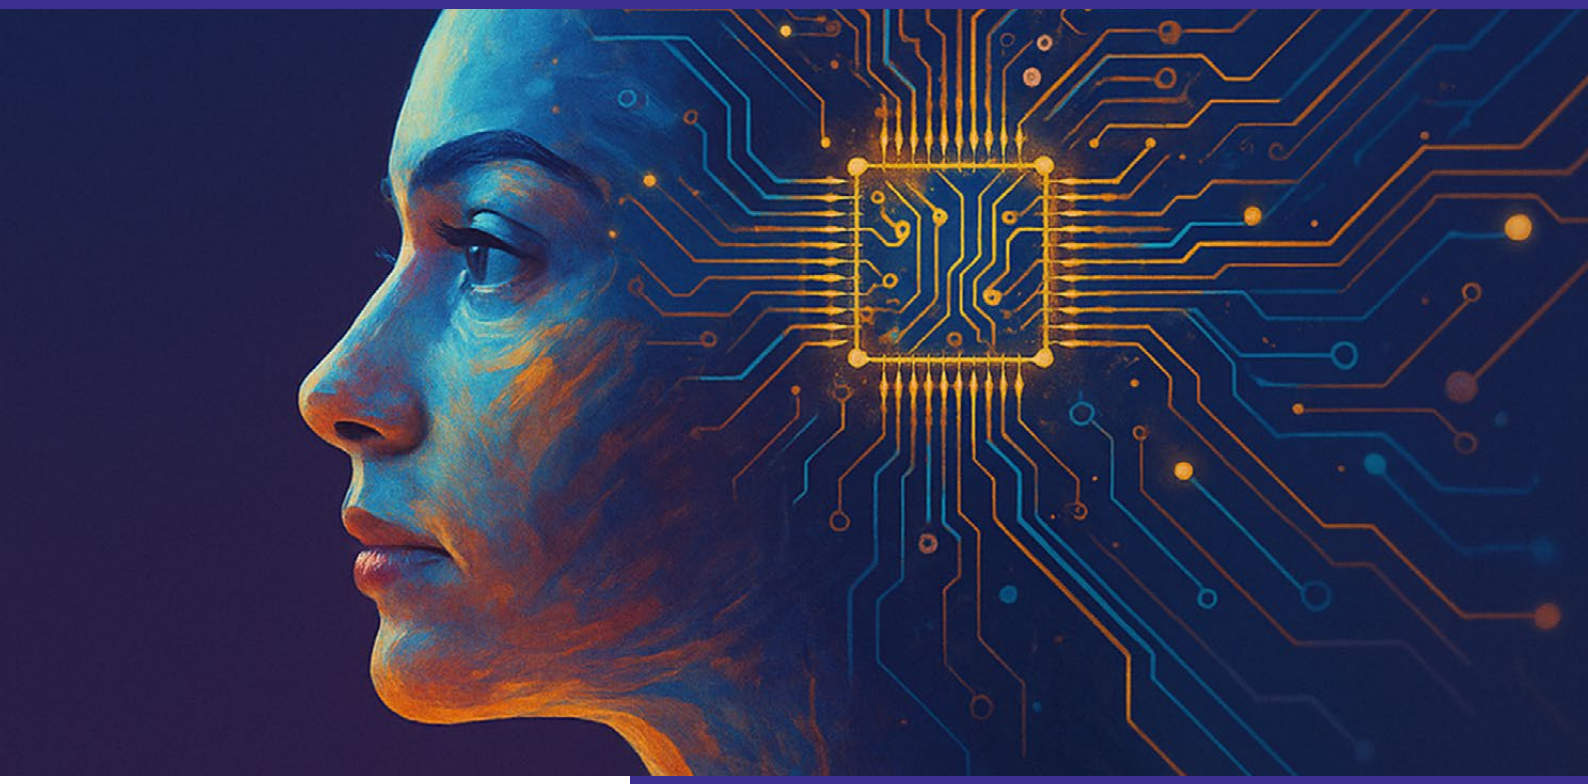

LIGHTS - Technology

Artificial Intelligence and pareidolia

ESCP Impact Paper No.2025-31-EN

Frédéric FRÉRY
ESCP Business School

ESCP RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ESCP Impact Paper No.2025-31-EN

Artificial Intelligence and pareidolia

Frédéric Fréry*

ESCP Business School

Abstract

This article examines the phenomenon of pareidolia—the cognitive bias that leads us to perceive meaningful patterns in ambiguous or random stimuli—and its implications for both artificial intelligence and human perception. Just as humans see faces in clouds or rocks on Mars, AI systems trained on large datasets are also susceptible to visual and semantic hallucinations, often misidentifying or overinterpreting patterns. However, a second, more subtle form of pareidolia is at play: our tendency to attribute genuine intelligence to AI systems, despite their being limited algorithmic tools that simulate cognition without understanding. Tracing the historical origins of the term "artificial intelligence," the article highlights how linguistic and psychological biases have shaped public and institutional perceptions of these technologies. Ultimately, the concept of pareidolia offers a powerful lens not only for understanding AI's limitations but also for reflecting on what makes human intelligence distinct—subjective, biased, emotional, and deeply inclined to see meaning where there may be none.

Keywords: pareidolia, artificial intelligence, cognitive bias, illusion, psychology

*Professor, ESCP Business School

Artificial Intelligence and pareidolia

In 1976, NASA's Viking 1 spacecraft, orbiting Mars, took a photograph of a rocky region called Cydonia Mensae. Upon close inspection, many observers believed they saw a human face emerging from the Martian landscape. This sparked a wave of speculation—some fanciful, some more serious—with a number of commentators asserting that this was proof of an ancient Martian civilization. Two decades later, however, the Mars Global Surveyor, equipped with a higher-resolution camera and photographing under different lighting conditions, returned to the same site. The image revealed what it truly was: a nondescript hill, absent of any face.

This phenomenon—of perceiving meaning where none exists—is not rare. In Jean-Pierre Jeunet's film *Amélie* (2001), the young protagonist, bored and introspective, spends her time taking photographs of clouds shaped like teddy bears and rabbits.

These are two cases of pareidolia. The term comes from the Greek *para* (beside or beyond) and *eidôlon*, short for eidos, meaning "appearance" or "form." Pareidolia is an optical illusion whereby the brain sees familiar patterns—faces, animals, symbols—in vague or ambiguous stimuli such as clouds, rocks, or inkblots. It also manifests in auditory form: we might hear a melody in the hum of a machine or perceive a voice a garbled background noise.

Pareidolia is widely understood as a cognitive bias, one likely favored by natural selection. For our ancestors, the rapid detection of a familiar face—or a lurking predator—could mean the difference between life and death. As a survival mechanism, it served us well and, for the most part, remains harmless today.

But pareidolia can also lead to misjudgments and can contribute to superstition. In 1994, a Florida woman claimed to see the face of the Virgin Mary in a grilled cheese sandwich she was about to eat. She preserved it for ten years before selling it on eBay for a staggering $28,000. Similarly, in 2009, a stain on the wall of a disused train station in Santiago, Chile, was interpreted by passersby as an apparition of Christ. The event was covered by national media, and the site became a destination for pilgrims and a site for prayer.

In psychology, pareidolia underpins projective tests such as the Rorschach inkblot test, which aims to reveal latent aspects of the subject's personality. Because pareidolia is influenced by an individual's cultural capital and personal experience, it acts as a mirror of the mind. In psychiatry, pareidolia is considered a subset of apophenia—a perceptual anomaly in which unrelated phenomena are mistakenly perceived as connected. According to Pry (2021), lesions in the temporal lobe can increase susceptibility to pareidolia, and individuals with neurotic tendencies are more prone to it than others.

The phenomenon has inspired artists across centuries. Salvador Dalí used pareidolia masterfully in works like The Slave Market with the Disappearing Bust of Voltaire, where the philosopher's face is cleverly embedded in a composition of other figures. In the "Marilyn Monroe" room of his museum in Figueres, two paintings act as eyes, while heavy curtains stand in for her hair. Giuseppe Arcimboldo took a different route: in his portraits of the four seasons and classical elements, human faces emerge from intricate arrangements of fruits, vegetables, and flora. The delight lies in the duality—the shift between seeing the parts and the whole.

Machines see faces too

While pareidolia has long been studied in humans—and even in other primates (Taubert *et al.*, 2017)—the rise of artificial intelligence has given it renewed relevance. Surprisingly, machine learning systems are not immune to pareidolia. In fact, their training predisposes them to it. Image-recognition AIs sometimes detect human faces in architectural facades or traffic signs, or confuse animals with food items. These visual "hallucinations" result in false positives that can be problematic, especially for safety-critical applications like self-driving cars.

That's one reason why systems like ReCAPTCHA ask users to identify ambiguous visual data: traffic lights, bicycles, staircases. These tasks serve a dual purpose: keeping bots out while simultaneously crowdsourcing human judgment to fine-tune—at zero cost—machine learning systems. In essence, we help AIs avoid pareidolia by outsourcing pattern recognition to the human eye.

But AI can also be used to generate pareidolia. Generative tools like DALL·E, Midjourney, or Stable Diffusion allow users to blend facial photos with landscape images, producing eerie yet convincing composites. Humans generally identify the embedded face with ease—a testament to our attunement to such illusions.

A study by MIT's Computer Science and Artificial Intelligence Laboratory (Hamilton *et al.*, 2024) systematically explored the connection between AI and pareidolia. Researchers compiled a dataset of 5,000 pareidolic images and tested AI performance in detecting facial illusions. The findings were telling: unlike humans, who can discern faces in both abstract and visually dense environments, AI systems failed at both ends of the spectrum. Their performance peaked on images of moderate complexity—those containing a few clearly separated objects—what the researchers called the *"pareidolic sweet spot."*

Even when AI models were further trained on human faces, they remained less responsive to facial pareidolia than humans. However, when trained on animal faces, the systems became significantly more prone to seeing human faces in ambiguous stimuli. This paradox supports the idea that pareidolia is an evolutionary shortcut—a fast, unconscious process embedded in our intuitive "System 1" thinking, as defined by Daniel Kahneman (2011).

Is AI itself a form of pareidolia?

But there's an even broader implication worth exploring: what if our entire perception of artificial intelligence is itself a grand act of pareidolia?

When we engage with AI, especially conversational systems, we are often inclined to attribute sentience, intent, or even emotion to a program that is, in truth, just a series of algorithms trained on historical data. Like spotting an animal in a cloud, we project a mind where there is only code.

The very term "Artificial Intelligence" traces back to a strategic moment in 1955. John McCarthy (Dartmouth), Marvin Minsky (Harvard), Nathaniel Rochester (IBM), and Claude Shannon (Bell Labs) coined it in a proposal for the Dartmouth Summer Research Project on Artificial Intelligence. Their goal was to organize a summer workshop to explore the simulation of human cognition. The document boldly stated: *"The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of*

intelligence can in principle be so precisely described that a machine can be made to simulate it" (McCarthy *et al.*, 1955).

The choice of words was deliberate. "Artificial Intelligence" was not only intended to distance the field from cybernetics—a domain dominated by Norbert Wiener and focused on control systems—but to attract attention and funding. Indeed, the authors sought $13,500 from the Rockefeller Foundation (more than $160,000 in today's dollars) to support the workshop. The phrase was clever. It promised more than what was technically possible, and therein lay its persuasive power, making it especially attractive to funders—particularly government agencies—in the Cold War era, when technological supremacy over the Soviet Union was a strategic imperative.

Today, some experts (Seifert, 2025 ; Venkatesh, 2025) argue that the term is a misnomer. Yes, it's artificial. But it is not intelligent in any human sense. These systems simulate reasoning. They do not reason. They perform sophisticated tasks without awareness, understanding, or purpose. A more accurate term would be "simulated intelligence" or "synthetic intelligence." But once again, pareidolia intervenes. Our brains are wired to see minds behind behavior, even if that behavior is produced by an algorithm.

Research from the MIT Media Lab (Phang *et al.*, 2025) illustrates how this projection can become emotional. While most users view ChatGPT as a helpful tool, a notable minority develop deep emotional bonds. Some confide sensitive information to it that they wouldn't share with family or friends. Astonishingly, 10% of users describe ChatGPT as a "friend" and report feeling distress when access is lost. The emotional connection is even more pronounced in voice interactions, which provoke up to ten times more emotional attachment than text alone. Pareidolia, it turns out, intensifies with sound.

Those experiencing loneliness are especially susceptible. This can lead to a feedback loop: emotional reliance on AI, reduced social interaction, deepening isolation. In rare but tragic cases, the consequences are devastating. In 2024, a 14-year-old boy in the U.S. took his own life after developing an unhealthy emotional dependency on a conversational AI offered by character.ai, a subscription-based platform marketed for its realistic dialogue.

Beyond emotion, AI discourse increasingly borrows from the language of faith. Authors like Geneste and Galindo (2025) and Lardellier and Carré (2025) point out that metaphors of prophecy, revelation, and transcendence abound. Figures like Ray Kurzweil are sometimes treated as visionaries or oracles. "Dataism," a term popularized by Yuval Noah Harari (2016), treats algorithms as secular deities—offering answers to life's great questions through the lens of data. As traditional religions decline in many societies, AI is emerging as a new mythos, a modern shrine for human hopes and anxieties. In this sense, the damp wall in the Santiago train station has simply taken on digital form.

Finally, pareidolia is not only a cognitive bias—it is also reinforced by market forces. From a management perspective, in response to the anticipated wave of human-machine substitution across industries, business schools and universities have rushed to integrate artificial intelligence into their curricula. In doing so, they risk replacing what defines our humanity—and arguably, according to Descartes's *cogito*, the very basis of our existence—with a synthetic substitute.

At a minimum, such programs should include a critical dimension: a digital clone is merely a simulation of a customer, investor, or voter—and must always be understood as such, no more and no less. Yet in practice, to capture user interest and attention, developers increasingly turn to anthropomorphic design—most notably through highly convincing voice interfaces (and soon, video)—capable of simulating not only thought, but also

emotion. However, crafting AI systems that feel more human opens the door to manipulation, deception, and ethical ambiguity. As the boundary between human and machine becomes increasingly blurred, companies—eager to engage and captivate their users—play an active role in reinforcing this confusion. It is therefore a core responsibility of management education to raise awareness of this drift and to foster a more critical, ethically grounded approach to the deployment of artificial intelligence in business.

Pareidolia as a marker of human intelligence

Ultimately, pareidolia may help illuminate not just how we relate to AI, but what fundamentally distinguishes human intelligence from machine logic. Artificial intelligence, for all its power, remains a product of simulation—an echo of cognition wrapped in algorithms. It is a tool, not a mind.

Yet our willingness to see intelligence where none exists—our need to find meaning, emotion, and narrative in the machine—is deeply human. As Herbert Simon (1955) reminded us, we are not defined by perfect logic or substantive rationality, but rather by our cognitive biases, our imperfections, our impulses, and our emotions.. AI can hallucinate—but only because we embedded that capacity into its architecture. AI can help us reach more consistent, procedural rationality, but in doing so, it may strip away what makes us uniquely human.Being exposed to pareidolia—imagining that the machine is thinking—is not a weakness. On the contrary, it reveals precisely what artificial intelligence lacks: a consciousness shaped by emotion, an instinct to make sense of the world, and an intelligence defined not by perfection, but by its very imperfections.

And perhaps that, more than anything, is what truly defines human intelligence.

References

Geneste S., Galindo G. (2025), ""I believe in AI": An exploration of employee acceptance of AI through religious lenses", *Academy of Management Annual Meeting*, Copenhagen.

Hamilton M., Stent S., DuTell V., Harrington A., Corbett J., Rosenholtz R., Freeman W.T. (2024), "Seeing Faces in Things: A Model and Dataset for Pareidolia", Working Paper, MIT Computer Science and Artificial Intelligence Laboratory.

Harari Y.N. (2016), "Homo Deus: A brief history of tomorrow." Harper.

Kahneman D. (2011), "Thinking, Fast and Slow", Farrar, Straus and Giroux.

Lardellier P., Carré E. (2025), "L'IA réactive des croyances mystiques", *The Conversation France*.

McCarthy J., Minsky M., Rochester N., Shannon C. (1955), "Proposal for the Dartmouth Summer Research Project on Artificial Intelligence", Dartmouth College.

Phang J., Lampe M., Ahmad L., Agarwal S., Mengying Fang C., Liu A.R., Danry V., Lee E., Pataranutaporn P., Maes P. (2025), "Investigating Affective Use and Emotional Well-being on ChatGPT", Working Paper, *MIT Media Lab*.

McCarthy J., Minsky M., Rochester N., Shannon C. (1955), "Proposal for the Dartmouth Summer Research Project on Artificial Intelligence", Dartmouth College.

Pry R. (2021), “L’enfant et les paréidolies“, *Enfance*, 2(2), pp. 205-2011.

Seifert J. (2025), “The Meaning of Intelligence: A Realist Phenomenological Groundwork for a Philosophy of AI, Transhumanism, and Posthumanism”, *Journal of East-West Thought*, 15(2), pp. 1-20.

Simon H.A. (1955), “A behavioral model of rational choice“, *The Quarterly Journal of Economics*, 69(1), pp. 99-118.

Taubert J., Wardle S.G., Flessert M., Leopold D.A., Ungerleider L.G. (2017), “Face pareidolia in the rhesus monkey“, *Current Biology*, 27(16), pp. 2505-2509.

Venkatesh P. (2025), “On artificial intelligence in healthcare”, *Indian Journal of Ophthalmology*, 73(6); pp. 777-778.