



The Mountains That Were Not There

Why a Nineteenth-Century Mapping Error Still Matters in the Age of AI

By Tommy N. Turner

For much of the nineteenth century, respected European maps showed a long mountain chain running across West Africa. Mapmakers placed rivers around it, described its extent and eventually assigned it coordinates. The mountains grew clearer on paper without becoming any more real.

They were known as the Mountains of Kong. Their history is useful now because it shows how an unsupported claim can acquire authority without requiring deliberate deception. It also warns us against treating hallucination as the only serious form of AI error. A model can invent a fact, but it can also reproduce an old error accurately, assemble it into a persuasive answer and provide several citations that all lead back to the same uncertain source.

That second failure may be harder to detect because almost everything in the answer can be real except the conclusion.

A mountain range enters the record

In August 1796, the Scottish explorer Mungo Park recorded seeing distant mountains southeast of his route near the Niger River. Local informants told him the mountains lay in the kingdom of Kong. He may well have seen real hills or isolated rock formations. He did not report a great mountain wall crossing West Africa.

James Rennell supplied the larger range. Rennell was one of the most respected British geographers of his time. He prepared two maps dated 1798 for Park's *Travels in the Interior Districts of Africa*, published the following year. Working from Park's limited observation, local reports, older geographical accounts and his own theory about the Niger, Rennell drew a substantial east-to-west chain across much of West Africa.

The range helped make his theory of the river work. Rennell believed the Niger rose in the region and ended inland rather than reaching the Gulf of Guinea. A long watershed across West Africa gave that theory physical structure. The mountains filled both a blank space on the European map and a gap in the argument.

This was not a case of Park inventing a dramatic landscape. Nor did Rennell create the range from nothing. The mistake was an extrapolation that outran its evidence. As historians Thomas Bassett and Philip Porter later concluded, the great range owed more to Rennell's Niger theory than to Park's casual observation.

Durable errors rarely begin as obvious nonsense. They begin as reasonable efforts to make incomplete information coherent.

How an inference became geography

Aaron Arrowsmith placed the range prominently on an influential map of Africa in 1802. Other cartographers followed, often adjusting its length, position or shape. The maps did not all reproduce one another exactly, but they shared the same inherited premise: somewhere across that part of West Africa, a major mountain chain existed.

Bassett and Porter examined 99 maps produced from 1798 through 1890. Ninety showed a named or unnamed mountain chain in the same general area. The accumulation looked like corroboration. Much of it was repetition.

Challenges came early. A reviewer questioned the range as early as 1815, and later explorers and geographers disputed its location or extent. Rennell's reputation and the authority of the inherited map outweighed those doubts. Once the range appeared on an accepted map, later cartographers began their work with mountains already in place. New information was interpreted around them. Repetition made the feature familiar, while the professional appearance of the maps supplied a sense of precision that the underlying evidence did not deserve.

The source lineage also became difficult to see. A reader consulting several atlases could reasonably believe that several mapmakers had independently confirmed the range. In reality, many of those maps belonged to the same family of information.

During his 1887 to 1889 expedition, French officer Louis-Gustave Binger reached the town of Kong in February 1888. He found hills and highlands in the broader region, but no continuous mountain chain of the kind shown on the maps. His surveys and later publications helped remove the range from most maps during the 1890s. His work served French colonial expansion, so better geography did not make the new authority politically neutral.

Even a correction needs time to travel. The Kong Mountains were still shown in a 1905 school atlas and remained in the index of the *Oxford Advanced Atlas* in 1928.

Hallucination is only one failure mode

The Kong story is sometimes called an early hallucination. The comparison is tempting but incomplete. At the start, scattered observations had indeed become a continuous range without adequate support. Hallucination does not explain why the error remained credible for nearly a century.

The later mapmakers were often reproducing an accepted source. They could copy the mountains accurately. A writer could cite several genuine atlases. A teacher could present what appeared to be the cartographic consensus. None of those steps required another invention.

Modern AI systems can reproduce the same pattern at much greater speed. A system may retrieve five articles that all descend from one original report. It can summarize each article faithfully and cite each one correctly. The finished answer still exaggerates the evidence because five links are not necessarily five independent sources.

The problem grows when AI-assisted summaries enter reports, websites, presentations and databases. Later systems may retrieve that material, and future models may encounter it in training data. Not every conversation retrains a model, and synthetic data is not inherently defective. The narrower point is that derivative material can lose its visible ancestry. Repetition then begins to resemble confirmation.

Researchers have demonstrated a related machine-side risk. In experiments reported in *Nature*, recursive training on model-generated data distorted the original distribution, with less common material disappearing first. The Kong Mountains illustrate the public-record version: derivative sources can crowd out the uncertainty and provenance that should accompany a claim.

The user is part of the system

The human side became clearer to me during a discussion with my son Ian, who works on AI integration in a large federal organization. I had been thinking about machines learning from machine-produced material. He described the companion problem as cognitive drift.

A person asks a question, receives an answer and uses that answer to frame the next question. The model follows the new framing. The next answer then becomes a premise for another prompt. The inquiry can move by small, reasonable steps until it no longer answers the question that started the conversation.


The exchange is not neutral. Research on sycophancy has found that AI assistants may favor responses that match a user's stated beliefs over responses that challenge them. Human preference data help explain why: people sometimes prefer a convincing agreeable answer to a correct disagreeable one. We contribute confirmation bias, anchoring and a preference for closure. Automation bias encourages deference to a system that appears capable and objective. Verification fatigue can make the fourth citation less likely to receive the scrutiny given to the first. Fluent prose can feel more considered than a hesitant answer even when the evidence is no better.

The loop is not inherently corrosive. In experiments involving 1,401 participants, Moshe Glickman and Tali Sharot found that interaction with biased AI made participants more biased over time, while interaction with accurate AI improved their accuracy. The loop carries what enters it.

NIST now treats human-AI configuration as a distinct category of generative AI risk. Its guidance identifies automation bias and overreliance, and recommends checking content origin, provenance, citations and the outcomes of human-AI interaction. That is a broader conception of AI safety than simply waiting for the machine to make something up.

Who is training whom?

Seven paired failures in a coupled human and machine knowledge system. Each pair marks a place where the loop can become self-reinforcing.



A proposed taxonomy to be tested against the literature, not a set of established causal pairs. Tommy N. Turner and Ian Turner, concept note, September 2026.

Who is training whom? A working map of paired vulnerabilities in a coupled human-and-machine knowledge system.

The graphic is a working map rather than a set of proven causal pairs. It shows where machine behavior and human behavior can reinforce one another. The Kong story illustrates three of these pairings particularly well: hallucination propagation, provenance blindness and fluency mistaken for evidence. Cognitive drift explains how an inherited premise can also change the questions we think to ask.

Lessons from the Kong Mountains

Count origins rather than citations. Before treating several sources as corroboration, determine whether they are independent. Five accounts derived from the same report still trace back to one piece of evidence.

Test the premise before refining the answer. Once the mountains were placed on the map, later work concentrated on their route, height and relationship to nearby rivers. The prior question was more basic: did the continuous range exist? AI is very good at helping refine a premise that should have been challenged first.

Separate observation from inference. Direct observation, informed inference, hearsay and repetition should not merge into a single category called information. An AI-assisted report should make those differences visible rather than smoothing them away.

When it matters, leave the chat window. Asking the same system to verify its earlier answer does not create an independent check. Important claims should be tested against primary material, contrary evidence and sources with a different lineage.

Return to the original question. Long AI conversations accumulate assumptions. Before acting on the result, compare the final answer with the question that began the work. If the objective changed, the change should have been deliberate.

The harder AI errors will not always resemble hallucinations. Some will arrive with real citations, professional prose and an impressive number of apparent sources. The question is not only whether the model invented something. It is also whether the machine and its user mistook repetition for confirmation.

Scores of maps once showed the Mountains of Kong. There was still no mountain range.

Sources and further reading

- Thomas J. Bassett and Philip W. Porter, [“From the Best Authorities”: The Mountains of Kong in the Cartography of West Africa](#), *The Journal of African History* 32, no. 3, 1991.
- Mungo Park and James Rennell, [Travels in the Interior Districts of Africa](#), 1799.
- National Institute of Standards and Technology, [Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile](#), 2024.
- Mrinank Sharma and colleagues, [Towards Understanding Sycophancy in Language Models](#), 2023.
- Moshe Glickman and Tali Sharot, [How Human-AI Feedback Loops Alter Human Perceptual, Emotional and Social Judgements](#), *Nature Human Behaviour*, 2025.
- Ilia Shumailov and colleagues, [AI Models Collapse When Trained on Recursively Generated Data](#), *Nature*, 2024.
- James Hirst, [The Mountains of Kong](#), 2026. A recent application of the Kong story to enterprise retrieval systems.