

Communications of the ACM

Opinion

[Artificial Intelligence and Machine Learning](#)

As Answers Get Cheaper, Questions Grow Dearer

Abundant machine answers imply power shifts to better human questions.

By [Sangeet Paul Choudary](#) and [Marshall Van Alstyne](#)

Posted Dec 10 2025

In 19th-century Paris, the Académie des Beaux-Arts defined what counted as legitimate art. Realism, the prevailing standard, emphasized precision and visual accuracy. Success was based on how well you aligned with these norms. The system rewarded consistency, not experimentation.

Photographic advances in the 1830s and 1840s began to challenge this standard. At first, photography seemed like a threat to painters. If a machine could record the world more precisely and more quickly than a human hand, what role did painting have? But over time, photography freed painting from its representational obligations. Painters no longer had to compete with the camera in copying reality. Instead, they could focus on the subtleties that early cameras could not capture—the play of light, the texture of perception, new interpretations of the familiar.

Without photography, art would have progressed—at least for some time—on a predictable path toward more of the same: More Realism, with improvements in accuracy. If Realism was the prevailing *answer* for what was asked, artists would have gone on to give better answers. Photography, ironically, collapsed the cost of generating answers. You could get the most realistic portraits without hours of effort on the part of an artist. Photography freed painting from Realism, but what became more interesting was what rose to take its place.

Impressionist painters including Monet and Degas began experimenting with subjective experiences of color and light. Instead of representing reality, which the camera could do with far less effort, they started interpreting it. Instead of providing better answers—more Realism—Impressionists redefined the question altogether. With Realism, art was judged based on its representation. With Impressionism, art had a new purpose, as a means for interpretation. The camera provided cheap replicas: abundant answers. The Impressionists changed the framework and positioned art as a basis for asking better

questions.

When what was previously scarce suddenly becomes abundant, look for the new scarcity because that is what creates leverage.

LLMs and Cheap Answers

LLMs are the latest step in a long arc of technologies that have made answers progressively cheaper. From pocket calculators to spreadsheets to data analytics to recommendation engines to GenAI, each wave has broadened access and driven the cost of answers toward zero. These systems specialize in answers—not necessarily correct ones, and rarely final ones, but answers that are immediate and, most importantly, plausible. And, as far as LLMs are concerned, their answers have confidence.

The problem is that plausible answers can be worse than clearly wrong answers. Facts illustrate. When questions have unambiguous answers, LLMs add real value. If you are a call center worker, LLMs might increase your productivity an average of 14%.² But when questions have ambiguous, context-dependent answers, LLMs can lead astray. If you are an inexperienced entrepreneur, their advice might cut your profits 10%.⁴ Then real judgment is required.

When an answer feels good enough, we tend to stop asking. In an environment overloaded with content and starved for attention, plausibility becomes a stand-in for truth. Search results that confirm our assumptions rise to the top, memes that reinforce ideas in our heads get shared further, language models that complete our thoughts reinforce existing narratives.

As costs of continuing the inquiry rise, good questions become more expensive than ever. Additionally, our environment becomes less stable as we move toward structural uncertainty:³ a world where the rules are no longer static. What worked yesterday may not apply tomorrow, not because facts-on-the-ground have changed, but because the terrain itself has shifted. In such an environment, answers that were once reliable quickly become outdated. Static knowledge has limited utility in dynamic systems.

What matters more is the capacity to stay curious and continue a line of inquiry. This is where good questions—even though expensive—become strategic. A good question expands the field of awareness. It reframes the problem. In systems marked by structural uncertainty, value is created not by declaring what is known, but by directing attention to what remains unresolved. Valuable answers today are not necessarily those that appear complete and articulate, but those that reveal where we must keep looking. In a system with structural uncertainty, the goal moves past understanding the present to navigating the future, to continuously find orientation in a moving landscape.

The Consulting Conundrum

You might choose to hire consultants for their answers but that too can fall into traps that

don't, initially, feel like traps at all. As a client, you might assume that because answers become abundant, naturally you can evaluate them like other commodities: by cost.

But this mindset quickly leads to a race to the bottom. The cheaper the answers, the greater the likelihood that no one has done the expensive work of asking the hard questions. You save money upfront, but what you receive in return is often superficial, or worse, misleading.

Alternatively, you might bring in external expertise to help navigate uncertainty. But that introduces a second trap: you pay for guidance, without necessarily having a way to evaluate its real worth. Unless you can assess not just the answers but, more importantly, the quality of the questions asked, you have no reliable way to judge the value of what you're buying.

Unless clients learn to recognize the value of good questions, they risk falling into one of these traps. To navigate this shift, clients must build a new muscle: the ability to evaluate thinking not by fluency or presentation, but by the structure and depth of the inquiry behind it.

As Clay Christensen said of consulting: the value is not in having all the answers, it is in teaching clients how to think.³ Christensen did not need to know more about microchips than his clients. He helped them see the broader patterns. A good theory, in that sense, is a strategic flashlight: it shows you where to look when the data alone is dark. In a world of cheap answers, the expensive mistake is mistaking them for good ones.

The Re-Skilling Cop-Out

The go-to response to AI job disruption is often “re-skilling.” Whenever the conversation turns to automation and employment, the default solution is to invest in upskilling displaced workers. But if re-skilling only teaches people how to generate more answers, it can offer diminishing returns. Continuous AI improvement only makes those answers cheaper and more abundant. The real differentiator lies elsewhere: in the ability to frame better questions. This is a fundamentally different skill from what our education and training systems have traditionally emphasized. For much of the 20th century, success was measured by domain mastery: how well one could provide correct answers. More knowledge meant more advantage. But that logic only works when the underlying rules stay constant. When conditions are stable, the same answers have value. When conditions shift, it is the ability to ask the right questions that becomes invaluable.

What matters now is not just knowledge accumulation, but the capacity to navigate complexity without becoming trapped by the illusion of understanding. The most effective knowledge workers will treat uncertainty not as a threat to be minimized, but as a terrain to be explored. They will look to construct partially but directionally correct maps, instead of falling for the trap of irrelevant answers.

Good Questions Need Strong Theoretical Foundations

Curiosity alone is not enough. Understanding requires better framing. In environments where data are missing or misleading, we need more than observation. We need tools to reason in the dark. That's where deep theories become essential. While most AI and data science projects start from available data, theory allows us to ask better questions before data exists. A good theory provides structure: it helps you imagine what should happen, anticipate second-order effects, and evaluate whether plausible answers are meaningful or misleading.

In the 1840s, Ignaz Semmelweis saved countless lives of mothers and newborns by asking doctors to wash their hands in obstetrics wards after performing autopsies. His theory of "cadaverous particles" prompted reframing decades before Pasteur provided evidence that germs cause disease. Repeating history, countries that leaned on theoretical models of disease spread had vastly better public health outcomes than those who waited for COVID-specific data. Early actors, like Taiwan, New Zealand, and South Korea averaged less than 15 deaths per million compared to wait-for-evidence countries, like the U.S., U.K., and Italy, who averaged more than 900 deaths per million.^b Exponential systems severely punish late actors who wait for further data.

Deep theory does three things well: it helps you frame what matters, it helps you design what does not yet exist, and it helps you sanity-check answers, especially in ambiguous domains. Theories make implicit assumptions explicit. They challenge implications that might not hold. They can reveal when we're hallucinating patterns that are not causal.

AI systems and data-driven models typically detect correlations, but without theoretical grounding, they can easily mistake correlation for causation. This can lead to spurious conclusions on the basis of statistical patterns that have no power to control or design better outcomes. A well-formed theory offers a sanity check. It helps determine whether an observed pattern makes sense within a broader causal model or is simply an artifact of overfitting.

Theories also help us anticipate second-order effects. In environments marked by structural uncertainty, ideas from behavioral economics and game theory can guide how we design systems and rules, well before all relevant data arrive. These theoretical tools help us anticipate behavior, design for cooperation, and create mechanisms for coordination when empirical evidence is not yet available. Theory offers not just a lens for understanding the world, but a tool for shaping it.

Good Questions Change System Framing

Much of the world we live in today traces back to a deceptively simple question asked in 1948. At the time, engineers at Bell Labs were trying to improve the clarity of telephone calls. They were tinkering with wires, amplifiers, and filters, optimizing for better answers. But Claude Shannon asked something else entirely. Instead of asking how to reduce noise on the line, he asked a more fundamental question: What is the information passing through it? Shannon's breakthrough was to show that the more uncertainty a message resolves, the more information it contains, and with that, the more potential it holds for distinguishing between possibilities.⁵

Shannon's insight reshaped how we think about communication, encoding, and uncertainty. He shifted the emphasis from noise reduction to information transmission. And his work gives us one very useful lens: a good answer is one that reduces uncertainty. If someone tells you there is a 90% chance of sun tomorrow, in the middle of summer, you might shrug. You already assumed that. But if they tell you there is a 90% chance of hail, that changes something. You prepare differently. The difference is not in the volume of information—which stays the same—but in closing the gap between what you believe and what is. It makes uncertainty actionable. Good answers revise beliefs, not just confirm them.

Modern LLMs dazzle with their fluency. They speak with such coherence that we often forget to ask whether they're actually telling us anything new. This is the first trap. As answers become cheaper and more abundant, we begin to confuse ease of access with quality of resolution. These systems do not encourage you to keep asking; they subtly persuade you that there is no need for further inquiry.

The second trap is assuming that more answers are better than fewer. That is not always true.¹ In a world saturated with knowledge and starved of attention, the limiting factor is focus, not facts. And the more data we accumulate, the more we risk misallocating our attention to what's abundant instead of what's unresolved. Our systems reward speed and verbosity. But good inquiry often requires slowing down, noticing what's missing, and tolerating what's unresolved.

A truly good answer, in a knowledge-dense world, must do two things: it reduces uncertainty, clarifying something that was previously ambiguous, and it uses attention wisely, delivering insight without demanding more focus than it deserves. Anything less is noise; no matter how well it is phrased.

Most answers live within the boundaries of existing frames. The real breakthroughs happen when we step outside those frames. Good questions do not just seek better answers. They widen the frame itself. Copernicus and Einstein asked questions that broke the prevailing assumptions. So did the researchers who turned CRISPR, once just an obscure bacterial immune system, into a gene-editing revolution. None of these were answers to existing questions. In a flood of answers, the rarest and most valuable act may be the ability to ask a question that reveals the limits of the frame and helps us see the world anew.

The Surprising Human + LLM Advantage

Structural uncertainty is often resolved not in following the most obvious path, but in the unexpected intersection between distant ideas. This is where good questions become critical, because they can bridge concepts that do not already belong to the same system. This is one place where a human asking good questions can get superpowers when paired with an LLM throwing out cheap answers. LLMs are particularly good at connecting unconnected domains, but only if they are asked the right questions. In the hands of someone looking for cheap answers, LLMs can be a liability. But in the hands of those asking great questions, LLMs could confer superpowers.

References

1. Ackoff, R.L. Management misinformation systems. *Management Science* 14, 4 (1967).
 2. Brynjolfsson, E., Li, D., and Raymond, L. *Generative AI at Work*. National Bureau of Economic Research (Apr. 2023).
 3. Christensen, C.M. *How Will You Measure Your Life?* Harvard Business Review Press (2017).
 4. Otis, R. et al. *The Uneven Impact of Generative AI on Entrepreneurial Performance: Evidence from a Field Experiment in Kenya*. UC Berkeley (Oct. 2025).
 5. Shannon, C.E. A mathematical theory of communication. *The Bell System Technical J.* 27, 3 (1948).
-

Footnotes

^a See <https://bit.ly/4qOJ3Jr>

^b Source: <https://bit.ly/4ovc8lp>

About the Authors

Sangeet Paul Choudary (sangeet@platformthinkinglabs.com) is a Senior Fellow at the Haas School of Business, University of California, Berkeley, CA, USA.

Marshall Van Alstyne (mva@bu.edu) is Allen & Kelli Questrom Professor of Information Economics at the Questrom School of Business, Boston University, Boston, MA, USA.

Submit an Article to CACM

CACM welcomes unsolicited submissions on topics of relevance and value to the computing community.

© 2026 Copyright held by the owner/author(s).

Join the Discussion (0)