

The Trusted Advisors

EXPLORING WISDOM BEYOND EGO

PERSPECTIVE

The Amplifier in the Machine

Why the most urgent work of the AI era isn't technical. It's learning, finally, to see ourselves.

By **Roland Hess**

FOUNDER & MANAGING DIRECTOR

THETRUSTEDADVISORS.GLOBAL

Here is the short version, before the argument: what an organization calls its shared operational truth is a constructed reality — a version of the world everyone inside has quietly agreed not to question. Agentic AI does not correct that construction. Built the way most organizations are building it, the AI learns the construction, acts on it, and multiplies it.

It was raining sideways in Amsterdam the morning it happened. One of those grey March Thursdays where the light never properly arrives and everyone turns up with wet shoulders and coffee that went cold on the tram. Sixteen of us in a fourth-floor room: two executives, the CFO, a head of data who had clearly been told to attend, and a woman named Marieke who ran operations and had — I would work this out by lunchtime — the sharpest instincts in the building. The agenda taped to the wall promised agentic AI. Use cases. A deployment roadmap by four o'clock. Somewhere between the second coffee and the roadmap, I asked them to close their laptops and sit with a stranger question:

What, exactly, do you think is real?

It is not the question anyone came for. The agenda promised systems that don't just execute instructions but infer intent, fill in gaps, and make the kind of judgment calls once reserved for people. The room expected use cases. It got neuroscience instead, and I could feel two or three people deciding whether to be annoyed.

Here is the finding that was waiting for them, and is waiting for the rest of us. None of us perceive reality directly. Our brains do not record the world; they predict it. What we experience as plain, unmediated fact is a best guess — assembled from memory, expectation and pattern, continuously revised but never fully corrected by what the senses actually report. Researchers call this a controlled hallucination, and the phrase does them no favours, because it sounds like an accusation. It isn't. It is not a flaw in perception; it is how perception works. What we call “the way things are” is a working model — a scaffold of values and assumptions laid down since childhood, so familiar that we mistake it for the world itself.

That alone is a decent dinner-party fact. What makes it urgent is the step I used to skip. The brain behaves this way because it is a predictive network: a mesh of connections that runs on priors, compresses ambiguity into a usable story, and acts before it has finished checking. An organization is also a network — of people, processes, relationships, incentives, and the things colleagues tell each other in the corridor rather than the meeting. Feed any network of that kind a stream of ambiguous input and a set of inherited priors, and it will do precisely what a brain does. It will predict first, verify later, and mistake the prediction for the world.

It is tempting to file this under culture. Culture is the visible layer: how we behave, what we reward, who gets promoted, what is safe to say out loud. What I am

pointing at sits underneath it. The question is not only how we react to the numbers on the dashboard — it is which numbers made it onto the dashboard at all. What we went looking for and what we passed over. What got compressed out of a messy situation on its way into a two-page summary. By the time the slide exists, the interpretation has already happened, invisibly, and we experience the result as data. Every company is, in this sense, already a belief system wearing an org chart. Nobody wrote it down. Everybody is fluent in it.

Now introduce the machine — and here I want to be precise, because the imprecise version of this argument is the one people rightly dismiss.

Almost no organization will train an AI. The frontier models — Claude, ChatGPT and their peers — are closed-weight and general-purpose: highly capable, broadly reasonable, and, out of the box, not yours. Nobody in that Amsterdam room was going anywhere near a model's internal weights. What organizations actually build is the harness around the model: the retrieval, the tools, the permissions, the policies, the worked examples, the objectives. The company context. That harness is where the organization lives, and it is entirely, unavoidably, the organization's own work.

This is the good news and the bad news in a single sentence. Good, because the direction is yours to set — the model is not driving. Bad, because a harness assembled from an organization's own history — its emails, its past decisions, its escalation patterns, its record of what got approved and what quietly died — reimports the true character of that organization at full fidelity. Unless it is built consciously, the harness will instruct a very capable general-purpose system to act on the bad habits alongside the good, and it will not flag the difference. It has no way to know there is one.

Which is why the mirror metaphor only takes you halfway. A mirror is passive; nobody was ever harmed by a reflection. This thing acts. Point a tireless, capable system at a constructed reality that nobody has examined and you have not built a mirror. You have built an amplifier.

Hiring is the example I keep returning to, because everyone recognises themselves in it. The old system was a manager skimming CVs at six in the evening — tired, hungry, and, intentionally or not, favouring a certain type of person: broadly, the one whose CV looked like his own. A real problem, but a slow-moving one, bounded by how many CVs a tired person can read. The new system does exactly the same thing. It

just never gets tired and never goes home. It acted on the behaviour it learned, then multiplied it — a thousand times a day, with perfect consistency and an audit trail that makes the whole thing look like rigour.

If, like me, you have spent the better part of two decades inside organizational belief systems — sitting in these rooms, noticing the gap between what people say in the workshop and what they do on the Tuesday afterwards — none of this lands as news. My colleagues and I have argued for years that culture is not the soft accompaniment to transformation; it is the mechanism of it. Leadership, method and skilled people compound each other. They multiply, they do not add. Leave one of the three out and the sustainable result is not merely reduced. It is zero.

What is new is where that argument now points. For years, “know your organization’s blind spots” was a leadership-offsite platitude, filed somewhere between the trust falls and the mission statement. It is considerably harder to file away once you accept that those blind spots are about to be encoded into a system that acts with autonomy, at speed, on your behalf, at a scale no group of humans has ever operated at.

One tool I use for surfacing these hidden operating truths comes from Richard Barrett, whose Seven Levels of Consciousness model maps how individuals and organizations move from survival and self-protection, through belonging and achievement, toward service and contribution. The most important thing to say about it is what it is not: it is not a ladder of virtue. The foundational levels — safety, relationship, self-esteem — are not immature. They are load-bearing. The work is never to skip them. It is to notice when an organization is operating from fear at a level it no longer needs to, and to have enough awareness to choose otherwise. That is the entire mechanism, and it is why the model earns its place in this argument: you cannot consciously change a truth you don’t know you are standing inside.

Which brings us back to something almost old-fashioned — authenticity, treated as a discipline rather than a personality trait. Being your fully conscious, examined self was always a reasonable thing to ask of a leader. It becomes load-bearing infrastructure in a world where machines get better by learning from you. If the amplifier is going to be built regardless, the only real choice left is what it amplifies. An organization that has done the work — named its actual values, distinguished them from its stated ones, got honest about which level of consciousness runs its Tuesday-morning decisions — hands its AI a truer picture to learn from. An organization that hasn’t hands it a better-produced version of its own blind spots.

There is a version of this story that reads as alarming, and it isn't wrong to feel that. But the people who do this work for a living tend to describe it, surprisingly, as hopeful — even joyful. Their language keeps circling back to something like delight: the joy of seeing what really is. Organizational self-examination turns out to be less a reckoning than a workout. Uncomfortable during, clarifying after.

And there is a second, harder-edged reason for hope, which is the one I would press on any executive team. An agentic system transforms an organization on the basis of what has already worked — a learned dataset of past successful cases. That is genuinely valuable, and it is also, by construction, backward-looking. It is very good at the last problem. The organizations that will matter are the ones already working on the challenge nobody has named yet, and that work is simply not available to a system trained on the past. It requires human wisdom: judgment exercised under real uncertainty, by people who can see their situation clearly enough to read it. Authentic clarity is not just the honest option. It is where genuine innovation starts. By the time your AI has finished optimising yesterday, the clear-eyed organizations have already cracked tomorrow.

Marieke found me at the end of that grey Rotterdam Thursday, coat already on, and asked whether the question I had opened with actually had an answer. It does, though not a comfortable one. Inside an organization, what is real is whatever the organization has agreed to see. That used to be a philosophical curiosity, the sort of thing you turn over on the train home. It is now a specification — and you are writing it either way, this quarter, whether or not anyone in the room has read it.

BEFORE THE ROADMAP GETS SIGNED

So here is what I would actually do about it, before the roadmap gets signed. First, write down the three values your organization genuinely rewards — not the ones on the wall — and hand that list to whoever is building the harness. Second, audit the history you are about to feed the machine as though it were a policy document, because that is exactly what it becomes. Third, before you automate any decision, make someone say out loud what a good version of that decision looks like and why; if the room can't agree, you have found the work.

The task in front of leaders is not to out-engineer the machine, or to slow it down until humans catch up. It is to become, deliberately, the kind of organization worth learning from — because the learning is happening either way, and it has already started.

The one-line version, the kind that ends up on a slide because it is the only sentence anyone still remembers by Friday, goes like this:

AI won't hallucinate randomly. It will hallucinate exactly like we do — and then act on it, tirelessly, at a scale we never could. You get the AI you deserve. The only question left is how honest we are willing to become before it finishes learning.

Roland Hess is Founder and Managing Director of The Trusted Advisors and an Executive Performance Coach.
The Trusted Advisors · CH-6300 Zug · 45276 Essen · thetrustedadvisors.global