

BEYOND THE ARTEFACT

Automating the Automators: Generative AI and the Future of the Transformation Profession

Why the real threat is not that the models replace practitioners, but that they reveal how much of the work was never craft to begin with

Giovanni Leonardi

October 2023



Contents

1. Executive Summary	3
2. The question every practitioner is quietly asking	3
3. What the models are genuinely good at	4
4. The ceremony problem	5
5. The help — and why organisations will probably waste it	6
6. The apprenticeship problem	7
7. So, help or threat?	8

Automating the Automators: Generative AI and the Future of the Transformation Profession

Transformation · No. 357

Originally written October 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The models have not come for our judgement; they have come for our ceremony — and in doing so they have shown how much of the trade was ceremony all along.”

1. Executive Summary

A year on from the arrival of the first genuinely capable general-purpose language models, the transformation profession is asking, quietly and not always out loud, whether the technology it is now being paid to deploy is a gift or a threat to its own trade. This essay argues that the question, as usually posed, is the wrong one.

The models are exceptional at producing the visible artefacts of transformation — the target operating model deck, the benefits map, the stakeholder narrative, the requirements catalogue — and consistently poor at the work those artefacts were only ever a proxy for: judgement under ambiguity, the reading of organisational politics, and the discipline of deciding what not to do. The threat is therefore real, but misplaced. It falls not on the profession but on the ceremony inside it — the large share of the trade that was always artefact-production dressed as insight.

The help is equally real, and equally conditional. The same collapse in the cost of the artefact layer can either free practitioners for the judgement that remains stubbornly human, or simply industrialise the ceremony — faster, glossier, and with even less thought behind it. Which of the two an organisation gets will be settled not by the model but by the profession's own honesty about what its work was ever for. And there is one problem underneath both readings that neither the enthusiasts nor the doom-mongers have answered: if the machine now does the apprentice's work, where will the next generation's judgement be formed?

2. The question every practitioner is quietly asking

Somewhere this week, in a programme you will never hear about, a transformation lead pasted a rough brief into a chat window and watched a forty-slide operating model

assemble itself in the time it took to fetch a coffee. Structure, headings, a phased roadmap, a RACI, a benefits logic, even a tidy set of risks — all coherent, all confidently expressed, all delivered before the kettle had boiled. Work that a two-person workstream would once have carried for the better part of three weeks now arrived, in draft, before lunch.

The lead felt two things at once, and most of us have now felt them too. The first was exhilaration: this changes everything I do. The second, arriving a beat later and harder to admit, was a colder thought: if a machine can produce this, what exactly was I being paid for?

That second thought is the real subject of this essay. The technology that appeared, seemingly from nowhere, at the end of last year — the conversational assistant that reached a hundred million people faster than any consumer tool before it — did not announce itself as a threat to consultants and change leaders. It announced itself as a threat to writers, coders, illustrators, customer-service teams. But the transformation profession has always occupied an awkward position: it is the trade whose entire purpose is to reshape other people's work, to automate, to restructure, to render roles redundant in the name of a better operating model. There is a particular vertigo in watching the tool turn, for the first time, toward the people who wield it. The automators are being automated.

The reflexive answers came quickly, and both are too neat. The vendors and the optimists offered the story of pure augmentation: a co-pilot, a productivity multiplier, drudgery removed so that the practitioner can rise to higher work. The anxious offered the story of hollowing: the decks write themselves, the analysis writes itself, the junior ranks empty, and the profession thins from the bottom up. This essay's claim is that both stories are describing the same event from opposite ends, and that neither has looked closely enough at what the work actually is.

3. What the models are genuinely good at

Begin with an honest inventory, because condescension toward the technology is its own kind of denial. The models are not a gimmick. On a wide band of tasks that make up the daily texture of transformation work, they are already better than competent.

- They draft. Give one a messy brief and it will return an operating-model narrative, a communications plan, a set of guiding principles, a first-cut benefits map — structured, literate, and usually a stronger starting point than a tired human produces at four in the afternoon.
- They restate. Feed one a two-hundred-page strategy and it will compress it to a page, re-pitch it for a different audience, or turn a wall of prose into a briefing without losing the thread.
- They generate at volume. Ask for two hundred user stories from a rough process description and they arrive overnight, plausibly formed and consistently phrased.
- They never tire, never sulk, and never tell you the request is beneath them.

Anyone who has run a workstream knows how much of its cost sat precisely here — in the production, formatting, and endless re-cutting of documents that had to exist before

anyone could decide anything. That cost is now collapsing. To pretend otherwise, to wave the models away as stochastic parrots unfit for serious work, is to repeat the mistake every profession makes at the edge of its own disruption. Something structural is happening.

But notice what every item on that list has in common. Each is a transformation of things that were already written down — a brief, a strategy, a process description, a template. And that is the exact boundary of the machine's competence.

What the models do well	What they cannot touch
Draft the operating-model deck, the benefits map, the comms	Decide which of five true problems matters this quarter
Turn a rough brief into a requirements catalogue	Read the room that will quietly veto it
Summarise, restructure, and restate what is documented	Weigh the history that was never documented
Produce a coherent plan at inhuman speed	Own the consequence when the plan meets the organisation

Return to that forty-slide operating model that assembled itself before lunch. It was coherent. It was literate. And it was wrong in the single way that mattered: it proposed a shared-services structure that the board had politically killed eighteen months earlier, after a bruising fight that left scar tissue in three directorates. Nobody had written that decision down in any document the model could have seen; it lived in the memory of about nine people and in the careful way they now avoided the subject. The machine produced forty flawless slides and was missing the only one that would have got the programme shot on sight.

This is not a defect to be patched in the next release. It is the shape of the thing.

“A model can only ever know what the world has written down. Transformation is the management of everything the organisation has decided not to say.”

The models interpolate brilliantly across the vast corpus of what humanity has documented. But the material of our trade is disproportionately the undocumented: the private veto, the sponsor's real motive, the reorganisation that failed so badly six years ago that no one will countenance its vocabulary again, the two executives who cannot be in the same governance forum. A practitioner's value has never been the deck. It has been knowing which deck would get someone fired.

4. The ceremony problem

If the story ended there — machines handle the written, humans handle the unwritten, everyone keeps their job — this would be a comfortable essay. It does not end there, because of an uncomfortable truth the technology has dragged into the light.

A great deal of transformation work was never insight. It was ceremony. The RAID log maintained so that a log would exist, consulted by no one. The status report engineered, week after week, to read a shade greener than the programme deserved. The two-hundred-page strategy whose length was its argument, produced because a serious effort was expected to weigh something. The workshop whose real output was the feeling of progress. We know this. We have always known this. We simply had a professional interest in not saying it too loudly, because the ceremony was billable and the hours were real.

The models are merciless with ceremony, precisely because ceremony is pure artefact — form that has detached from function. A machine that produces the artefact directly, in an afternoon, for almost nothing, does not merely speed the ceremony up. It asks, implicitly and unavoidably, what the ceremony was for.

The artefact was never the point. It was evidence that thinking had occurred. The danger of a machine that produces the evidence directly is that it lets us skip the thinking and keep the evidence.

This is why the second thought our transformation lead had — what was I being paid for? — is the honest one, and why it should be welcomed rather than suppressed. For the fraction of the trade that was genuine judgement, the machine is no threat at all; it cannot do that work and shows no sign of learning how. For the fraction that was ceremony, the machine is not so much a threat as an exposure. The models have not come for our judgement; they have come for our ceremony — and in doing so they have shown how much of the trade was ceremony all along.

A profession confident in its craft would find this clarifying. A profession that has quietly depended on the ceremony being indistinguishable from the craft will find it frightening. Both reactions are, in their way, correct.

5. The help — and why organisations will probably waste it

Suppose we take the optimistic reading seriously, as intellectual honesty demands. The cost of the artefact layer falls to near zero. Practitioners are liberated from the drudgery of production and redirected toward the irreducible work: the judgement, the politics, the deciding. Programmes get cheaper and better at once. Is this not obviously the future?

It is a genuine possibility, and I do not want to argue it away. But everything I have watched organisations do with a productivity windfall argues against it. When a technology makes an output cheaper, organisations do not generally bank the saving as thought. They spend it as volume. Give a team a machine that produces decks ten times faster and the overwhelmingly likely result is not ten times the reflection per deck; it is ten times the decks. The strategy refresh that used to happen annually because it was

expensive now happens quarterly because it is cheap, and each iteration carries less conviction than the last. This is the productivity paradox in its newest dress: faster outputs, unchanged decisions — or worse decisions, made more often, with more supporting material and less actual deliberation behind each one.

Here the strongest objection deserves a hearing. You are overcomplicating it, a sensible colleague might say. This is just a better tool. The spreadsheet did not end financial analysis; it made analysts more productive. The word processor did not end writing. This is the same story, and the profession will absorb it as it absorbed those. The analogy is comforting and, I think, wrong in a specific and important way. The spreadsheet automated calculation — a task nobody mistook for judgement. What these models automate is the appearance of judgement: fluent, structured, confident output that looks exactly like the product of thought. That is a far more corrosive thing to make cheap, because in our profession the appearance of judgement has always been disturbingly easy to mistake for the real article — by clients, by sponsors, and, most dangerously, by ourselves. The spreadsheet never once fooled an analyst into thinking the analysis was done. This technology fools us daily.

So the help is real, but it is not automatic. It is a fork. Down one path, cheap artefacts free scarce human attention for the judgement that matters. Down the other — the better-trodden path, given how organisations actually behave — cheap artefacts simply flood the organisation with more confident, less considered material, and the practitioner's hardest new job becomes resisting the temptation to generate when they should be thinking. The technology does not choose the path. We do, and our track record is not encouraging.

6. The apprenticeship problem

There is one objection to the optimistic case that I cannot answer, and I distrust any account of this moment that pretends to. It is the strongest thing the pessimists have, and it deserves to be stated at full strength rather than in caricature.

For a generation, the profession has manufactured judgement in exactly one reliable way: by making junior people do the artefact work. You learned to see the organisation by drafting the hundredth status report and noticing, finally, which risks were real and which were theatre. You learned requirements by writing two hundred user stories and discovering, painfully, that the bottleneck was never the writing — it was deciding which thirty were worth building, and that discovery only came after the writing. You learned political judgement by producing the deck that got torn apart in the steering committee, and by watching why. The grind was not waste. The grind was the training. Judgement was the residue left behind by artefact work honestly done.

Now hand that grind to the machine. The two hundred user stories arrive overnight; the junior analyst never writes them, and so never has the discovery on the far side of writing them. The status reports assemble themselves; the graduate never sits with the discomfort of the number that will not go green. We will have automated not merely the drafts, but the apprenticeship that drafting was secretly providing.

We built judgement in our juniors by making them do the very work we are now handing to the machine. Automate the apprenticeship and you do not lose a generation of drafts; you lose a generation of judgement.

I have watched a capable analyst generate in an afternoon what would have taken her a fortnight, and I have watched the fortnight's incidental education simply fail to happen — the false starts she never made, the dead ends she never had to reason her way out of, the sponsor conversation the delay would have forced. The output was better and faster. The analyst, a year on, was not. This is the movement of the argument I am least able to resolve, and I will not insult the reader by resolving it in a paragraph. The optimists have no answer to it. The pessimists have no answer either, beyond nostalgia. The profession that works out how to form judgement without the drudgery that used to form it — how to teach the reading of a room when the room no longer needs the junior to draft the papers for it — will have solved something the technology's makers are not even thinking about.

7. So, help or threat?

The honest answer to the question this essay began with is that it depends entirely on what you believed the job was.

If you believed the job was producing artefacts — decks, documents, plans, catalogues, the tangible weight of a programme — then the machine is a threat, and a genuine one, because it produces those things faster and cheaper than you ever will, and it is improving while you read this. If you believed the job was carrying judgement and consequence — deciding what matters under ambiguity, reading the politics no document records, owning the outcome when the plan meets the organisation and the organisation pushes back — then the machine is help, and something better than help. It is a clarifier. It strips away the ceremony that let the two definitions of the job blur together, and it forces the profession to say which one it was practising all along.

*“Ask not whether the machine can do your job. Ask
what you believed your job was.”*

We have seen this shape of moment before, even if we have not seen this technology before. Boards are already asking "what is our AI strategy?" in the identical anxious register with which they asked "what is our digital strategy?" and, before that, "what is our cloud strategy?" — the same reaching for a noun to make the disquiet manageable. The hype will overshoot; a European regulator has already pulled the tool offline once over data questions, enterprise editions are arriving with promises that your data stays yours, and legislators are negotiating the first serious attempt to draw legal lines around it. Some of the loudest present claims will look foolish within eighteen months. That, too, is a pattern we know.

But underneath the cycle, something real has shifted, and it will not shift back. For the first time, the trade whose business is transforming other people's work has met a technology that transforms its own. The threat it poses is not extinction. It is exposure — of how much of the work was ever craft, and how much was only ever the performance of craft. The practitioners who thrive will not be the ones who resist the tool, nor the ones who surrender their judgement to it, but the ones honest enough to let it burn away the ceremony and stand on what remains.

Whether what remains is enough to sustain a profession — and whether we can still grow the next generation's judgement once the machine has taken the apprentice's pen — is the question we should be arguing about. It is a better question than help or threat, and it is the one the answer to help or threat keeps handing us.