

BEYOND DOCUMENT WORK

Generative AI Will Not Replace Transformation Professionals—But It Will Expose the Weak Ones

In 2023, the real contest is not human versus machine, but judgement versus administrative competence

Giovanni Leonardi

October 2023



Generative AI Will Not Replace Transformation Professionals—But It Will Expose the Weak Ones

Transformation · No. 356

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

“Generative AI does not remove the need for judgement; it raises the cost of pretending to have it.”

1. The scene behind the green status

At 8.30 on a Monday morning, a transformation director opens the weekly steering pack. Eleven workstreams have supplied updates. The pack runs to forty-seven pages, carries sixty-three risks and issues, and still fails to answer the question the executive committee will ask at noon: what has materially changed since last week?

A business analyst copies the workstream narratives into a generative AI tool. Seven minutes later, the machine has grouped the recurring concerns, exposed three contradictory claims about readiness, and drafted a sharper summary. The first reaction is relief. The second is unease. If a machine can do in minutes what occupied a capable professional for half a day, what exactly is the professional for?

That question has become difficult to avoid in 2023. Large language models have moved from technical curiosity to browser-accessible instruments at astonishing speed. Demonstrations are persuasive: draft a communication, summarise an interview, generate a workshop agenda, rewrite a business case. The temptation is to frame the development as a contest between human and machine.

It is the wrong contest. Generative AI is neither the saviour of transformation nor its executioner. It is a solvent. It dissolves the administrative work around transformation and reveals the quality of judgement underneath. That makes it an extraordinary help to a strong profession and a genuine threat to a profession that has mistaken production for value.

2. The profession we say we practise

Transformation professionals describe their work in the language of outcomes, choices and organisational change. Much of the working week, however, is spent producing representations of those things: plans, packs, minutes, process maps, requirement catalogues, risk narratives and communication drafts. These artefacts matter. The

problem is that their production has gradually become confused with the exercise of judgement they are meant to support.

The distinction is easy to miss because good administration can look like control. A complete risk register conveys diligence. A polished status report conveys confidence. A workshop with an impressive volume of captured notes conveys engagement. Yet none of these proves that the difficult decision was made, that the operating model is viable, or that the promised benefit can survive contact with line management.

Generative AI attacks precisely this zone of ambiguity. It is already competent at first drafts, synthesis, reformulation and pattern extraction. It can turn rough notes into structured prose, compare two versions of a process, propose categories for a risk set, and generate alternative explanations for a delayed milestone. The output requires checking, but so does work delegated to any junior team member. The economic question is not whether the machine is infallible. It is whether the combined cost of generation and review is lower than the old cost of manual production.

For many routine artefacts, the answer will be yes.

Generative AI does not remove the need for judgement; it raises the cost of pretending to have it.

3. Help, in the places where friction has become habit

The immediate opportunity is not autonomous transformation. It is the removal of avoidable friction from professional work. Three uses are already sufficiently plausible to change how a team should operate.

- Compression: reducing long interviews, workshop notes and workstream updates into candidate themes for human verification.
- Expansion: turning a clear decision or outline into draft communications, test scenarios, stakeholder questions or alternative formulations.
- Challenge: asking for contradictions, missing assumptions, adverse consequences and arguments against a proposed course.

The third use is the most important and the least celebrated. Most public demonstrations treat the model as a rapid copywriter. Transformation teams should treat it as an inexpensive dissenter. A programme office can ask the tool to compare the benefits register with the delivery plan and list benefits that have no credible enabling milestone. A change lead can ask it to identify groups implied by a new process but absent from the stakeholder map. A business analyst can ask for the strongest case against a requirement that everyone has stopped questioning.

Consider the steering pack again. The weak application is to produce a smoother executive summary from the same workstream prose. The stronger application is to force a chain of tests:

1. Extract every claim of completion or readiness.

2. Identify the evidence offered for each claim.
3. Mark claims supported only by self-report.
4. Compare dependencies named by different workstreams.
5. Draft the three decisions that the evidence appears to require.

The machine may make errors at every step. But it changes the economics of asking the questions. Work that was previously omitted because the team lacked time can now be attempted, reviewed and improved. The benefit is not that the model knows the programme. It is that it makes disciplined scrutiny cheaper.

4. Threat, where professional identity rests on output

The strongest opposing view deserves serious treatment. Generative models can invent facts, flatten nuance, reproduce bias and expose confidential information when used carelessly. They do not understand organisational history, informal authority or the personal cost of a restructuring decision. Their fluent prose can make a weak answer appear settled. On this view, transformation is too contextual and consequential to entrust to probabilistic text generation.

Much of that is correct. It does not follow that the profession is protected.

The same limitations apply unevenly across the work. No responsible leader should delegate a workforce decision, an investment recommendation or an assurance conclusion to a model. But refusing to delegate the drafting of interview questions because the model cannot own the final decision is like refusing a calculator because it cannot understand the business case. The relevant question is always: which part of the task is being transferred, and where does accountability remain?

Work	Suitable machine contribution	Irreducible professional contribution
Status reporting	Extract themes, compare statements, draft summaries	Test evidence, read incentives, frame the decision
Process design	Generate variants, identify common exceptions	Judge feasibility, authority and behavioural consequence
Change communication	Produce drafts for different audiences	Decide what can honestly be promised
Benefits management	Cross-check measures, milestones and owners	Challenge causality and secure operational ownership
Workshop preparation	Suggest questions, scenarios and structures	Read the room, change course and resolve conflict

The threat therefore falls first on roles whose value has been expressed mainly through document production and coordination. A professional who receives information, reformats it and passes it upward is operating in the model's strongest territory. A professional who determines which information is reliable, notices what is being withheld, and converts ambiguity into a decision remains in far stronger territory.

This is not a comforting distinction. Many organisations have trained transformation staff to be excellent at compliance with a delivery method and cautious about exercising judgement beyond it. Standardisation made large programmes easier to govern, but it also encouraged a class of work in which the template became the product. Generative AI will expose that bargain.

5. The return of judgement

The profession has spent years seeking legitimacy through method. We created clearer stage gates, richer reporting disciplines and more elaborate bodies of knowledge. These advances corrected real failures. They also invited an unintended belief: that repeatable process could substitute for experienced judgement.

Generative AI may reverse the emphasis. When a competent first draft becomes abundant, the scarce contribution is no longer composition. It is selection. Which assumption deserves attention? Which stakeholder is being treated as a communication audience when they are actually a source of operational veto? Which measure is a proxy that can be improved while the underlying outcome deteriorates? Which apparently technical dependency conceals a dispute over authority?

These are not mystical human qualities. They are developed through exposure to consequences, explicit reflection and the willingness to be accountable. Nor are they beyond assistance. A model can generate possible interpretations; it cannot decide which interpretation the organisation should act upon. That decision depends on evidence, context, ethics and appetite for risk.

The practical consequence is that professional development must change. Teaching people to produce the standard artefacts will remain necessary, but it will no longer be sufficient preparation for a career. Teams should cultivate four capabilities deliberately:

- Problem framing: defining the decision before commissioning the analysis.
- Evidence discrimination: separating observation, inference, assertion and generated possibility.
- Organisational reading: understanding incentives, informal power and the distance between declared and actual behaviour.
- Responsible challenge: making dissent specific enough to be useful and safe enough to be heard.

These capabilities were always important. The difference is that the surrounding production work can no longer hide their absence.

6. A new division of labour

The sensible response in 2023 is neither wholesale adoption nor defensive prohibition. It is controlled redesign of work. Start with a bounded activity whose inputs can be handled safely and whose output is already subject to human review. Measure elapsed time, correction effort and error type. Then decide whether the result should be adopted,

adapted or rejected.

A transformation office might choose the weekly synthesis of non-confidential workstream updates. For four weeks, it can run the existing process and an assisted process in parallel. Suppose the manual route takes six hours each week. The assisted route takes forty minutes to prepare inputs, ten minutes to generate candidate themes and ninety minutes for a senior reviewer to check source statements and rewrite the decision requests. Even if the machine draft is unusable in places, the team has recovered nearly four hours. The question then becomes what those hours purchase.

If they purchase another layer of formatting, little has changed. If they purchase two conversations with benefit owners, a review of an unresolved dependency and better preparation for the steering meeting, the technology has improved transformation practice. The value appears only when saved production time is deliberately reinvested in judgement and intervention.

The control model should be equally concrete:

- Classify information before it enters any external service.
- Keep source material and generated text visibly distinguishable.
- Require named human ownership for every consequential output.
- Test claims against primary evidence, not against the model's confidence.
- Record recurring error patterns so that controls evolve with use.

This is not glamorous governance, but it is the difference between experimentation and negligence. A blanket ban drives use into private, unobservable behaviour. Unbounded enthusiasm turns fluency into false authority. Controlled use creates evidence about where the tools help and where they mislead.

7. What the machine cannot carry

Transformation is ultimately an alteration of commitments. Budgets move. Roles change. Measures become consequential. A leader agrees to stop protecting one priority in order to fund another. A line manager accepts disruption now for a benefit that may arrive later. These movements depend on trust, legitimacy and the capacity to remain present when the answer is unpopular.

A language model can help prepare the conversation. It cannot bear its consequence.

That boundary matters because the hardest transformation failures are rarely failures of prose. They are failures to confront a trade-off, assign authority, remove an obstacle or admit that the original case no longer holds. Faster documents may even worsen such failures if they allow the organisation to circulate increasingly persuasive descriptions of a decision it continues to avoid.

The profession should therefore resist two forms of vanity. The first says our work is too subtle for machines, when much of it is repetitive synthesis. The second says the machine is already a strategist, when it has no stake in the result and no obligation to those affected. Both positions protect an identity. Neither improves the work.

8. Help or threat is a choice about the profession

Generative AI will help transformation professionals who use it to reduce clerical effort, widen challenge and spend more time at the point where evidence becomes decision. It will threaten those whose authority depends on controlling the production of information or on making familiar artefacts appear more valuable than they are.

The dividing line will not be technical fluency alone. Prompt-writing tricks will evolve, tools will change, and today's impressive demonstration may become tomorrow's routine feature. The durable advantage lies in knowing what to ask, what to distrust, what matters in context and what one is prepared to recommend.

We should welcome the pressure. A profession that claims to transform organisations cannot reasonably demand immunity from the same forces. The task is not to defend every activity we currently perform. It is to preserve and strengthen the contribution that only becomes visible when the activity is stripped away.

The machine will produce more words than we need. Our value will be measured by whether we can turn that abundance into fewer evasions, better choices and consequences honestly owned.