


BEYOND THE IDEA LIST

Use-Case Hunting Is How AI Programmes Avoid Choosing What Must Change

Behind the crowded idea lists, isolated pilots and the unresolved argument between central ambition and local adoption

Giovanni Leonardi

June 2022



Use-Case Hunting Is How AI Programmes Avoid Choosing What Must Change

Programme · No. 211

Originally written June 2022 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“AI value is not the value of a prediction. It is the value of a changed decision multiplied by the organisation’s ability to act on it.”

1. The Wall of Possibilities

The artificial-intelligence workshop ends with forty-seven use cases on the wall.

There are forecasts, recommendations, document classifiers, fraud alerts, maintenance predictions, customer prompts and automated decisions. Each has a coloured score for value, feasibility and data availability. Six are selected for further analysis. The executive sponsor is pleased: an abstract commitment to AI has become a visible portfolio.

Three months later, four of the six have no operational owner. One depends on historical data whose meaning changed during the pandemic. Another would optimise a process the organisation has already agreed to replace. The remaining use case becomes a pilot because a capable manager, an accessible dataset and an interested supplier happen to coincide.

The AI programme appears to have moved from strategy to delivery. In reality, it has moved from a technology ambition to a list of places where technology might fit.

This is the recurring weakness of “use-case hunting”. It looks practical, democratic and evidence-led. It invites the business to identify problems and prevents the central technology team from imposing solutions. Yet as a starting point for an AI programme, it often allows the organisation to avoid the harder decision: which part of the operating model is important enough to redesign around better prediction, classification or judgement?

A use case is not a strategy in miniature. It is a hypothesis about a decision. Without a change in ownership, process and consequence, it remains an interesting technical possibility.

2. Why the Hunt Became So Attractive

The attraction is understandable in 2022. Artificial intelligence has moved beyond research and specialist technology companies. Cloud platforms, packaged tools and a

growing supplier market have made machine learning, natural-language processing and computer vision more accessible. Boards want evidence that their organisations are not being left behind. Delivery teams need bounded work. Operational managers are more likely to engage with a specific problem than with a general data-and-AI strategy.

The use-case workshop resolves all these pressures at once.

It gives executives a pipeline, technology teams a mandate, suppliers an opportunity and business units a route to propose improvements. It also creates the appearance of bottom-up adoption: ideas emerge from people close to the work rather than from a remote centre.

But that appearance conceals the selection mechanism. The ideas most likely to survive are not necessarily those most important to strategy. They are those that can be described quickly, scored confidently and demonstrated with available data.

The hunt therefore favours:

- processes already measured, whether or not they matter most
- managers willing to sponsor experiments, whether or not they own the outcome
- datasets that are accessible, whether or not they represent the future process
- benefits that can be modelled, whether or not the organisation can realise them
- narrow decisions that avoid cross-functional change, whether or not they scale

The portfolio becomes a map of organisational convenience.

This is why so many AI programmes produce a curious mismatch: impressive local proofs and weak enterprise consequence. The pilots are not random. They have been selected by the path of least institutional resistance.

3. Top-Down Ambition and Bottom-Up Reality

The structural argument is often described as top-down versus bottom-up.

The top-down case says AI investment should follow strategy. Leaders identify the capabilities that matter, establish shared data and technology, and direct resources towards high-value outcomes. This creates coherence and avoids scattered experimentation.

The bottom-up case says useful applications cannot be designed from the boardroom. Operational experts know where judgement is slow, repetitive or inconsistent. Small experiments reveal what the technology can actually do. Adoption grows through demonstrated value rather than executive instruction.

Both cases are right about the other's weakness.

Top-down programmes can produce grand architectures, central platforms and priority themes with no route into daily work. They underestimate local knowledge and may force AI into decisions that do not need it. Bottom-up programmes can accumulate pilots that share neither infrastructure nor purpose. They optimise fragments while leaving the operating model untouched.

Use-case hunting is attractive because it appears to reconcile the two. Leaders set the ambition; teams supply the ideas. Yet the difficult connection between them is postponed. Nobody has to explain how a local prediction changes a strategic capability, or who will alter the process when the prediction is available.

The gap is filled with scoring.

A three-by-three matrix of value, feasibility and readiness gives unlike ideas a common language. But it also creates false precision. A forecast that could reduce inventory is compared with a classifier that could save handling time. Their estimates depend on different assumptions, time horizons and ownership. Multiplying scores does not make them commensurable.

The exercise converts judgement into arithmetic precisely where judgement is most needed.

4. The Pilot That Won the Scorecard

Consider a composite organisation with an AI opportunity register of 62 ideas. A central team reduces the list to eight using weighted scores: 40 per cent value, 30 per cent technical feasibility, 20 per cent data readiness and 10 per cent sponsor commitment.

The highest-ranked idea is a model to predict late deliveries. Historical records are plentiful, the outcome is clear and the potential cost reduction is substantial. The pilot achieves 81 per cent accuracy in identifying high-risk orders five days earlier than the existing process.

The result is celebrated. Deployment proves more complicated.

Operations already know that certain orders are at risk; the problem is that expediting capacity is limited and allocated by customer pressure. Procurement controls supplier contact. Customer service owns communications. Finance measures working capital. No single leader owns the decision the model is intended to improve.

The prediction enters a daily report. Teams discuss it, but the allocation rule remains unchanged. After six months, on-time delivery has improved by less than one percentage point. The model worked. The operating decision did not move.

A lower-ranked idea had proposed classifying service requests at entry. Its technical value appeared modest, but the service director owned the end-to-end process and was already redesigning team roles. That pilot reduced transfers between teams from 28 per cent to 17 per cent because the prediction arrived inside a process with authority to respond.

The scorecard selected the stronger model. The organisation gained more from the stronger change context.

This is the mechanism use-case hunting misses. AI value is not the value of a prediction. It is the value of a changed decision multiplied by the organisation's ability to act on it.

5. The Reasonable Defence of Discovery

There is a serious defence of the hunt. Organisations do not yet know where AI will prove useful. Requiring every experiment to trace directly to strategy can suppress curiosity and favour familiar problems. Early models often teach more about data and process than the original use case. A small local success can build confidence, capability and demand. In an uncertain field, exploration is not waste; it is how the organisation learns.

This is true. The alternative cannot be a central committee choosing “strategic AI use cases” from a distance and demanding certainty before evidence exists.

Nor should every pilot carry the full burden of enterprise transformation. Some experiments should be cheap, short and allowed to fail. Their value may be knowledge rather than deployment.

The problem is not discovery. It is discovery without a declared learning question or a decision about what follows.

A hunt asks, “Where could we use AI?” Governed discovery asks more demanding questions:

- Which decision or constraint are we trying to understand?
- What would we learn that changes an investment or operating choice?
- Who owns that choice?
- What evidence would justify scaling, narrowing or stopping?
- Which capability—data, process, skill or control—would be reusable afterwards?

The distinction appears subtle at the workshop and decisive six months later. One produces a backlog of possibilities. The other produces evidence tied to choices.

An AI pilot should earn the right to change a decision, not merely the right to become another project.

6. The Portfolio as Theatre

Use-case hunting persists because it serves a political function.

AI ambition is difficult to govern. Executives fear both overinvestment and inaction. Business units want access to funding without surrendering control. Technology leaders want common standards but cannot own every operational outcome. Suppliers need bounded opportunities. A use-case portfolio allows each party to participate without resolving these tensions.

The central team can report demand. Business units can demonstrate innovation. Executives can approve pilots rather than commit to operating-model change. When results disappoint, responsibility is dispersed: the data was not ready, adoption was weak, the model needed refinement, or the use case was not quite right.

The portfolio is busy enough to sustain belief and fragmented enough to avoid accountability.

This does not make the participants cynical. Most are responding rationally to the incentives around them. The business case funds a pilot more easily than a shared data capability. A manager can sponsor an experiment without accepting a change to service measures. A central team is rewarded for the number of ideas progressed. The structure produces the behaviour.

That is why better ideation methods do not solve the problem. The organisation does not lack ideas. It lacks a governing connection between strategic intent, operational ownership and technical evidence.

7. Start with Decisions, Not Ideas

A more useful starting point is not a list of use cases but a small number of consequential decisions or constraints.

Where does the organisation repeatedly lose value because it acts too late, applies inconsistent judgement, cannot see a pattern or spends scarce expertise on routine classification? Which of those decisions sits inside a capability the strategy genuinely needs? Who owns the operational response? What would improve if the prediction became available—and what would have to change for that improvement to occur?

Only then should teams explore whether AI is an appropriate mechanism.

This reverses the usual sequence:

Use-case hunt	Decision-led discovery
Starts with what AI could do	Starts with which decision must improve
Scores estimated benefits	Tests the mechanism from evidence to action
Selects available data	Examines whether data represents the decision
Finds a sponsor for the pilot	Identifies an owner for the operating consequence
Treats adoption as implementation	Designs process and authority with the model
Counts pilots progressed	Counts decisions improved or assumptions disproved

The model may not be the answer. Better rules, clearer data, process redesign or ordinary automation may solve the problem more cheaply. That is not failure. It is the discipline of refusing to turn AI ambition into a requirement that every search ends with AI.

8. What the Hunt Tells Us About Transformation

The deeper lesson extends beyond technology.

Transformation programmes often begin with solution-shaped ambition and then search for local justification. “Become digital” produces channel projects. “Move to the cloud” produces migration lists. “Adopt AI” produces use-case backlogs. The organisation mistakes the breadth of opportunity for clarity of purpose.

Use-case hunting reveals the gap between transformation intent and transformation reality because it shows where authority actually sits. Strategy may declare the destination, but local managers control data, process, measures and adoption. Central teams can fund experiments, but they cannot manufacture ownership. Suppliers can demonstrate capability, but they cannot decide which institutional trade-off matters.

The work of transformation is connecting these levels without pretending either can substitute for the other.

In 2022, experimentation is necessary. AI capability is developing quickly, and organisations need practical learning. But experimentation should not become an indefinite holding pattern in which the enterprise accumulates demonstrations while avoiding a decision about where it will change.

The right question is no longer, “How many AI use cases have we found?”

It is, “Which important decisions are we now prepared to redesign—and what have our experiments taught us about doing so?”

Until that question is answered, the crowded wall is not a strategy. It is evidence that the organisation has many ideas and has not yet chosen what transformation means.