

BEYOND THE PREFIX

Everyone Is an AI Company Now — Which Is Exactly Why the Label Is Worthless

Separating genuine machine intelligence from a marketing prefix — before you buy a transformation on it.

Giovanni Leonardi

February 2022



Everyone Is an AI Company Now — Which Is Exactly Why the Label Is Worthless

Transformation · No. 275

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

"The capability was real. The transferability was the fiction."

1. The demonstration that always lands the same way

Somewhere in the second half of the pitch — after the roadmap slide, before the pricing — the account team clicks to a screen where the software does something that looks like magic. A forecast that revises itself overnight. A support queue that sorts itself by intent. A form field that fills itself in from three systems the buyer did not know were connected. The word arrives a beat later, and it is always the same word. Powered by AI. The room leans in. The procurement lead writes it down. And somewhere near the back, the person who will actually have to run the thing after go-live feels a small, familiar unease they cannot quite name and, in that room, cannot afford to voice.

I want to give that unease a name, because it is correct. What the buyer is watching is real software doing real work. What they are hearing is a marketing decision. The two are no longer the same thing, and the whole difficulty of buying enterprise technology this year lives in the gap between them.

We have arrived at the moment when every software company has become an AI company. Not through invention — through renaming. The forecasting vendor who called itself an analytics platform eighteen months ago, and a cloud platform two years before that, now calls itself AI-native, and the underlying code has changed far less than the noun on the box. This is not fraud. It is something more ordinary and, for the buyer, more dangerous: it is a market in which the most valuable word in the language has been attached to almost everything, and has therefore stopped carrying information.

2. When a signal is on everything, it stops being a signal

The purpose of a label is to let a buyer discriminate — to tell one thing from another without opening both. A label works precisely because it is not on every product. The instant it is, it collapses into decoration.

That is where "AI" now sits. Walk any exhibition floor, read any category on any analyst's grid, open any inbound deck, and the term is present on the finance tool, the HR tool, the ticketing tool, the procurement tool, and the tool that schedules the meeting rooms.

When the descriptor is universal, its information content is zero. The buyer who screens for "AI capability" in a requirements document is, this year, screening for nothing at all — every serious vendor will tick the box, and most of them will be telling a version of the truth, because the definition has been stretched wide enough to admit a rules engine, a regression line, and a decision tree with a confident colour scheme.

The label tells you the vendor's marketing budget. It tells you almost nothing about whether the capability behind it will survive contact with your data, your process, and the people you have left to operate it.

The analysts saw this coming. Three years ago the now-familiar audit of European "AI startups" found that a striking share — roughly two in five — showed no meaningful evidence of machine learning in their actual products. The phrase AI washing entered the vocabulary shortly after. What has happened since is not that the washing stopped; it is that it moved upmarket, out of the startups and into the incumbents, who have every reason to reach for the word and the balance sheets to make it stick.

3. What actually changed, and what only got renamed

None of this means nothing changed. It means the change and the noise are running on separate tracks, and the buyer's job is to tell which track a given product is on.

Genuinely new, and genuinely useful: the cost of training and serving a model has fallen sharply as the cloud platforms turned machine learning into a metered service. Pretrained models for vision and language can now be reached through an API by an engineer who is not a research scientist. The operational discipline of keeping models alive in production — the plumbing we have started calling MLOps — is maturing from craft into practice. At the frontier, the striking research is in models that generate fluent text from a prompt, but that work still sits largely in the labs and the API sandboxes, not in the enterprise procurement catalogue. These are real advances, and a transformation programme is right to want them.

Re-badged, and mostly unchanged: the substantial majority of the "AI" features arriving in enterprise contracts this year are things the field has had for a decade or more. A threshold alert is now anomaly detection. A linear regression is now predictive AI. A scripted software robot clicking through a legacy screen is now intelligent automation. A branching decision tree answering the ten most common questions is now a conversational AI agent. Each of these can be worth buying. None of them is what the word in the demo implied, and the price has quietly been reset to match the word rather than the mechanism.

Consider the shape this takes in a real procurement. A programme replacing its demand-forecasting process is shown an AI-powered forecasting module and a headline: nine per cent more accurate than the incumbent. The number is true — in the pilot. What the slide does not say is that the pilot was tuned by two of the vendor's own data scientists over six weeks against a clean historical extract; that the "AI" underneath is a gradient-boosted model no more exotic than what the buyer's own analytics team could

stand up; and that when the specialists rotate off the account and the model meets live data with its gaps and its seasonal breaks, the nine per cent erodes to something closer to two — which the buyer discovers eleven months later, after the case study has already been written. The capability was real. The transferability was the fiction.

4. Why the goldrush recurs

It would be comfortable to blame the vendors, and lazy. The more useful account is that everyone in the transaction is responding rationally to the incentive in front of them, and the collective result is a distortion no single actor intends.

The mechanism is a chain, and every link is behaving sensibly on its own terms:

- The analysts create a category, because a named category is what their subscribers pay to be ranked within.
- Capital flows to the label, because the label is what commands the valuation multiple this year, and a company that does not claim it is punished for the omission.
- The incumbents must defend their position, so a cloud company that has not become an AI company looks, to the market, like it is falling behind — whatever its software actually does.
- The buyers reward the word, because "AI capability" has found its way into the requirements template, and so the safest thing a vendor can do is supply the word, whether or not it supplies the capability.

Read the chain in one direction and it manufactures hype out of ordinary self-interest, with no villain required. Read it in the other and you find the escape: the distortion is sustained by buyers rewarding the word. That is the one link the buyer controls.

5. The strongest case for the goldrush — and where it fails the buyer

The honest objection to everything above is that hype is not merely waste — it is a pump. The flood of capital and attention that a goldrush unleashes pulls genuine capability forward faster than sober demand ever would. It forces cautious incumbents to actually invest rather than talk. It funds the research whose by-products become next year's ordinary tools. And it drags useful machine learning into products that buyers, left to write their own specifications, would never have thought to ask for. Cynicism, on this view, is not shrewdness; it is a way of missing the upside while congratulating yourself on your realism.

Every part of that is true, and none of it helps the person signing the contract. The rising-tide argument is an argument about the market — about aggregate progress averaged across a whole sector over years. The buyer does not live in the aggregate. They live in one contract, for one tool, against one process, on one timeline. They do not get the sector's averaged-out benefit; they get the specific thing they bought, and the label they bought it on tells them nothing about whether this one, here, will work. That the goldrush is good for progress in general is entirely compatible with it being bad for you in particular — and confusing the two is exactly the error the demo is engineered to

produce.

*"The tide rising is a fact about the ocean. You are
buying a boat."*

6. How to buy through the noise

The discipline is not scepticism for its own sake — a programme that refuses every "AI" claim will miss the real advances alongside the re-badged ones. The discipline is discrimination: a small set of questions that the marketing prefix cannot answer for the vendor, and that separate the mechanism from the label without requiring the buyer to become a data scientist.

1. Ask what it did before it was AI. Every re-badged feature has a prior name. If the honest answer is "it was our reporting module," you now know what you are buying and can price it as such.
2. Ask to see it fail. Request a demonstration on your own messy data, not the vendor's clean extract — and specifically ask what happens at the edges, with the missing fields and the seasonal breaks. Confidence about the failure modes is the surest sign of a real capability; a vendor who cannot describe how their model degrades has not watched it run.
3. Ask who has to be in the room for it to keep working. The forecasting number that needed two of the vendor's specialists to achieve is a consulting engagement wearing a product's clothes. Find out whether the capability is in the software or in the people the software arrived with.
4. Ask for the baseline, and insist the pilot measures against it. "Nine per cent better" is meaningless without "than what, measured how, over what period." A vendor confident in the mechanism will welcome the baseline. A vendor selling the label will find reasons to avoid it.

None of these questions require you to understand the model. They require you to refuse to let the word stand in for the evidence — which is the one thing the goldrush is structured to make you do.

7. The honest account

The label will keep moving. A few years ago the universal prefix was cloud; before that it was digital; a while before that, e-. Each began as a genuine distinction, spread until it described everything, and then dissolved into the background noise from which the next distinction had to be carved. AI is further along that road than its champions admit and less far than its cynics claim, and by the time this particular word has worn smooth, another will have taken its place at the top of the deck.

What does not move is the buyer's task. Underneath every renaming, the question a transformation programme must answer has not changed at all: what does this system actually do, how will it behave when my reality meets it, and what will it still be worth when the people who sold it have moved on? The vendors have become AI companies. Our job is to stay buyers of capability — and to keep asking, of every dazzling demonstration and every confident prefix, the plain and slightly deflating question the whole performance is designed to make us forget to ask: yes, but what is it, really?