What Is AI?

AI Is Not One Thing: why the capability stack matters - and where generative AI fits

LinkedIn and kjerside.com companion essay | Companion 1 of 3 | 29 July 2026

In many business conversations, AI has become shorthand for generative AI: ChatGPT, Claude, Gemini or another language-based assistant. That is understandable, because generative AI made artificial intelligence visible to almost everyone. It is also misleading.

AI is not one thing because the work of intelligence is not one thing. A fraud model, a recommendation engine, a visual-inspection system and a language model all sit under the AI umbrella. They do not depend on the same inputs, fail in the same ways or require the same governance.

Bad categories lead to bad management. Treat every AI opportunity as a chatbot problem and the organisation misses prediction, optimisation, perception and automation. Call every dashboard or data lake AI and the word stops distinguishing anything.

AI is a category. Capability is the unit of decision.

The practical question is not: are we using AI? It is: which capability are we trying to make more scalable, in which workflow, with what evidence, verification and human ownership?

> *The practical question is not: are we using AI? It is: which capability are we trying to make more scalable, in which workflow, with what evidence, verification and human ownership?*

The five-part AI series makes the economic and organisational argument. This first companion supplies the technical map underneath: what AI is, which levels are being collapsed into one word, where generative AI fits, and how to forecast where capability moves next.

1. AI is machine-based inference

AI, at its broadest useful level, is software that works things out rather than merely executing instructions. Given inputs, an AI system infers outputs: predictions, classifications, recommendations, generated content, plans, decisions, tool calls or physical actions. Anything that does not infer is ordinary software, however sophisticated.

That definition is deliberately dry. It avoids two opposite mistakes: anthropomorphism, assuming AI must think as a person thinks; and dilution, treating any impressive software as AI. It is also how the OECD, the EU AI Act and NIST frame AI systems: machine-based systems that infer from inputs to generate outputs that may influence physical or virtual environments.[1-3]

The management implication is direct: define AI by what it does, not by the mystique around it. Big data, dashboards, algorithms and automated workflows may supply data, rules, infrastructure or execution. They are not machine inference. Appendix A works through the common boundary cases.

The first discipline is to separate inference from ordinary software - and capability from the model, product and workflow in which it appears.

2. Capability, model, application and workflow are different levels

The word AI is used to describe four different things at once. A capability describes what inferential work becomes scalable. A model describes how that capability is produced. An application packages models for a use case. A workflow arranges applications, evidence, people and controls into recurring work.

The distinction matters because decisions happen at different levels. A leader may need a forecasting capability, not "a language model". A developer may need a retrieval component, not "an AI strategy". A risk owner may need an approval gate around an action, not a better prompt. A language model is not a use case.

AI history is often told as waves: rules, machine learning, deep learning, generative AI, agents and robotics. In practice the waves accumulate rather than replace one another; yesterday's frontier becomes today's cheap, reliable component. The useful business map is therefore not a timeline but a capability stack. A fuller reference taxonomy appears in Appendix B.

Capability	What AI makes more scalable	Business examples	Management question
Perception	Turning images, speech, video, documents or sensor data into usable information.	Inspection, transcription, OCR, safety monitoring, chart and PDF interpretation.	What is captured, and what context is lost when reality becomes data?
Classification	Assigning cases, documents, customers or events to categories.	Risk flags, document tagging, claims routing, compliance triage.	Who defines the categories, and what happens to borderline cases?
Prediction	Estimating future or unknown outcomes.	Demand, churn, fraud, credit risk, maintenance.	Is the target measurable, stable and causally meaningful?
Optimisation	Choosing actions under objectives and constraints.	Scheduling, logistics, pricing, capital allocation.	Is the objective function the right one?

Recommendation	Ranking options or next-best actions.	Product suggestions, sales next-best action, content ranking, expert finding.	What incentives does the ranking system create?
Generation	Producing or transforming artefacts.	Drafts, code, summaries, images, synthetic data, structured outputs.	How are outputs verified, and who owns provenance?
Reasoning and planning	Structuring arguments, solving multi-step problems and sequencing actions.	Issue trees, scenario analysis, legal reasoning, project plans, task decomposition.	Has the problem been framed correctly, and are constraints explicit?
Interaction	Communicating with people through language, speech or multimodal interfaces.	Assistants, customer support, training, sales support.	Does the interaction build trust or create false confidence?
Action	Executing digital or physical steps.	API calls, coding and computer-use agents, workflow automation, robotics.	What permissions, audit and rollback mechanisms exist?

The table turns "we need AI" into a design question. A company rarely needs AI in the abstract. It needs better forecasting, faster classification, more scalable drafting or more reliable execution. Start with the work. Identify the capability. Define the evidence and quality threshold. Then decide which model, product or system pattern belongs in the solution.

The capability tells you what to test. The consequence tells you how hard to test it.

3. Generative AI is one layer - and the visible one

> *The capability tells you what to test. The consequence tells you how hard to test it.*

Generative AI is the capability to produce or transform artefacts: text, code, images, audio, video or structured outputs.[4] It is one row of the capability stack, not the whole of AI. What sets it apart is visibility: it is the first AI most people could see and use directly.

Earlier business AI produced useful but unobtrusive outputs: scores, forecasts, labels, rankings and alerts. The model sat inside a pricing engine or fraud system; most users never met it. Generative AI changed both the interface and the artefact: the interface became conversational, the output recognisable.

Foundation models amplified the shift: broad pretrained models can be adapted or instructed for many downstream uses.[5] Language models became the most visible branch because language is the operating surface of knowledge work. Organisations store themselves in emails, contracts, policies, transcripts, filings, reports, models, code and meeting notes.

Generative AI therefore matters for two reasons. It cheapens the production and transformation of represented artefacts. And it provides a general interface through which other capabilities can be composed. The visible layer is important; the system around it is what turns visibility into work.

4. The forecasting rule: capability grows where output can be checked

AI capability does not advance evenly across all work. It advances fastest where the work is written down and where the result can be tested. That single rule predicts more about where capability moves next than any vendor roadmap.

Four conditions make work receptive. Represented: the work exists in language, code, documents, tables, images or models. Pattern-rich: recurring structures can be inferred. Repeatable: similar work appears often. Reviewable: outputs can be checked, challenged or verified.

The last condition matters most. Frontier models increasingly improve through training on tasks with checkable outcomes: code that passes tests, mathematics with verifiable answers. The same logic applies in deployment. Capability is easiest to trust and scale when the surrounding system can test what the model produced.

That is why code and mathematics moved early, and why consulting, law and finance are moving now: issue trees, contracts, filings, models and memos are represented, pattern-rich and reviewable, even where the final judgement remains contested.

The rule is more discriminating than asking whether work is prestigious, creative or complex. Work with clear evaluation standards will often be automated sooner than lower-status work whose quality is difficult to observe. Use AI where cognition is represented and reviewable; keep people where salience, trust, legitimacy and consequence-bearing remain load-bearing.

5. From capability map to system design

A capability map is not yet an operating model. Real value rarely comes from one capability used alone: serious workflows combine several rows of the stack, with

> *Use AI where cognition is represented and reviewable; keep people where salience, trust, legitimacy and consequence-bearing remain load-bearing.*

rules, evidence and human decision rights around them. The map is the vocabulary; system design is the work.

A demand-planning process may use prediction to forecast, classification to segment, retrieval to supply context, generation to explain and rules to enforce policy. Generation may be the visible interface; the value lies in the composition.

The same is true of agents. An agent is not a new species of intelligence. It is a system pattern: a model inside a software loop that plans, calls tools, observes results, revises and continues.[6] When a system only drafts, a category error creates a poor answer. When it can act, the same error can change a record, commit money or expose data. The design question moves from what the model can say to what the system is allowed to do.

A practical discipline follows: before naming a provider, write down five things. The input: what enters, and what representation loses. The capability: which row of the stack is being scaled. The acceptance test: what counts as good enough. The failure consequence. And the authority boundary: inform, recommend or act. Those five lines expose most category errors early - and they explain why model comparisons usually arrive too soon.

The next companion follows how a model becomes an application and then a governed system; the third explains how that system should choose among models and providers as the frontier moves.

The strategic question is no longer "which model should we use?" It is: which capability are we augmenting, what evidence does it need, what must be verified, what may be automated and what must remain human-owned?

AI is not one thing. The next question is how the pieces become a system.

Appendix A: Boundary cases - what is not AI?

A useful taxonomy needs boundaries. Otherwise AI becomes a label for every data product, dashboard, algorithm or automated workflow.

Category	Not AI by itself	When it becomes AI-enabled
Big data	Large datasets, data lakes and extensive information stores.	When used by systems that infer, classify, predict, generate, recommend or act.
Analytics and BI	Dashboards, reporting and descriptive analytics.	When AI models generate forecasts, classifications, explanations or recommendations.
Algorithms	Clearly specified computational procedures or rules.	When they form part of systems that infer, learn, generate, plan or act.

Automation / RPA	Scripted process execution or task routing.	When automation uses classification, NLP, vision, LLMs or adaptive decisioning.
Cloud computing and databases	Compute, storage, networking and data retrieval infrastructure.	When used to train, serve, monitor, ground or govern AI systems.
Knowledge graphs / simulations	Structured representations of entities, relationships or system behaviour.	When used for reasoning, recommendation, semantic search, prediction or optimisation.
Workflows and robotics hardware	Process routing, approvals, sensors, actuators and physical machinery.	When AI performs perception, planning, extraction, reasoning, prediction or action inside them.

The point is not to narrow AI artificially. It is to keep the word useful. A dataset is not intelligence. A dashboard is not inference. A workflow is not judgement.

Appendix B: Detailed taxonomy for reference

The main article uses a capability lens. The reference taxonomy below separates techniques, model classes, application fields and system patterns. The categories overlap and are not meant to be mutually exclusive.

Segment	What it does	Type
Machine learning	Learns patterns from data and applies them to new cases.	Technique / paradigm.
Deep learning	Uses multi-layer neural networks to learn complex representations.	Model class / technique.
Natural language processing	Processes, classifies, extracts, translates, summarises or generates language.	Application field.
Computer vision / speech AI	Processes images, video or speech.	Application field / capability.
Robotics and embodied AI	Combines perception, planning, control and actuation.	System / application area.
Generative AI	Creates or transforms content.	Capability / model-application class.
Foundation models	Large pretrained models adaptable to many downstream tasks.	Model class / infrastructure layer.
LLMs, reasoning and multimodal models	Operate over language, code and other modalities; some are optimised for multi-step problem solving.	Model class / product category.

Agents, RAG and external memory	Combine models with tools, retrieval, APIs, state, permissions and workflows.	System pattern / infrastructure.

The discipline is to label the level correctly. Generative AI, language models, retrieval-augmented generation and agents are all important. They are not the same kind of thing. That distinction is the difference between an AI strategy and a vocabulary accident.

Selected evidence and references

1. OECD, What is AI? The updated definition of an AI system, reflecting the definition approved in November 2023 and the Principles updated in May 2024.

2. Regulation (EU) 2024/1689, Article 3(1), definition of an AI system.

3. NIST, Artificial Intelligence Risk Management Framework 1.0, NIST AI 100-1, January 2023.

4. NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, July 2024.

5. Bommasani et al., On the Opportunities and Risks of Foundation Models, Stanford Center for Research on Foundation Models, 2021.

6. Yao et al., ReAct: Synergizing Reasoning and Acting in Language Models, ICLR 2023.

Evidence and inference. The official definitions establish the breadth of the AI umbrella and the generative-AI category. The capability stack, boundary cases and forecasting rule are the author's managerial synthesis: useful because they expose different tests, decisions and consequences, not because the categories are exhaustive or mutually exclusive.

Disclosure

This article was AI-assisted. The argument, judgement and conclusions are my own.

Jesper Kjerside is co-founder of *taonta*, an AI-native advisory firm, and *fastland*, the software company behind it.