

Artificial intelligence: what actually *contains* what

The popular version of this diagram draws four concentric rings — AI & ML, generative AI, AI agents, agentic AI — and calls them layers. Three of those rings are genuine subsets. One is not a subset at all. **Containment and maturity are different claims, and the ring version quietly swaps one for the other.**

READ OUTSIDE IN. EACH BOX IS **FULLY INSIDE** THE BOX AROUND IT.

Artificial Intelligence

THE SUPERSET

Systems that perform tasks we associate with human reasoning. **Not everything in here learns.** A large share of AI running in production today is a rules engine or a search algorithm.

Symbolic & rules-based AI

OMITTED FROM THE RING VERSION

Behaviour written by humans as logic, not learned from data. Sits beside machine learning, not beneath it.

ALGORITHMS	A* and Dijkstra search · minimax with alpha-beta pruning · constraint satisfaction · forward and backward chaining · PDDL planners
USE CASES	Route and fleet optimisation · shift scheduling · tax and benefits determination · underwriting rules · chess and game engines · clinical decision support
TOOLS	Google OR-Tools · Drools · CLIPS · Neo4j and other knowledge graphs

Machine Learning

SUBSET OF AI

Learns parameters from data instead of following written rules. **Still where most enterprise value sits.**

SUPERVISED	Linear and logistic regression · decision trees · random forest · gradient boosting (XGBoost, LightGBM) · support vector machines · k-nearest neighbours · naive Bayes
UNSUPERVISED	k-means · DBSCAN · hierarchical clustering · PCA · t-SNE · isolation forest · association rules
REINFORCEMENT	Q-learning · SARSA · DQN · PPO
USE CASES	Credit scoring · churn and demand forecasting · fraud and anomaly detection · recommendations · predictive maintenance · pricing
TOOLS	scikit-learn · XGBoost · Databricks · SageMaker · Vertex AI · Azure ML

Deep Learning

SUBSET OF ML

Multi-layer neural networks that learn their own features rather than being handed them.

ARCHITECTURES	Multilayer perceptron · convolutional networks (ResNet, YOLO) · recurrent networks (LSTM, GRU) · graph neural networks · autoencoders · transformers
USE CASES	Defect detection on a production line · medical imaging · OCR and document extraction · speech recognition · translation
TOOLS	PyTorch · TensorFlow · JAX · Hugging Face

Generative AI

SUBSET OF DEEP LEARNING

Learns a distribution, then samples new output from it. **The innermost box, not the outer one.**

ALGORITHMS	Transformer decoders (the LLM family) · latent diffusion · denoising diffusion · generative adversarial networks · variational autoencoders · flow matching
USE CASES	Drafting and summarising · code generation · image, video and audio synthesis · retrieval-augmented question answering · synthetic test data
MODELS	Claude · GPT · Gemini · Llama · Mistral · Stable Diffusion · Midjourney · Whisper

SEQUENCES ANY BOX ON THE LEFT

Workflows

NOT A SUBSET

A model called at fixed points along **a path a person wrote in advance**. The sequence is decided by code; the model fills in the steps. Most of what gets sold as agentic today are workflows.

MADE OF	A model · a predefined code path · triggers and schedules · validation and approval gates · retries and fallbacks
PATTERNS	Prompt chaining · routing and classification · parallelisation (sectioning, voting) · orchestrator and worker · evaluator and optimiser · retrieval pipelines
USE CASES	Invoice extraction into validation into posting · document classification and routing · drafting behind a review gate · alert enrichment · nightly reporting · KYC checks
TOOLS	n8n · Zapier · Make · Airflow · Temporal · AWS Step Functions · Power Automate · LangGraph

A workflow fails *where you put the gate*. An agent fails somewhere nobody wrote down.

WRAPS ANY BOX ON THE LEFT

Agents & agentic systems

NOT A SUBSET

An **architecture built around a model**, not a smarter class of model. The same generative model becomes an agent when you give it tools, memory, a planning loop and write access. Nothing about the model changed.

MADE OF	A model · tool and API access · memory or state · a planning loop · a permission set · an audit trail
PATTERNS	Function and tool calling · ReAct · plan-and-execute · reflection and critique loops · single agent versus multi-agent orchestration · Model Context Protocol
USE CASES	Ticket triage and resolution · CRM and record updates · contract review pipelines · code migration · SOC alert triage · outbound sales sequencing
TOOLS	LangChain and LangGraph · CrewAI · AutoGen · Semantic Kernel · n8n · Claude Code · Codex · Copilot Studio

Autonomy is a permissions decision, *not a capability tier*. A weak model with write access is more dangerous than a strong model with read access.

What the four-ring version gets wrong

- 1. It puts generative AI outside machine learning.** Generative AI is deep learning, which is machine learning. Drawing it as a wider ring reverses the containment.
- 2. It implies all AI learns.** The rules engine deciding your insurance excess is AI and has no training data anywhere in it.
- 3. It treats agents as a fourth stage of intelligence.** Same models. More permissions. The jump is in blast radius, not in reasoning.
- 4. It files governance as a term inside the outer ring.** Governance is not the last thing you reach on the way out. It runs through every box, at every level.

Human capability

Named owners who hold launch-hold, budget and escalation authority — not a control point in a diagram.

The ground: Baseline AI literacy sits under all three columns and cuts across all of them. Without it the other three become paperwork.

Governance

Decision rights, oversight and measurement. Regulatory alignment is a sub-obligation inside this, never the point of it.

Supply chain

Models, data, weights, tooling and agent identity — everything you did not build and still answer for.

FOLLOW FREDRIK LINDSTROM FOR MORE

Fredrik Lindstrom

AI and cybersecurity governance for boards, CISOs and senior leaders. 25+ years in enterprise security.

A practitioner to guide you.

ARTICLES & BOARD ADVISORY

fredriklindstrom.info

Long-form governance analysis and toolkits

LINKEDIN

in/fredriklindstrom

Daily commentary and article launches

YOUTUBE

@Promise-and-Risk-AI

AI education without the hype or fear

COMPANION SITE

promiseandrisk.ai

Episode notes, sources and explainers

NEWSLETTER

The Governance Memo

Monthly, board-facing. Subscribe at fredriklindstrom.info