cprime AI LEARNING SERIES

AI Cheat sheet

Your plain-English guide to AI — from the fundamentals

to agentic systems, data strategy, and responsible deployment.

Evolution of AI

Models vs Systems

Data & RAG

Fine-Tuning

AI Risks

Agentic AI

01 THE EVOLUTION OF AI

From Rules to Reasoning

ARTIFICIAL INTELLIGENCE

1950s+

The broadest category: any computer system that mimics human intelligence to perform tasks — including rule-based systems, vision, robotics, and voice recognition. Early AI used hand-crafted rules; modern AI learns from data. Think of AI as the umbrella — everything else lives beneath it.

e.g. Spam filters, chess engines, smart assistants, fraud detection

MACHINE LEARNING (ML)

1980s+

ML is a subset of AI where systems learn patterns from data rather than explicit rules. Train it by feeding examples — it figures out the patterns itself. The more quality data you provide, the more accurate the predictions. No human needs to hand-code every rule.

e.g. Demand forecasting, churn prediction, anomaly detection, recommendations

LARGE LANGUAGE MODELS (LLMs)

2017+

LLMs are the specific type of GenAI trained on language. They understand intent, context, tone, and nuance — enabling natural conversation. They work by predicting the most probable next word (token). They don't 'know' facts — they predict fluent responses.

e.g. ChatGPT, Claude, Gemini, Copilot — Q&A, summarisation, drafting, translation

GENERATIVE AI (GenAI)

2020s

A leap forward: instead of classifying or predicting, GenAI creates new content — text, images, code, audio, and video — by learning patterns from billions of examples. The shift is from 'what will happen?' to 'generate something new.'

e.g. Drafting documents, writing code, generating images, summarising reports

AGENTIC AI

Emerging

The frontier: AI that doesn't just answer questions but plans and executes multi-step tasks autonomously. Agentic systems use tools, call APIs, browse the web, write and run code — working toward a goal with minimal human intervention. Governance is critical.

e.g. Automated research workflows, code deployment agents, multi-system orchestration

KEY INSIGHT:

Small / specialized models (SLMs) — many enterprises now deploy smaller models for cost, speed, and privacy (e.g., Llama variants, Mistral, Phi).

02 AI MODELS vs AI SYSTEMS

Understanding the building blocks

When people say 'AI', are they talking about the model itself or the full application built around it? These are very different things — and the distinction matters for deployment, governance, and results.

THE AI MODEL	THE AI SYSTEM
WHAT IT IS	
The trained neural network — a mathematical system that has learned patterns from training data. Takes an input (your prompt) and generates an output (text, code, images).	The full application: a model plus surrounding infrastructure — APIs, data connections, retrieval layers, access controls, user interfaces, and guardrails.
WHAT IT KNOWS	
Only what was in its training data, up to its knowledge cut-off. No awareness of your company, files, or events after training unless provided to it.	Whatever you connect to it. A well-built system integrates your documents, databases, CRM, and policies — giving the model the context it needs.
HOW IT REASONS	
LLMs don't 'think' — they predict statistically likely next tokens. They can sound confident while being factually wrong. Pattern-matching, not truth-checking.	Systems include logic layers: routing prompts to the right model, injecting context via RAG, enforcing permissions, filtering harmful content, and audit logging.
ANALOGY	
A brilliantly educated person locked in a room with no phone, no internet, and no knowledge of your organisation.	The same person — given access to your documents, a clear job description, and a manager who reviews and overrides their decisions.

KEY INSIGHT:

Choosing a model is only the first 10% of the work. Building the system — with the right data, guardrails, and governance — is the other 90%.

03 DATA, RAG & FINE-TUNING

How AI gets smarter and more relevant for your context

TRAINING DATA

Foundation layer

Training data is the raw material of every AI model. A model learns by processing billions of examples — books, websites, code, scientific papers — adjusting its internal parameters to capture patterns. The quality, breadth, and recency of training data directly determines what the model can do.

Rule of thumb: garbage in, garbage out. Biased, outdated, or narrow training data produces biased, outdated, or narrow models.

RAG — RETRIEVAL-AUGMENTED GENERATION

Most common approach

RAG is the most widely used technique to make a general AI model useful in your specific context — without retraining it. The system first searches your knowledge base (documents, emails, databases) to find relevant content, then feeds that as context into the model's prompt. The model generates a response grounded in your actual data.

Why RAG matters: keeps your AI current without retraining, respects access permissions, and provides citations so users can verify answers.

VECTOR DATABASES:

The engine behind RAG. Documents are converted into numerical representations (embeddings) that capture meaning — not just keywords — enabling semantic search: finding content by meaning rather than exact words.

FINE-TUNING

Advanced technique

Fine-tuning takes a pre-trained model and continues training it on a smaller, specialised dataset — teaching it your terminology, tone, style, and domain knowledge at a deeper level than RAG. Like sending a knowledgeable employee on an intensive industry course. The result: a model that has internalised your way of working.

Use RAG when:

Content changes frequently, you need citations, or want to stay current without retraining.

Use Fine-Tuning when:

Style/tone consistency is critical, domain terminology is specialised, or RAG alone doesn't deliver quality.

04 AI CHALLENGES & RISKS

What every leader needs to understand before deploying AI

HALLUCINATION

AI models confidently generate text that sounds authoritative but is factually wrong. LLMs predict probable word sequences — they don't verify facts. An AI may invent statistics, misattribute quotes, or describe non-existent products.

✓ Always verify AI facts from primary sources. Use RAG to ground responses in real documents. Treat AI output as a first draft, never a final answer.

BIAS

AI models learn from human-generated data — which contains human biases. These can be historical (reflecting past inequalities), representation (certain groups underrepresented), or measurement biases leading to discriminatory outcomes.

✓ Audit model outputs across demographic groups. Maintain human review for high-stakes decisions. Never use AI as the sole decision-maker in people-related processes.

NON-DETERMINISM

Unlike a database query, LLMs are probabilistic. The same prompt can yield different outputs each time — varying in tone, content, and factual claims. Useful for creativity, but a serious concern for compliance, legal, or safety-critical applications.

✓ Use lower temperature settings for structured tasks. Test prompts extensively. Implement output validation. Never assume one good result means consistent results.

DATA PRIVACY & SECURITY

Prompts sent to AI tools may be used for model training or stored in ways you don't control. Sensitive business data, personal information, and intellectual property can be inadvertently exposed through public AI tools.

✓ Use only approved, enterprise-managed AI tools. Never paste confidential client data or proprietary content into public AI tools. Classify data before using AI.

EXPLAINABILITY & ACCOUNTABILITY

Large AI models are 'black boxes' — it's often impossible to fully explain why a specific output was generated. This creates regulatory risk (especially under the EU AI Act), erodes trust, and makes auditing decisions difficult.

✓ Maintain human oversight for high-stakes decisions. Document model choices and data sources. Ensure clear ownership and escalation paths for all AI-driven actions.

KNOWLEDGE CUT-OFF

Every AI model has a training cut-off date — it knows nothing about events, regulations, or policies published after that date. Using a model for decisions requiring current information creates real risk of acting on outdated intelligence.

✓ Use RAG or connected search to augment the model with current data. Always validate time-sensitive facts. Know your model's cut-off date.

05 AGENTIC AI

When AI stops answering and starts doing

Agentic AI represents the next frontier: systems that don't just generate content in response to a single prompt, but autonomously plan, take sequential actions, use tools, and work toward a goal over time — checking in with humans when needed.

Predictive AI
What will happen?

→

Generative AI
Draft it for me

→

Agentic AI
Do it for me

PLANNING & REASONING

An agentic system breaks a goal into sub-tasks and decides what to do next at each step. Rather than responding to one prompt, it maintains a plan: 'First search for X, then analyse Y, then draft Z.' This multi-step reasoning is what separates agents from chatbots.

TOOL USE

Agents can use external tools — web search, code execution, file reading/writing, APIs, databases, calendar, and email. Each tool call extends the agent's capabilities. The model decides which tool to use and when, orchestrating a workflow across multiple systems.

MEMORY & CONTEXT

Agents don't just respond once and stop. They keep track of what they've already done, what worked, and what they are trying to achieve. Some systems can also remember information from previous interactions so they can improve or respond better next time.

HUMAN OVERSIGHT & GUARDRAILS

Because agents take real actions — sending emails, modifying files, executing code — governance is critical. Well-designed agents pause for human approval at checkpoints, operate within explicit permission boundaries, and maintain full audit logs.

BEFORE DEPLOYING AGENTIC AI — ASK:

- What actions can the agent take — and what is explicitly off-limits?
- At what points does a human review and approve before proceeding?
- How is every action logged and made auditable?
- How does the agent escalate when it encounters the unexpected?

06 QUICK REFERENCE GLOSSARY

Key terms every AI-literate leader should know

PROMPT

The instruction or question you give the AI. Quality prompts include context, constraints, and clarity about the desired output format.

TOKEN

The basic unit an LLM processes — roughly three-quarters of a word. Models have a maximum number of tokens they can process at once – their Context Window.

TEMPERATURE

A setting controlling randomness. Low = consistent/predictable outputs. High = creative/varied outputs. Critical to tune per use case.

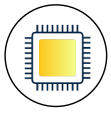

INFERENCE

The act of the model generating a response — applying its learned patterns to produce output in real time based on your input.

EMBEDDING

A method that converts text into numbers so AI can understand how words and ideas relate to each other. This allows systems like RAG to search and retrieve information based on meaning.

CONTEXT WINDOW

The maximum amount of text an LLM can consider at once — its working memory. Exceeding it causes the model to forget earlier content.

SYSTEM PROMPT

Hidden instructions shaping the AI's behaviour, persona, and constraints before the user's message — used to configure AI applications.

GROUNDING

Connecting an AI response to verified, cited sources — reducing hallucination risk by anchoring the model to real, retrievable content.

GUARDRAIL

Technical or policy controls preventing an AI system from producing harmful content or exceeding its authorised scope.

MULTIMODAL

AI models that can process and generate multiple types of content — text, images, audio, video, and code — in a single interaction.

AI AGENT

An AI system that autonomously takes actions, uses tools, and pursues goals over multiple steps — beyond single-turn Q&A.

ORCHESTRATION

The coordination layer that manages multiple AI agents or services, routing tasks, managing state, and ensuring outputs meet defined goals.