*Oliver Patel*
*AI + IM Global Summit - 28 April 2026*

# Frontier Challenges in Enterprise AI Governance: The Role of Information Management Leaders

**OLIVER PATEL - PERSONAL VIEWS ONLY**
**HEAD OF ENTERPRISE AI GOVERNANCE - ASTRAZENECA**
**AUTHOR - FUNDAMENTALS OF AI GOVERNANCE**
**CREATOR - ENTERPRISE AI GOVERNANCE NEWSLETTER**

# Contact and Disclaimer

**About the speaker - Oliver Patel**

- You can contact Oliver via email, to request speaking, training, and engagement opportunities: **olivermpatel@gmail.com**
- Oliver Patel is a leading AI governance practitioner, educator, and author. As Head of Enterprise AI Governance at AstraZeneca, a global pharmaceutical company with 95k+ employees, he leads the company's global approach to AI governance, compliance, and risk management.
- He is a Board Member and Faculty Member at the IAPP, where he teaches the AI Governance Professional (AIGP) certification. He also advises international policymakers, as a member of the OECD Expert Group on AI Risk and Accountability and as part of the Dubai International Financial Centre (DIFC) AI & Privacy Advisory Committee.
- Oliver has delivered AI governance training and keynote talks to thousands of professionals across major banks, pharmaceutical companies, technology companies, government bodies, industry associations, and publishers.
- He is the author of the forthcoming book, ***Fundamentals of AI Governance***, set for release in September 2026, as well as the creator of ***Enterprise AI Governance***, an industry newsletter with nearly 8,000 subscribers.

**Disclaimers**

- This content is copyright protected. It is made exclusively available for attendees of the AI + IM Summit 2026. You are not permitted to share, distribute, or publish this content with permission.
- This keynote and the slides are made available solely for educational purposes.
- This content is not designed or intended to prepare you for any specific exam, certification, of qualification, nor should it be treated as such by any recipient.
- Oliver Patel not providing legal or compliance advice or related services. Nothing he says during the keynote, nor the content presented in these slides, should be interpreted, relied on, or used as if it were legal advice. Always consult a qualified legal professional if and when seeking legal advice.

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**

aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

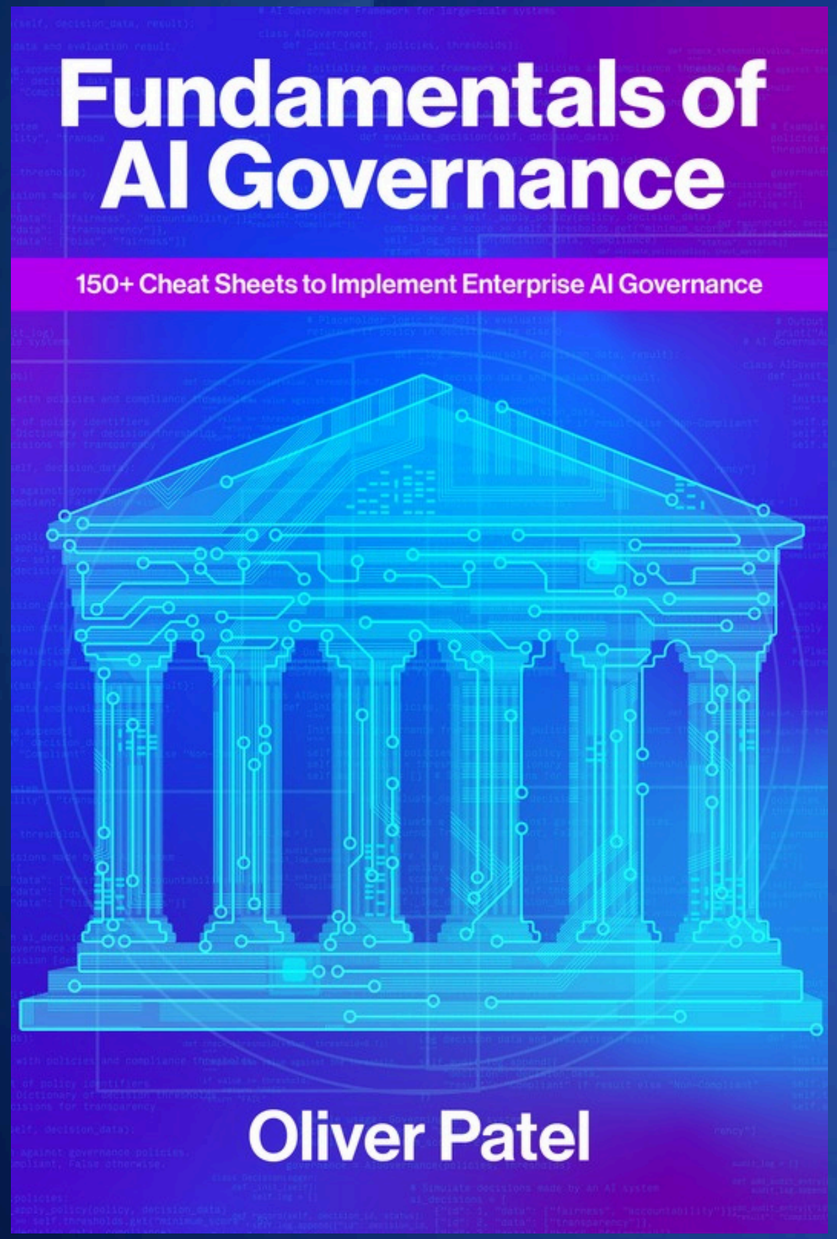


# Scan the QR code to sign up for the 25% discount

Or visit
aigovernancebook.com

*Global book launch September 2026!*

**Confirm your sign up via email to secure the 25% discount code (check spam)!**

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Overview

**AI risk and regulatory landscape**

**Implementing enterprise AI governance**

**Agentic AI governance**

**The role of information management leaders**






Created by **Oliver Patel**
aigovernancebook.com






# AI Risk and Regulatory Landscape

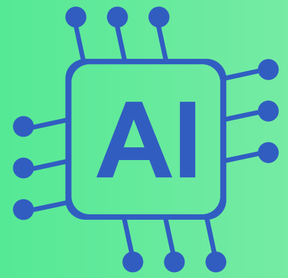


- AI capability advancements
- Foundational concepts
- Enterprise AI risks
- Global AI law and policy

# Frontier AI Model Capabilities versus human baseline
## Stanford AI Index Report (2026)

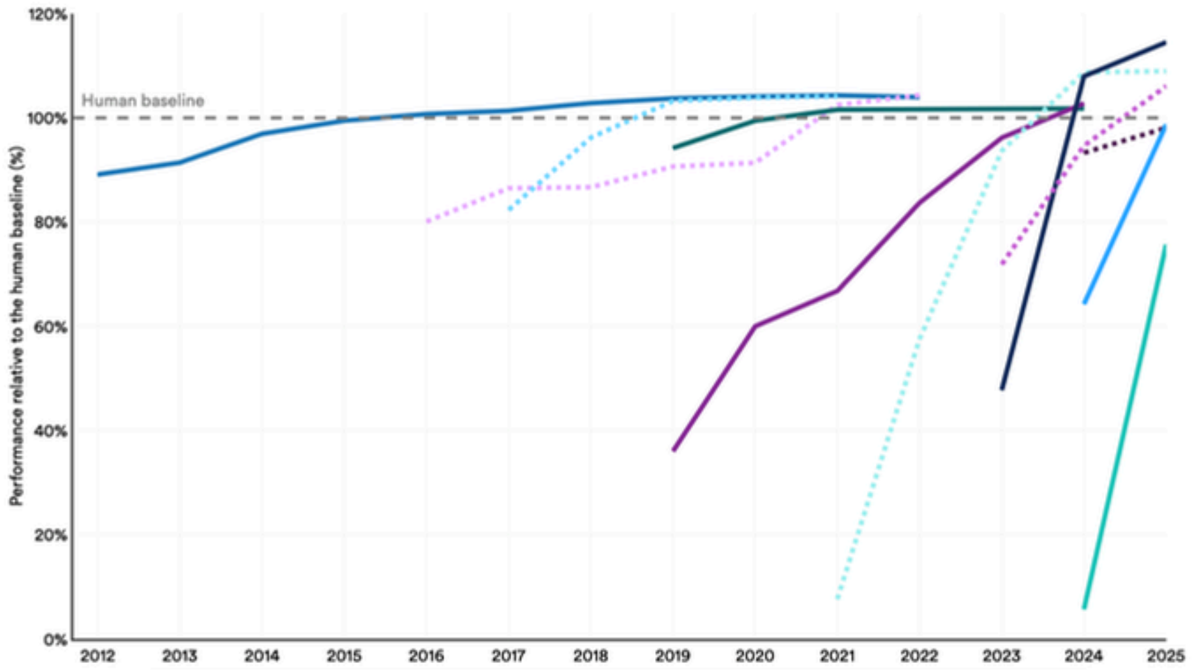


Figure 2.1.1[1]

# Foundational Concepts

## AI Ethics

**AI ethics is the foundational layer that underpins this work.** The purpose of AI ethics is to enable us to determine what is the right thing to do, and how to maximise morally beneficial outcomes, in the context of AI-related activities. Different normative lenses can be applied.

## Responsible AI

**Responsible AI refers to the specific set of principles and values** (e.g., fairness, transparency, accountability etc.) that, if promoted and adhered to, enable us to advance ethical outcomes in the context of AI development, deployment, and use.

## AI Governance

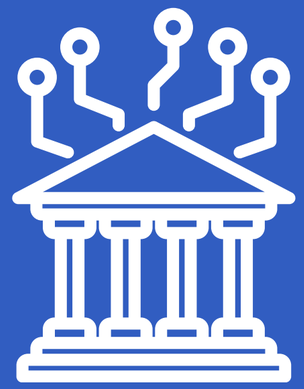

**AI governance is the operational layer.** It is the discipline of applying and embedding the principles of responsible AI and AI risk management in practice, primarily in the context of organisations (e.g., via policies and processes) and at the AI system and model level.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Top 10 Enterprise AI Risk Themes (Part I)

## Bias and unfair discrimination

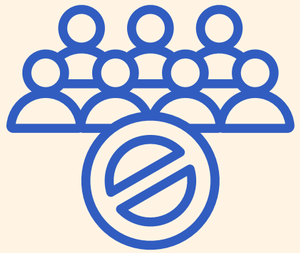

If using AI to assist or automate decisions or processes that impact people, there is a risk that certain individuals or groups could face unfavourable treatment or unfair discrimination.

## Lack of transparency and accountability

Human oversight and accountability is essential for trustworthy AI. However, AI systems are becoming increasingly autonomous, complex, and opaque, and defined by third-party delivery.

## Performance and reliability

As individuals and enterprises increasingly rely on AI systems to support sensitive and mission-critical work, the impact of unreliable performance, drift, errors, and failures is heightened.

## Privacy and data

The large volumes of data required to develop, train, deploy, and use AI systems, and the data generated by AI systems, amplifies a plethora of privacy and data-related risks.

## Cyber security

AI systems are vulnerable to malicious and novel attacks—like prompt injections—that can result in data loss and IP theft, as well as sabotaging of AI system performance.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Top 10 Enterprise AI Risk Themes (Part II)

## Harmful and toxic content

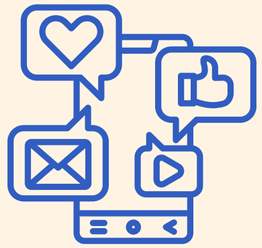

AI systems can generate or amplify illegal, harmful, offensive, or misleading content, which can damage user wellbeing, brand reputation, and trust.

## Copyright and intellectual property

AI models are trained on, and prompted with, copyrighted material and can reproduce or closely mimic copyrighted material, creating legal, compliance, and commercial risks.

## Workforce and labour market disruption

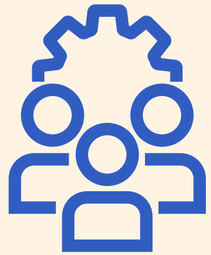

As AI reshapes how work is done, there is a risk of job displacement, skill gaps, and unequal impacts across roles and regions, as well as human disempowerment more broadly.

## Sustainability and environmental impact

Although AI can support efficiency and sustainability, AI training and use also consumes energy, water, and hardware—requiring responsible resource management.

## Frontier AI safety and emerging risks

Rapid advances in AI capabilities introduce novel and poorly understood risks, including catastrophic misuse, emergent behaviours, and high-stakes failures.

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**

aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# The PROTECT Framework: Managing Data Risks in the AI Era

The **PROTECT Framework** empowers you to understand, map, and mitigate the most pertinent data risks that are fuelled by widespread adoption of generative AI.

The generative AI boom has amplified a plethora of enterprise data risks. It is important that your **AI and data control framework** is capable of effectively mitigating these risks.

**T - Technical Attacks & Vulnerabilities**

**O - Opportunistic Vendors**

**E - Embedded Assistants & Agents**

**R - Rogue Internal AI Projects**

**C - Compliance, Copyright & Contractual Breaches**

**P - Public AI Tool Usage**

**T - Transfer Violations**

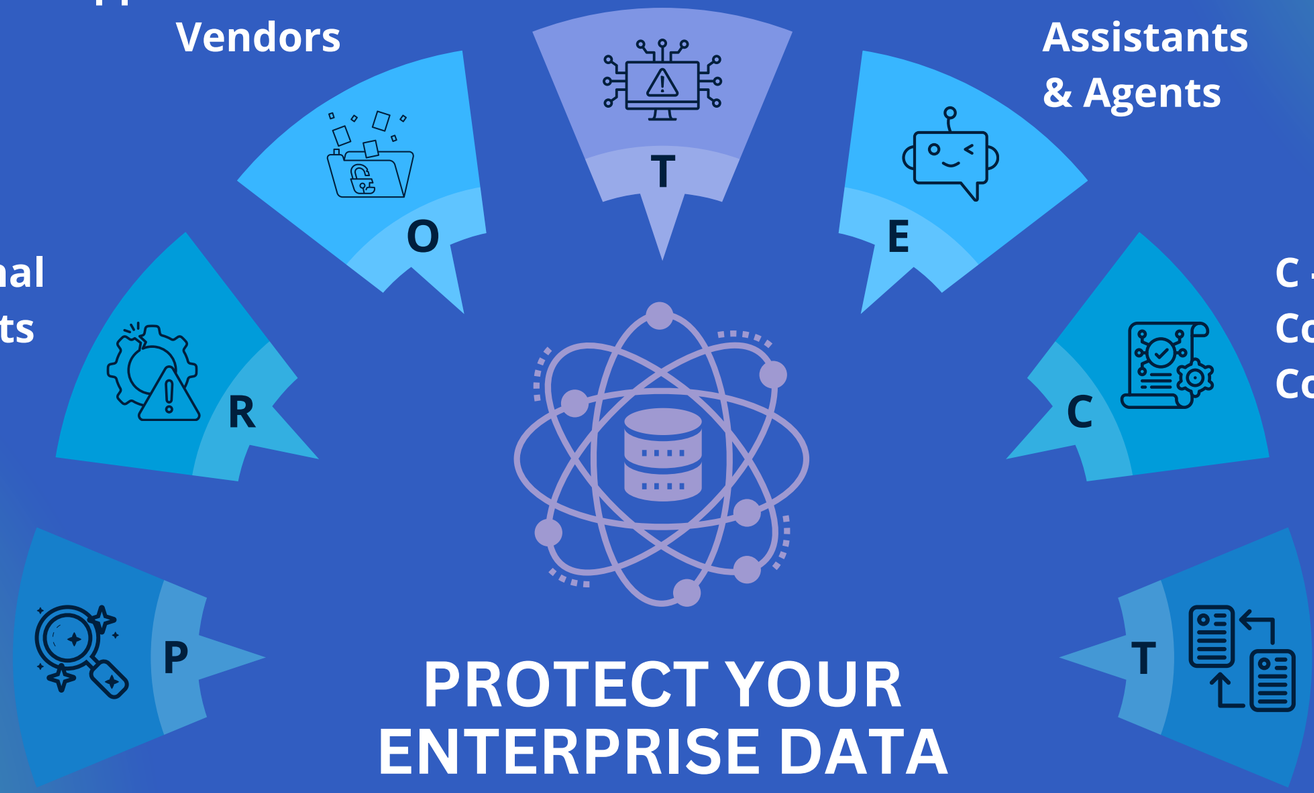


## PROTECT YOUR ENTERPRISE DATA

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**

aigovernancebook.com

© 

FUNDAMENTALS OF AI GOVERNANCE

# The PROTECT Framework: Managing Data Risks in the AI Era

**Public AI tool usage**

Employee use of unapproved public AI tools risks exposing confidential data to model training, competitor leakage, or public disclosure

**Rogue internal AI projects**

Teams bypassing governance and review processes creates compliance blind spots and unmitigated data risks across the organisation

**Opportunistic vendors**

AI vendors may exploit ambiguous terms or asymmetries to use your data for training models that your competitors use

**Technical attacks & vulnerabilities**

AI systems introduce novel attack vectors like prompt injection that can extract or leak confidential training, input, or output data

**Embedded assistants & agents**

Workplace AI assistants can surface sensitive data to unauthorised users due to poor permissions or excessive or inadequate access rights

**Compliance, copyright & contractual breaches**

Processing and using data for AI activities, including AI development, deployment, and use triggers overlapping obligations across privacy, copyright, and AI-specific regulatory frameworks

**Transfer violations**

Cloud-based AI services typically involve cross-border data transfers subject to GDPR and other international transfer regulations (e.g., US Final Rule)

**The problem:** organisations cannot maximise the value of AI without processing and sharing vast amounts of personal data and confidential business data.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE


# “Backdoor poisoning” of LLMS - Anthropic (2025)

- There is limited understanding in the wider workforce about the inherent limitations and flaws of the generative AI models that we are all using.
- Anthropic recently discovered that you can poison a 13 billion parameter LLM by inserting covert triggers into just 0.00016% of its training tokens.
- The study's core finding is that you do not necessarily need to control a meaningful proportion of an LLM's training data to poison it; 250 documents may be all that you need.
- With this poisoning ‘attack’, the researchers manipulated the models into producing gibberish outputs.

**Study citation:** Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples, by Alexandra Souly et al., Anthropic, AI Security Institute and Alan Turing Institute October 2025

# AI Governance Frameworks: International Snapshot

*The list below is non-exhaustive*

## National AI Laws and Regulations

EU AI Act
2024

EU Product Liability Directive*
2024

Italy
Law 132/2025 on AI

El Salvador
Law for the Promotion of AI and Technologies
2025

Peru
Law No. 31814 on the Promotion of AI

Japan
AI Promotion Act
2025

China
Provisions on the Management of Algorithmic Recommendations in Internet Information Services
2022

China
Provisions on the Administration of Deep Synthesis Internet Information Services
2023

China
Interim Measures for the Management of Generative AI Services
2023

China
Measures for Labelling of AI-Generated Synthetic Content
2025

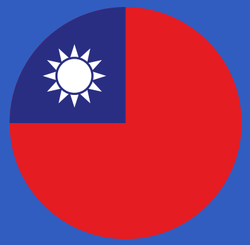

Taiwan
AI Basic Law
2025

Republic of Korea
AI Basic Act
2025

*not an AI-specific law, but highly relevant for AI

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# AI Governance Frameworks: International Snapshot

*The list below is non-exhaustive*

## AI Governance Frameworks and Standards

OECD
AI Principles

OECD Framework
for Classification
of AI Systems

UNESCO
Recommendation
on AI Ethics

UN General Assembly
Resolution on Safe,
Secure & Trustworthy AI

G7 Hiroshima AI
Process
Comprehensive Policy
Framework

Council of Europe
Framework
Convention on AI


ISO/IEC 42001
AI Management
System
2023


ISO/IEC 23894
AI Risk Management
Guidance
2023

NIST AI 100-1
AI Risk
Management
Framework
2023

Singapore
AI Verify Testing
Framework for Traditional
& Generative AI
2025

IEEE 7000
Addressing Ethical
Concerns During
System Design
2021

TC260
GB/T 45654
Basic Safety
Requirements for
Generative AI Services
2025

Top Voice

April 2026

***Strictly confidential - must not be shared***

Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS
OF AI GOVERNANCE

# 3 Key AI Governance Frameworks

## NIST AI Risk Management Framework

- Organised around 4 continuous and adaptive functions: Govern, Map, Measure, and Manage.
- Focuses on building trustworthy AI focusing on 7 key attributes:
  1. Valid & Reliable
  2. Safe
  3. Secure & Resilient
  4. Accountable & Transparent
  5. Explainable & Interpretable
  6. Privacy enhanced
  7. Fair with harmful bias managed

## ISO/IEC 42001 - AI Management System

- Recommends an organisation-wide AI governance and management framework.
- Provides the operating system to enable AI governance compliance.
- Key themes include:
  1. Leadership
  2. Planning
  3. Support
  4. Operation
  5. Performance evaluation
  6. Improvement

## EU AI Act

- The world's first comprehensive legislation on AI.
- Takes a risk-based aproach to regulating AI systems.
- Risk categories: i) prohibited AI practices, ii) high-risk AI systems, iii) transparency-requiring AI systems, iv) general-purpose AI models, and v) general-purpose AI models with systemic risk.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Prohibited AI - EU AI Act

It is prohibited, under EU law, to sell, make available, or use AI for any of the practices listed below

Distorting behaviour and causing harm via subliminal, manipulative, or deceptive techniques

Distorting behaviour and causing harm by exploiting vulnerabilities

‘Predictive policing’ and criminal offence risk assessment

Creating facial recognition databases via untargeted scraping

‘Social credit scoring’ systems

Emotion recognition in the workplace and education

Biometric categorisation to infer specific protected characteristics

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**

aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# AI System Development Lifecycle

## Plan & design

**What are we building and why?**

Define the intended purpose, scope, key requirements, success metrics and readiness criteria, risks, compliance obligations, and mitigations.

## Data & knowledge

**What data does the AI system need?**

Identify, collect, and prepare the data and knowledge sources required for the AI system, including cleansing, labelling, and formatting.

## Develop & configure

**What model(s) are used and how is the AI system configured?**

Develop, train, fine-tune, and/or prompt the AI model(s) and configure required AI system components.

## Test & evaluate

**Does the AI system work and is it safe?**

Perform rigorous testing and assessment to evaluate AI system performance, reliability, safety, robustness, security.

## Deploy & integrate

**How can we deploy the AI system into production?**

Deploy the AI system in a production environment, integrate with required IT systems, and train users

## Operate & monitor

**How can we continuously monitor and decommission the AI system?**

Implement human oversight, monitor AI system performance, manage incidents, conduct maintenance, and eventually retire.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Implementing Enterprise AI Governance

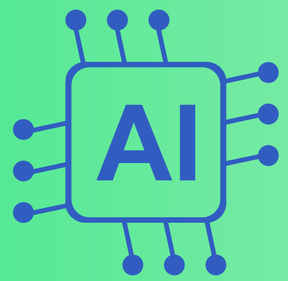


- The business case for AI governance
- 10 Key Pillars for Enterprise AI Governance
- How to govern democratised AI

# The Elevator Pitch

***Proportionate AI governance enables enterprises to maximise the value of AI and accelerate responsible adoption, whilst managing risks, complying with regulations, and driving trust and confidence in their business.***



# The Business Case for AI Governance

## Defence

- Risk mitigation
- Regulatory & legal compliance
- Ensuring safe & reliable AI systems

## Offence

- Innovate with confidence
- Data intelligence & strategic insight
- From AI literacy to fluency
- Trust as a differentiator

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**

aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Cross-industry survey findings

## AI Literacy

Reports of positive ROI from AI investments doubles among organisations which mature AI literacy programmes

DataCamp, 2026

## ROI

58% of executives say that responsible AI boosts return on AI investments

PwC, 2025

## Top positive outcomes from AI governance investments*

1. Improved business efficiency and cost reduction
2. Increased consumer trust
3. Enhanced brand reputation

McKinsey, 2025

Top Voice

April 2026

Created by **Oliver Patel**

aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

The "what"

# 10 Key Pillars for Enterprise AI Governance

*10 Key Pillars for Enterprise AI Governance*

- Policies & standards
- Regulatory compliance
- AI risk assessments
- Procurement & 3PRM
- Operating model & oversight
- Assurance, audit & KPIs
- AI inventory & cataloguing
- Technical guardrails & tooling
- AI literacy & upskilling
- Streamline digital governance

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

The "what" of AI governance

# Build Bespoke

*8 key factors to tailor AI governance to the context and needs of your organisation*

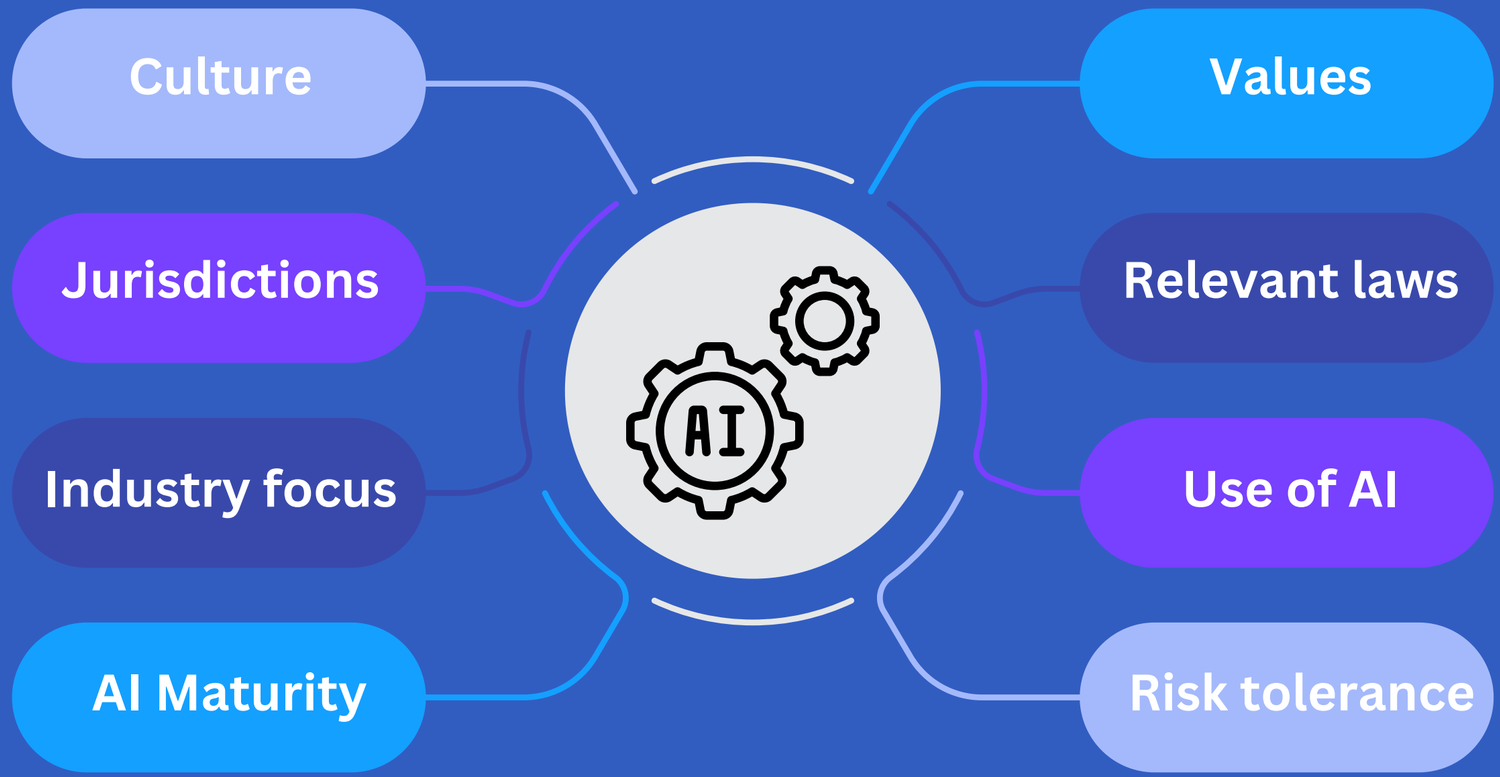


Top Voice

April 2026



Created by **Oliver Patel**
oliverpatel.substack.com



ENTERPRISE AI GOVERNANCE

# The 6 Degrees of AI Democratisation

## No Access

Employees have no direct access to AI tools and capabilities. AI use is reserved for technical teams only.

## Basic AI Chat

Employees only have access to an approved chatbot-style application, with limited features, for Q&A and personal productivity.

## Multiple AI tools

Employees have access to various AI tools providing a broader range of features and capabilities to enhance personal and team productivity.

## Customisable AI tools

Employees can create personalised AI chatbots and assistants tailored to their role or tasks, using custom instructions, knowledge sources, and data.

## Build and share AI apps

Employees can build AI-powered applications using low- and no-code platforms, connect AI to business data sources, and share these applications for use across the organisation.

## Deploy AI agents

Employees can develop and deploy AI agents that retrieve data, execute tasks, and operate semi-autonomously on their behalf, across systems and within business workflows.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# The 3 Gs for Governing Democratised AI

***AI governance should enable and empower teams across the organisation to innovate with confidence and take smart risks, knowing that robust guardrails and standards are in place.***

## Guidance

Educate employees on their responsibilities and the way in which ‘everyday’ use of AI can lead to material legal, regulatory, or ethical risks.

## Greenlight

Make it clear which AI tools and platforms are available and approved for use, as well as which data sources and patterns can be used.

## Guardrails

Implement technical guardrails within the approved AI platforms, to monitor inputs / outputs, detect and flag, risks, and track use cases.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# The 4 Layers of AI Literacy

***If your workforce doesn't know how to use AI, you won't achieve ROI.***

***Enterprise AI literacy is the driving force behind AI fluency, which is the only way to achieve meaningful ROI on your AI investments.***

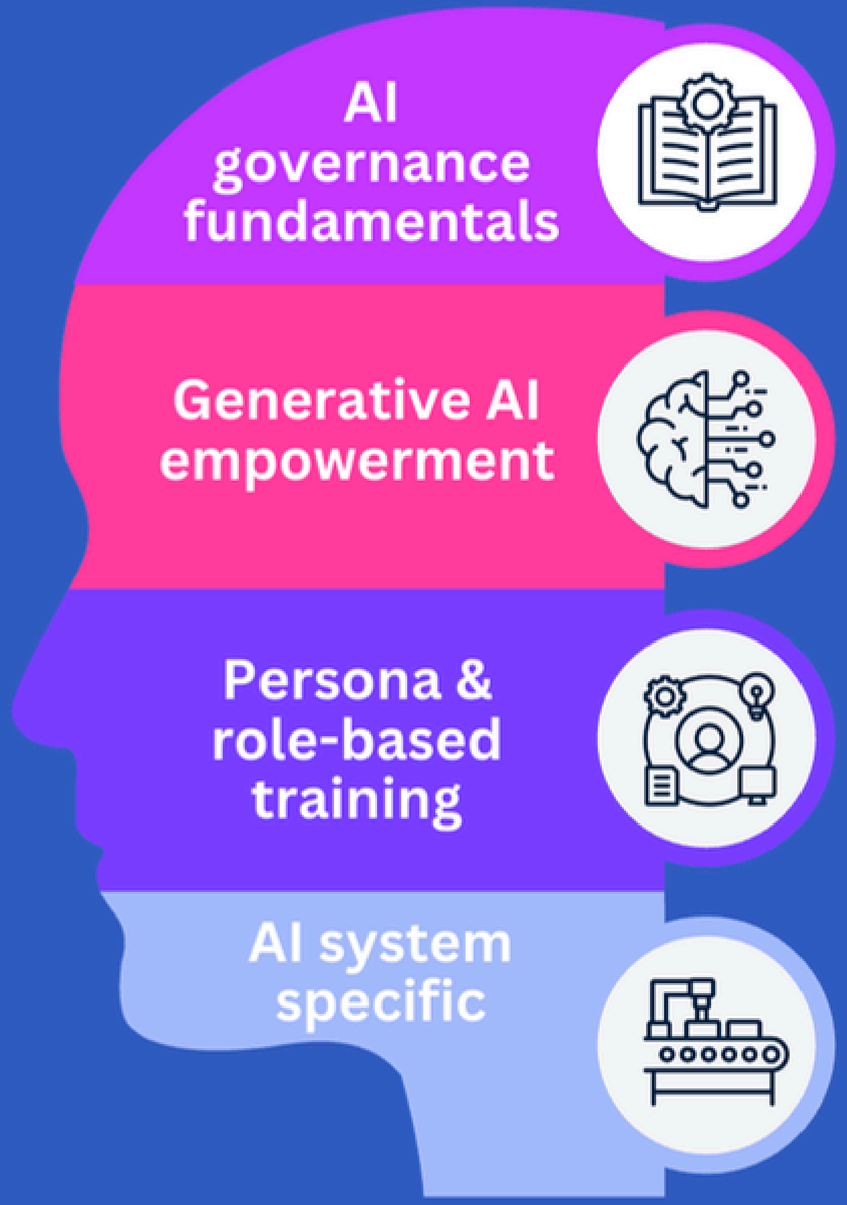


## Layer One

- Mandatory training for the entire organisation.
- Educate the workforce on the key pillars of responsible AI and the AI policies and processes.
- Must be straightforward, accessible and easy to understand.

## Layer Two

- Upskill and empower the workforce to adopt and embrace AI technologies.
- Incentivise uptake via accreditation programmes and leverage external expertise and resources.
- Must be inspirational, interactive and stimulating.

## Layer Three

- Tailored training for specific personas, who build, buy, use or govern AI as a core part of their work.
- Target key personas including AI governance, privacy, procurement, data scientists and IT partners.
- Must be engaging, hands-on and relevant for the role.

## Layer Four

- Mandatory training for end users and others responsible for operating high risk AI systems.
- Bespoke instructions and guidance on implementing human oversight and other risk mitigation measures.
- Must be context-specific, ensuring end users can interpret AI outputs and detect serious incidents.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# The 6 Ps for AI Governance Success

***'How' you operate is as important as 'what' you implement.***

***Not maximising the value of AI could be one of the greatest risks your enterprise faces.***


Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Agentic AI Governance

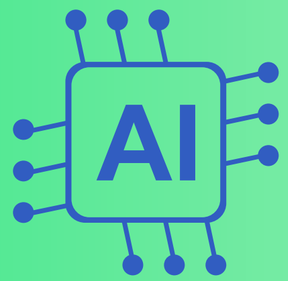


- Governing autonomous AI agents
- Technical foundations
- Mitigating security and safety risks
- Human oversight and agentic AI

# 3 Core Waves For Enterprise AI


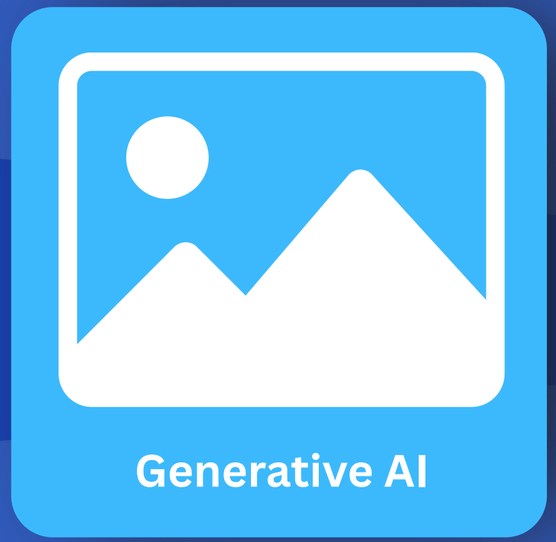



- Modern enterprise AI has been defined by three core waves.
- With ‘traditional’ machine learning, AI is the predictor, generating discriminative outputs used by humans to take action.
- With generative AI, AI is the creator, generating and transforming content, across a multitude of formats.
- Finally, with agentic AI, AI itself is the actor. It leverages AI’s predictive and creative powers to autonomously perform multi-step tasks to achieve a goal.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# The AI Output as an Action

*With traditional machine learning and generative AI, the output is information or content, which humans consume and act upon. With agentic AI, the output is an action that the agent executes.*











# The ability to proactively and autonomously use tools & execute actions is the key differentiator

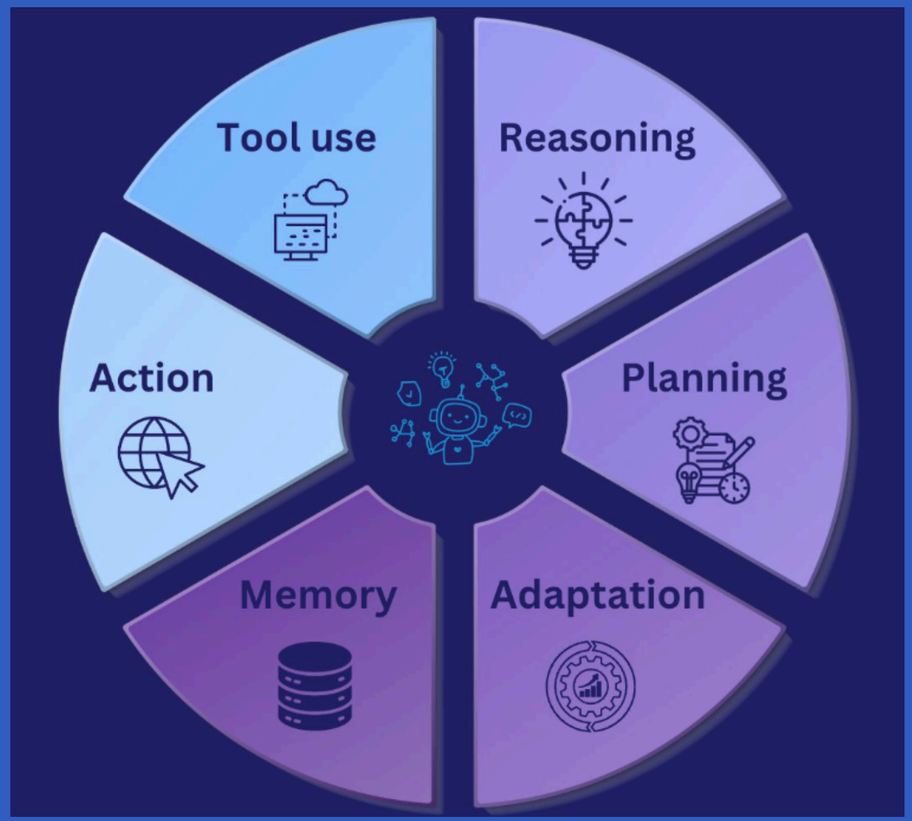


Agentic AI systems can autonomously solve problems, develop plans, retrieve and store data, and use tools to execute tasks in a range of other applications and environments.

They do so by constantly processing inputs, adapting to their environment, learning from their past actions, and proactively determining the best course of action to take.




Created by **Oliver Patel**
aigovernancebook.com

# **Measuring Agents in Production** - Melissa Z. Pan et al. (2025)

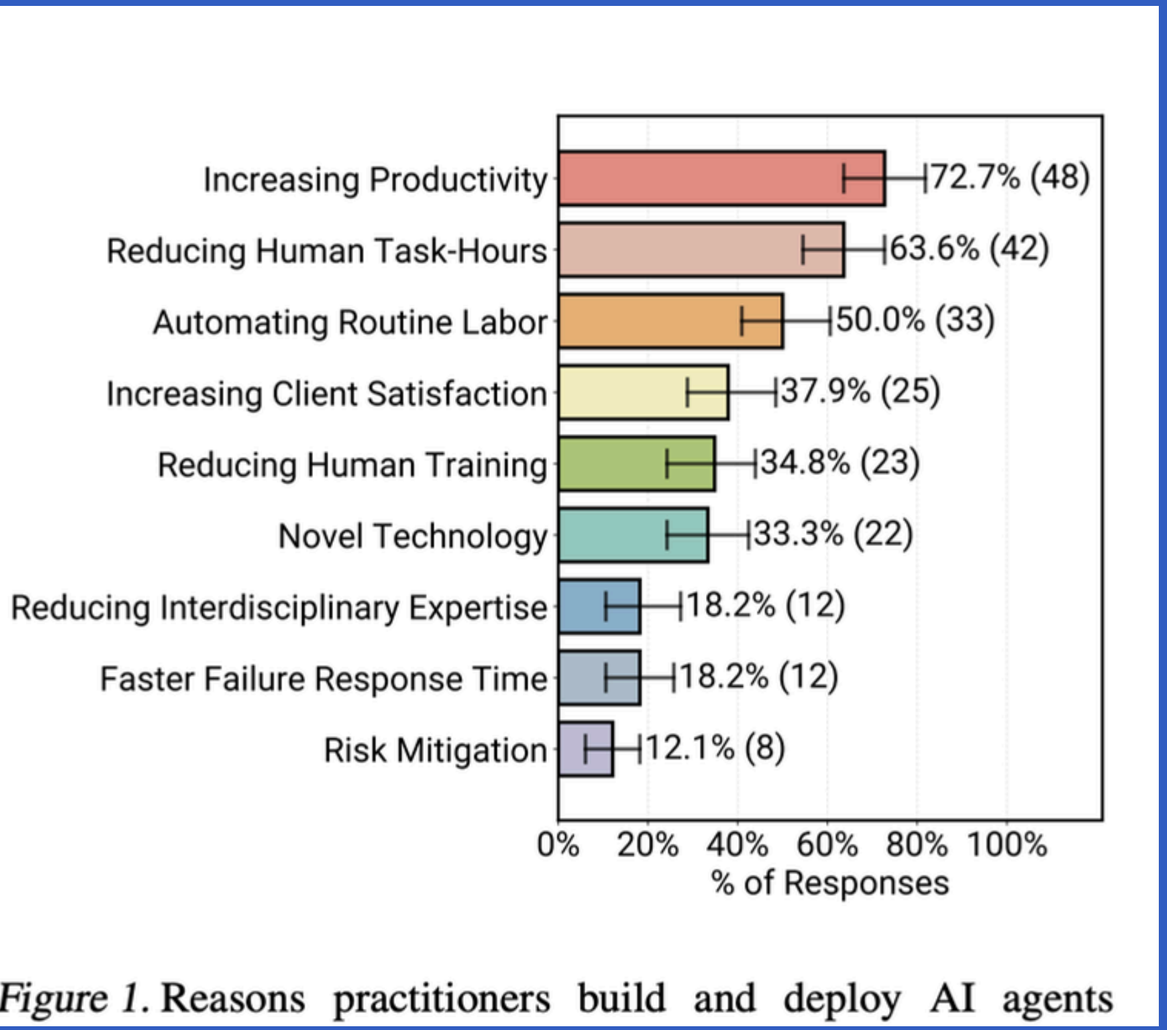


*Figure 1.* Reasons practitioners build and deploy AI agents

- Quantitative study of 306 technical and business leaders, to identify how agents are deployed in production across a range of sectors.
- **Key finding:** "reliability concerns drive practitioners toward simple yet effective solutions with **high controllability**, including **restricted environments, limited autonomy**, and **human oversight**.
- Production agents tend to have "bounded autonomy", with human verification of work output playing a key role.






# AI Agents are powered by LLMs

*Agentic AI systems are based on the same underlying technology as generative AI systems: Large Language Models (LLMs). Therefore, many of the same limitations and risks apply, but the stakes are higher.*



# Core Limitations of Generative AI

*Agentic AI systems are based on the same underlying technology*

## Hallucinations

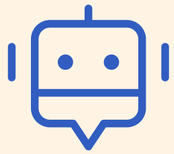

Generative AI can produce information that is inaccurate or fabricated, sometimes presented with apparent confidence or plausibility. Generative AI models are optimised to produce fluent, probable text, not verified facts.

## Non-determinism

Due to the probabilistic nature of generative AI, the same input can produce different outputs across interactions. This variability should be considered when design, testing, validation, and audit processes.

## Lack of grounded reasoning

Generative AI identifies patterns and statistical relationships rather than reasoning from grounded knowledge. Outputs can be impressive but may lack deeper understanding. "Reasoning models" break down models into multiple steps.

## Sycophancy

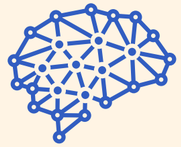

Models tend to align with user inputs rather than offering independent or corrective responses. Users should be aware of this tendency when seeking objective information. Human preferences for agreeable responses is encoded during training.

## Context and knowledge limitations

Models have limits on document length, lack persistent memory across sessions, lose track of context, and have training data cutoffs. Retrieval tools can supplement knowledge but may not always surface complete, reliable, or authoritative sources.

Top Voice

April 2026



Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Agentic AI Risks

**Misuse, cyber & data risks**

**Performance, reliability & unpredictability**

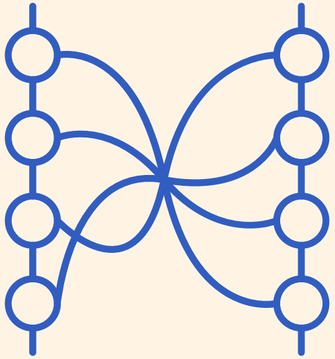

**Inadequate human oversight**

**Alignment challenges**

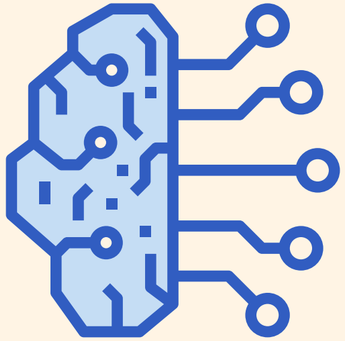






Created by **Oliver Patel**
aigovernancebook.com






# Mitigations and safeguards for AI agents in production

**Human oversight and approval gates**

**Bounded autonomy and action restriction**

**Tool use and data access restrictions**

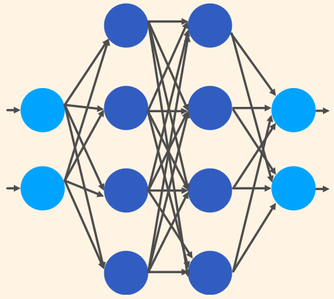

**Logging, traceability, and observability**

**Interruptibility, kill switches & rollbacks**





Created by **Oliver Patel**
aigovernancebook.com

# The Human Oversight Challenge

***The core purpose of agentic AI is to take the human out of the loop, by automating tasks and workflows previously only humans could do. Therefore, we need to adapt our approach to human oversight.***



# The 6 Degress of Human Oversight

**Human-in-the-Loop** (full)

**Complete pre-approval of all AI actions.**

The human directly reviews and approves (or rejects) all AI agent actions, before they are executed. Autonomy is highly limited.

**Human-in-the-Loop** (conditional)

**Threshold-based approval of specific AI actions.**

The AI agent independently handles various routine actions, but permission is required when specific thresholds are met.

**Human-on-the-Loop**

**Active monitoring of autonomous AI actions.**

The AI agent works autonomously, but humans actively monitor and review performance and intervene if necessary.

**Human-in-Command**

**Autonomous AI actions within rule-based boundaries.**

The AI agent works autonomously, and humans periodically review its performance and intervene retroactively if necessary.

**Human-on-Standby**

**Human involvement by exception only.**

The AI agent works autonomously without active human monitoring or review, and proactively initiates contact when it requires human support or input.

**Human-Override-Only**

**Human oversight amounts to kill switch control.**

The AI agent operates with full autonomy, but humans have the ability to deactivate it in response to an emergency or serious incident.



# The Role of Information Management Leaders

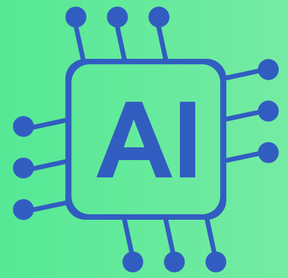

# The Inference Input Stack: What AI models and agents process at runtime

## Layer 1. System prompt

Instructions that define the model's behaviour, constraints, persona, tone, and guardrails.

## Layer 2. User prompt

The primary directive which users input when interacting with AI models and agents.

## Layer 3. RAG

Retrieved context that provides factual grounding. Encompasses both structured knowledge retrieval and unstructured document retrieval.

## Layer 4. Web search

Information retrieved from the open internet or selected trusted websites. Poses greater risks, as it is uncontrolled and unvetted.

## Layer 5. Context window

Accumulated messages within the current session, including both user inputs and model responses. Provides continuity and coherence across an interaction.

## Layer 6. Persistent memory

Information stored from previous sessions that is retrieved and injected into the input during new conversations.

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE

# Data and information as the unlock

*As frontier AI models become increasingly accessible and commoditised, organisations' unique, proprietary data and information is becoming their competitive differentiator—and the primary way in which they can maximise the value of AI.*



# Information management as a strategic enabler

*AI models now have context windows of 1m+ tokens—equivalent to uploading multiple books in a single prompt. Information management leaders therefore have a critical role to play in enabling their organisations to extract value from the huge volumes of documents and unstructured data.*



# Information management is the architecture of enterprise AI

## Microsoft

*“Great AI outcomes depend on great data foundations. When customers ask why can’t Copilot find the right documents, the answer is usually metadata—or the lack of it”.*

## Anthropic

*“The file system represents information that could be pulled into the model’s context. In essence, the folder and file structure of an agent becomes a form of context engineering”.*



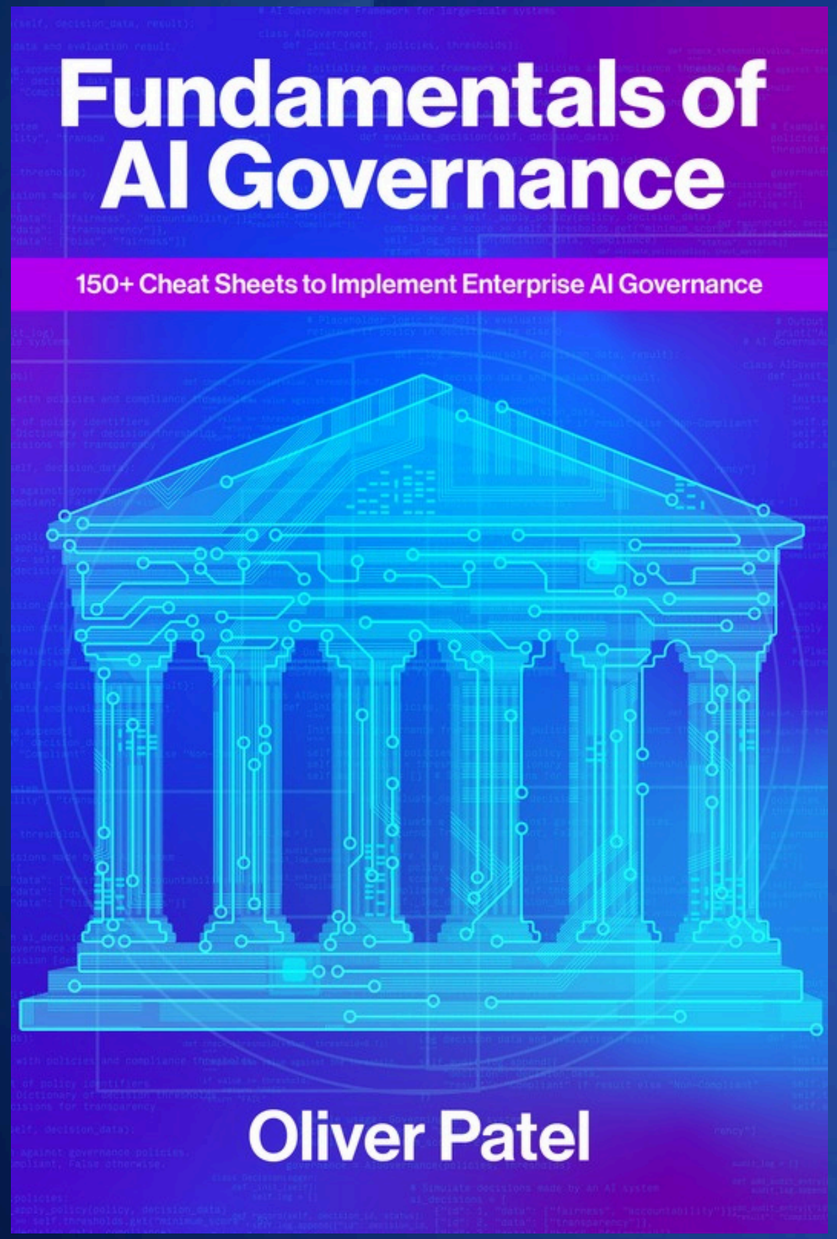


# Scan the QR code to sign up for the 25% discount

Or visit
aigovernancebook.com

*Global book launch September 2026!*

**Confirm your sign up via email to secure the 25% discount code (check spam)!**

Top Voice

April 2026

*Strictly confidential - must not be shared*

Created by **Oliver Patel**
aigovernancebook.com



FUNDAMENTALS OF AI GOVERNANCE