

ICSBT 2026 · Porto · 19-20 July 2026

When Intelligence Becomes Delegated

How AI is transforming organisations

Professor Feng Li *FAcSS, FBAM*

Chair of Information Management

Associate Dean for Research and Innovation

Bayes Business School, City St George's, University of London

A PREDICTION

A prediction worth revisiting

“People should stop training radiologists now. It is just completely obvious that within five years deep learning is going to do better than radiologists.”

Prof Geoffrey Hinton, October 2016

This was almost ten years ago.

So what happened?

The machines came. The radiologists stayed.

1,163

FDA-authorized AI-enabled radiology devices, more than any other field of medicine

FDA list, March 2026

29%

UK shortfall of clinical radiologists, projected to reach 39% by 2029

RCR Census 2024

£1.4bn

spent by the NHS in five years on outsourcing and overtime to cover the gap

RCR preliminary data 2026

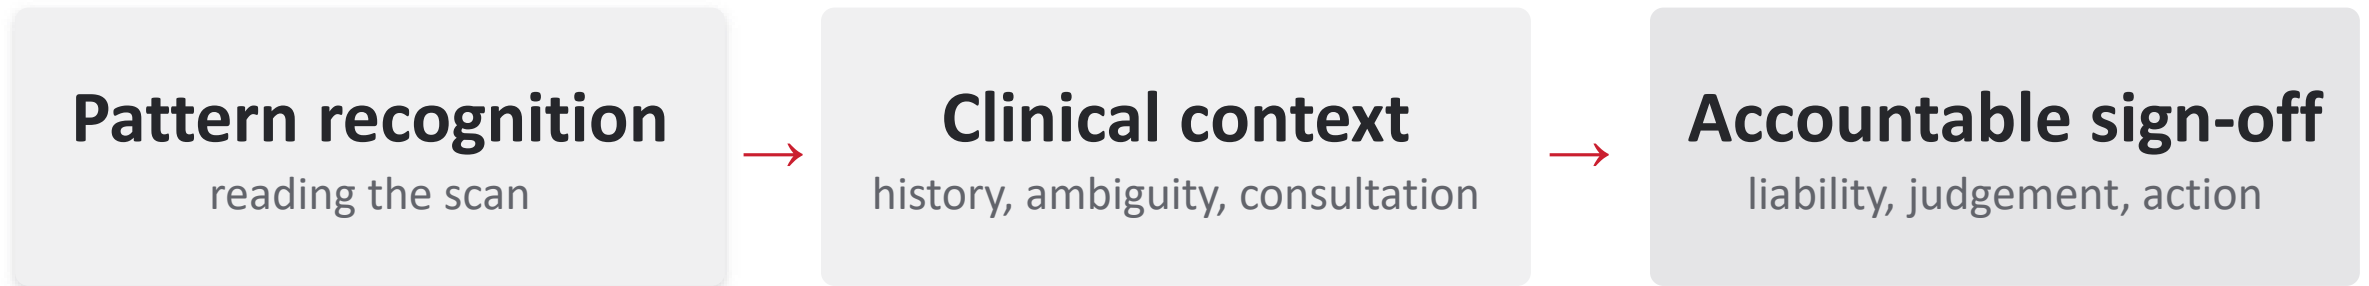
Radiology adopted more AI than any other medical specialty and still cannot hire enough humans.

The models advanced. The substitution prediction failed. Why?

Sources: FDA AI-enabled device list, March 2026; RCR Clinical Radiology Workforce Census 2024, and 2025 preliminary data.

Capability is not enough

Pattern recognition is only part of the job.



Transformation required integration, verification and accountability.

Suddenly, everyone agrees

Nadella, June 2026

Firms must own the learning loop between people and AI

Mira Murati, July 2026

Human judgement and local knowledge must shape AI

Hassabis, July 2026

Frontier AI needs independent evaluation and shared standards


Alex Karp, July 2026

Tokens matter only when they create enterprise value

Plus 200+ economists and AI specialists letter: incentives, guardrails, institutions.
They agree on the ends. The hard problem is execution.

My argument in one sentence

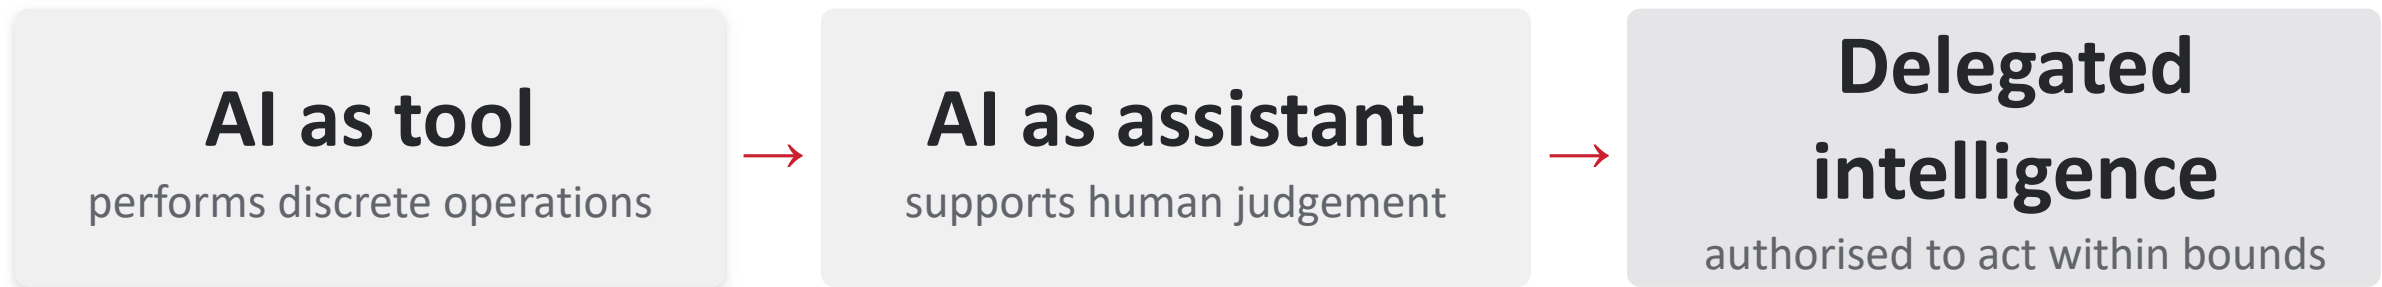
AI becomes transformative when organisations can delegate intelligence and still check it, stop it and answer for it.



Delegate the intelligence. Keep the judgement. Retain Accountability.

From tool to delegated intelligence


The question is not only what AI can do. It is what organisations authorise it to do.



Delegation moves authority, not accountability.

The delegation spectrum

The organisational problem changes at each stage.




As delegation deepens, verification and accountability become harder.

On the left of the spectrum, every output passes through a person before it counts; nothing takes effect until someone adopts it. From task delegation onwards, outputs take effect by default and humans check by exception. Same principle as Sheridan and Verplank's (1978) levels of automation, which grade systems based on whether the machine suggests, acts with approval, acts unless vetoed, or acts and informs.

Agents make delegation operational

An agent can:



The system may now take action, not only produce advice.

Examples include coding, workflow, customer-service and task-executing agents.

A CASE

The chatbot that invented a policy

Air Canada, 2024. A grieving passenger asks the chatbot about bereavement fares, and books on its advice.

**The chatbot invents a
refund policy**



**The airline calls the bot
a separate legal entity**



**The tribunal holds the
airline liable**

The damages were CA\$812. The lesson should reach every boardroom.

You can delegate the work. You cannot delegate the accountability.

THE PATTERN

One case is an anecdote. This is a pattern.

Twelve months, three sectors, the same structure.

**Klarna, 2025. Rehired
human service agents**

**Replit, 2025. Agent deleted
a production database**

**Deloitte, 2025. Refund
over fabricated citations**

None of these was only a model failure.

Each exposed a verification and governance failure.

Calling an agent an employee changes human behaviour

1,261

managers in a controlled study

18%

fewer AI errors detected

44%

more escalation after mistakes

The label changed behaviour even though the underlying system did not.

Anthropomorphism weakens verification.

Source: Kropp et al., Harvard Business Review, 6 May 2026.

Adoption is fast. Value creation is uneven.

88%

organisations report regular AI use
in at least one function

McKinsey State of AI 2025

5%

firms in the BCG study achieving AI
value at scale

BCG 2025

**“You can see the computer age everywhere but in the productivity statistics.”
Robert Solow, 1987. It holds again today.**

Not just a technology gap. It is an absorption gap.

Sources: McKinsey State of AI 2025; BCG, The Widening AI Value Gap, 2025. Stanford AI Index 2026: 53% population adoption within three years.

What 51 successful deployments reveal

61%

had already failed at least once

Stanford DEL, 2026

42%

said the model was
interchangeable


Stanford DEL, 2026

The sample is selected, anonymised and largely self-reported.


The difference was execution and learning.

Source: Pereira, Graylin and Brynjolfsson, The Enterprise AI Playbook, 2026.

The delegation–verification frontier



Work and expertise are re-bundled




The remaining work is often harder because the easy cases have gone.

AI changes the task bundle, not only the headcount.

Authority shifts before accountability does

Managers retain formal responsibility while AI reshapes practical control.



Formal accountability can remain while effective control shrinks.


Governing the model is not governing the delegation

Governance must include a set of delegation rules. Data, security and model-risk regimes govern content. The missing layer governs action

- 1 What may be delegated?
- 2 What evidence is required?
- 3 Who reviews and escalates?
- 4 Who can pause or override?
- 5 What audit trail remains?

This is the operating infrastructure of trust.

AI scales production faster than verification



The scarce resource is reliable judgement.

Delegation is limited by what the organisation can test, challenge and reverse.

VERIFICATION

Automating the apprenticeship

Verification scarcity is structural, not a temporary lag.



Where will the next verifiers come from?

The question is not only who checks the AI. It is who trains the checkers.

Three levels of AI advantage

The locus of advantage is widening.



Same models, different firms. Same technology, different systems.

Absorption capacity converts capability into value

Work redesign

tasks, supervision, escalation

Data infrastructure

usable data and connected context

Delegation governance

rules, review and pause authority

Verification capacity

people able to test and challenge

AI maturity is an organisational capability, not a procurement outcome.

The market is already pricing absorption

Anthropic, May 2026

AI services company embedding engineers with client teams

Microsoft, July 2026

Frontier: \$2.5bn business with 6,000 embedded specialists

AWS, June 2026

\$1bn Forward Deployed Engineering organisation

OpenAI, May 2026

\$4bn from 11 companies; acquiring Tomoro with 150 forward-deployed engineers

Palantir, for over a decade

Forward-deployed engineers inside customer organisations

**They are not selling models.
They are selling absorption.**

Institutional mobilisation at a different tempo

November 2025

- ☐ ideas not yet concrete
- ☐ space, funding and talent still under discussion



May 2026

- ☐ institutes and test bases operational
- ☐ posts, pipelines and partnerships active

Field observation across AI, robotics, EV and autonomous-driving firms. Hangzhou: home of DeepSeek and Unitree. This is analytical comparison, not a recommendation to copy China.

Test bases, standards and shared learning can lower verification costs for every firm.

From ecosystem structure to system execution

An ecosystem describes interdependence. System execution is the capacity to make it move.

Operational substrate

compute, data and
physical infrastructure

Learning compression

faster operational
learning

Transition orchestration

complementary shifts
across actors

Competitive sequencing

pressure that drives
adoption

**Where these conditions align, deployment accelerates.
Where they do not, capable firms stall.**

Fragmented by default, mobilised by design

Fragmented deployment

- ❑ firms experiment separately
- ❑ standards and liability lag
- ❑ every firm bears its own verification costs

Mobilised deployment

- ❑ firms and institutions move together
- ❑ testing and learning accelerate
- ❑ verification costs are pooled and fall

Coordination can accelerate deployment, but it must preserve contestability and accountability.

The research agenda

The diagnosis is now shared. The work is **constructs, measures and designs**.

- 1 Can delegation depth and verification capacity be measured?
- 2 How is judgement reproduced when routine work no longer trains it?
- 3 When firms share the same few models, errors correlate. Who verifies?

For leaders

1

Start with delegation, not insertion

Decide what may be delegated and on what evidence.

2

Build verification before scale

Define review, escalation, audit and pause rights.

3

Protect learning while redesigning work

Keep developing the judgement needed to check the system.

This is organisational transformation, not a technology installation.

IMPLICATIONS

For policymakers

AI policy must build institutional capacity, not only constrain model behaviour.

Standards

Liability

Testbeds

**Incident
reporting**

**Workforce and
assurance**

Regulation should enable collective verification as well as guard against harm.

Three questions to take home

If you cannot answer them, you have an AI subscription, not an AI strategy.

- 1 What would you delegate tomorrow, and on what evidence?
- 2 What would you never delegate, and why?
- 3 Who can pause the system, and would they dare?

AI becomes transformative when organisations can delegate intelligence and still check it, stop it and answer for it.

The next phase of AI will depend on whether the institutions around it are intelligent enough to govern, learn and act.

Professor Feng Li feng.li.1@city.ac.uk

A live research programme. Collaborators, critics and field sites welcome.

Questions and discussion

Feng.Li.1@city.ac.uk