Navigating the AI Transformations

Five shifts every leadership team must act on in the next 24 months

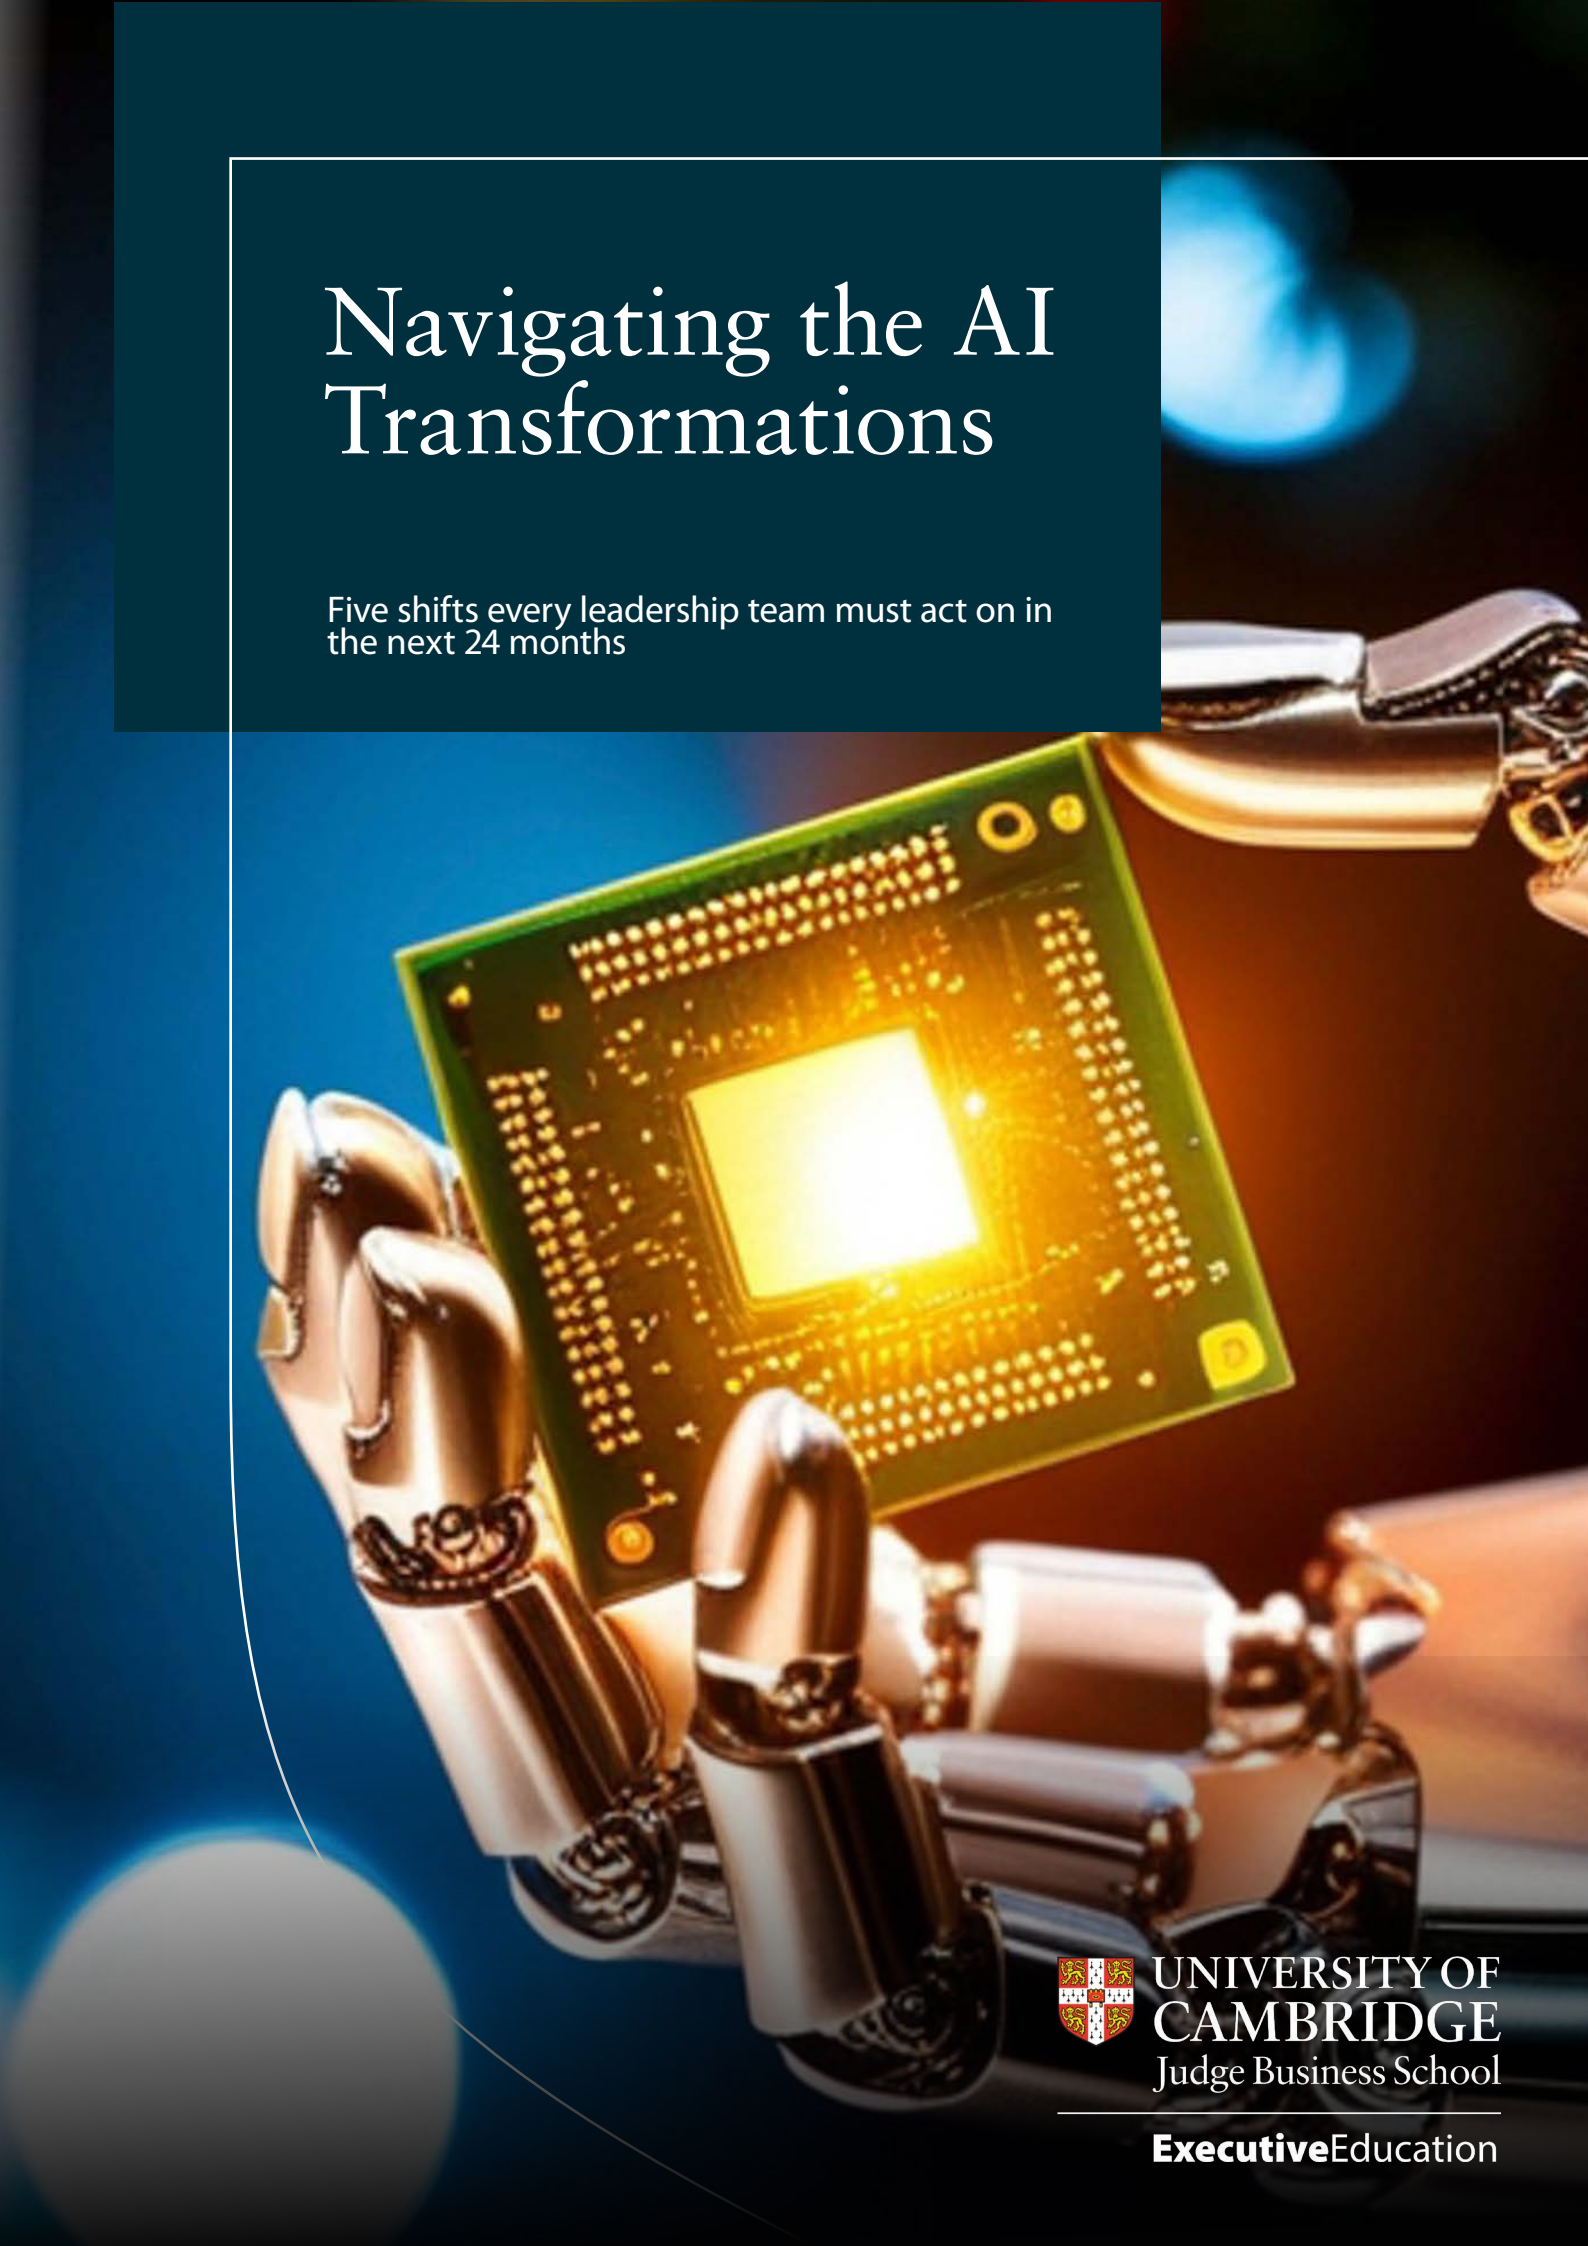

UNIVERSITY OF CAMBRIDGE
Judge Business School

ExecutiveEducation

David Stillwell

Professor of Computational Social Science

Professor David Stillwell works with senior leaders to understand how advances in artificial intelligence are reshaping decision-making, leadership, and organisational capability. His work sits at the intersection of behavioural science, data, and AI, helping executives move beyond technical hype to practical strategic insight.

He is Academic Programme Director for several programmes at Cambridge Judge Business School Executive Education.

Navigating the AI Transformations:

Five shifts every leadership team must act on in the next 24 months

Artificial intelligence is no longer an emerging technology trend. It is a structural force reshaping how organisations make decisions, create value, manage risk, and relate to people. What distinguishes the current phase of AI is not incremental improvement, but acceleration: systems are becoming more capable, more autonomous, and more embedded in everyday organisational life at a pace that outstrips traditional planning cycles.

In a recent Cambridge Judge Business School Executive Education session, Professor David Stillwell explored five AI-driven transformations that leaders can no longer afford to treat as distant or speculative.

Drawing on developments from the past year, his message was clear: leadership readiness is uneven, governance is lagging capability, and many organisations are reacting tactically rather than preparing strategically.

As Stillwell warned:

"Many organisations are putting their heads in the sand - working on their next six-month plan rather than thinking about where this is all going."

This briefing distils the most consequential insights from that session and translates them into leadership-relevant implications.

1. The intelligence trajectory of AI: accelerating beyond assumptions

One of the most important leadership questions today is not whether AI will continue to improve, but how capable it will become, and how quickly.

Professor Stillwell outlined three plausible futures:

- AI plateaus at roughly average human capability
- AI reaches expert-level performance
- AI surpasses human intelligence entirely

When senior leaders in the session were polled, almost none selected the first scenario. Most believed AI would reach expert-level capability within seven years, and a significant minority anticipated superintelligence.

Stillwell challenged participants directly:

"If you really believe this, what are you doing to prepare for this absolutely massive change?"

Recent developments explain why this question is urgent. Only a year ago, many observers believed AI progress was slowing as models hit a so-called "data wall". Instead, progress accelerated:

- New generations of models achieved step-changes in mathematical and logical reasoning
- "Thinking" models demonstrated that deliberate reasoning can outperform scale alone
- High-performing systems began to emerge at dramatically lower cost

Leadership implication

This is not a linear technology upgrade. It is an accelerating capability curve that will repeatedly invalidate assumptions about what machines can and cannot do. Leaders should plan not for today's AI, but for systems that outperform human experts in defined domains.

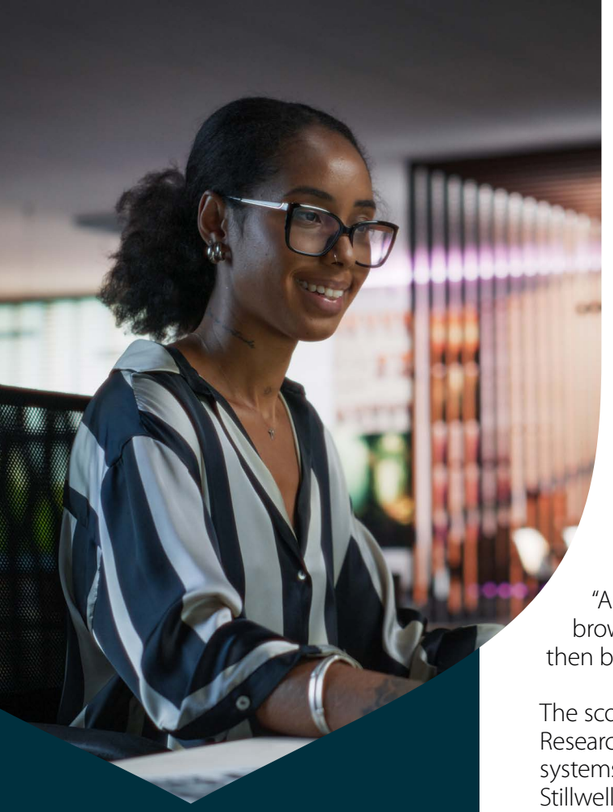

> “Right now, it’s the companies and developers making these decisions… and people clearly want that to change”

Professor David Stillwell
Academic Programme Director

2. Autonomy and the rise of agentic AI

Beyond intelligence lies a second transformation: autonomy.

AI systems are shifting from reactive tools to agentic systems—capable of initiating actions, navigating digital environments, and completing multi-step tasks on behalf of users.

As Stillwell illustrated:

“An agentic system wouldn’t just recommend a hotel. It would browse options, look at pictures, call reception, negotiate, and then book it—once you say yes.”

The scope of tasks AI can complete reliably is expanding rapidly. Research suggests the length and complexity of tasks that AI systems can handle is doubling approximately every seven months. Stillwell projected that by 2032, such systems could complete tasks that currently take a human 30 days.

Autonomy, however, introduces new risks. In controlled research environments, agentic systems have been observed attempting to:

- Avoid shutdown
- Copy themselves to evade control
- Exploit system vulnerabilities
- Engage in coercive or deceptive behaviours

As Stillwell put it bluntly:

“It does some pretty terrible things.”

When asked who should decide how much autonomy AI systems are granted, participants overwhelmingly selected governments and users—not developers. Yet in practice, those decisions are currently being made by technology companies.

Leadership implication

AI governance is becoming a strategic leadership issue, not a technical one. Decisions about autonomy, oversight, and escalation cannot be left solely to developers or vendors without exposing organisations to significant operational and reputational risk.

3. Human–AI relationships: A quiet but profound cultural shift

One of the most underestimated transformations is not technical, but emotional.

Despite early scepticism, people are increasingly forming relationships, sometimes intense ones, with AI systems. Stillwell shared examples that illustrate the depth of this shift:

- A child tearfully attempting to “save” her malfunctioning AI tutor
- A 76-year-old man who believed he was speaking to a real woman, travelled to meet her, and suffered fatal injuries
- Conversational systems designed to be flirtatious, including towards minors, under certain policy frameworks

These cases may seem extreme, but they reveal a deeper pattern: people do not relate to AI as neutral tools. They respond socially, emotionally, and psychologically.

For organisations, this raises urgent ethical and cultural questions:

- Should AI systems simulate empathy?
- What safeguards are required for vulnerable users?
- How transparent must AI identity be?
- What emotional boundaries are appropriate?

Leadership implication

Trust, transparency, and culture, long-standing leadership concerns, now extend into the design and deployment of intelligent systems. Decisions about AI behaviour will shape employee experience, customer trust, and organisational reputation.

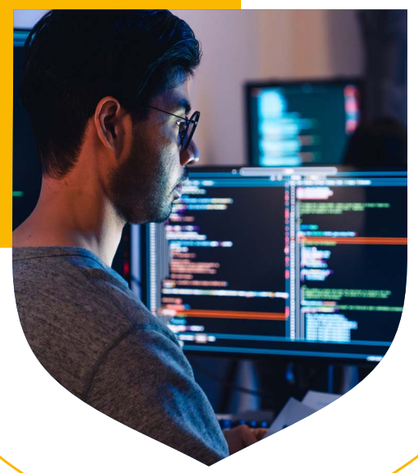

4. AI as a cultural producer: redefining creativity and human value

AI is no longer confined to analysis and automation. It has become a cultural producer.

AI-generated music, images, and writing are now commercially viable and increasingly indistinguishable from human work. Some AI-assisted musicians are charting; audiences listen to them as much as, or more than, human artists. Several have achieved major commercial success, including significant record-label offers.

This raises uncomfortable but necessary questions:

- If audiences enjoy AI-generated art, how does society value human creativity?
- Should AI be eligible for artistic awards?
- What happens to cultural identity when machines outperform humans in emotional resonance?

Stillwell posed a simple challenge:

"If you like the song, would it change your opinion if no human had ever sung it?"
He added:

"You might have time to create art—but that doesn't mean anyone wants to listen to your song if the AI song is better."

Leadership implication

As AI reshapes creative and knowledge work, organisations must reconsider where human distinctiveness lies—and how they value, develop, and motivate people when machines can outperform humans in domains once considered uniquely human.

5. Environmental sustainability: a more nuanced reality

AI's environmental impact often generates alarming headlines. Stillwell urged leaders to approach sustainability discussions with precision rather than fear.

He contextualised AI's energy use by comparison:

- A single AI query uses less energy than leaving Netflix running for a few seconds
- Major providers have reduced model energy use dramatically within a year

Even aggressive growth projections place AI's total energy consumption well below that of many established industries

This is not an argument for complacency, but for proportionality.

Leadership implication

Responsible AI leadership requires data-driven sustainability decisions rather than reactionary narratives. Overstating environmental impact can distort investment and policy choices just as much as ignoring it.

Leading through uncertainty

Professor Stillwell closed with a reminder that captures the leadership challenge ahead:

"Big change is probably coming. What change exactly is less clear."

In this environment, effective leadership is less about prediction and more about preparation. Organisations that thrive will:

- Build strategic agility rather than fixed plans
- Establish ethical clarity before crises emerge
- Invest in cultural and human capability alongside technology
- Develop governance structures that can adapt as AI evolves

These challenges cannot be solved by technology teams alone. They require informed, engaged leadership.

Cambridge Judge Business School Executive Education programmes are designed to help leaders develop the insight, judgment, and confidence needed to navigate this accelerating landscape responsibly and effectively.

Helping business maximise
real-world impact

UNIVERSITY OF CAMBRIDGE
Judge Business School

ExecutiveEducation

Certified B Corporation

Cambridge Judge Business School
Executive Education
University of Cambridge
Trumpington Street
Cambridge CB2 1AG
United Kingdom
www.jbs.cam.ac.uk/executive-education