UNIVERSITY OF CAMBRIDGE | Cambridge Advance Online

Five ways to create sustainable AI-driven growth

Ray Eitel-Porter

Global AI expert and Senior Research Associate, The Intellectual Forum, Jesus College, University of Cambridge

About Ray Eitel-Porter

Ray is an expert in AI safety and ethics and a Senior Research Associate at the Intellectual Forum, Jesus College, University of Cambridge. As a pioneer in the field of responsible AI, he worked on the first AI bias and fairness tool and designed and built Accenture's internal AI compliance programme.

He advises organisations across industries on how to use AI safely and responsibly; past projects include multi-year programmes at a global bank, a global retailer and a major healthcare brand. Ray has led ethical AI research collaborations with Stanford, MIT, the Alan Turing Institute and the Institute for Ethics in AI at the University of Oxford.

About Dr Professor Magda Osman

Dr Professor Magda Osman is a cognitive psychologist whose research focuses on judgement and decision-making (both basic and applied), as well as human agency and control. She has published three academic books and authored over 180 articles across psychology, management, economics, philosophy, linguistics, engineering, computer science and cognitive science.

She is currently a Research Fellow at Cambridge Judge Business School, a Research Fellow at the El-Erian Institute, University of Cambridge, and a Research Affiliate at the Centre for the Study of Existential Risk (CSER), University of Cambridge. She is also Visiting Professor of Policy Impact at Leeds Business School, University of Leeds.

About Dr Alexandru Marcoci

Dr Alexandru Marcoci is an Assistant Professor of Global Risk and Resilience in the Centre for the Study of Existential Risk (CSER) at the University of Cambridge and Course Director of the MPhil in Global Risk and Resilience. From 2025–26, he was a UKRI Policy Fellow in the Security and Online Harms Directorate at the Department for Science, Innovation and Technology.

Dr Marcoci's research portfolio includes significant contributions to methods for long-term risk management, particularly in AI, emerging technologies and systemic risks. He has published over 30 articles, including in publications such as Proceedings of the National Academy of Sciences, Nature Human Behaviour, and Artificial Intelligence.

Successful AI adoption takes a whole organisation

Too many organisations focus first and foremost on the systems rather than the people. But it's your people and how they deploy AI that is your advantage, not the AI itself.

Companies roll out systems like Copilot and are disappointed when they don't see a jump in productivity – I have heard this from so many senior executives.

But there are a number of understandable reasons this can happen:

- People are busy with their existing workload and struggle to find time to learn something new.
- They may lack the imagination to recognise how AI could transform the way they do their work.
- They may be concerned about contravening an AI policy which they haven't had time to read.
- They may be nervous about giving the impression their role can be eliminated because AI can replace them.

> Many organisations are chasing tools, pilots and productivity wins using AI at the moment. But overall I am seeing very little organisational change. What I am seeing is that AI pilots are rarely developed in collaboration with the end users, taking into account their preferences, workflows and needs. **If AI is not built with the people who will actually use it, it will not be used – no matter how impressive the technology is.**

Alexandru Marcoci
Assistant Professor of Global Risk and Resilience, Centre for the Study of Existential Risk, University of Cambridge

1 Decide what you need AI to solve

While organisations typically follow a number of AI principles, the single most important principle is to ensure that AI supports humans in what they do and that humans remain accountable for how AI works.

This means asking the question 'just because AI could do this, do we think we should use AI in this context?' at the very start of considering an AI product or service.

Then the relevant business colleagues, not just technologists, should be involved throughout the AI design and development process. They need to be convinced of the value and be active supporters, not reluctant victims. This will also help ensure that the AI system is used to best effect.

> An approach that considers a human-technological integration of AI is going to lead to a functioning and productive organisation. **The alternative will have the appearance of a very expensive vanity project.**"

Magda Osman
Research Fellow, Cambridge Judge Business School, University of Cambridge.

2 | Design the right operating model for AI

As AI becomes increasingly widespread across organisations it needs to be managed, and the risks assessed and mitigated.

AI won't deliver business value and meet its business case if it behaves in unintended ways. Establishing the right governance processes will ensure that the right questions are asked and the right tests undertaken to minimise the chances of the AI malfunctioning and harming business value.

You'll need a clearly appointed leader in the C-suite or immediately below who is accountable for implementing AI governance, has the political clout to get collaboration across functions and has the budget to support it.

> A common misconception is that AI strategy is mainly about choosing the right model or platform. In practice, competitive advantage rarely comes from the model itself. It comes from whether an organisation has the data discipline, governance, culture, processes and leadership alignment to use AI well. **The question is not 'Which AI tool should we buy?' but 'What operating model do we need for AI to create value safely and repeatedly?'**
>
> **Alexandru Marcoci**
> Assistant Professor of Global Risk and Resilience, Centre for the Study of Existential Risk, University of Cambridge

3 Give risk management responsibility to the project leader

The AI risks should be owned by the person commissioning the AI product or service, not by technology.

Most AI systems will also carry some element of residual risk, even after careful risk mitigation steps and controls have been applied.

This individual needs to satisfy themselves that enough has been done to reduce the risk to a level where the benefits significantly outweigh the remaining risks and that plans are in place to rapidly deal with any negative fallout should it occur.

> Leaders don't need to be exceptionally technically gifted. It isn't so much a capability issue when it comes to transforming an organisation through the adoption of AI. **Rather, leaders need to develop the qualities that enable them to make the best use of the people around them. Those that are technically gifted and pragmatic.**

Magda Osman
Research Fellow, Cambridge Judge Business School, University of Cambridge.

4 | Prioritise hands-on training for your organisation

A big reason for AI strategies failing is organisations not providing enough personalised training for how to use AI and how to do so safely.

Reports consistently show that over 90% of AI budgets are spent on the technology with less than 10% on training and enablement.

However, successful business leaders have recognised that someone with AI skills actually needs to sit down with business colleagues, watch how they work, and suggest where AI can help. Only then do they see significant adoption and return on investment.

Early investment in contextual, hands-on training is often the difference between AI experimentation and meaningful adoption. A good example of this is the rollout of Copilot at Paramount. Despite providing online training, it wasn't until the CIO team sat down with business colleagues and helped them to understand where AI could be used and how to avoid the risks that they saw real uptake.

> What's needed isn't unique to AI, every technological innovation that is massively disruptive to an industry or sector still needs to address the fundamentals of how to enable a functioning workforce. We aren't well calibrated to what AI systems can do. **The fall out is an under appreciation of human values that come with an overestimation of what AI can do**"

Magda Osman
Research Fellow, Cambridge Judge Business School, University of Cambridge.

5 | Create the confidence to adopt AI at scale

Drafting AI principles and policies is important but easy. The hard part which really matters is how to turn those into robust processes and controls which ensure that your use of AI at scale will deliver the intended results.

Responsible AI is often wrongly thought of as a compliance burden which will slow down innovation and AI adoption. In fact, it is what gives an organisation the confidence to adopt AI at scale, knowing it has taken all reasonable steps to ensure the AI delivers what is expected.

Yes, you need expert AI governance or a responsible AI team acting as the second line of defence to shape policies, tools, processes and so forth but you also need to equip frontline business managers with responsible AI skills.

> "AI risk governance needs to be continuous because the monitoring of AI in its adoption in an organisation will constantly be evolving because of how people co-opt it, and because the tools themselves are frequently updated. **Both will demand regular evaluation in terms of implementation, use, and ethics.**"

Magda Osman
Research Fellow, Cambridge Judge Business School, University of Cambridge.

The biggest risks of relying on AI

Placing too much trust in AI output

Automation bias is the phenomenon whereby most people naturally tend to be overly trusting of AI results and, like any scientific or technical output, assume it to be correct. Paradoxically, this risk increases as AI becomes more accurate.

As long as AI makes regular mistakes, people remain on their guard and check the outputs thoroughly but let's say the AI only makes a mistake 1% of the time, it is far easier to become complacent and assume it is always correct. Yet that 1% error could, in a medical context for example, be fatal.

Loss of key skills

Cognitive atrophy is another concern as people lose the capability to perform certain activities which are always delegated to AI. In the classroom, for example, some teachers are allowing students to use AI to write a first draft of an essay and then ask the student to critique and improve it.

While this is indeed a useful skill, it is much harder to sit in front of a blank sheet of paper and come up with a first draft – a daunting task if you have never done it before.

> People are gradually treating AI outputs as if they are inherently reliable because they are fast, fluent and confident. That creates a subtle but serious drift in accountability.
>
> Teams stop interrogating outputs, junior employees stop building foundational judgement, and poor decisions become harder to trace because responsibility is diffused between human and machine. **This is not just a technical risk – it is a cultural and managerial one.**

Alexandru Marcoci
Assistant Professor of Global Risk and Resilience, Centre for the Study of Existential Risk, University of Cambridge

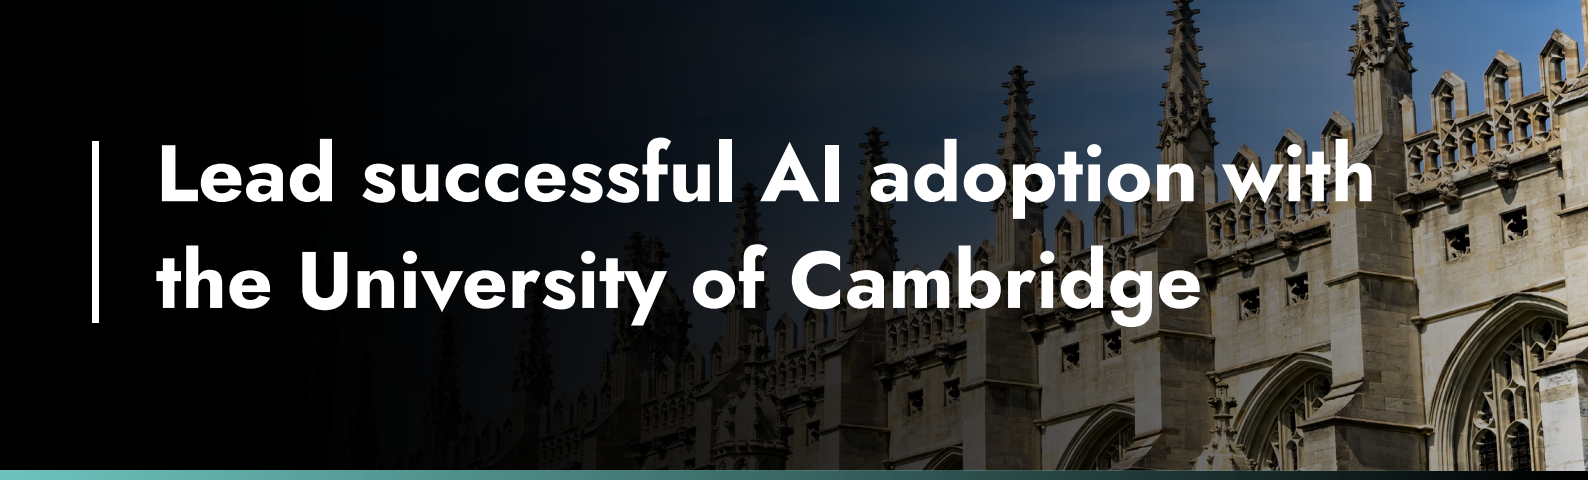

Lead successful AI adoption with the University of Cambridge

The challenge for leaders is no longer access to ideas, but the judgement to know which ones will matter over time.

Shaped by University of Cambridge research and engagement, our courses and programmes reflect how ideas are debated, tested and applied within organisations, helping leaders navigate change with confidence.

The AI Leader: From Strategy to Transformation

This six-month programme gives senior leaders the skills and confidence to lead AI transformation — from shaping AI strategy to guiding organisational change. Gain advanced insight into AI opportunities and risks from University of Cambridge academics whose research is shaping how AI is understood, governed and applied across industries, and design an AI strategy tailored to your organisation.

Skills-focused learning for your workforce

Our courses and programmes support professionals across your organisation to build practical, job-relevant capability - strengthening skills, confidence, critical thinking and ways of working as technologies and tools continue to evolve.

Create a bespoke learning programme

Drawing on expertise from across the University of Cambridge, we develop tailored, blended and custom learning programmes designed to meet your organisation's unique needs and skill gaps - from agile micro-learning for large-scale rollouts to programmes for boards and senior leaders.

UNIVERSITY OF CAMBRIDGE | Cambridge Advance Online

Find out more at

executiveprofessional.cam.ac.uk/ai-leadership