

Insights
from CACTUS
Dialogues
London 2026

CACTUS[®]

Executive Summary

Artificial intelligence has entered a new phase of organisational life. What began as experimentation at the margins of organisations has rapidly become a board-level concern, shaping decisions around investment, governance, risk and long-term competitiveness. At CACTUS Dialogues London, senior leaders from publishing, finance, insurance, life sciences, tech and pharma came together to discuss a challenge that increasingly transcends industries: how organisations can embrace the opportunities presented by AI while preserving trust, quality and integrity.

The conversations throughout the afternoon made one thing clear: the challenge of AI is no longer primarily technological. It is organisational, cultural and deeply human. Discussions repeatedly returned to questions of leadership, judgement, accountability and trust. While industries may differ in their operating realities, many of the concerns raised around the room were remarkably similar. How should leaders make decisions when technology is evolving faster than organisational norms? What role should human expertise continue to play? How can organisations innovate responsibly while maintaining confidence in the quality of their outputs?

Five themes emerged from the forum.

01

AI has moved decisively from experimentation to executive accountability.

02

Trust remains the principal barrier to adoption, not technological capability.

03

Human expertise becomes more valuable as AI scales.

04

Organisations must move beyond pilots and towards disciplined implementation and governance.

05

Competitive advantage may belong not to the fastest adopters, but to those that become trusted navigators of uncertainty.

This paper brings together the key ideas, questions and insights that emerged during CACTUS Dialogues London and explores what they may mean for leaders preparing their organisations for an increasingly AI-enabled future.

Table of Contents

Introduction: Why This Conversation Matters	04
Different Industries, One Question	05
Opening Perspective Akhilesh Ayer on the Leadership Challenge Ahead	06
The Real Bottleneck Is Trust	07
Keynote Katie King on Leading Through AI Transformation	08
Industry Perspective Peter de Jong on Building Trustworthy AI in Publishing	09
Panel Discussion Nikesh Gosalia on Leadership, Trust and Accountability in an AI-Enabled World	10
Conclusion: The Advantage of Being a Trusted Navigator	12



Introduction

Why This Conversation Matters

The conversations surrounding artificial intelligence have changed dramatically over the past two years. Not long ago, AI was largely viewed as an experimental technology, explored through isolated pilots and championed by innovation teams operating at the edges of organisations. Today, the conversation looks very different. AI increasingly appears on board agendas, features prominently in discussions around risk and competitiveness, and influences decisions concerning long-term organisational strategy. In many ways, AI has ceased to be a technology conversation and has become a leadership one.

Our London chapter was designed against this backdrop. The forum brought together leaders from publishing, finance, insurance, life sciences, tech and pharma to step back from day-to-day operations and engage in candid discussions about how AI is reshaping executive decision-making not only operationally, but strategically, culturally and ethically. The emphasis was not on products or technical capabilities. It was on insight, perspective and peer-level dialogue.

Throughout the day, one idea surfaced repeatedly: the questions facing organisations may differ in language and context, but they are increasingly shared in substance.

Whether organisations are grappling with research integrity, decision-making under risk or balancing speed with rigour, they are all ultimately asking the same question: how do we embrace the opportunities presented by AI while preserving trust, quality and integrity?

Different Industries, One Question



One of the defining characteristics of the forum was the diversity of industries represented in the room. Publishing leaders discussed integrity at scale and the growing strain on peer review. Financial institutions explored decision-making in increasingly uncertain environments. Insurance organisations considered the challenge of pricing the unknown, while life sciences companies examined the tension between speed and scientific rigour.

Despite these differences, the conversations revealed a remarkable convergence in leadership concerns. Organisations across sectors are grappling with how to deploy AI responsibly, how to preserve trust in environments of increasing complexity, and how to make decisions when both the technology and its implications continue to evolve rapidly.

The value of bringing together leaders from different sectors lay precisely in recognising that the challenge of AI is not unique to any one industry. It is increasingly a cross-sector leadership challenge that demands new forms of judgement, governance and organisational readiness.

Akhilesh Ayer, a man with dark hair, is wearing a dark blue suit jacket over a light-colored shirt. He is holding a silver microphone in his right hand and a white paper in his left hand. He is looking slightly to the right of the camera with a slight smile. The background is a blurred indoor setting with large windows and curtains.

Opening Perspective

Akhilesh Ayer on the Leadership Challenge Ahead

In his opening remarks, Akhilesh framed the central challenge facing organisations today: how can leaders harness the opportunities presented by AI while preserving trust, quality and integrity? His message was clear—the challenge lies in determining when and how organisations can trust the technology enough to embed it into consequential decisions.

The forum repeatedly returned to the speed with which AI has moved from experimentation to executive accountability. Only a few years ago, AI projects often existed at the periphery of organisations, led by small teams conducting isolated pilots. Today, AI has become a board-level concern, appearing on strategic agendas and influencing conversations around investment, governance and organisational resilience. The journey from experiment to strategic bet has happened remarkably quickly.

One statistic shared during the forum captured the scale of this shift particularly well: half of chief executives now say that their own job security depends on getting AI right. The question has moved decisively up the organization from “Should we try this?” to “How do we make decisions our organisations can confidently stand behind?” The consequences of getting these decisions wrong are no longer operational inconveniences. They are strategic risks with implications for trust, reputation and competitive position.

The Real Bottleneck Is Trust



For all the extraordinary capabilities demonstrated by modern AI systems, the discussions repeatedly returned to a more fundamental issue: trust.

Throughout the afternoon, trust emerged in four distinct forms. Leaders must trust the data that underpins AI systems and understand where it originates. They must trust the quality and reliability of AI-generated outputs and determine whether those outputs can withstand expert scrutiny. They must trust the organisations and vendors providing AI capabilities. Most importantly, they must trust their own judgement, knowing when to rely on machine recommendations and when to challenge them.

Audience polling reinforced the significance of this issue. Among the various challenges associated with AI adoption, participants identified protecting trust, quality and integrity as the issue that resonated most strongly with their own organisations. The finding is significant because it fundamentally changes the nature of the AI conversation. The barriers to adoption are no longer primarily technological. They are organisational, cultural and human.

The implications are profound. Organisations that succeed in an AI-enabled future are unlikely to be those that simply move the fastest. They may instead be those that establish the governance, accountability and transparency necessary to build trust among employees, customers, partners and regulators.

A photograph of Katie King, a woman with blonde hair, wearing a white jacket, speaking into a microphone. The image is partially obscured by text on the left side.

Keynote

Katie King on Leading Through AI Transformation

Katie's keynote challenged leaders to think beyond the excitement surrounding AI and towards the practical realities of implementation. Her central argument was that sustainable value does not come from adopting every new tool or pursuing automation for its own sake. Meaningful value is created when organisations integrate AI thoughtfully into their people, processes and decision-making structures.

Several principles emerged from the keynote. AI should augment expert judgement rather than replace it. Governance should be viewed as an enabler rather than an obstacle. Organisations should focus on real problems rather than technology in the abstract. And transformation depends far more on organisational readiness than on technology itself.

The now-familiar 10-20-70 rule reinforced this message.

10% of successful transformation lies in algorithms. **20%** lies in technology and data. **70%** depends on people and process.

The technology, as Katie noted, is often the easy part. The more difficult work lies in leadership alignment, change management, culture and creating environments in which experimentation can occur safely and responsibly.

Her keynote also emphasised that organisations are entering a new phase in their AI journey. The era of fascination and experimentation is giving way to one of discipline and implementation. The organisations that derive sustained value from AI will likely be those that start with one real problem, establish clear ownership, build AI literacy and create environments in which it is safe to try and safe to learn. The challenge, ultimately, is not keeping pace with technology, it is preparing organisations to change with it.

A photograph of Peter de Jong, a man with long hair and a beard, wearing a dark suit jacket over a black shirt. He is holding a microphone in his right hand and gesturing with his left hand. The background is blurred, showing what appears to be a conference or meeting setting.

Industry Perspective

Peter de Jong **on Building Trustworthy AI in Publishing**

Drawing on six months of conversations with more than twenty publishers across scholarly and medical publishing, Peter's research explored how senior leaders are thinking about AI in the submission and peer-review process. His findings suggested that organisations are increasingly motivated by three factors: cost, capacity and the need for more consistent quality checks. At the same time, concerns around accuracy, output quality and review capacity continue to dominate leadership discussions.

One of the most compelling findings from the research was the absence of a shared playbook. Organisations continue to wrestle with questions around whether to build, buy or partner; how to balance rigour with speed; whether governance should be centralised or distributed; and how to move from pilots to scalable implementation.

The research also highlighted the growing strain on the scholarly record itself. Submission volumes continue to rise, reviewer capacity remains constrained, and publishers face increasing pressure to maintain trust and integrity at scale. As AI becomes more deeply embedded across workflows, validating outputs and maintaining accountability become just as important as generating them.

Perhaps the strongest message from Peter's research was that while AI is already firmly embedded in many workflows, human oversight remains essential.

The question is no longer whether there should be a human in the loop, but where that human intervention should sit and how it can scale. Increasingly, proof becomes the new standard: the ability to demonstrate what the AI did, who checked it, and how accountability was maintained throughout the process.

A photograph of Nikesh Gosalia, a man with dark hair and glasses, wearing a dark blue suit jacket over a light blue shirt. He is holding a black microphone in his right hand and a smartphone in his left hand. He is sitting in a light-colored armchair. The background is a bright, out-of-focus window with greenery visible outside.

Panel Discussion

Nikesh Gosalia **on Leadership, Trust and Accountability in an AI-Enabled World**

The panel discussion, led by Nikesh and featuring Katie and Peter, brought together many of the themes that had surfaced throughout the day and translated them into a series of practical leadership questions. Rather than focusing on technology alone, the conversation examined the decisions, behaviours and organisational realities that determine whether AI initiatives succeed or fail.

A recurring theme throughout the discussion was the relationship between ambition and validation. Organisations are moving quickly to embed AI into their workflows, yet many are still determining how to trust the outputs these systems produce. This tension gave rise to a broader question that resonated across industries: are organisations moving faster than their ability to govern and validate the technologies they are deploying?

The panel also challenged assumptions around leadership and accountability.

If, as Katie's framework suggests, the majority of successful transformation depends on people and process rather than technology, are most AI failures ultimately leadership failures wearing a technology costume?

Are organisations too quick to blame the model when they should instead be examining the decisions, incentives and governance structures that surrounded its deployment?

Questions of trust surfaced repeatedly. Is the growing pressure on chief executives to get AI right creating better decision-making, or is it encouraging defensive behaviour and risk aversion? What should leaders make of deliberate non-adopters who choose to move cautiously rather than quickly? Are they falling behind, or are they exercising sound judgement in the face of uncertainty?

The discussion also examined the increasingly popular idea of keeping a human in the loop. Every industry represented in the room claimed to have meaningful human oversight of AI-assisted decisions. Yet the panel questioned whether, in practice, human reviewers are genuinely exercising judgement or simply approving recommendations generated by machines. At what point does human oversight become oversight theatre? And if organisations cannot answer that question clearly, what does it mean for accountability and trust?

Perhaps the most valuable aspect of the discussion was that it produced few definitive answers. Instead, it surfaced a set of questions that are likely to shape leadership conversations for years to come:

- Can organisations achieve scale without compromising trust and quality?
- Will innovation itself change if discoveries fail to fit established models?
- What work belongs to the machine, and what work should remain distinctly human?
- Should organisations begin to redefine themselves in response to changing technological realities?
- Is deliberate caution a sign of falling behind, or a sign of thoughtful leadership?



The panel concluded with one final question that lingered long after the discussion ended:

In two years, will any of our industries be talking about AI trust as a solved problem, or will we still be having a version of this exact same conversation?

An impromptu contribution from Tejas Joshi added another dimension to the afternoon, prompting further discussion and reinforcing one of the forum's defining characteristics: some of the most valuable insights emerge when leaders are willing to challenge assumptions and learn from perspectives beyond their own.



Conclusion

The Advantage of Being a Trusted Navigator

CACTUS Dialogues London demonstrated that AI is no longer simply a technology issue. It is a question of leadership, governance, trust and judgement.

The technology itself will continue to evolve, and new capabilities will continue to emerge. The more enduring challenge lies in how organisations choose to respond. The discussions throughout the day suggested that the organisations best positioned to succeed may not be those that adopt AI first or move fastest. They may instead be those that distinguish signal from noise, preserve integrity while innovating, and make decisions that their organisations can defend with confidence.

As Akhilesh observed in his opening remarks, and as the conversations throughout the day repeatedly reinforced:

The advantage no longer belongs to the fastest adopter. It belongs to the most trusted navigator.

CACTUS[®]

CACTUS offers expert services and AI-driven products that help improve how research gets funded, published, communicated, and discovered. For 20+ years, we have helped researchers, publishers, universities, government bodies, pharmaceutical companies, biotech companies, and other key players communicate their impact to the world.

To know more about our solutions and explore how we can support you, visit www.cactusglobal.com or reach out to us at partnership@cactusglobal.com.



www.cactusglobal.com



[@Cactusglobal](https://twitter.com/Cactusglobal)



[@Cactus Communications](https://www.linkedin.com/company/Cactus-Communications)



[@cactus.communications](https://www.instagram.com/cactus.communications)



[@cactusglobal.bsky.social](https://bsky.app/profile/cactusglobal.bsky.social)



[@CactusCommunications](https://www.facebook.com/CactusCommunications)



partnership@cactusglobal.com