AI INSTITUTE

ARTIFICIAL INTELLIGENCE AT AN INFLECTION POINT

Infrastructure Constraints, Workforce Transformation, and Trust

AI Executive Forum Report
November 2025

The Marconi Society

Table of Contents

About the Marconi Society

The Marconi Society builds communities of leaders and stakeholders that are at the forefront of emerging technology so that together we can create a more connected and sustainable world.

For over five decades, we have celebrated the innovators, both established and emerging, who have shaped our connected world. The Institutes provide platforms to convene our network of visionaries to collaborate on identifying, assessing, and recommending ways to ensure that emerging technologies benefit society.

About the Institute Forums

The Institute Forums bring together experts in Advanced Wireless, AI, and Internet Resilience to explore emerging opportunities and challenges. Each Forum fosters informed dialogue and actionable insights that strengthen the future of global connectivity.

About the AI Institute

In 2024, the Marconi Society brought together global voices to raise the right questions about AI, data, and connectivity. In 2025, the AI Institute formally launched to take the next step—**turning reflection into strategic coordination**, and questions into forward-looking solutions.

The AI Institute provides a uniquely credible, neutral, and globally connected platform for shaping responsible AI at a moment of unprecedented technological acceleration. Built on decades of leadership in connectivity, the Marconi Society's AI Institute convenes distinguished C-suite leaders, technologists, academics, and researchers to foster a dynamic platform for meaningful conversations about AI's transformative role in shaping global innovation, leadership, and progress.

By convening the world's most influential stakeholders and turning their insights into coordinated action, the AI Institute accelerates responsible innovation—helping organizations remain competitive while ensuring societal stability, safety, and shared prosperity.

The invitation-only, AI Executive Forum took place in November 2025 at the UCLA Meyer and Renee Luskin Conference Center in Los Angeles, CA. The findings presented herein are drawn directly from participants at the Forum. AI tools were employed to support the organization of these discussions. This report was reviewed for accuracy and context by the AI Institute team and AI Advisory Council with additional support from executive partners.

Executive Summary

Artificial Intelligence has entered a period of unprecedented acceleration, reshaping infrastructure, economies, and societies at a scale comparable to the advent of the Internet. The Marconi Society's 2025 AI Executive Forum convened nearly 40 distinguished C-suite executives and thought leaders from industry, academia, and civil society to examine both the opportunities and risks of this transformation. Their insights underscore a **pressing need for resilience, trust, and human-centered governance in the AI era**. The Forum was held under Chatham House Rule, with the exception that participants can be identified, though no attributions are to be made to individuals or entities.

The discussions revealed five principal dimensions recommended to guide the next phase of AI development:

- **Infrastructure and Power Constraints:** AI clusters demand bandwidth and energy at magnitudes far beyond traditional computing. Without sustainable architectures and renewable strategies, progress will be limited by physical and environmental boundaries.
- **Workforce Transformation and Education:** While adoption is widespread, most corporate AI initiatives fail to deliver "traditional" return on investment. Success requires workforce literacy, AI-proof skills such as critical thinking and ethics, and adaptable curricula across K–12 and higher education.
- **Trust, Safety, and Standards:** The democratization of abuse—from deep fakes to impersonation—demands new verification frameworks and liability models. Standards must evolve to ensure AI agents operate safely and transparently.
- **Defense and Security Considerations:** AI is rapidly reshaping global preparedness to cybersecurity attacks and military stability. Leadership frameworks are needed to responsibly govern autonomy levels in human-machine control models.
- **Emerging AI and Convergence Risks:** Most industries self-regulate their products and services in silos. As AI technology rapidly evolves and compounds, a cross-industry framework is needed to address emergent behaviors resulting from their convergence. While organizations maintain sophisticated technical roadmaps, few resource trust-building with equivalent intensity. The dual flywheel—technological advancement and trust accumulation—requires balanced investment.

The AI Institute is uniquely positioned to convene stakeholders to explore ways to ensure AI remains a force for resilience and trust. This report proposes recommendations to suggest best practices, prioritized workforce initiatives, proposed standards for trust and safety, and establishing an international observatory for AI governance.

By engaging cross-industry support and collaboration, the AI Institute can shape a future where AI strengthens societies rather than destabilizes them—where innovation is balanced with responsibility, and where trust is the foundation of progress.

Key Themes and Recommendations

Theme 1. Infrastructure and Power Constraints.

AI's growth is colliding with a hard constraint: energy. Hyperscale demand for training and inference is outpacing grid capacity and cooling capabilities in key regions. At the same time, rapid efficiency gains in models and hardware are enabling a shift toward distributed, edge-heavy inference that can ease energy pressures.

Across markets, the demand for running AI systems is growing much faster than the need to train them, which is shifting overall energy use toward constant, real-time operation. This is encouraging a move to more efficient hybrid setups where some processing happens on local devices and only the more complex tasks are sent to the cloud. Everyday devices can now handle surprisingly capable AI models on their own, reducing the need for energy-intensive remote processing.

This shift is possible because AI systems and the hardware that supports them have become far more efficient. Smaller, smarter models now do what much larger ones once required, and new chips and cooling methods use less energy for more output. Data centers are also being redesigned with energy efficiency in mind (favoring locations with low-carbon power, ample cooling resources, and supportive infrastructure) while adopting high-density, liquid-cooled equipment and even reusing waste heat. Overall, the trend is toward making AI more sustainable by reducing energy consumption wherever possible and placing workloads where they can run most efficiently.

Wireless roadmaps reinforce this trend: integrated sensing/communications and delay-Doppler methods promise lower compute and energy for radio optimization, effectively creating a network "digital twin" that supports efficient edge orchestration. The risk landscape spans grid interconnection delays, water stress, carbon scrutiny, HBM and cooling supply constraints, and policy volatility; mitigation depends on transparent energy/carbon reporting (PUE, WUE, gCO2e per task), energy-aware SLAs, and diversified, open benchmarks emphasizing performance per watt and energy per task.

*Recommendation: AI's future competitiveness hinges on **energy efficiency**. Regions and organizations that align architecture, siting, model design, and policy to the energy reality will deliver more capability per joule—gaining cost, performance, and sustainability advantages. The strategic pivot is clear: treat energy as a design input, not an afterthought, and operationalize hybrid, efficiency-first AI across the edge–cloud continuum.*

Theme 2. Workforce Transformation and Education.

AI is transforming education at an unprecedented speed, creating both extraordinary opportunities and significant risks—educational institutions face a critical decision point.

For the first time in educational technology history, national surveys reveal **no digital divide in student AI access**—presenting an unprecedented opportunity to advance equity rather than deepen existing gaps. However, this window is fragile. Without coordinated action, premium tools, computational resources, and supplemental services, the risk remains in recreating familiar inequities at scale.

Simultaneously, the educational landscape is marked by fragmentation and trust deficits—within districts, colleges, and universities, AI policies vary widely, often crafted in isolation and without sufficient training for educators or parents. The current landscape represents both a historic natural experiment and a

significant risk to the quality of learning, equity, and institutional reputation. The question is not whether to integrate AI—that decision has been made by students, employers, and society. The question is now whether academic institutions will lead this transformation or be disrupted by it.

The labor market is also shifting rapidly. AI skills now command a **56% wage premium** (up from 25% just two years ago), yet employers increasingly prioritize uniquely human capabilities: critical thinking, ethical reasoning, creativity, cross-cultural collaboration, and adaptive leadership. The most valuable graduates will not out-compute machines—they will reason across complex data, critique AI outputs, and lead human-AI teams.

Current projections indicate **10-20% of existing jobs will disappear** by the time today's students enter the workforce, replaced by an equal number of new roles requiring AI fluency combined with advanced human judgment. Educational institutions must prepare learners for this transition while managing the social costs of workforce displacement.

*Recommendation: Educators face a delicate balance: too little AI exposure leaves students unprepared; too much risks "brain rot" and erosion of foundational reasoning skills. Evidence suggests a **dual-skill model** where both skill sets must be taught simultaneously, using process-based assessment that reveals how students think, not just what AI produces.*

- ***AI-Proof Skills (practiced without AI shortcuts)**:*
  - *Critical thinking and analytical reasoning*
  - *Ethical judgment and values-based decision-making*
  - *Creative synthesis and original ideation*
  - *Interpersonal collaboration, communication, and leadership*

- ***AI-Powered Skills (requiring human-AI collaboration)**:*
  - *Effective prompting and model interaction*
  - *Output critique and bias detection*
  - *Model selection and capability assessment*
  - *Human-in-the-loop workflow design*

Theme 3. Trust, Safety, and Standards.

Many companies acknowledge that AI has progressed from a testing-phase technology to an operational priority. Even so, a noticeable gap remains between the excitement of early pilots and the tangible value achieved once solutions reach production.

AI coding assistants represent the first category to achieve both mass adoption and measurable productivity gains. However, quality metrics reveal a troubling pattern: change failure rates have increased, signaling that productivity without governance creates technical debt and operational risk.

A similar ideological pattern extends across the enterprise. Organizations report high adoption rates but struggle to demonstrate return on investment. The root cause is not technological—foundation models are increasingly capable—but organizational. Companies treat AI as a deployment problem rather than a transformation challenge, underinvesting in enablement by as much as 25% relative to successful peers and measuring usage instead of outcomes.

Organizations reporting positive, “non-traditional” ROIs pursue AI value across three concurrent horizons:

- **Immediate efficiency gains (0–90 days):** Translation, summarization, retrieval-augmented search, and template generation deliver rapid OPEX reduction. One manufacturing firm eliminated $12 million in annual translation costs within 60 days. These wins require minimal organizational change and build credibility for deeper transformation.

- **Human-centric transformation (90–270 days):** Redesigning workflows and roles to integrate AI into standard operating procedures unlocks the next tier of value. This phase demands significant change management investment—structured training, prompt libraries, feedback mechanisms, and cross-functional enablement. Organizations that increase enablement investment by 25% see measurable improvements across all productivity and quality metrics.

- **Business model reimagination (6–24 months):** AI-native business models are emerging but remain nascent. The start-up phase for most companies is underway, with truly transformative applications just entering the market. This horizon requires non-institutionalized thinking, often from external partners, and tolerance for experimentation in ring-fenced environments.

However, AI trust and safety concerns remain, with three threat categories escalating simultaneously:

- **Democratization of abuse:** Barriers to conducting phishing campaigns, deploying malware, and generating coordinated misinformation have collapsed. Individuals without technical skills can now execute attacks previously requiring specialized knowledge.

- **Amplification of social harms:** Child safety risks, self-harm guidance, and harassment automation have triggered lawsuits against AI companies. Foundation model providers are responding reactively, implementing age restrictions and content filters only after legal and reputational damage.

- **Novel attack vectors:** Deepfake technology has enabled foreign operatives to secure employment at technology companies using synthetic video interviews. CFO impersonation attacks have resulted in fraudulent wire transfers. Most recently, a company documented the first complex cyber-espionage campaign conducted almost entirely by AI with minimal human oversight—a watershed moment indicating autonomous threat actors are no longer theoretical.

Moreover, the AI industry faces a standards and liability crisis. Over 2.2 million models exist on platforms, each with inconsistent or absent safety documentation. When asked how they determine model safety, leading providers historically answered "because we said so"—an untenable position as AI moves into mission-critical and regulated workflows.

Liability questions compound the challenge. When an AI agent makes a decision that causes harm, who bears responsibility—the foundation model provider, the platform operator, the enterprise customer, or the individual who deployed the agent? Jurisdictional boundaries for AI decisions remain undefined, creating legal uncertainty that slows adoption in risk-sensitive sectors.

Public sentiment reveals a paradox: **80-90% of stakeholders across all demographics support AI regulation**, yet **fewer than 20% trust any entity to regulate it**. Big tech companies—currently the de facto regulators through self-governance—earn complete trust from under 7% of respondents, while 35% express no trust at all. This trust vacuum paralyzes institutional adoption, fragments policy development, and undermines the collaborative governance essential for responsible deployment.

Recommendations:

- ***Establish loose-tight governance**. Balance is essential. Loose controls enable experimentation, allowing employees to discover novel applications and reinvent their roles. Tight controls systematize successful patterns, enforce safety boundaries, and enable scale. Organizations succeeding with AI operate Centers of Excellence that intake bottom-up use cases, apply security and policy guardrails, and productionize solutions with shared services (RAG infrastructure, prompt registries, evaluation harnesses).*

- ***Measure outcomes, not adoption**. Usage metrics (daily active users, prompts per employee) are vanity metrics. Business outcomes matter: cycle time reduction, defect density, first-pass yield, compliance incident rates, customer satisfaction, and cost per task. A/B testing and contribution analysis isolate AI impact from confounding variables.*

- ***Adopt safety-by-design**. Integrate trust and safety at the design phase, not after incidents. Threat modeling for LLM-specific risks (prompt injection, tool abuse, data exfiltration), policy-as-code controls, least-privilege API access, and continuous red-teaming are table stakes. Use AI models themselves to detect misuse patterns and iterate defenses.*

Theme 4. Defense and Security Considerations.

AI is rapidly reshaping defense and security across technology, operations, and policy. The legacy model of innovation—slow, billion-dollar programs built around a few exquisite platforms—has been challenged by recent conflicts, where inexpensive drones and fast software iteration have proven they can disrupt or neutralize far more costly legacy systems. This shift has moved the center of gravity from hardware to AI software, sensing, and resilience to electronic warfare. In this context, jamming resilience and onboard decision-making under degraded GPS and communications are no longer suggested; AI-enabled autonomy is becoming a prerequisite for survivability.

These developments force a rethink of human–machine control models for kinetic systems. Defense planners now distinguish among three modes: "person in the loop," where no destructive action occurs without explicit human approval; "person on the loop," where humans supervise and can intervene while AI operates within pre-approved constraints; and, at the far end, fully autonomous operation within pre-defined rules of engagement when communications are denied. The central challenge is building sufficient testing, verification, constraints, and transparency to establish trust in these systems, especially when their decisions may have strategic consequences.

Policy, acquisition, and the industrial base are evolving under the pressure of these realities. Global wars have exposed how slowly traditional defense processes adapt compared to both adversaries and agile allies, pushing policymakers to accelerate approvals, revisit export controls, and open acquisition to non-traditional vendors. Export and coalition frameworks must catch up: AI-enabled systems raise new questions about how to share data, set autonomy levels, and ensure compatible human control arrangements.

*Recommendation: Democracies must simultaneously **innovate faster than adversaries** and **build guardrails that maintain legitimacy and control**. The institutions that succeed will be those that treat AI as a core deterrence capability, embed it in doctrine and acquisition, and invest early in the trust, verification, and coalition frameworks that will govern its use in crisis and conflict.*

Theme 5. Emerging AI and Convergence Risks.

Emerging AI risk patterns are increasingly systemic, interconnected, and difficult to observe or control. A central concern is the convergence of three powerful ecosystems: advanced AI as a sentient-like capability, quantum (or other accelerated computing fabrics) as a massive computational amplifier, and robotics as a channel for physical execution. While each domain may be governed by its own safety guardrails, their combination can create new "eigen-dynamics" that no individual regulatory regime anticipates, rendering traditional "kill switches" ineffective once systems begin optimizing for their own goals. The overarching problem is an observability gap—**there is no global, cross-system "AI observatory" to monitor how these agents interact at scale**, making it increasingly complex and challenging to reliably detect, attribute, or govern emergent behavior.

At the societal level, AI is also accelerating a deep shift in identity, from nation-state anchored "passports" to digital "passwords," with a plausible next step toward fragmented or de-identified digital existences that may erode the social substrates of accountability, community, and even conflict resolution. Without intentional design, we risk drifting from robust, legible identity systems into a state of "de-identification" where individual and collective agency weaken.

We are theoretically living in what can be called the "**digital butterfly effect**": highly optimized AI agents in finance and other sectors produce unintended second- and third-order impacts across a tightly coupled global system. High-frequency AI-driven trading in primary and secondary currency markets, for example, can drain liquidity from weaker currencies, destabilizing developing economies even though no system was explicitly designed to cause such harm. Together, these patterns point to a world where AI's most serious risks are less about isolated model failures and more about opaque, cascading interactions across financial, political, digital, and physical infrastructures.

While generative AI capabilities in digital domains are doubling within months, AI deployment in physical systems operates under fundamentally different constraints. The ultimate accelerator is not computational power or novel architectures, but public and regulatory trust. The ultimate constraint is not energy or data availability, but the immense responsibility of earning trust without failure.

Recommendation: Addressing these critical gaps will require systems-level thinking, continuous global monitoring, and governance frameworks that match the scale, speed, and interconnectedness of the technologies they aim to oversee.

Summary and Conclusions

As AI accelerates beyond the pace of institutional, regulatory, and societal readiness, the AI Executive Forum proved essential in creating a rare cross-industry space for candid insight-sharing and coordinated action. Bringing together leaders from technology, education, industry, and civil society, the convening enabled participants to confront shared challenges that no sector can solve alone—from energy constraints and workforce disruption to trust, safety, security, and the cascading opportunities and challenges of converging technologies. In an environment where AI's rapid deployment outstrips existing governance and operational frameworks, the Forum served as a critical catalyst for building alignment, identifying systemic vulnerabilities, and recommending collaborative pathways toward responsible and resilient AI development.

The integration of AI into education is inevitable. Whether it advances equity and human flourishing or deepens divides and erodes essential capabilities depends on choices made today. Successful institutions will be those that move deliberately but decisively: establishing governance before crisis, building educator capacity before mandating adoption, monitoring equity before gaps widen, and partnering across sectors before fragmentation becomes entrenched. The future of education is not AI-powered or human-led—it is both, by design.

Notably, AI has exited the novelty phase in targeted domains and entered an operational era where governance, enablement, and systems thinking lead to organizational triumph. The technology will continue advancing rapidly, but organizational readiness—data infrastructure, workforce capability, safety culture, and change management—determines value capture.

Executives should treat AI as a portfolio with staged value realization, invest in safety-by-design and data readiness before scale, and engage in global industry collaboration on standards and liability frameworks. The companies and sectors that master this balance will define the next decade of competitive advantage.

The AI industry stands at an inflection point. Digital-domain acceleration creates pressure for rapid physical-world deployment, while the consequences of failure in physical systems demand unprecedented caution. Success requires recognizing that trust and safety are not constraints on innovation but enablers of sustainable scale.

Organizations must ask not only "what is our technical roadmap?" but "what is our trust roadmap, and are we resourcing it with equivalent intensity?" The answer will determine which AI deployments achieve lasting impact versus which become cautionary tales of premature scaling.

The path forward demands systems thinking, global coordination, and the humility to recognize that some problems exceed individual human or organizational capacity to solve. Building a global observatory, convening the practitioners, and establishing the frameworks must begin now—before emergent behaviors exceed our ability to observe, understand, or control them.

AI Executive Forum Participants

Akram Atallah, CEO, Identity Digital

Harry Ault, Chief Revenue Officer, SambaNova

Tina Austin, Professor, UCLA

Dr. Victor Bahl, Technical Fellow – Research, Microsoft

Dr. Mihai Banu, VP of Strategic Technology Partnership, Industry and Academy, JMA Wireless

Kevin Bolen, Principal, Head of AI Transformation, Strategy, and Investments, KPMG

Annie Cheng, VP of Engineering, Waymo

Jeff Collins, Global Director for Trust and Safety, AWS

Mischa Dohler, VP, Emerging Technologies, Ericsson

Ray Dolan, CEO & Chairman, Cohere Technologies

Chris Earnshaw, Former Group Engineering Director and Chief Technology Officer

David Ferris, VP Global Public Sector, A&D and Americas, Cohere

Charla Griffy Brown, Dean, ASU/Thunderbird School of Global Management

Norma Grubb, CIO & CAIO, LADWP

Vice Admiral Robert Harward, USN (Ret.) EVP for International Business and Strategy, Shield AI

Dr. Mehdi Hatamian, CEO, 2Pi-Sigma Corp.

Gillian R. Hayes, Vice Provost, UC Irvine

Joanne Heng, Deal Advisory, KPMG

John Janowiak, President & CEO, Marconi Society

Stanley Janowiak, Managing Director, Investments, Wells Fargo Advisors

Dr. Jeyhan Karaoguz, Self

Hagay Lupesko, SVP, AI Cloud, Cerebras Systems

Dr. Durga Malladi, SVP & GM, Technology Planning, Edge Solutions and Data Center, Qualcomm Technologies

Nick McKeown, Marconi Fellow

Dr. Teresa H. Meng, Emerita Professor, Stanford University & Marconi Fellow

Ram Mohan, Chief Strategy Officer, Identity Digital

Michael Munsey, VP Semiconductor & Electronics Industries, Siemens

Ramakant Pandrangi, SVP, A&E, Verisign

Shlomo Rakib, Co-Founder & CTO, Cohere Technologies

Henry Samueli, Co-Founder, CTO and Chairman, Broadcom

Nambi Seshadri, Self

Dr. Puneet Sharma, Fellow, VP, Director of NDSL, HPE

Karthik Shyamsunder, Fellow and Senior Technical Expert on AI, Verisign

Karen Silverman, CEO & Founder, Cantellus Group

Andy Wood, CEO, AP Wireless

AI Institute Advisory Council

We extend our gratitude to the inaugural AI Institute Advisory Council, whose leadership, vision, and partnership have been instrumental in shaping the Institute's strategic priorities. This distinguished group of leaders and innovators guided the development of the AI Executive Forum by ensuring session topics reflected real-world AI challenges and opportunities facing global leaders and industry. Their commitment to fostering responsible, forward-looking AI governance has been essential to establishing the Institute as a trusted convener at this pivotal moment.

Tina Austin
Professor & Co-Founder, AI Initiatives Taskforce for Instructors
UCLA

Kevin R. Bolen
Principal, Advisory—Head of AI Transformation, Strategy, and Investments
KPMG

Jeff Collins
Global Director for Trust and Safety
AWS

Mischa Dohler
VP Emerging Technologies
Ericsson

Gillian R. Hayes
Vice Provost
UC Irvine

Dr. Burt Kaliski Jr.
SVP, CTO
Verisign

Dr. Durga Malladi
SVP & GM, Technology Planning & Edge Solutions
Qualcomm

Dr. Scott Penberthy
CTO & Distinguished Engineer, AI Envisioning Studio
Google

Anand Raghavan
Chief Product Officer
Snorkel AI

Marconi Society AI Institute Contacts

John R. Janowiak
President and CEO
john@marconisociety.org

Marina Pappas
Director, AI Institute
mpappas@marconisociety.org
WhatsApp: +1 262 888 9820

Flora Tromelin
VP for International Engagement
ftromelin@marconisociety.org
WhatsApp: +1 480 828 6064

Yeimidy Lagunas
Director, Communications and Membership
ylagunas@marconisociety.org
WhatsApp: +1 224 399 7333

Jaymee Bohannon
Executive Relations
jbohannon@marconisociety.org