

Ai4 Conference: AI and Innovation

August 4-6, 2026

Venetian Expo Center, Las Vegas, Nevada

America's largest AI conference drew 12,000 attendees from over 100 countries and showcased 400 companies capturing the strange mix of excitement and anxiety defining today's AI climate. A technology once confined to decades of science fiction is now reshaping medicine, food distribution, and global strategy. In this era of technological progress, even credible experts are offering sharply conflicting predictions ranging from massive prosperity and universal basic income to extinction-level risks. I heard both predictions being discussed at Ai4. The scale of the event reflects both the promise and the uncertainty of a world where AI is no longer theoretical but built daily by researchers and developers, prompting questions about whether we are constructing a dream, a threat, or something in between. Ai4 positions itself as a space for good-faith dialogue, encouraging participants to think for themselves and work toward a responsible future amid today's intense and sometimes absurd narratives about machines taking over.

AI Race to Reinvent Medicine

Leading figures in biotechnology and artificial intelligence examined how rapidly advancing AI systems are transforming drug discovery, DNA design, and the broader landscape of human health. Gevorg Grigoryan of Generate Biomedicines emphasized that biology's complexity makes traditional drug development slow and unpredictable, with each medicine essentially a one-off effort. His team aims to make biology programmable and modular, allowing researchers to reuse components and design therapies more systematically. By combining machine-learning models with modular biological parts, he argued, scientists may eventually predict side effects more reliably and accelerate the path from concept to clinical reality. He also stressed the need for thoughtful boundaries that protect against misuse without slowing beneficial innovation.

Moderator Alice Park of *TIME* highlighted that medicine has always evolved through breakthroughs that redefine what is possible. AI now promises to reduce the cost and failure rate of clinical trials—helping more drugs cross the “Valley of Death” where many candidates fail for financial reasons. Alex Zhavoronkov of Insilico Medicine described how AI accelerates the thousands of steps required to bring small-molecule drugs to human testing. His company focuses on aging, arguing that even one additional year of healthy life would generate enormous global benefit. He noted that AI can shorten development timelines, improve molecule quality, and help navigate the challenge that animals and humans often respond differently to treatments.

Eric Nguyen of Radical Numerics explored the profound implications of AI-generated DNA, including the ability to design biological systems nature has never produced. He talked about how his company developed a new virus through manipulation of DNA. He acknowledged the dual-use risks. Thus, the same tools that create life-saving therapies could also generate harmful agents. Nguyen called for stronger biodefense, clearer guidelines, and proactive testing frameworks for entirely new biological constructs, which are now possible because of AI. Across the panel, speakers agreed that AI is reshaping

every stage of drug development, outpacing existing rules and demanding responsible governance to ensure these powerful technologies serve human health rather than threaten it.

Using Agentic AI in Business

Companies are rapidly adopting agentic AI, often without formal approval, because employees feel an urgent need to use these tools to keep pace with rising expectations and competitive pressure. Surveys show that 96% of CEOs believe their employees are already using AI without authorization, revealing that this is not a matter of a few rogue actors but a widespread organizational shift. The demand for AI is legitimate, yet the foundational infrastructure to support safe, accountable use is missing. This mirrors the rise of spreadsheets in the 1980s, when business teams adopted Excel long before IT sanctioned it; for decades, critical decisions were made on tools never designed for enterprise-grade reliability. With AI, however, the risks are far greater: an agent can make mistakes faster, at scale, and with consequences that are harder to detect.

The core problem is organizational design. Traditional centralized IT models assume scarcity and gatekeeping, but AI breaks that model because employees build and deploy agents as if their jobs depend on it. The more effective approach is to treat AI agents as labor—hired, reviewed, budgeted, and held accountable just like human workers. Business units should own ROI for the agents they deploy, while IT shifts from gatekeeper to a platform and oversight function, providing identity, access, safety, and observability. Companies need visibility into how many agents exist, what they are doing, and how they perform, just as they would for human employees. As confidence erodes and competitors advance, the organizations that succeed will be those that build strong foundations for managing agentic AI, ensuring that the people who understand the work are the ones empowered to build and oversee these new digital workers.

Chips, Computation, and Constraints

Computation is the true foundation of modern AI, and the global race to build more of it is reshaping both technology and economics. As AI models scale, demand for chips, memory, and data-center capacity has become insatiable, making semiconductors the “oil” of a token-driven AI economy. Even highly sophisticated AI leaders often lack deep knowledge of chip engineering, yet the future of AI capability is inseparable from advances in hardware. More computation consistently enables more complex problem-solving, and the industry is now confronting whether intelligence itself is becoming a commodity tied to massive pools of computation. This raises questions about who captures value. Is it algorithm developers or chip manufacturers?

The bottlenecks are severe because data centers take years to build, energy capacity is limited, and water shortages constrain cooling. Also, data centers are not popular. There are four critical pillars in increasing computing power are (1) infrastructure, (2) efficiency, (3) energy, and (4) economics. All four must improve dramatically, potentially by factors of 10,000, to support future AI systems. Investments are already enormous, but they remain concentrated and insufficient for the scale required. Lowering computation costs is essential to ensure AI benefits are widely accessible rather than restricted to a few dominant players. With better hardware, the size of data centers and the pollution and the water used by these centers will hopefully become much less.

Future acceleration will depend on breakthroughs in chip specialization, memory-bound architectures, cryogenic cooling, and robotics-driven construction that can shorten data-center build times. Faster chip design cycles and more efficient inference could fundamentally change the economics of AI. Ultimately, the industry must convince governments and markets to invest across the entire compute stack, such as chips, energy, data centers, and manufacturing, so that AI's expanding capabilities can be unlocked safely, affordably, and at global scale for both the Global North and Global South.

Where AI Is Headed

Across venture capital, community leadership, and enterprise innovation, an AI future will be defined by rapid adoption, new organizational structures, and a shift toward agent-driven workflows. Ray Wu, a venture capitalist, emphasized that AI startups are becoming more compelling for investments because end users are retaining and expanding their use of AI services, even though traditional hard metrics are still emerging about the actual consumer need for these services. The strongest companies demonstrate sustained customer engagement and clear evidence that AI meaningfully improves their operations. Claire Xie, Women in AI, highlighted that founders must continually ask whether they are solving real customer needs and whether their solutions can be repeated across many users, including women and children. Momentum in AI comes from scalable problem-solving rather than novelty.

AI agents are beginning to alternate work between humans and machines, reshaping trust, workflow design, and team building. The fastest-growing companies are those that integrate AI into every function from day one, requiring founders to understand new processes and help customers navigate the transition. Trust and strategic alignment will become central as AI becomes embedded in government, industry, and product development. Bioengineering, high-performance computing, and specialized hardware will define the next wave of innovation. While large companies struggle to integrate AI quickly, smaller startups can adopt and iterate faster, making AI a competitive battlefield where agility matters more than size. Together, the panelists suggested that AI is headed toward deeper integration, more agentic workflows, and entirely new industry structures built around speed, specialization, and continuous adoption.

AI and University Education

A panel of university presidents and provosts described how AI has dramatically shortened the cycle of knowledge creation, forcing professors and students to learn together rather than follow traditional lecture-and-quiz models. Instead of relying on fixed authoritative sources, faculty now receive real-time feedback and even real-time fact-checking from AI systems, which requires faculty to upskill, adopt best practices, and share effective methods for teaching and productivity. Globally, AI adoption varies. U.S. institutions tend to rely on commercialized models, while many international partners use open models that are more affordable and accessible. Universities must balance cost, data control, and long-term reliability as they choose which AI platforms to integrate.

Panelists also emphasized the need for international collaboration, especially in areas such as geospatial technologies and supply-chain analytics, where AI can help communities lacking resources. Countries without heavy legacy systems may even be better positioned to adopt new AI-driven educational models. Looking ahead to 2030, they anticipate prototypes of robotic systems and more abundant intelligence supporting community health and global research. Amid this transformation, university presidents acknowledged a "ball of confusion" where institutions must prepare students to use AI

responsibly, yet companies themselves do not fully understand what skills their students need. Students are comfortable with AI, but universities must ensure students understand how to use it wisely and ethically.

The Future of Unmanned War Drones

Samantha Hamilton at Heven AeroTech talked about how today's drone systems can perceive, navigate, and execute pre-planned routes reliably, but they cannot yet make independent judgments or understand intent. Most drones operate at the autopilot or basic autonomy level, where onboard perception allows limited replanning, but true decision-making remains far out of reach. Autonomy is difficult because drones have tight power constraints, millisecond-level control loops, and high-risk failure modes. Small aircraft cannot "phone home" for help, and larger drones pose greater danger when something goes wrong, requiring human oversight to provide context and direction.

Despite these challenges, progress is accelerating. Over-the-air model updates allow fleets to learn from every flight, turning drones into software-defined platforms. Advances such as neuromorphic chips, edge-based learning models, and new hardware designs promise more capable systems. Autonomy can reduce human exposure to dangerous environments, provide persistent surveillance, and compress detection-to-decision timelines from hours to seconds.

However, the risks are significant. Autonomous systems can make mistakes at machine speed, and every AI capability introduces new attack surfaces, including adversarial inputs and model poisoning. The industry must avoid misplacing authority by delegating critical decisions to machines prematurely. As drones become more AI-enabled, society must decide how autonomous they should be in both wartime and peacetime, recognizing that these choices will shape the coming decade.

The AI Bubble

Ed Zitron argues that the AI industry is inflating a dangerous economic bubble driven largely by OpenAI and Anthropic, which together account for the overwhelming majority of AI-related revenue. Smaller AI companies are struggling, and many hyperscalers mask weak demand by using misleading run-rate figures. As model labs begin charging customers the true cost of token usage, businesses are discovering that AI is far more expensive and far less predictable than expected. Token-based billing (1) makes budgeting nearly impossible, (2) forces users to pay for AI mistakes, and (3) exposes the lack of clear ROI for many AI deployments. Despite industry pressure to adopt AI, companies are increasingly questioning whether the benefits justify the costs.

Zitron contends that common counterarguments—such as comparing AI to early Amazon Web Services—misunderstand the scale of losses and the absence of sustainable demand. Energy and chip costs continue to rise, and most announced data-center expansions remain theoretical. Nearly all meaningful computational consumption comes from OpenAI and Anthropic, whose own financials reveal massive losses. Private equity and credit have poured hundreds of billions into AI infrastructure, much of it backed by pensions and insurance funds, meaning ordinary Americans are financially exposed. Zitron warns that if the AI bubble bursts, the economic fallout could spread across the tech sector, major U.S. companies, and household retirement savings.

AI and Job Security

The rise of AI represents a profound shift in how economic value is created, with implications for labor that echo (and may exceed) the disruptions of the Industrial Revolution. Historically, technological progress moved bottlenecks from land to labor, raising the value of human work and lifting living standards. In an AI-driven economy, however, both capital and labor become reproducible. Machines can perform cognitive and physical tasks, potentially at scale. As AI systems grow more capable through increased compute and recursive self-improvement, output could rise astronomically while the share of value flowing to human labor declines. This creates the possibility of two parallel economies—one human, one AI—where wages fall even as productivity surges.

Panelists noted that AI will eliminate repetitive work and shift human roles toward management, decision-making, and critical thinking. Some job functions will disappear, just as past technologies displaced telephone operators or in-house hardware designers. Yet rapid change raises questions about whether society can absorb the transition. Companies adopting comprehensive AI often grow and hire more people, but current tech-sector layoffs show that firms may also use AI as justification to reduce headcount. Entry-level workers are especially vulnerable, and individuals face anxiety about reskilling, income stability, and the adequacy of social safety nets.

Critical thinking, inquisitiveness, and the ability to evaluate AI-generated work will become essential skills. Organizations may flatten, with fewer lower-management roles as agentic AI handles routine tasks. Workers must be willing to learn and unlearn, adapting quickly as models evolve. While new technologies historically create new jobs, the pace of AI advancement may require stronger oversight, international coordination, and policies that support lifelong learning. The challenge is to ensure that the benefits of abundant machine intelligence enhance human well-being rather than deepen inequality or destabilize labor markets.