

 COLUMBIA UNIVERSITY
DATA SCIENCE INSTITUTE

Presented by the DSI Industry Catalyst Network

DATA SCIENCE AND AI EXCHANGE

2026 EVENT SUMMARY





This document provides a high-level summary of the key sessions and themes from this year's event. For additional details, to follow up with a specific presenter or researcher, or to find out how you can engage more deeply with the DSI Industry Catalyst Network (ICN), please contact:

Lori Glover

Chief of Strategic Alliances

The Data Science Institute at Columbia University

lg3478@columbia.edu

We look forward to welcoming you back next year during AI Week New York in Spring 2027!

Sign up for ICN newsletter below for all upcoming event information.

Sign Up →

industry.datascience.columbia.edu



TOPICS COVERED

Overview: Data Science and AI Exchange (DAX)

Strategic Priorities

Key Takeaways

Keynote Address with Bob Hedges

Session 1: Privacy & Security

Session 2: Safety & Agentic AI

Session 3: Governance: Explainability & Compliance

Fireside Chat with Francesca Rossi

Innovation Expo



Overview: Data Science and AI Exchange

About DAX

The DAX Conference, formerly known as Data Science Day, is the flagship annual event at the Data Science Institute (DSI). Designed as an exchange between industry, government, and university leaders on pressing AI challenges, DAX offers attendees an inside look at emerging work in domains that matter most to business. Through cross-sector dialogue and applied examples, DAX connects cutting-edge developments with real-world strategic opportunities. Additionally, the Innovation Expo features a glimpse of what is to come with research presentations and demos from Columbia AI startups.

DAX 2026

- **419** registered attendees, participants from Canada and UK
- **5** sessions, **57** presenters, **24** innovation demos
- **22** topic discussions, **69** session questions posed, including “Which CLI agent are you using and why?”



Overview: Industry Partnerships

About the Data Science Institute and DSI Industry Catalyst Network

The Data Science Institute (DSI) at Columbia University enables interdisciplinary connections in data science and artificial intelligence across the breadth of Columbia's 5 campuses, 20 schools, and 69 departments and divisions. Breakthrough innovations are born at the intersection of disciplines, where subject matter expertise accelerates technical progress. From engineering to medicine, law to chemistry, international policy to journalism, AI is impacting every sector, and DSI connects you to the leading researchers shaping the space and the talent to power your organization.

Through DSI, you can tap into cutting-edge research, new technologies, top talent, and startups, all within the leading research institution in the heart of New York City that generates over 400 new inventions each year.

The DSI Industry Catalyst Network — Columbia DSI's industry engagement program — provides structured pathways to align these capabilities with your business needs to solve priority challenges and accelerate growth.

Companies can also engage with DSI through targeted opportunities, such as sponsored research and capstone projects, depending on your goals and timeline.



“

AI's impact today is being defined by the choices we make now and the partnerships we build. DAX is an effort to advance this at the Data Science Institute and Columbia University.”

Garud Iyengar

Avanessians Director, Data Science Institute, and Professor of Industrial Engineering and Operations Research, Columbia Engineering





STRATEGIC PRIORITIES

Advance Responsible and Ethical AI as a Core Institutional Principle

Provost Angela V. Olinto and Garud Iyengar, Avanessians Director of the Data Science Institute, emphasize that responsible AI is not optional. It is foundational. Columbia's strategy prioritizes integrating ethics (fairness, transparency, safety) directly into AI development and deployment. The goal is to ensure that technological progress aligns with societal trust, regulatory readiness, and long-term sustainability.

Position the University as a Cross-Sector Hub for AI Collaboration

A central priority is to bring together academia, industry, and government to move from theory to real-world implementation. Through initiatives from the Data Science Institute and forums, such as the DAX Annual Conference, Columbia aims to act as a convening platform that fosters partnerships and shared practices to address AI's complex, system-wide impacts.

Drive Transformative Partnerships to Shape the Future of AI

Columbia seeks to drive foundational research and train the next generation of AI leaders. This includes advancing technical innovation while embedding governance, ethics, and societal considerations into education, ensuring graduates are equipped to build AI systems, and also guide how they should be used and governed.





Artificial intelligence is no longer confined to a single domain. It is shaping every sector of the economy from science,...healthcare and finance, to media, infrastructure, and public policy. The defining challenge is not whether AI will transform these fields, but how responsibly and equitably that transformation unfolds.

Responsible AI is not just an ethical aspiration, it is a necessity for all of us.”

Angela V. Olinto

Provost of Columbia University, and Rutherford Professor of Astronomy and Professor of Physics, Columbia University

RESPONSIBLE AI: 3 KEY TAKEAWAYS



1

Responsible AI Builds Long-Term Business Value

2

Meaningful AI Governance Means Examining Systems – And Expectations

3

Human Expertise is an Essential Part of Successful AI Systems

RESPONSIBLE AI: 3 KEY TAKEAWAYS

1

Responsible AI Builds Long-Term Business Value

As AI becomes embedded in business workflows and products, neglecting core consumer concerns like security, privacy, and trust can quickly turn short-term gains into long-term liabilities, especially as AI systems relying on external inputs challenge traditional security models. Barnard computer science professor Rebecca Wright warns that leaders must assume anything AI can access may be misused, while Bob Hedges, retired Global Chief Data Officer at Visa, emphasizes that AI's success ultimately depends on consumer trust—currently low, with only 26% of Americans viewing it positively—arguing companies should innovate in ways that benefit users, not just shareholders. Building that trust requires embedding ethical responsibility across organizations, as IBM Global Leader for responsible AI & AI Governance Francesca Rossi explains, through shared accountability, values-driven risk assessment, and a culture that prioritizes societal impact, which in turn reduces business risk and supports sustainable long-term value.

RESPONSIBLE AI: 3 KEY TAKEAWAYS

2

Meaningful AI Governance Means Examining Systems – And Expectations


Governing modern AI systems is challenging because their “black box” nature and unpredictable real-world use make it difficult to fully test or anticipate outcomes, as Columbia computer science professor Zhou Yu notes. This uncertainty means governance must go beyond model performance to include a clear understanding of what systems do, what data they access, and what actions they can take, according to Buka Gurgendidze-Steinau, Head of Centralized Intelligence at Memorial Sloan Kettering. At the same time, Columbia law professor Talia Gillis points out that many governance challenges are not new, emphasizing the need to clarify what information is needed, who it serves, and its intended impact. While AI can generate explanations for its behavior, Columbia electrical engineering professor Micah Goldblum cautions these may be unreliable—much like human reasoning—highlighting the limits of accountability and the importance of evolving governance approaches that can keep pace with both innovation and misuse.

RESPONSIBLE AI: 3 KEY TAKEAWAYS

3

Human Expertise is an Essential Part of Successful AI Systems

Despite rapid advances in AI, human judgment remains critical to using these tools effectively, as educators and professionals adapt in real time to integrate them responsibly into learning and work, with Rachel Cummings, Columbia engineering and operations research professor, emphasizing the shift from fearing misuse to teaching practical, ethical application. In business contexts, Columbia computer science professor Eugene Wu notes AI excels in low-stakes scenarios but still falls short where high-quality standards are required, reinforcing the need for human oversight, while Omar Santos, Distinguished Engineer at Cisco, highlights the risks of unchecked experimentation. Beyond oversight, AI's complexity demands interdisciplinary collaboration, as Rossi explains, making it essential for experts across fields to communicate and share responsibility. Ultimately, human expertise provides the structure, judgment, and ethical grounding needed to ensure AI systems deliver sustainable business value and align with organizational and societal expectations.

A photograph of a large, multi-story building with a green roof, likely a part of MIT, at night. The building is illuminated from within, and the sky is dark with some clouds. The image is partially obscured by a dark blue overlay on the right side.

“

What are the things we would do if we were really serious about gaining consumers’ trust and earning it, so we could fully capitalize on the potential for AI to be transformative in our society? What are the things that we would do?”

(referencing Steve Jobs’ “Show Them” philosophy)

Bob Hedges

Retired Global Chief Data Officer, Visa, and Research Fellow,
MIT’s Initiative for the Digital Economy

KEYNOTE ADDRESS

How Consumers Will Drive the Success of AI and How We Can Help Them

AI has the transformational power to change work, improve lives, and support society-level gains. Consumers' embracing of AI, however, is dependent on establishing a level of trust and confidence that data scientists and technologists have carelessly taken for granted. Achieving the required level of consumer trust will come not through regulation, but a collective commitment to keep consumers at the center of our data and AI innovation efforts.

Here is what we can do.



Bob Hedges

Retired Global Chief Data Officer, Visa

Research Fellow, MIT's Initiative for the
Digital Economy

KEYNOTE ADDRESS

How Consumers Will Drive the Success of AI and How We Can Help Them



In his keynote, Bob Hedges, highlighted the ways that some companies have been using data and AI that has eroded the trust of the consumer.

As a result, many consumers are skeptical of the incredible technical innovations of the last few years, and how companies and industries are using these new technologies and tools. According to Hedges, when companies use consumer data without explicit permission, track across apps, aggressively target messaging, or apply some aspects of dynamic pricing, it diminishes consumer trust in both the companies and AI systems in general.

In this high-energy talk, Bob made the case for putting consumers back at the center of innovation, and recognizing that it can be better business for companies to think about how to add value for consumers, not extract value from them.



SESSION 1



[Responsible AI] comes back to the fundamentals of what is a safe environment: having proper sandboxes in place where you can run experiments and then roll into production at a really fast pace. This requires a strong foundation of technology being truly implemented with just-in-time permissions for identity, inventory tracking, and visibility into who is doing what with which resources. It also involves understanding how the technology may impact safety, security, and related concerns. So it's a mix. It is not just about building the next step for upskilling; it is the infrastructure around it that you have to keep in mind."

Omar Santos

Distinguished Engineer, Cisco Product Security Incident Response Team (PSIRT) Security Research and Operations

The logo features a stylized gear or fan shape composed of several blue and orange segments. An orange shield icon with a white padlock is positioned at the top of the gear.

SESSION 1

Privacy & Security

Trust begins with how data is protected.

As AI systems rely on complex and sensitive data, organizations must safeguard privacy while securing models, pipelines, and infrastructure. What are emerging and practical approaches to protecting data, mitigating risk, and building AI systems that earn and maintain trust?



Rachel Cummings

Associate Professor, Industrial Engineering and Operations Research; Affiliate, Department of Computer Science; Co-chair of Cybersecurity Research Center, Data Science Institute, Columbia University



Kelly Moan

Chief Information Security Officer, Cyber Command



Omar Santos

Distinguished Engineer Cisco Product Security Incident Response Team (PSIRT) Security Research and Operations



Rebecca Wright

Druckenmiller Professor of Computer Science and Director of the Vagelos Computational Science Center, Barnard College



Moderated by:

Kara Miller

Host, It Turns Out Podcast

The logo features a circular graphic composed of several blue and orange segments. An orange shield icon with a white padlock is positioned at the top of the circle.

SESSION 1

Privacy & Security

The rush of investment in agentic AI requires ensuring these tools still maintain essential privacy and security safeguards. Both companies and researchers need to be sure they're carving out time to be proactive about creating innovative security and privacy features that increase trust and can preserve the advantages of AI tools, not only reacting to after-the-fact harms.

Because keeping up with every aspect of AI's development in real time is unrealistic, the panel addressed the need to be strategic and also collaborative in working towards security solutions. Perhaps most importantly, the increasing need for interdisciplinary work means that even experts in a given field need to get comfortable saying, "I don't know," when facing new technology, new security risks, or new challenges.





SESSION 2



We're now in full-blown production environments, where agents are part of core workflows and scaling, even just beyond workflows, into a wide array of use cases.”

Brittney Cantor

Agentic Commerce Lead, Accenture's Center for Advanced AI





SESSION 2

Safety & Agentic AI

The deployment of AI Agents requires demand new approaches to **control and accountability**.

As agentic AI moves from experimentation to deployment, questions of safety, oversight, and human-in-the-loop design become critical. How are organizations managing risk, defining guardrails, and ensuring responsible behavior? How is this new technology helping with productivity and cost reduction? What are some of the workplace challenges surfacing?



Brittany Cantor

Agentic Commerce Lead,
Accenture's Center for Advanced AI



Ranjan Roy

VP Industry Lead - Retail, Writer



Eugene Wu

Associate Professor of Computer Science and Co-director of the Data, Agents, and Processes Lab (DAPLab), Columbia University



Zhou Yu

Associate Professor, Computer Science Department, and Co-director of the Data, Agents, and Processes Lab (DAPLab), Columbia University



Moderated by:

Kara Miller

Host, It Turns Out Podcast



SESSION 2

Safety & Agentic AI

We don't know what AI agents are going to do when we set them loose – but we also don't know how humans are going to put those agents to use. Despite the exciting tasks that agentic AI can accomplish in demos, it's crucial to recognize that such systems still face challenges for truly safe use in most professional or personal environments.

This panel's robust conversation explored how to balance the excitement about the future potential of AI agents with the real challenges of making sure they actually do what we intend. Panelists also pointed out that even agents designed with explicitly beneficial intent can create real risk – and that they can also be deployed maliciously. Not only do we not know exactly what will happen when we put AI agents to work, we can't know all the ways humans will choose to use them – underscoring the need makes safety testing and guardrail development a core focus.





SESSION 3

“ Executives need to know, first of all, what AI systems are running.

What are they doing? What do they have access to? And what actions they can take? So now governance is about governing the actions from those AI agents. And that's where explainability and safe governance comes in.”

Buka Gurgenedze-Steinau

Head of Centralized Intelligence and Automation for Enterprise Infrastructure, Memorial Sloan Kettering Cancer Center



SESSION 3

Governance: Explainability & Compliance

AI cannot scale without governance that executives can stand behind.

Regulatory expectations and stakeholder expectations are reshaping how AI systems are designed and deployed. What does it take to build governance frameworks that support explainability, compliance, and confident decision-making?



Buka Gurgenedze-Steinau

Head of Centralized Intelligence and Automation for Enterprise Infrastructure, Memorial Sloan Kettering Cancer Center



Alessandro Petroni

Head of Quality Engineering - Fraud Payment Services, The Clearing House



Talia Gillis

Daniel G. Ross Professor of Law, Columbia Law School



Micah Goldblum

Assistant Professor, Department of Electrical Engineering, Columbia University



Moderated by:

Kara Miller

Host, It Turns Out Podcast





SESSION 3

Governance: Explainability & Compliance

How can we create effective governance systems for AI technologies that we don't fully understand? On the one hand, we often can't fully understand why humans behave the way they do. On the other hand, we have much clearer laws and regulations around what happens when a human breaks the rules.

This panel zoomed in on the problem of how to regulate technology when both the general public and lawmakers often lack sufficient understanding to design governance that is both feasible and appropriate. One key approach the panel identified is making sure that data and AI literacy is foregrounded in both education and professional development, while still making progress towards real governance that ensures the viability of these technologies in the long-term.





#DAX2026

“Trust is really fundamental in having people and society accept technology, and also having companies, institutions adopt technology. So if we care about the adoption of AI and all the benefits that it can bring, we should care about not just building a system of trust around the technology, but also around those that build and deploy the technology.”

Francesca Rossi

IBM Global Leader for Responsible AI and AI Governance, T.J. Watson IBM Research Lab, New York, and Lecturer, MS in Technology, Columbia University School of Professional Studies

FIRESIDE CHAT

Responsible AI

The conversation examined how AI has evolved into a foundational force across industries and how that shifts the demands of ethics and governance today. Rossi outlined building effective AI governance in practice, from risk assessment and technical tools to integrated processes, education, and culture change. She explored the growing role of global regulation and multi-stakeholder collaboration, and why investing in AI ethics delivers both immediate and long-term value.



Francesca Rossi

IBM Fellow and the IBM Global Leader for Responsible AI and AI Governance, T.J. Watson IBM Research Lab, New York

Lecturer, MS in Technology, Columbia University School of Professional Studies

FIRESIDE CHAT

Responsible AI



In this thoughtful conversation, Francesca Rossi emphasized that responsible AI development is not just a "good thing" for companies to do, but is actually a core requirement for long-term business success. While some critics contend that responsible tech threatens to slow down innovation, Rossi emphasized that truly responsible technology development ensures innovation is sustainable. This allows AI innovations to build durable value for companies – not short-term returns followed by new or unexpected costs. Globally, Rossi said, companies that devote more resources and people to responsible tech are also more successful. Rather than treating ethical considerations as an afterthought, Rossi highlighted that developing a clear ethical culture towards AI development throughout your company is key to building more effective, efficient, and higher-value businesses.

INNOVATION EXPO

CUTTING-EDGE RESEARCH AND STARTUPS PARTICIPATING IN THE EXPO CONNECTING RESEARCH TO REALITY

- Tactorum Inc. →
- Fair Aggregation in Virtual Power Plants
- AI-Driven Clinical Decision Support for Diabetes Treatment Selection
- Mirror AI
- AskAI Projects (Explainable Asset Allocation and Portfolio Construction; and AI Deployment Framework for Consumer Credit Models)
- RateMate Energy →
- Checkpoint-lite →
- Metlas →
- Data Flow Control →
- LakeAgent: Search and Reasoning over Large Heterogeneous Data Lakes
- A Census-Tract Analysis of Data Center Climate Vulnerability in the United States
- Curiously AI →
- Sky Valley →
- Reliable Vibe Coding
- Doppel Market →
- Inside Outside Health →
- nMonica →
- Generalizing Risk Parity via Optimal Risk Budgeting with Target Returns: Exact and Tight Approximation Algorithms
- Arlex.AI →
- Quantifying Video Storyline: A Theory-Guided Multimodal LLM Approach For Creative Insights
- Linking Data Centers, Energy, and Weather Risks via AI
- AI and the Tapestry of Ethnic Placemaking in American Suburbs
- Delegated Control Under Attack: A Measurement-Driven Threat Model for Exposed AI Agents
- Socrat →

INNOVATION EXPO

Emerging Technologies from Columbia University Spanning 6 Distinctive Industries

The Innovation Expo highlighted cutting-edge projects from **24 research projects and emerging startups** across Columbia University's research, technology, and innovation ecosystem. Attendees explored live demonstrations from Columbia-affiliated startups and heard from graduate students and Data Science Institute (DSI) scholars as they presented their latest research breakthroughs.

Chemistry • Engineering • Healthcare • International Policy • Journalism • Law



INNOVATION EXPO

AI-Driven Ideas Shaping the Future


The event concluded with an Innovation Expo featuring the latest research from several of Columbia's labs, as well as AI-focused startups emerging from the University.

Some participants included startups addressing:

- Overlooked mental health issues of athletes
- The challenge in assessing pain in lab rats
- Tax preparation
- Improving care for patients with diabetes
- Tracking climate change at the hyper-local level

To kick off the Expo, students and startups delivered a one-minute-pitch to attendees on what they are up to, why it's important, and how to find out more. Followed by a free-form networking session where attendees could meet with early-stage founders and researchers to discuss their work, the Expo was a unique glimpse at what the next phase of AI-enabled problem-solving is going to look like. For anyone interested in an inside look at what's next for these exciting tools and technologies, there was no need to look further than these 24 promising new projects.

For a full list of the featured research and projects, please visit: [DAX 2026 | Responsible AI](#).

A nighttime photograph of a city skyline, likely New York City, with illuminated buildings and a bridge visible in the background. The image is dark blue with a grid pattern.

#DAX2026

“ This is a chance to rethink how we engage across institutions and sectors. Connecting industry and academia provides the fuel for the Innovation Supply Chain that enables economic growth and positive societal impact.”

Lori Glover
Chief of Strategic Alliances, Data Science Institute, Columbia University



COLUMBIA UNIVERSITY DATA SCIENCE INSTITUTE

INDUSTRY CATALYST NETWORK

Unlock the potential of collaboration with
Columbia University's Data Science Institute
through the DSI Industry Catalyst Network.

industry.datascience.columbia.edu