



Limor Ziv

FUTURIST

Boards are pouring resources into AI and seeing thinner returns than promised. Regulatory scrutiny is rising in parallel. The two pressures converge at the same operational layer, and that is where most deployments quietly fail.

Dr. Limor Ziv, founder and CEO of Humane AI, helps organisations turn AI investments into operating value that performs commercially and holds up under regulatory scrutiny.

Limor Ziv's 2026 Biography

Why organisations work with Limor Ziv

- She closes the gap between AI governance theory and AI commercial performance, with named client work at Google, KPMG, Bright Data, SolarEdge, and Medison Pharma.
- Humane AI was selected by Meta among 2025's most promising AI startups; she sits on the Scientific Council of the Israeli Association for Ethics in AI.
- Her peer-reviewed research with 74 senior AI and data professionals (MDPI, 2025) gives boards a working diagnostic for the breakdowns that destroy AI value, from data and governance to context and business alignment.
- She works the practitioner seam between academic credibility and commercial deployment, the same person teaches the lecture and runs the workshop afterwards.
- Operator credibility on the women-in-AI question, with public roles in Israel's data and ethics communities, without converting that into a separate keynote category.

Biography highlights

- Founder and CEO, Humane AI, selected by Meta among 2025's most promising AI startups.
- Postdoctoral fellowship in Human-AI Interaction; Ph.D., Tel Aviv University, where she was named to the "100 Outstanding Lecturers" list.
- Co-author, "Behind the Algorithm: International Insights into Data-Driven AI Model Development" (MDPI, 2025), based on interviews with 74 senior AI and data professionals.
- Scientific Council member, Israeli Association for Ethics in AI (IAEAI).
- Named keynote and consulting work for Google, KPMG, Bright Data, SolarEdge, Medison Pharma, government ministries, and municipalities including Tel Aviv.

AVAILABLE FOR

- Consultancy
- Speaking
- Training and Workshops

LIMOR'S SPEAKING THEMES

- AI Ethics & Responsible Technology
- Artificial Intelligence & Generative AI
- Cybersecurity
- Data Analytics
- Digital Transformation
- Risk Management

LANGUAGES: English

Biography

AI investments are returning less than executive committees expected, and regulatory scrutiny is rising in parallel. The two pressures meet at the same operational layer, where most AI deployments quietly fail. Dr. Limor Ziv built Humane AI to work at that seam. The company was selected by Meta among 2025's most promising AI startups.

In a keynote, she takes apart the hidden risks built into everyday algorithms and gives senior audiences practical ways to govern AI systems and to lower the exposure they create. She works with organisations to capture the upside of AI and avoid the failure modes that come with deploying it.

The research underwrites the consulting. Her 2025 paper with Maayan Nakash in MDPI's Machine Learning and Knowledge Extraction draws on interviews with 74 senior AI and data professionals across countries and sectors. The recurring breakdowns it surfaces are the same ones that destroy commercial performance: weak data quality, missing context, governance gaps, and business misalignment.

Her client base spans government ministries and municipalities, universities, and technology companies, including Google, KPMG, Bright Data, SolarEdge, and Medison Pharma. She lectures in graduate programmes at universities and was named to Tel Aviv University's list of 100 Outstanding Lecturers. The brief is consistent across them: make AI deployments that earn their investment and hold up when someone audits the decisions.

Key speaking topics

- Responsible AI and AI ethics
- AI risk management and compliance
- AI governance and regulatory readiness
- Data quality and data-centric AI
- Human-AI interaction and human-centred design
- Realising value from AI investment
- AI in regulated sectors, including healthcare and finance
- Operationalising responsible AI in the enterprise

Ideal for

- Boards and executive committees accountable for both the return on AI investment and the defensibility of the AI systems they sign off.
- Chief AI officers, chief data officers, and CISOs operationalising responsible AI policy while showing it earns its cost.
- Regulated-industry leadership in healthcare, financial services, and public sector wrestling with AI deployment under compliance pressure.
- Heads of AI, innovation, and transformation under pressure to show a return on AI spend without creating regulatory exposure.

Audience outcomes

- A working definition of AI risk that distinguishes data risk, model risk, and decision risk, and a sense of where each lives in the organisation.
- A view of what AI deployment actually demands of operating teams, drawn from her empirical research with 74 senior practitioners across countries and sectors.
- A grounded read on the live regulatory and ethical debates, from a

Scientific Council member of the Israeli Association for Ethics in AI.

- An honest read on the gap between what AI investments cost and what they return, and on the operating work that decides whether the gap closes.

Limor Ziv's 2026 talks & topics

Opportunities and Risks in AI Development

A senior-level orientation to where AI creates measurable value and where it carries the most material organisational risk.

Key takeaways:

- A practitioner's map of AI risk categories, from data and model risk to deployment and reputational risk.
 - The decision points where executive sponsors typically lose control of an AI programme.
 - A framework for prioritising AI investment against organisational risk appetite.
-

Risk Management in AI

A working session on identifying, governing, and mitigating AI risk inside operating organisations.

Key takeaways:

- How AI risk differs from conventional technology risk and why existing risk functions struggle to absorb it.
 - Practical guardrails for AI development teams that survive contact with regulators and auditors.
 - The data hygiene and documentation practices that turn AI risk policy into AI risk practice.
-

Responsible AI Practices

What responsible AI looks like when it is operationalised across product, data, and governance functions.

Key takeaways:

- The translation problem between AI ethics principles and shippable system requirements.
 - Roles, accountabilities, and review mechanisms that make responsible AI auditable.
 - Where responsible AI overlaps with existing privacy, security, and compliance regimes, and where it does not.
-

AI Regulation and Compliance

A read on the current and emerging regulatory landscape, with implications for AI strategy.

Key takeaways:

- The architecture of the EU AI Act and adjacent regimes, in plain operating terms.
 - What regulators are likely to ask first, and what documentation organisations need ready.
 - How to set internal policy that anticipates regulation rather than reacts to it.
-

AI Governance and Management

How to design AI governance that integrates with existing corporate governance, not parallel to it.

Key takeaways:

- Governance structures for AI portfolios, from steering committees to model registries.
 - Decision rights and escalation paths between data, AI, risk, and legal functions.
 - Governance failure patterns in early enterprise AI programmes and how to avoid them.
-

Principles and Practices of Human-Centered Technology Development

A grounding in human-AI interaction research applied to product and workplace design.

Key takeaways:

- The evidence on how users actually engage with AI systems, and where trust breaks.
 - Design choices that protect human judgement rather than displace it.
 - Cross-disciplinary inputs, from psychology to policy, that strengthen AI product decisions.
-

Building Synergistic Human-AI Ecosystems

A view of how organisations should structure work, roles, and capability around AI rather than alongside it.

Key takeaways:

- The shift from automation thinking to augmentation thinking inside operating teams.
- Capability investments leaders consistently underestimate during AI rollout.
- A model for designing human-AI workflows that hold under operational pressure.

ARTICLES AND PODCASTS



AI Risk Management in Healthcare & Biotechnology: A Practical Approach