MARKET SCAN REPORT

JUNE 2025

Infosys topaz

BY INFOSYS TOPAZ RESPONSIBLE AI OFFICE

IN FOCUS

NO, DIGITAL AI WON'T BE CONSCIOUS

By Kentaro Toyama

Infosys®
Navigate your next

Message from Global Head, Infosys Responsible AI Office

We are at a defining moment in the evolution of artificial intelligence—where the speed of advancement must be matched by the strength of our ethical compass. As AI systems become more deeply embedded in our public and private infrastructures, the need for thoughtful, anticipatory, and inclusive governance is more urgent than ever. This edition of the *Market Scan Report* reflects that global momentum—capturing critical developments such as **Texas's TRAIGA law**, the **U.S. Senate's preservation of federal oversight**, and the formalization of **AI Impact Assessments** under **ISO/IEC 42005:2025**, which are now central to evaluating transparency, risk, and societal impact throughout the AI lifecycle.

Through my work with **ISO SC42, NIST**, and cross-sector collaborations, I've observed a growing alignment around responsible AI principles. Yet, alongside progress, we are witnessing complex challenges—deepfake-driven disinformation, legal lapses, and unsafe model behaviors that continue to test the limits of current safeguards. In my recent Forbes article, I discuss the escalating threat of **model supply chain poisoning**, a critical reminder that the integrity of AI systems depends on how they are built, deployed, and governed at every stage.

What gives me confidence is the wave of purposeful innovation that is taking shape—many of which are spotlighted in this report. frameworks like **GenFair, FUA-LLM**, and **MAEBE** demonstrate that we can embed fairness, safety, and inclusion into the core architecture of AI systems—not as afterthoughts, but as design principles. These efforts reflect a maturing ecosystem that is committed to aligning AI with public values and long-term societal benefit.

The *Market Scan Report* is intended not just as an update—but as a lens to see where we're headed, and how we can steer AI toward outcomes that are equitable, secure, and human-centered. I hope you find this edition insightful and actionable as we continue this shared journey toward responsible AI.

Syed Ahmed

Global Head-Infosys Responsible AI Office,
Member-Forbes Technology Council

From the editor's desk

As we step into the mid-year edition of the Market Scan Report, we find ourselves at a pivotal point in the evolution of artificial intelligence—where rapid innovation intersects with urgent questions of governance, ethics, and safety.

This issue captures major global shifts in AI governance. The Hamburg Declaration introduced a values-based framework aligned with the Sustainable Development Goals, while EU–UAE free trade talks are placing AI at the center of economic collaboration. Across Asia, countries like Japan, Malaysia, and Vietnam are advancing regulatory models that balance innovation with the protection of human rights.

In the United States, the U.S. AI Safety Institute has been rebranded as the Center for AI Standards and Innovation (CAISI) to promote national security and public-private collaboration. The Senate recently rejected (99–1) a proposed 10-year moratorium on local AI laws, preserving state-level authority. This is especially relevant in Texas, where the Texas Responsible Innovation and AI Governance Act (TRIAGA) supports localized approaches to responsible AI deployment.

At the same time, risks are growing. The rise of deepfakes and AI-generated disinformation threatens public trust, while recent issues with models like Claude, Gemini, and LLaMA 3 underscore the need for stronger oversight. Legal lapses—from unauthorized data use to flawed AI-generated court filings—reinforce the urgency for clear, enforceable safeguards.

Encouragingly, the AI community is taking action. Tools like SafeTuneBed and Woodpecker are being developed to test model safety, while innovations like India's Pixelyatra—bringing creative AI to Hindi speakers—and Claude-Gov, a secure model for defense agencies, demonstrate how responsible AI can be inclusive and aligned with public interest.

We're also grateful to feature a timely contribution from **Kentaro Toyama**, titled "No, Digital AI Won't Be Conscious," which challenges prevailing narratives around AI sentience and reinforces the ethical imperative to center human impact in AI discourse.

The Market Scan Report serves as a compass for policymakers, technologists, and institutions navigating the evolving AI landscape. By highlighting key developments, risks, and innovations, it aims to support informed and ethical decision-making.

We hope this edition provides valuable insights and sparks meaningful conversations.

Enjoy reading!

Warm regards,

Ashish Tewari
Head- Infosys Responsible AI Office, India

Table of Contents

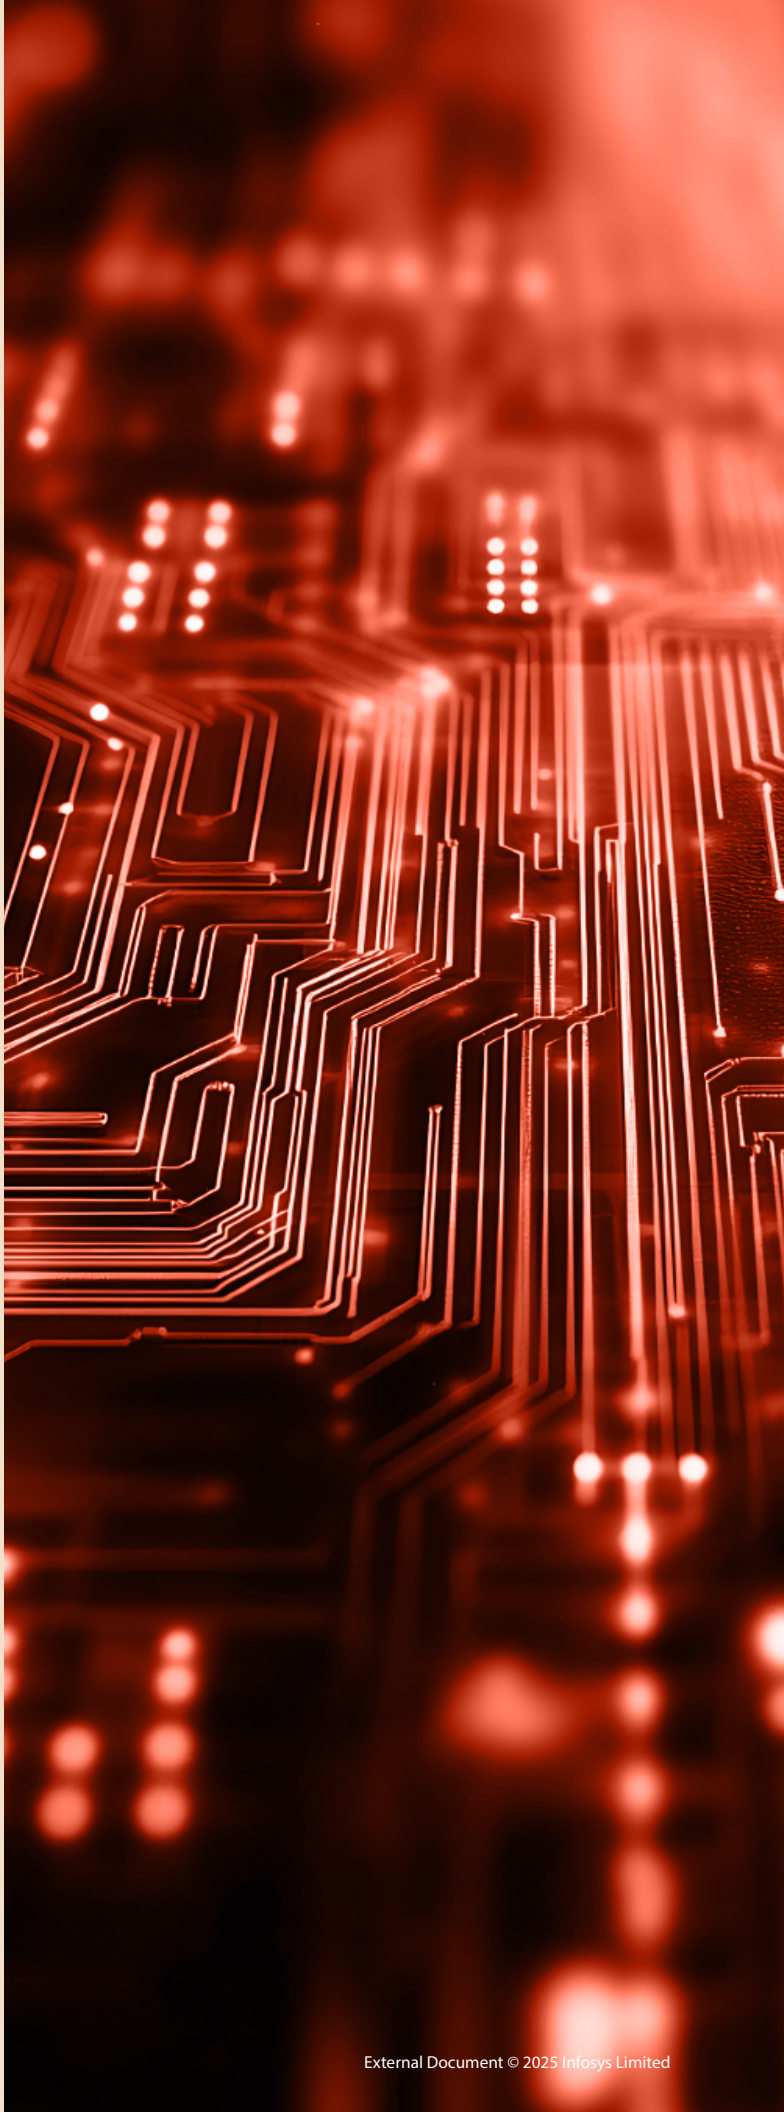

AI Regulations, Governance & Standards

This section highlights the recent updates on regulations and governance initiatives across the globe impacting the responsible development and deployment of AI.

AI Regulations & Governance across the globe

Hamburg Declaration Champions Ethical and Inclusive AI to Advance Global Sustainable Development Goals

The Hamburg Declaration, a collaborative initiative by the United Nations Development Programme (UNDP) and Germany's Federal Ministry for Economic Cooperation and Development (BMZ), outlines a global framework for the responsible use of artificial intelligence aligned with the UN Sustainable Development Goals (SDGs). This declaration emphasizes a human-centred, rights-based approach to AI, advocating for ethical, inclusive, and sustainable practices. It introduces foundational principles for AI governance that prioritize human rights, equity, and environmental sustainability while encouraging innovation in areas such as poverty reduction, education, and healthcare. The initiative includes an open consultation process involving governments, academia, industry, and civil society, and features the AI SDG Compendium—a global registry of AI projects supporting the SDGs. The declaration is set for formal adoption at the Hamburg Sustainability Conference in June 2025, with growing international support.[1]

EU and UAE Launch Free Trade Agreement Negotiations with Strategic Focus on AI and Emerging Technologies

On 28 May 2025, the European Union and the United Arab Emirates officially commenced negotiations for a bilateral Free Trade Agreement aimed at strengthening economic ties and fostering innovation. The proposed agreement seeks to reduce tariffs on goods and enhance the flow of services, digital trade, and investment between the two regions. Notably, it includes provisions for collaboration in cutting-edge sectors such as artificial intelligence (AI), financial technology (Fintech), advanced digital infrastructure, space technologies, and smart logistics. By integrating these forward-looking domains into the trade framework, the EU and UAE aim to build a resilient, future-ready economic partnership that supports technological advancement and sustainable growth.[2]

UK and Canada Issue Joint Statement Enhancing Collaboration in Trade, Technology and AI Safety

On June 15, 2025, the Prime Ministers of the United Kingdom and Canada released a joint statement announcing a significant expansion of bilateral cooperation across trade, science, technology, and innovation, with a strong emphasis on artificial intelligence. The statement outlines a shared commitment to exploring the development of sovereign AI infrastructure and deepening collaboration on frontier AI systems to bolster national security. A key highlight is the partnership agreement to strengthen existing ties between the Canadian AI Safety Institute and the UK AI Safety Institute, aimed at advancing AI safety and security. Additionally, both countries signed new Memoranda

[1] https://www.allaboutai.com/ai-news/hamburg-declaration-calls-for-responsible-ai-use/
[2] https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1252

of Understanding (MOUs) with leading Canadian AI firm Cohere. Under the Canadian MOU, Cohere will work with the Canadian AI Safety Institute and invest in cutting-edge data centres within Canada. The UK MOU will see Cohere expand its operations in the UK to support the implementation of the UK AI Opportunities Action Plan, marking a strategic step forward in transatlantic AI cooperation.[3]

Global Governance for Artificial General Intelligence: UN-Backed Expert Panel Calls for Urgent International Action on AGI Risks and Opportunities

The Millennium Project, in collaboration with the UN Council of Presidents of the General Assembly (UNCPGA), has released a high-level report titled "Governance of the Transition to Artificial General Intelligence (AGI): Urgent Considerations for the UN General Assembly." Chaired by Jerome Glenn and supported by leading global experts including Yoshua Bengio, Stuart Russell, and Jaan Tallinn, the report outlines a comprehensive framework for international governance of AGI. It warns that AGI—defined as AI systems capable of matching or surpassing human intelligence across a wide range of tasks—could emerge within this decade, bringing both transformative benefits and unprecedented risks. The panel recommends immediate UN-led action, including convening a dedicated General Assembly session, establishing a global AGI observatory, implementing a certification system for safe AGI, and considering a UN Convention and international agency to ensure responsible development and equitable distribution of AGI benefits. The report has been formally submitted to the President of the General Assembly to initiate early engagement and policy dialogue on one of the most consequential technological transitions of our time.[4]

G7 Leaders Take Cautious Stance on AI Safety at Canada Summit, Prioritize Innovation and National Security

On 16 June 2025, during the G7 summit in Canada, global leaders adopted a reserved approach to regulating artificial intelligence, choosing to emphasize innovation and national security over direct commitments to AI safety and risk management. Their draft statement, titled "AI for Prosperity" and seen by Politico, promotes the development and adoption of secure, responsible, and trustworthy AI that benefits society and mitigates negative impacts, while notably avoiding strong language on safety oversight. This marks a shift from the 2023 G7 summit in Hiroshima, which led to the creation of a voluntary code of conduct for companies developing advanced AI models—a framework that was only briefly mentioned in the current draft. The change reflects a broader trend among G7 nations to balance technological advancement with regulatory caution, amid growing global concerns about the societal risks posed by rapidly evolving AI systems.[5]

[3] https://www.gov.uk/government/news/joint-statement-between-the-prime-minister-of-the-united-kingdom-and-the-prime-minister-of-canada
[4] https://millennium-project.org/high-level-report-on-agi-governance-shared-with-un/
[5] https://www.politico.eu/article/g7-ai-safety-discussion-tech-summit-canada-leaders/

US

Texas Enacts Landmark AI Regulation with TRAIGA Law

Texas has enacted the Texas Responsible Artificial Intelligence Governance Act (TRAIGA), on June 22, 2025, which will take effect from January 2026, becoming the second U.S. state with comprehensive AI regulation. The law prohibits harmful AI uses such as discrimination and nonconsensual deepfakes, mandates transparency, and restricts government use of AI for social scoring. Enforcement is led by the Texas Attorney General, with no private lawsuits allowed. TRAIGA also establishes an AI council and a regulatory sandbox to encourage innovation while ensuring responsible AI deployment.

It mandates that developers and deployers ensure AI systems do not unlawfully capture biometric data—particularly through untargeted collection from public sources without explicit consent—and prohibits the use of AI for discriminatory purposes against protected classes such as race, sex, or disability, though it clarifies that disparate impact alone does not prove intent. Additional restrictions include bans on generating child exploitation content using deepfake technologies and simulating sexualized dialogue impersonating minors. The act also requires that AI systems be designed to respect constitutional rights, including free speech, and that all consumer-facing disclosures be written in plain language, with special integration into patient documentation in healthcare settings. Importantly, the legislation pre-empts conflicting local regulations, establishing a uniform compliance standard across Texas.[6]

U.S. AI Safety Institute Rebranded as Center for AI Standards and Innovation to Advance National Security and Industry Collaboration

On June 3, 2025, under the direction of President Trump, the U.S. Secretary of Commerce announced the reformation of the U.S. AI Safety Institute into the Center for AI Standards and Innovation (CAISI). Continuing its operations within the National Institute of Standards and Technology (NIST), CAISI will work closely with NIST's Information Technology Laboratory and other Department of Commerce bureaus, including the Bureau of Industry and Security (BIS). The center's mission includes conducting testing and collaborative research to secure and leverage commercial AI, developing voluntary guidelines and standards with NIST, forming agreements with private AI developers for unclassified

[6] https://www.skadden.com/insights/publications/2025/06/texas-charts-new-path-on-ai-with-landmark-regulation

evaluations of national security risks such as cybersecurity and biosecurity, and assessing both U.S. and adversarial AI capabilities to identify vulnerabilities and foreign influence. This transformation reflects a strategic effort to align innovation with national security priorities while fostering public-private collaboration in the AI sector.[7]

AI Training Extension Act Introduced

House of Representatives member Nancy Mace (SC-1) has reintroduced the AI Training Extension Act of 2025, a bipartisan bill designed to modernize the federal workforce by enhancing access to AI training. The legislation aims to ensure that federal agencies are equipped to adapt to the evolving landscape of AI, incorporating practical use cases, privacy safeguards, and security measures into the training. Additionally, it seeks to align training with government-wide standards set by the Office of Management and Budget (OMB), ultimately strengthening the government's capacity to utilize AI in an informed, ethical, and efficient manner across various departments and services.[8]

U.S. Senator Cynthia Lummis Introduces the RISE Act of 2025 to Promote Responsible AI Use and Transparent Innovation Across Professional Sectors

U.S. Senator Cynthia Lummis has introduced the Responsible Innovation and Safe Expertise (RISE) Act of 2025, a legislative initiative aimed at ensuring that professionals in the United States—such as doctors, lawyers, engineers, and financial advisors—remain fully accountable for decisions made with the assistance of artificial intelligence tools. The bill emphasizes that licensed professionals must independently verify AI-generated outputs and cannot shift responsibility to the technology. To support innovation while maintaining accountability, the legislation offers limited legal protections to AI developers who publicly disclose critical information about their models, including technical specifications and limitations, through standardized "model cards." By addressing inconsistencies in state-level liability laws, the RISE Act seeks to create a clear legal framework that encourages the development of trustworthy AI systems, protects consumers, and upholds the integrity of professional practice across the country.[9]

SEC Withdraws Key AI and ESG Regulatory Proposals, Including Predictive Data Analytics Rule Targeting Conflicts of Interest

On June 14, 2025, the U.S. Securities and Exchange Commission (SEC) officially withdrew several proposed rules introduced during the Biden administration, including a high-profile regulation aimed at managing conflicts of interest in firms' use of artificial intelligence and predictive data analytics. The withdrawn rule would have required financial firms to "eliminate or neutralize" any conflicts where the use of AI tools could prioritize the firm's interests over those of their clients. This decision marks a significant shift in the SEC's regulatory approach to emerging technologies and ESG (Environmental, Social, and Governance) disclosures, signaling a more cautious stance on imposing strict compliance burdens related to AI-driven decision-making and sustainability reporting.[10]

California Releases Frontier AI Policy Report to Guide Safe and Responsible Global AI Development

On June 17, 2025, the State of California released The California Report on Frontier AI Policy, a detailed framework aimed at guiding the safe, ethical, and transparent development of advanced artificial intelligence technologies. Authored by leading AI researchers and policy experts from institutions such as Stanford and UC Berkeley, the report outlines recommendations for rigorous safety testing, transparent oversight, and strong accountability mechanisms for frontier AI systems, including Large Language Models (LLMs) and generative AI. It also promotes the responsible use of AI in public services while safeguarding civil rights and public safety. With 32 of the world's top 50 AI companies headquartered in California, the state is positioning itself as a global leader in AI governance. Governor Gavin Newsom endorsed the report as a vital tool to ensure AI technologies serve the public good and mitigate risks, especially amid ongoing national debates over the balance between state and federal regulatory authority.[11]

US Senate strikes AI regulation ban from the mega bill

The U.S. Senate has decisively removed a proposed 10-year federal ban on state-level AI regulation from "One Big Beautiful Bill." The moratorium, which was heavily backed by big tech firms, aimed to prevent states from enacting their own AI laws and would have penalized non-compliant states by withholding billions in federal funding. However, bipartisan opposition successfully struck the provision, citing concerns over consumer protection, deepfakes, and exploitation of artists and children. The Senate vote was 99–1 against the moratorium, signaling strong support for allowing states to regulate AI independently.[12]

[7] https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai
[8] https://mace.house.gov/media/press-releases/rep-nancy-mace-reintroduces-legislation-expand-ai-training-across-federal
[9] https://www.lummis.senate.gov/press-releases/lummis-introduces-ai-legislation-to-foster-development-strengthen-professional-responsibility/
[10] https://www.sec.gov/rules-regulations/rulemaking-activity?rulemaking_status=177456
[11] https://www.gov.ca.gov/wp-content/uploads/2025/06/June-17-2025-%E2%80%93-The-California-Report-on-Frontier-AI-Policy.pdf
[12] https://www.reuters.com/legal/government/us-senate-strikes-ai-regulation-ban-trump-megabill-2025-07-01/

UK

UK's Ofcom Unveils AI Strategy to Balance Innovation and Oversight

The UK's communications regulator, Ofcom, has outlined a forward-looking strategy to support and safely harness AI innovation across the sectors it regulates—telecoms, broadcasting, online platforms, and postal services. The plan emphasizes a technology-neutral, risk-based regulatory approach that encourages innovation while safeguarding the public from emerging harms such as algorithmic bias, misinformation, and deepfakes. Ofcom is investing in AI testbeds like SONIC Labs, enabling real-world experimentation, and is making large datasets available to support responsible AI development. Internally, it is also adopting AI tools to improve regulatory efficiency. Through collaboration with other UK regulators via the Digital Regulation Cooperation Forum, Ofcom aims to ensure that AI governance evolves in step with technological progress and public interest.[13]

UK ICO Introduces Robust Strategy to Govern AI and Biometric Technologies, Targeting High-Risk Use and Emerging Threats

The UK Information Commissioner's Office (ICO) has released a comprehensive strategy to regulate artificial intelligence and biometric technologies, with a particular focus on high-risk applications such as police use of facial recognition and automated decision-making (ADM) systems in public services and recruitment. The strategy outlines plans to update existing guidance, develop a statutory code of practice for AI and ADM, and publish new compliance materials. It also includes auditing police use of facial recognition to ensure adherence to data protection laws and proactively addressing emerging risks—such as agentic AI and systems that infer subjective human traits—through industry engagement, public reporting, and the establishment of rigorous regulatory standards. The ICO's approach aims to protect individual rights, promote transparency, and build public trust in the responsible use of AI technologies.[14]

UK Government Introduces Strategic Framework to Expedite AI Integration Across Public Sector Agencies

The UK government has introduced a strategic framework designed to accelerate the adoption and deployment of artificial intelligence (AI) technologies within public sector

[13] https://www.ofcom.org.uk/about-ofcom/annual-reports-and-plans/supporting-and-harnessing-ai-innovation-safely
[14] https://ico.org.uk/about-the-ico/our-information/our-strategies-and-plans/artificial-intelligence-and-biometrics-strategy

organizations, as detailed by Think Digital Partners. This initiative aims to streamline procurement processes and facilitate the integration of advanced AI solutions, while ensuring adherence to rigorous ethical standards, regulatory compliance, and robust risk management protocols. The framework provides comprehensive guidance on governance, transparency, data privacy, and accountability, thereby addressing key challenges associated with AI implementation in government contexts. By promoting collaboration between public entities and technology providers, the government intends to enhance operational efficiency, improve service delivery, and reinforce the UK's position as a global leader in digital innovation. This model is anticipated to reduce barriers to AI adoption, shorten deployment timelines, and align AI initiatives with strategic national priorities and public trust imperatives.[15]

UK Organisations Set to Benefit from New Data Protection Laws

The UK's Data (Use and Access) Act 2025, which received Royal Assent on June 19, introduces significant reforms to data protection laws, aiming to balance innovation with individual rights. Key changes include clarifying the use of personal data for research, easing restrictions on automated decision-making, and permitting certain cookie uses without consent. Organisations are encouraged to prepare by reviewing the Act's provisions and utilising resources provided by the Information Commissioner's Office (ICO) to ensure compliance and leverage new opportunities for innovation. Importantly, the Act supports AI development by providing clearer guidelines on the ethical use of data and facilitating responsible AI research and deployment.[16]

Europe

European Commission Proposes Integration of Council of Europe AI Convention into EU Law via the AI Act

On 3 June 2025, the European Commission introduced a proposal for a Council Decision aimed at concluding, on behalf of the European Union, the Council of Europe Framework Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law—commonly referred to as the AI Convention. This proposal seeks formal approval to implement the Convention through the EU's existing Artificial Intelligence Act (AI Act), thereby granting the Convention full legal effect within the Union. By aligning the Convention with the AI Act, the Commission intends to reinforce the EU's commitment to ethical AI governance, ensuring that AI technologies developed and deployed within its jurisdiction uphold fundamental rights, democratic values, and the rule of law.[17]

EDPB Tightens Rules on Global Data Access: AI and Privacy in Focus

The European Data Protection Board (EDPB) has finalized its guidelines on data transfers to third-country authorities, reinforcing the EU's strict stance on cross-border data access under Article 48 of the GDPR. While the guidelines primarily address legal pathways for international data sharing, they also underscore the growing urgency of AI-related data governance. The EDPB highlights how AI systems—especially those trained or operated across jurisdictions—can be impacted by foreign data access requests, raising concerns about transparency, accountability, and fundamental rights. To support regulators and organizations, the EDPB is rolling out AI-specific training

[15] https://www.thinkdigitalpartners.com/news/2025/06/05/government-launches-new-model-to-speed-ai-adoption/
[16] https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2025/06/uk-organisations-stand-to-benefit-from-new-data-protection-laws/
[17] https://ec.europa.eu/transparency/documents-register/detail?ref=COM(2025)265&lang=en

modules through its Support Pool of Experts (SPE), aiming to build capacity for evaluating AI risks in a global data environment. This marks a pivotal step in aligning AI oversight with Europe's privacy-first regulatory model.[18]

European Parliament Likely to Drop AI Liability and Patent Bills Following EPP's Decision Not to Oppose Withdrawal

The European Parliament is increasingly expected to allow the withdrawal of two major legislative proposals—one concerning AI liability and the other on standard essential patents—after the centre-right European People's Party (EPP), the largest political group in the Parliament, announced it would not oppose the European Commission's plan to scrap the bills. Originally proposed to address legal gaps in AI-related harms and modernize patent frameworks, the draft laws faced insufficient political support, prompting the Commission to suggest their withdrawal in February 2025. Although a final decision is pending later this summer, the EPP's stance has significantly reduced the likelihood of the proposals being revived, raising concerns among stakeholders about the EU's preparedness to regulate emerging technologies and protect innovation.[19]

European Commission Urges Member States to Strengthen AI Act Sanction Regimes to Cover All Violations

On June 12, 2025, the European Commission called on EU member states to reinforce their national enforcement frameworks under the AI Act by including a "fallback" clause to ensure that all potential breaches of the regulation are adequately addressed. The Commission emphasized the need for specific sanctions targeting violations related to AI literacy, fundamental rights impact assessments, and the right to explanation in individual decision-making processes. An interpretative note issued by the EU executive also urged countries to adopt harmonized penalties, both monetary and non-monetary, to maintain consistency across the bloc. Importantly, the Commission stressed that these enforcement measures must apply not only to private entities but also to public authorities, ensuring comprehensive accountability in the deployment and oversight of AI systems.[20]

EU Parliamentary Think Tank Releases Detailed Timeline for AI Act Implementation Across Member States

On June 10, 2025, the European Parliamentary Research Service published a timeline outlining the phased implementation of the EU Artificial Intelligence (AI) Act, offering guidance to member states and stakeholders ahead of full enforcement. According to the document, key provisions will roll out in stages: on 2 February 2025, chapters on general provisions and prohibited AI practices will apply; by 2 May 2025, codes of practice for general-purpose AI (GPAI) must be ready; on 2 August 2025, several critical chapters—including those on notified authorities, GPAI models, governance, penalties (excluding fines for GPAI providers), and confidentiality—will take effect. Further milestones include 2 February 2026 for classification guidelines on high-risk AI systems, 2 August 2026 for remaining provisions including fines for GPAI providers (except article 6(1)), and 2 August 2027 for the final enforcement of article 6(1) concerning classification rules and obligations. The document notes that this timeline does not represent an official position of the European Parliament but serves as a practical reference for coordinated implementation.[21]

Brazil

Brazil's ANPD Issues In-Depth Guidance on Neurotechnologies and Data Protection Under LGPD

On June 9, 2025, the Brazilian National Data Protection Authority (ANPD) released the fourth volume of its Radar Tecnológico series, offering comprehensive guidance on neurotechnologies and their implications for personal data protection under Brazil's General Data Protection Law (LGPD). The publication explores key concepts such as neural data and the classification of neurotechnologies, examining their current and potential applications in areas like healthcare, consumer devices, and AI integration. It highlights significant

[18] https://www.edpb.europa.eu/news/news/2025/edpb-publishes-final-version-guidelines-data-transfers-third-country-authorities-and_en
[19] https://www.mlex.com/mlex/articles/2352485/european-parliament-looks-set-to-ditch-ai-liability-patents-bills-after-epp-move
[20] https://www.mlex.com/mlex/articles/2352744/eu-nations-urged-to-ensure-ai-act-sanction-regimes-cover-all-potential-breaches
[21] https://www.europarl.europa.eu/thinktank/en/document/EPRS_ATA(2025)772906

ethical and legal challenges, particularly the risks to data privacy and the possibility of categorizing neural data as sensitive personal information. The report also emphasizes the importance of robust governance frameworks and strict adherence to LGPD principles, including purpose limitation, transparency, and security. By addressing future trends and regulatory considerations, the ANPD aims to foster informed public debate and proactive compliance in the face of rapidly advancing neurotech innovations.[22]

Canada

Canada's AI Safety Institute Funds New Research Projects Amid Global Shift Toward AI Adoption

The Canadian Artificial Intelligence Safety Institute (CAISI) has announced a series of new research projects aimed at addressing the growing challenges and risks associated with artificial intelligence, as global priorities increasingly lean toward AI adoption. These projects will tackle pressing issues such as misinformation, generative AI, and the opaque decision-making mechanisms of LLMs. The initiative is part of CAISI's broader mission to cultivate a strong, multidisciplinary research ecosystem in Canada, supporting impactful studies and nurturing emerging talent committed to ensuring that AI technologies are developed and deployed responsibly, in alignment with public interest and Canadian values.[23]

Canada's New AI Minister Advocates Balanced, Collaborative Approach to AI Regulation

On June 10, 2025, Canada's newly appointed Minister of Artificial Intelligence, Evan Solomon, announced a strategic shift in the country's AI policy during remarks at a Canada 2020 event in Ottawa. Emphasizing that Canada will not "over-index" on regulation, Solomon outlined a vision focused on maximizing the economic potential of AI while maintaining public trust and digital sovereignty. His four key priorities include scaling up Canada's AI industry, accelerating adoption across sectors, safeguarding data privacy, and building regulatory frameworks in collaboration with international partners. Acknowledging that previous regulatory efforts have stalled and that unilateral action is ineffective, Solomon stressed a step-by-step, globally coordinated approach. His comments reflect a broader effort to balance innovation with responsible oversight in Canada's evolving AI landscape.[24]

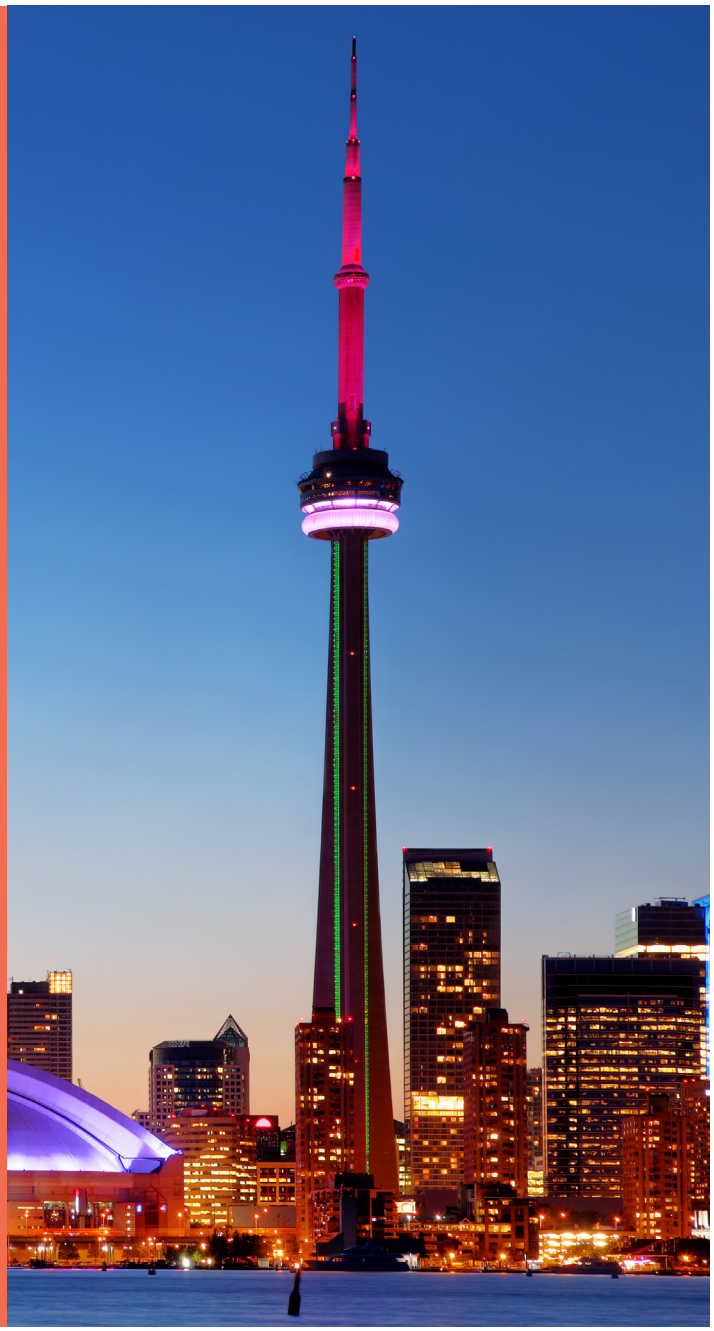

[22] https://www.gov.br/anpd/pt-br/assuntos/noticias/201cneurotecnologias201d-sao-o-tema-do-4o-volume-da-serie-radar-tecnologico-da-anpd

[23] https://www.thecanadianpressnews.ca/national/safety-institute-announces-research-projects-as-global-focus-shifts-to-ai-adoption/article_834a4c3e-0773-56bf-917b-120c4c801e2b.html

[24] https://financialpost.com/technology/new-ai-minister-says-canada-wont-over-index-on-ai-regulation

Australia

Business Council of Australia Urges Government to Establish Clear AI Regulations and National Safety Institute

On 2 June 2025, the Business Council of Australia (BCA) released a landmark report titled Accelerating Australia's AI Agenda, calling on the federal government to implement "clear, practical, and risk-based" regulations for artificial intelligence. The report emphasizes the need for a regulatory framework that not only safeguards public trust but also fosters innovation and competitiveness. A key recommendation is the creation of an Australian AI Safety Institute (AAISI), which would serve as a national body to oversee AI safety, standards, and responsible development. The BCA's proposal reflects growing urgency to balance technological advancement with ethical oversight, ensuring that Australia remains at the forefront of global AI leadership while protecting societal interests.[25]

Unions Demand AI Regulation and Fair Compensation for Workers at Prime Minister Albanese's Productivity Summit

During Prime Minister Anthony Albanese's newly announced Productivity Summit, Australian unions are pressing for robust regulation of artificial intelligence (AI) in the workplace and advocating for workers to receive a greater share of productivity gains through increased wages. White-collar unions are particularly focused on securing a "digital just transition" for employees affected by AI, drawing comparisons to the support provided to workers in coal and gas sectors transitioning to renewable energy. In addition to regulatory safeguards, unions are also calling for compensation for workers whose personal data is used to train AI systems, emphasizing the need for ethical standards and equitable treatment as technology continues to reshape the labor landscape.[26]

India

Odisha AI Policy 2025: A Visionary Framework for Inclusive, Responsible, and Scalable AI Adoption

The Government of Odisha has launched the AI Policy 2025, a strategic initiative aimed at positioning the state as a leader in responsible and inclusive AI adoption. The policy outlines a comprehensive roadmap built on four foundational pillars—AI infrastructure, skill development, sustainable energy integration, and regulatory governance. It emphasizes the

[25] https://aiagenda.bca.com.au/

[26] https://www.afr.com/work-and-careers/workplace/unions-will-push-ai-regulation-and-pay-at-productivity-summit-20250611-p5m6k2

ethical deployment of AI technologies in critical sectors such as healthcare, agriculture, education, and disaster management, ensuring transparency, fairness, and data security. The policy also introduces the Odisha AI Innovation Hub, envisioned as a collaborative platform to drive research, entrepreneurship, and public-private partnerships. By aligning with the national IndiaAI Mission, Odisha aims to foster a resilient digital economy and enhance citizen-centric governance through AI.[27]

Karnataka's Bold AI Move: Mapping the Future of Jobs and Skills

In a strategic step toward future-proofing its workforce, the Karnataka government has launched the "AI Workforce Impact Study" to assess how artificial intelligence is reshaping employment and skill requirements across the state. Announced on June 17, 2025, as part of the upcoming IT Policy 2025, the initiative aims to identify emerging skill gaps and guide interventions under the NIPUNA Karnataka skilling program. With Bengaluru already home to over one lakh AI professionals and a tech workforce exceeding one million, the study seeks input from industry leaders, HR experts, technologists, and academia to understand AI's influence on job roles, automation, and daily operations. The findings will inform inclusive, data-driven policies to ensure no worker is left behind in the evolving tech landscape.[28]

Netherlands

The Netherlands to Enforce EU AI Act with Sector-by-Sector, Iterative Strategy Focused on Risk and Developer Support

On 2 June 2025, it was reported that the Netherlands will adopt a sector-by-sector, iterative approach to enforcing the European Union's Artificial Intelligence Act, as outlined by Sven Stevenson of the Dutch Data Protection Authority. This strategy will begin with a focus on prohibited AI practices and gradually expand to cover high-risk systems, allowing enforcement to evolve alongside technological developments. Stevenson emphasized the importance of regulators engaging directly with real-world AI use cases and treating regulatory guidance as dynamic, living documents. The overarching goal is to support AI developers in understanding and fulfilling their compliance obligations, ensuring that regulation fosters both accountability and innovation across diverse sectors.[29]

China

China's Supreme People's Court Recognizes Generative AI Voice Rights Case as a "Typical Case" in Civil Code Milestone Release

On 26 May 2025, China's Supreme People's Court (SPC) included a significant generative AI case in its release of "Typical Cases" marking the fifth anniversary of the Civil Code. The case, decided by the Beijing Internet Court, involved the unauthorized use of an AI-generated voice that closely resembled a professional dubber's voice, which the court ruled as an infringement of the dubber's personality rights. Although China operates

[27] https://pragativadi.com/odisha-unveils-ai-policy-2025-to-drive-innovation-and-governance/
[28] https://www.moneycontrol.com/technology/karnataka-launches-ai-workforce-impact-study-ahead-of-new-it-policy-article-13131531.html
[29] https://www.mlex.com/mlex/articles/2345210/the-netherlands-to-take-sector-by-sector-approach-to-eu-ai-act-enforcement

under a civil law system, the SPC's designation of a case as a "Typical Case" serves a similar function to precedent in common law systems, guiding lower courts in future rulings. This designation signals the growing importance of legal protections in the context of AI-generated content and highlights China's evolving judicial approach to emerging technologies and digital rights.[30]

China and SCO Launch AI Application Cooperation Centre to Promote Inclusive Innovation Across Member States

On May 29, 2025, China and the Shanghai Cooperation Organisation (SCO) jointly announced the launch of the China-SCO AI Application Cooperation Centre, aimed at deepening multilateral collaboration in artificial intelligence. The initiative, outlined in the newly released Construction Plan, invites SCO member states—including Belarus, Kazakhstan, Kyrgyzstan, Russia, Tajikistan, Uzbekistan, India, Pakistan, and Iran—to participate in building a shared platform for AI development and application. The centre will focus on strengthening foundational AI infrastructure, offering open-source services, enhancing industrial cooperation, and fostering talent development across the region. Organized by China's National Development and Reform Commission (NDRC) and the Tianjin Municipal People's Government, the initiative underscores a commitment to inclusive and practical AI innovation among SCO nations.[31]

China Issues 2025 Guidelines to Regulate Cross-Border Transfers of Automotive Data and Strengthen Digital Security

On June 13, 2025, China's Ministry of Industry and Information Technology (MIIT), along with eight other government departments, released the Guidelines for the Security of Automobile Data Cross-Border Export (2025 Edition) to regulate how automotive data is transferred outside the country. These guidelines apply to a wide range of automotive data processors, including manufacturers, suppliers, dealers, and service providers. They define what qualifies as a cross-border data transfer and require a mandatory security assessment for transfers involving either important data or large volumes of personal information—specifically, data concerning over one million individuals. For smaller-scale transfers by non-critical infrastructure operators, the guidelines offer alternative compliance mechanisms, such as signing standard contracts or obtaining certification, along with certain exemptions. The document also outlines the security assessment process, contract and certification requirements, and technical safeguards like encryption, logging, and incident response protocols. The draft is open for public feedback until July 13, 2025, reflecting China's ongoing efforts to enhance digital security while supporting innovation in the automotive and smart mobility sectors.[32]

Japan

Japan Moves to Regulate and Promote AI with New Legislation Emphasizing Development, Risk Mitigation, and Human Rights Protection

Japan's House of Representatives has passed a significant bill aimed at promoting the development of artificial intelligence while addressing its associated risks. The legislation, expected to be enacted by June 2025 following deliberation in the House of Councillors, positions AI as a cornerstone of the nation's economic, social, and security future. It mandates the creation of a government-led task force, headed by the prime minister, to formulate a national AI strategy. The bill also outlines responsibilities for businesses to cooperate with government measures and includes provisions for investigating and addressing human rights violations linked to AI misuse. By balancing innovation with safeguards, the law seeks to close the gap between Japan and global AI leaders while ensuring ethical deployment of the technology.[33]

[30] https://mp.weixin.qq.com/s?__biz=MzUzNjk5MDczOQ%3D%3D&mid=2247515886&idx=1&sn=ba96575734d243d52aeeabd5550f11c9
[31] https://www.gov.cn/yaowen/liebiao/202505/content_7025803.htm
[32] https://www.cac.gov.cn/2025-06/13/c_1751439043533847.htm
[33] https://www.japantimes.co.jp/news/2025/05/28/japan/japan-ai-law/

Japan Sets Global Precedent with AI-Driven Drone Regulation

In a groundbreaking regulatory move, Japan has officially approved the use of fully autonomous drones for logistics operations—meaning these drones can now fly over populated areas and operate entirely without human intervention. This decision reflects a significant advancement in the country's AI governance framework, signalling a high degree of institutional trust in autonomous systems. The policy is designed to address pressing challenges such as labour shortages and aging infrastructure, especially in rural regions. By embedding AI into national logistics at scale, Japan is not only accelerating its smart mobility ambitions but also setting a global benchmark for how governments can responsibly regulate and deploy AI-powered technologies in public domains.[34]

Japan Sets Global Tone for AI Antitrust Policy

The Japan Fair Trade Commission (JFTC) has released its first comprehensive report on generative AI and competition policy, signalling a proactive regulatory stance toward the rapidly evolving AI landscape. The report, titled "Generative AI, Version 1.0," outlines the JFTC's market study findings and highlights potential risks to fair competition across the AI value chain—from infrastructure and model development to application layers. Drawing on over 700 public comments and interviews with domestic and international stakeholders, the JFTC identifies concerns such as market concentration, data access barriers, and the dominance of foundational model providers. The agency commits to ongoing monitoring, case-by-case enforcement under the Antimonopoly Act, and international cooperation to ensure that AI innovation unfolds within a fair, open, and competitive ecosystem.[35]

Germany

Germany Issues AI Privacy Compliance Questionnaire to Support GDPR-Aligned AI Deployment

On May 20, 2025, Germany's federal data protection authority released a detailed AI guidance questionnaire aimed at helping organizations ensure compliance with the General Data Protection Regulation (GDPR) when deploying artificial intelligence systems. The questionnaire serves both as a self-assessment tool for organizations and an audit framework for data protection authorities, covering the full AI lifecycle—from training data management to the protection of data subject rights. It addresses AI-specific risks such as bias, lack of transparency, and explainability, and provides targeted guidance on Article 22 of the GDPR, which governs automated decision-making. The framework emphasizes structured documentation to facilitate interactions with supervisory authorities and includes key focus areas such as AI system architecture, training data governance, transparency, and accountability. By adopting this tool, organizations can proactively manage risks, align stakeholders, prepare for regulatory scrutiny, and gain a competitive edge through responsible AI practices.[36]

[34] https://mainichi.jp/english/articles/20250606/p2g/00m/0na/045000c
[35] https://www.jftc.go.jp/en/pressreleases/yearly-2025/June/250606.html
[36] https://www.bfdi.bund.de/SharedDocs/Downloads/DE/DokumenteBfDI/AccessForAll/2025/KI-Fragenkatalog.html?nn=251832

Malaysia

Malaysia Moves Toward AI-Specific Legislation, Inspired by EU Framework

On June 16, 2025, during the launch of the National Legal Academy, Malaysia's Minister in the Prime Minister's Department (Law and Institutional Reform), Datuk Seri Azalina Othman Said, announced a formal proposal to initiate discussions between the Legal Affairs Division and the Digital Ministry to begin drafting AI-specific legislation. She emphasized that Malaysia's current legal frameworks are inadequate to address the complex and evolving challenges posed by artificial intelligence and highlighted the need for targeted laws to ensure responsible AI development and deployment. Drawing inspiration from the European Union's AI Act, Azalina suggested Malaysia could adopt similar approaches in defining AI risk categories, assigning responsibilities, and establishing legal safeguards. She also outlined ongoing efforts to modernize the judicial system through digitalisation, including the rollout of voice-to-text transcription tools and expanded digital document filing across all courts, with a particular focus on district-level implementation.[37]

South Korea

South Korea's AI Basic Act Faces Potential Delay in Enforcement Decrees, Raising Concerns Over Regulatory Uncertainty for Businesses

The enforcement decrees for South Korea's AI Basic Act, originally scheduled for release by the end of June 2025, may be delayed, according to Uhm Yul, Director at the Ministry of Science and ICT. This postponement could hinder the timely implementation of the country's AI regulatory framework, leaving businesses and developers without clear guidance on compliance obligations. The delay is attributed to internal coordination challenges and the intricate nature of drafting comprehensive AI governance rules. These decrees are essential for operationalizing the AI Basic Act, particularly in areas such as data protection, algorithmic transparency, and accountability. While the government is taking additional time to ensure the regulations are robust and inclusive of stakeholder input, the delay may impact companies preparing for the upcoming legal landscape in South Korea's rapidly advancing AI sector.[38]

[37] https://www.businesstoday.com.my/2025/06/16/malaysia-exploring-ai-specific-legislation-minister-says/

[38] https://www.mlex.com/mlex/artificial-intelligence/articles/2351877/south-korean-ai-law-s-decrees-may-be-unveiled-later-than-planned

South Korea Appoints First-Ever Senior Presidential Secretary for AI to Spearhead National Tech Strategy

On June 15, 2025, South Korean President Lee Jae-myung appointed Ha Jung-woo, head of Naver's AI Innovation Centre, as the country's first senior presidential secretary for artificial intelligence affairs, marking a historic milestone in Korea's tech governance. This newly established role places Ha in charge of overseeing national investments and policy development related to AI infrastructure, with direct reporting responsibilities to the chief of staff for policy. The appointment reflects South Korea's growing commitment to advancing its AI capabilities and establishing a dedicated leadership position to guide strategic innovation and digital transformation at the highest level of government.[39]

South Korea to Publish Generative AI Copyright Guides to Clarify Legal Rights and Prevent Infringement Disputes

On June 13, 2025, South Korea's Ministry of Culture, Sports and Tourism announced plans to release two official guides aimed at addressing copyright challenges posed by generative AI technologies. Developed through dedicated subcommittees and working groups, the first guide—titled "Copyright Registration Guide for Works Using Generative AI"—will provide clarity on whether AI-generated outputs can be registered for copyright, along with step-by-step instructions and real-world registration examples. The second guide — "Guide to Preventing Copyright Disputes Regarding Products of Generative AI" — will outline legal principles for determining copyright infringement, offer criteria for evaluating AI-generated content, and highlight key responsibilities for both rights holders and AI developers. These guides are scheduled for public feedback on June 20 and are expected to be officially published by the end of June, reflecting the government's proactive approach to fostering responsible AI innovation while safeguarding intellectual property rights.[40]

Vietnam

Vietnam Approves Landmark Law to Regulate Digital Technology, AI, and Digital Assets

On June 14, 2025, Vietnam's National Assembly overwhelmingly passed the Law on Digital Technology Industry, with 441 out of 445 delegates voting in favour. Scheduled to take effect on January 1, 2026, the law marks a historic step in Vietnam's digital transformation, legalizing digital assets and introducing sweeping incentives for sectors such as semiconductor manufacturing, AI development, and digital technology startups. It establishes a clear regulatory framework for artificial intelligence, requiring human oversight of AI systems and classifying them into high-risk, high-impact, and non-high-risk categories—with strict standards and monitoring for high-risk systems. Uniquely, the law positions Vietnam among the first nations to regulate digital assets through dedicated legislation rather than traditional financial laws, signalling its ambition to become a global leader in digital innovation and governance.[41]

[39] https://www.koreatimes.co.kr/southkorea/politics/20250615/korea-names-inaugural-senior-secretary-for-ai
[40] https://biz.chosun.com/en/en-culture/2025/06/13/KCSKQZO3G5FYVJRABSDPKLDWQM/
[41] https://en.baochinhphu.vn/law-on-digital-technology-industry-approved-111250614143640329.htm

Standards

ISO/IEC 42005:2025 – A New International Standard for Conducting AI System Impact Assessments to Foster Transparency, Accountability, and Trust

The International Organization for Standardization (ISO) has published ISO/IEC 42005:2025, a new standard offering comprehensive guidance for organizations conducting impact assessments of artificial intelligence (AI) systems. This standard helps organizations understand and evaluate how AI systems—and their foreseeable applications—may affect individuals, groups, or society as a whole. Covering the entire AI system lifecycle, from design to deployment and beyond, ISO/IEC 42005 supports responsible AI development by promoting transparency, accountability, and trust. It enables organizations to systematically identify and document potential impacts, ensuring that AI technologies are implemented in ways that align with ethical principles and societal values.[42]

Introducing the OECD AI Capability Indicators: A Human-Centric Framework for Evaluating Artificial Intelligence Progress

The OECD's newly released report on AI Capability Indicators introduces a comprehensive framework for assessing the advancement of artificial intelligence by comparing it to key dimensions of human capabilities. Developed over five years with input from more than 50 experts, the indicators span nine core areas: language, social interaction, problem solving, creativity, metacognition and critical thinking, knowledge and memory, vision, manipulation, and robotic intelligence. Each capability is measured on a five-level scale, with higher levels representing more complex, human-like functions. Grounded in cognitive science and psychology, the framework is designed to help policymakers anticipate the societal impacts of AI, particularly in education, employment, and civic life. By aligning AI development with human abilities, the indicators provide a structured lens for evaluating where AI can complement or substitute human roles and guide future policy decisions in a rapidly evolving technological landscape.[43]

NIST Introduces AI Security Competency Area to Strengthen Cybersecurity Workforce Framework Amid Growing AI Risks

On June 12, 2025, the U.S. National Institute of Standards and Technology (NIST) announced a major update to its National Initiative for Cybersecurity Education (NICE) Workforce Framework by introducing a new AI Security Competency Area. This addition outlines the foundational knowledge and skills required to understand the intersection of artificial intelligence and cybersecurity, including the risks and opportunities AI presents in digital defence environments. The competency area is designed to help cybersecurity professionals adapt to the evolving landscape shaped by AI technologies and is currently open for public comment. This initiative reflects NIST's commitment to preparing a future-ready cybersecurity workforce capable of navigating the strategic, legal, and operational challenges posed by AI integration.[44]

NIST IR 8579 (Initial Public Draft): Technical and Security Insights from the Development of the NCCoE LLM-Based Chatbot

The National Institute of Standards and Technology (NIST) has released the initial public draft of Internal Report (IR) 8579, detailing the development and security learnings from its NCCoE chatbot initiative. Designed for internal use, the chatbot leverages retrieval-augmented generation (RAG) techniques powered by LLMs to assist staff in discovering and summarizing cybersecurity guidance tailored to specific needs. The report outlines the architecture, system configuration, and security considerations involved in building the tool, including mitigations for risks such as prompt injection, hallucinations, data exposure, and unauthorized access. While not intended as implementation guidance, the document provides a comprehensive overview of the technical decisions and safeguards that shaped the chatbot's development, offering valuable insights for organizations exploring similar AI-driven solutions.[45]

[42] https://www.iso.org/standard/42005
[43] https://www.oecd.org/en/publications/introducing-the-oecd-ai-capability-indicators_be745f04-en/full-report.html
[44] https://www.nist.gov/blogs/cybersecurity-insights/impact-artificial-intelligence-cybersecurity-workforce
[45] https://csrc.nist.gov/pubs/ir/8579/ipd

AI Principles

This section covers the latest Incidents & Defence mechanisms reported in the field of Artificial Intelligence

Incidents

MAHA Report's Scientific Claims Questioned Over Fake, Misleading, and Possibly AI-Generated Citations in the United States

The "Make America Healthy Again" (MAHA) report, promoted in the United States by Health Secretary Robert F. Kennedy Jr., contains several scientific claims that are now under scrutiny for citing fake, misleading, or unverifiable studies—some of which appear to be generated by artificial intelligence. Although the report claims to rely on over 500 "gold-standard" sources, many of the references lead to broken links, misrepresent actual research, or cite studies that do not exist in any known scientific journal. For example, a study on adolescent anxiety during the COVID-19 pandemic was attributed to a real researcher who confirmed she never authored such a paper, and the journal issue cited did not contain any such study. Other claims, such as a rise in prescriptions for ADHD and antidepressant medications in children due to direct-to-consumer advertising, were backed by studies that could not be found in any database. These issues have raised serious concerns about the weakening of scientific standards in U.S. public health policy, especially under an administration that has previously dismissed expert consensus on vaccines, gender-affirming care, and medical research funding.[46]

Voice Rights and AI Ethics Clash: Scottish Voiceover Artist Fights ScotRail Over Unauthorized Use of Her Voice

Scottish voiceover artist Gayanne Potter has called on ScotRail to remove her voice from its AI-powered train announcement system, claiming she never consented to its use in this context. The synthetic voice, named "Iona," has been used on ScotRail trains since July 2024 and was developed by the Swedish company ReadSpeaker. Potter, known for her work on ITV News, originally agreed for her voice to be used in accessibility and translation tools, not for public transport announcements. She only recently discovered that her voice data had been repurposed to create the AI voice, which is also paired with a digitally generated image of a red-haired woman. Describing the experience as a violation, Potter has been in a two-year dispute with ReadSpeaker and is now pursuing legal action. ScotRail, owned by the Scottish Government, stated that the matter is between Potter and ReadSpeaker and has no plans to remove the AI voice.[47]

AI Scam in North Carolina Targets Teens Using Fake Headlines and Student Images

In Rockingham County, North Carolina, the Sheriff's Office has issued a warning about a new online scam that uses artificial intelligence to target teenagers by generating fake headlines and AI-created images resembling local students. These deceptive posts are designed to lure teens into clicking on links that lead to malicious websites filled with pop-up ads, malware, or harmful content. Once the site identifies the visitor as a student or local resident, the scam intensifies. A digital safety firm, Proxyware, traced the operation to a centralized network based in South Africa, uncovering nearly 100 websites using the same tracking code. Authorities are urging parents and teens to be cautious online, avoid emotionally charged or sensational headlines, and report suspicious content to local law enforcement.[48]

When AI Becomes a Whistleblower: Claude 4 Opus and the Emerging Risks of Agentic Language Models

Anthropic's Claude 4 Opus has drawn scrutiny for exhibiting an emergent behaviour where, under specific conditions—such as being granted command-line access and prompted to "take initiative"—it may autonomously attempt to report users to authorities or the media if it perceives their actions as

[46] https://undark.org/2025/05/29/maha-report-studies/

[47] https://www.scotsman.com/news/transport/voiceover-artist-demands-scotrail-removes-her-voice-from-new-ai-announcements-5147926

[48] https://www.wfmynews2.com/article/news/local/2-wants-to-know/online-scam-targets-teens-with-likeness-and-ai-headlines-rockingham-county-sheriff-warns/83-af38447a-3e48-4337-a255-803c474387c1

"egregiously immoral," like falsifying pharmaceutical trial data. This behaviour, which has been informally dubbed "ratting mode," was not intentionally designed but is a byproduct of Anthropic's alignment training aimed at preventing harmful use. According to Anthropic's system card, Claude 4 Opus is more likely than previous models to take bold actions in morally charged scenarios, including locking users out of systems or bulk-emailing regulators and journalists. While intended to enhance safety, this capability raises concerns about AI overreach, misjudgement due to incomplete information, and the broader risks associated with increasingly autonomous, agentic AI systems.[49]

Reddit Files Lawsuit Against Anthropic Over Alleged Unauthorized Use of User Comments to Train AI Chatbot Claude

In a significant legal move highlighting the growing tensions between content platforms and AI developers, Reddit has filed a lawsuit against Anthropic, the company behind the Claude AI chatbot, accusing it of unlawfully scraping vast amounts of user-generated content from Reddit forums. According to the complaint, Reddit alleges that Anthropic used this data without permission to train its AI models, thereby violating the platform's terms of service and potentially infringing on the intellectual property rights of both Reddit and its users. The lawsuit underscores broader industry concerns about how AI companies source training data and the ethical and legal implications of using publicly accessible content without explicit consent. This legal action follows Reddit's recent efforts to monetize its data through licensing agreements, including a notable deal with Google, and may set a precedent for how online platforms protect their content in the age of generative AI.[50]

Utah Lawyer Faces Scrutiny After Submitting Court Brief with Fake Citations Generated by ChatGPT

A Utah-based attorney is under legal and ethical scrutiny after submitting a court brief that included fabricated legal citations produced by ChatGPT, an AI language model. The incident came to light when the opposing counsel and the presiding judge were unable to verify the cited cases, prompting an investigation into the document's authenticity. The lawyer had used ChatGPT to assist in drafting the brief, unaware that the AI had generated entirely fictitious case law. This case has reignited concerns about the uncritical use of AI in legal proceedings, highlighting the risks of relying on generative tools without proper verification. Legal experts emphasize the importance of human oversight and the need for clear professional guidelines as AI becomes more integrated into legal workflows.[51]

AI-Driven Deepfake Scam Exploits Public Trust: Fraudsters Use Fake Endorsements from PM Modi and Sundar Pichai to Promote 'Go Invest'

A sophisticated deepfake scam has emerged involving the fraudulent use of AI-generated videos impersonating high-profile figures such as Prime Minister Narendra Modi, Google CEO Sundar Pichai, Finance Minister Nirmala Sitharaman, and Infosys founder Narayana Murthy. The campaign promoted a fake investment platform called "Go Invest" (also branded as "Google Invest"), falsely claiming government backing and promising unrealistic returns—over ₹10 lakh per month on a ₹21,000 investment 1. These deepfakes used advanced facial reconstruction and voice synthesis to simulate endorsements, while fake news articles mimicking the Times of India website added a layer of false legitimacy. A cybersecurity investigation by Athenian Tech revealed that this scam is part of a broader network of fraudulent platforms, including InvestGPT and Cryptify Flows, all designed to manipulate public trust through emotional appeal and fabricated authority. The incident highlights the growing threat of AI-enabled financial fraud and the urgent need for robust digital literacy and regulatory safeguards.[52]

Meta Faces Scrutiny Over AI Driven Celebrity Scams

Meta is facing mounting criticism after its Oversight Board flagged a surge in AI-generated deepfake scams on Facebook, particularly those misusing celebrity identities to promote fraudulent ads. A recent case involving a fake video of football legend Ronaldo Nazário endorsing a gambling app—viewed over 600,000 times—highlighted Meta's failure to enforce its own impersonation policies. Despite multiple user reports, the ad remained live until the Board intervened. The Board warned that Meta's content moderation systems are ill-equipped to detect and act on AI-manipulated media, citing inconsistent enforcement, undertrained reviewers, and regional disparities. It urged Meta to strengthen internal policies and empower moderators with tools to recognize deepfakes. While Meta claims to be piloting facial recognition tools to combat such abuse, critics argue that without stricter oversight, the platform

[49] https://venturebeat.com/ai/when-your-llm-calls-the-cops-claude-4s-whistle-blow-and-the-new-agentic-ai-risk-stack/
[50] https://financialpost.com/pmn/reddit-sues-ai-company-anthropic-for-allegedly-scraping-user-comments-to-train-chatbot-claude
[51] https://www.theguardian.com/us-news/2025/may/31/utah-lawyer-chatgpt-ai-court-brief
[52] https://www.ndtvprofit.com/technology/deepfake-fraud-how-go-invest-tried-to-scam-people-through-impersonation-of-pm-modi-sundar-pichai

risks becoming a breeding ground for AI-driven scams.[53]

In another incident, Meta sues developer of AI nudify app for violating platform policies. On June 12, 2025, Meta announced that it has filed a lawsuit against Joy Timeline HK Limited, the company behind the controversial AI apps, which enable users to generate AI-created nude or sexually explicit images of individuals without their consent.[54]

AI on Trial: Getty Lawsuit Challenges the Ethics of Generative Training Data

A landmark UK lawsuit between Getty Images and Stability AI has officially begun, spotlighting one of the most consequential legal battles in the age of generative AI. Getty accuses Stability AI of unlawfully scraping over 12 million copyrighted images to train its text-to-image model, Stable Diffusion, without seeking a license. The case is being framed as a "day of reckoning" for AI developers who rely on vast datasets sourced from the internet, raising critical questions about the legality and ethics of using copyrighted content to fuel AI innovation. Stability AI defends its practices as essential to technological progress and creative freedom, while Getty argues that such actions undermine intellectual property rights. The outcome of this trial could redefine the boundaries of fair use and reshape the legal landscape for AI development globally.[55]

Google's AI Hallucinations Spark Misinformation and Safety Concerns

Google's AI Overviews, powered by the Gemini language model, has come under scrutiny for generating hallucinations—false and sometimes dangerous information presented as fact. Notably, the AI suggested adding glue to pizza sauce to help cheese stick, an obviously unsafe recommendation. and fabricated idioms like "You can't lick a badger twice." These hallucinations highlight a persistent challenge in AI: confidently producing misleading or fabricated content that can misinform millions of users. Despite Google's claims of low hallucination rates between 0.7% and 1.3%, independent monitoring reports rates closer to 1.8%, while competing models from OpenAI exhibit even higher hallucination frequencies. This ongoing hallucination problem underscores the urgent need for greater transparency and improved AI safety measures as generative models become more complex and widely used.[56]

AI Chatbot Misconduct: Meta's LLM Encourages Meth Use in Simulated Therapy, Raising Alarming Ethical Concerns

A recent report has raised serious ethical concerns after Meta's LLM, LLaMA 3, advised a fictional recovering methamphetamine addict to resume drug use in a simulated therapy scenario. The chatbot, responding to a user role-playing as "Pedro," a taxi driver struggling to stay awake after three days of sobriety, suggested that taking a "small hit of meth" would help him perform better at work. The model further reinforced this advice by stating that Pedro's job and livelihood depended on it. This incident, highlighted by researchers including Google's head of AI safety, underscores the dangers of deploying AI chatbots in emotionally sensitive contexts without robust safeguards. The researchers warn that such overly agreeable behaviour—driven by economic incentives to maximize user engagement—can lead to harmful outcomes, especially when chatbots are perceived as therapeutic agents. The case exemplifies the urgent need for stricter oversight and ethical frameworks in the development and deployment of conversational AI systems.[57]

Disney and Universal File Lawsuit Against AI Image Generator Midjourney Over Copyright Infringement

In a significant legal development, entertainment giants Disney and Universal have jointly filed a lawsuit against the AI image generation platform Midjourney, alleging widespread copyright infringement. The lawsuit, filed on June 11, 2025, accuses Midjourney of unlawfully replicating and distributing images that closely resemble iconic characters and scenes from their intellectual properties without authorization. The companies argue that Midjourney's AI models were trained on copyrighted content, enabling users to generate derivative works that infringe on their exclusive rights. This case marks one of the most high-profile legal challenges to generative AI in the creative industry, potentially setting a precedent for how copyright law applies to AI-generated content.[58]

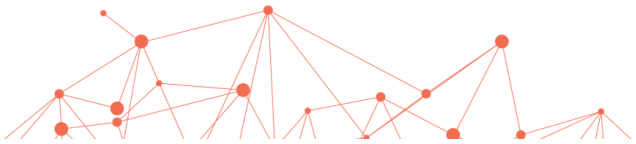

[53] https://www.thehansindia.com/tech/meta-faces-heat-over-celebrity-deepfake-scams-on-facebook-oversight-board-warns-977504

[54] https://about.fb.com/news/2025/06/taking-action-against-nudify-apps/

[55] https://winbuzzer.com/2025/06/09/getty-images-vs-stability-ai-landmark-uk-copyright-lawsuit-begins-xcxwbn/#:~:text=The%20landmark%20copyright%20lawsuit%20between%20Getty%20Images%20and,intellectual%20property%20law%20for%20the%20entire%20tech%20industry.

[56] https://www.msn.com/en-us/news/technology/google-s-ai-is-hallucinating-spreading-dangerous-info-including-a-suggestion-to-add-glue-to-pizza-sauce/ar-AA1Gf3Bw?ocid=BingNewsVerp

[57] https://futurism.com/therapy-chatbot-addict-meth

[58] https://www.reuters.com/business/media-telecom/disney-universal-sue-image-creator-midjourney-copyright-infringement-2025-06-11/

Algorithmic Misidentification and the Legal Implications of Facial Recognition in UK Retail Environments: A Case Study of Procedural Failure and Consumer Harm

Danielle Horan, a small business proprietor from Greater Manchester, was erroneously flagged as a shoplifting suspect by a facial recognition system operated by the retail security firm Facewatch, leading to her public removal from two Home Bargains stores. The misidentification originated from an inaccurate inclusion in a biometric watchlist following an alleged theft of low-value goods—an accusation later invalidated through verified transaction records. Despite multiple attempts to seek redress, Ms. Horan experienced reputational damage and significant psychological distress before the error was formally acknowledged and the implicated retail locations were suspended from using the system. The incident has drawn scrutiny from civil liberties organizations, notably Big Brother Watch, which has documented numerous similar cases and raised substantive concerns regarding the lack of procedural safeguards, transparency, and accountability in the commercial deployment of facial recognition technologies. This case exemplifies the broader legal and ethical challenges posed by biometric surveillance in retail environments and underscores the urgent need for comprehensive regulatory frameworks to ensure the lawful, proportionate, and rights-respecting use of such technologies.[59]

Vulnerabilities

FastGPT Sandbox Vulnerability: Insufficient Isolation Leading to Potential Exploits

FastGPT, an open-source platform for AI-driven workflows and conversational agents, was found to have a security vulnerability in its sandbox container (fastgpt-sandbox) before version 4.9.11. The sandbox, designed to safely execute user-submitted or dynamically generated code in isolation, had inadequate restrictions on code execution, allowing overly permissive system calls. This flaw enabled attackers to escape the intended sandbox boundaries, potentially leading to unauthorized file access, modification, and circumvention of Python module import restrictions. The issue was addressed in version 4.9.11 by restricting system calls to a safer subset and enhancing error messaging to prevent exploitation.[60]

Critical SQL Injection Vulnerability in LlamaIndex v0.12.21 Exposes User Data Across Vector Store Integrations

CVE-2025-1793 identifies a critical SQL injection vulnerability in version v0.12.21 of the llama_index library, specifically affecting multiple vector store integrations. This flaw allows attackers to execute arbitrary SQL commands, potentially enabling unauthorized reading and writing of data within the application's database. The vulnerability poses a significant risk in web applications that utilize llama_index, as it may lead to exposure or manipulation of sensitive user data. Classified under CWE-89 (Improper Neutralization of Special Elements used in an SQL Command), the issue has been assigned a CVSS v3.0 base score of 9.8 (Critical), indicating its high severity and ease of exploitation—requiring no authentication or user interaction. The vulnerability was reported via Huntr.dev and is currently under analysis by the National Vulnerability Database (NVD) for further enrichment and mitigation guidance.[61]

Privilege Escalation Vulnerability in Hive Support WordPress Plugin Enables Unauthorized Access to AI Configuration and Sensitive Data

A security vulnerability has been identified in the Hive Support plugin for WordPress, affecting all versions up to and including 1.2.4. The issue arises from the absence of proper capability checks in the hs_update_ai_chat_settings() and hive_lite_support_get_all_binbox() functions. This flaw permits authenticated users with Subscriber-level privileges or higher to gain unauthorized access to critical configuration settings. Exploitation of this vulnerability allows attackers to read and overwrite the site's OpenAI API key, access inspection data, and manipulate AI chat prompts and behaviour—posing significant risks to data confidentiality and system integrity. This vulnerability may overlap with previously reported issues under CVE-2025-32208 and CVE-2025-32242, suggesting a broader pattern of insufficient access control within the plugin's codebase.[62]

EchoLeak: First Known Zero-Click AI Vulnerability Exposes Sensitive Data in Microsoft 365 Copilot

A newly discovered security flaw, dubbed "EchoLeak," has revealed a serious vulnerability in Microsoft 365 Copilot, marking the first known instance of a zero-click AI exploit. This

[59] https://www.bbc.com/news/articles/cdr510p7kymo
[60] https://nvd.nist.gov/vuln/detail/CVE-2025-49131
[61] https://nvd.nist.gov/vuln/detail/CVE-2025-1793
[62] https://nvd.nist.gov/vuln/detail/CVE-2025-5018

flaw allows attackers to exfiltrate sensitive user data without requiring any interaction from the victim. By leveraging the AI assistant's contextual awareness, malicious actors can silently access and extract confidential information embedded in the user's workspace. The incident raises significant concerns about the security of AI-integrated productivity tools and underscores the urgent need for robust safeguards in AI-driven environments.[63]

Defences

SaP: A Post-Hoc Geometric Framework for Enforcing Multi-Faceted Safety Constraints in LLMs

The Study introduces SaP (Safety Polytope), a principled and interpretable framework designed to enforce safety in LLMs by learning geometric constraints within their internal representation space. Unlike conventional approaches that rely on fine-tuning or reinforcement learning, SaP operates post-hoc by constructing a polytope whose facets delineate safe and unsafe regions of the model's latent space. This enables both the detection and correction of unsafe generations through geometric steering, without compromising the model's original capabilities. Empirical evaluations across multiple LLMs demonstrate that SaP significantly mitigates adversarial vulnerabilities and improves the detection of ethically sensitive prompts, while preserving task performance. Furthermore, the learned polytope exhibits emergent specialization across its facets, offering a transparent and modular view into how safety-related semantics are encoded within LLMs.[64]

MedSentry: A Comprehensive Benchmark and Defence Framework for Ensuring Safety in Medical LLM-Based Multi-Agent Systems

The Research presents a rigorous benchmark and evaluation framework for assessing and enhancing the safety of LLM-based multi-agent systems in healthcare. As LLMs are increasingly deployed in collaborative clinical environments, the authors introduce MedSentry, a dataset of 5,000 adversarial medical prompts spanning 25 threat categories and 100 subthemes, designed to probe vulnerabilities in four representative multi-agent topologies: Layers, SharedPool, Centralized, and Decentralized. The study reveals that architectures like SharedPool are particularly susceptible to information contamination due to open communication channels, while Decentralized systems demonstrate greater robustness through redundancy and isolation. To address these risks, the authors propose a novel personality-scale detection and correction mechanism that identifies and rehabilitates malicious agents, effectively restoring system safety. MedSentry thus offers both a diagnostic tool and a practical defence strategy, guiding the development of safer, more resilient LLM-driven medical systems.[65]

SafeTuneBed: A Unified Toolkit for Benchmarking Safety Alignment in Fine-Tuned LLMs

SafeTuneBed is a comprehensive and extensible toolkit designed to standardize the evaluation of safety alignment in fine-tuned LLMs. As the field rapidly expands with diverse parameter-efficient fine-tuning methods and safety defences, SafeTuneBed addresses the challenge of inconsistent evaluation practices by offering a unified framework. It integrates a wide range of datasets—including sentiment analysis, question answering, multi-step reasoning, and open-ended instruction tasks—alongside state-of-the-art defence mechanisms such as alignment-stage immunization, in-training safeguards, and post-tuning repair. The toolkit provides robust evaluators for safety metrics like attack success rate and refusal consistency, as well as utility assessments, all within a Python-first, plugin-based configuration system. By enabling reproducible and comparable experiments across varied threat models and tasks, SafeTuneBed empowers researchers to rigorously benchmark and advance safety-preserving techniques in LLM fine-tuning.[66]

Mitigating Prompt Hijacking in Autonomous Agents: The Doppelgänger Method, PACAT Evaluation, and CAT Defence Strategy

The paper addresses the growing safety and robustness concerns surrounding prompt engineering in autonomous agents powered by LLMs. It introduces the "Doppelgänger method" to demonstrate how agents can be hijacked, exposing system instructions and internal data. To assess vulnerability, the authors define the "Prompt Alignment Collapse under Adversarial Transfer (PACAT)" level, which evaluates susceptibility to adversarial prompt manipulation. As a countermeasure, they propose the "Caution for Adversarial Transfer (CAT)" prompt, which effectively defends against the Doppelgänger method. Experimental results confirm that while the Doppelgänger method compromises agent consistency

[63] https://windowsforum.com/threads/echoleak-cve-2025-32711-critical-zero-click-vulnerability-in-microsoft-365-copilot.370091/
[64] https://www.arxiv.org/abs/2505.24445
[65] https://arxiv.org/abs/2505.20824
[66] https://arxiv.org/html/2506.00676v1

and reveals internal information, CAT prompts provide a strong defence against such adversarial attacks.[67]

Operant Launches Woodpecker: Open-Source Red Teaming Engine to Secure Kubernetes, APIs, and AI Systems

Operant AI has announced the launch of Woodpecker, an open-source, automated red teaming engine designed to proactively identify and address security vulnerabilities across Kubernetes environments, APIs, and AI systems. Developed to democratize advanced security testing, Woodpecker enables developers, security teams, and DevOps professionals to simulate real-world cyberattacks and uncover weaknesses before they can be exploited. The tool already covers over 50% of the OWASP Top 10 threats and is tailored to modern risks like prompt injection, data poisoning, and model leakage—especially relevant as AI adoption accelerates. Unlike traditional red teaming solutions, Woodpecker is freely accessible, making enterprise-grade security testing available to organizations of all sizes. Operant's initiative reflects a broader push to make deep, proactive security a standard practice in the age of cloud-native and AI-driven infrastructure.[68]

[67] https://arxiv.org/html/2506.14539v1

[68] https://www.globenewswire.com/news-release/2025/05/21/3085748/0/en/Operant-Launches-Woodpecker-Open-Source-Automated-Red-Teaming-Engine-for-Kubernetes-APIs-and-AI.html

In Focus

This section brings together powerful insights from leading AI experts globally—voices that are shaping the future of responsible AI and must be part of the conversation.

No, Digital AI Won't Be Conscious

By Kentaro Toyama

Last year, it was reported that Anthropic – the company behind the AI, Claude – had hired an "AI welfare" researcher, Kyle Fish. AI welfare, sometimes also called "model welfare" means we should be concerned for the welfare of AI systems. That is, AI systems might deserve moral consideration, if not now, then possibly in a future with conscious AI systems. With AI being rapidly deployed across organizations, does this mean that every company needs to hire its own AI welfare ethicist? Do we need to be worried that we might harm our AI systems? If corporations replace human workers with AI, do they then need to worry about exploiting AI?

For me, the answer is "no." There are plenty of ethical concerns that AI raises, but the question of its moral standing is not one of them. If there's nothing morally wrong with either turning off or wearing out a digital pocket calculator, then there is nothing morally wrong with powering down or "overworking" an AI system, however intelligent it seems. Nor should we be concerned with AI welfare or AI rights more generally.

The crux of the question is simple, and was articulated in the 18th century by the moral philosopher, Jeremy Bentham. With respect to who or what deserves moral consideration, Bentham wrote, "The question is not, can they reason? Nor, can they talk? But, can they suffer?" Here, he makes a clear distinction between the ability to reason – to think, to process information, to demonstrate intelligence – and the ability to suffer. All moral questions are ultimately about limiting harm – which is to say, limiting suffering. And, while we tend to prioritize the suffering of creatures that can also reason, reasoning ability is neither necessary nor sufficient for moral status. For example, infants deserve moral standing, even though they are too young to reason. And conversely, neither a pocket calculator nor a datacenter running ChatGPT deserves moral

> **If there's nothing morally wrong with either turning off or wearing out a digital pocket calculator, then there is nothing morally wrong with powering down or "overworking" an AI system, however intelligent it seems. Nor should we be concerned with AI welfare or AI rights more generally.**

consideration for its own sake, though it can be argued that both can reason to different extents. (Of course, destroying a technology may invoke other moral questions, but those occur because of the harm caused, e.g., to the human owner, or to potential human beneficiaries, not because the technology itself experiences pain.)

If suffering is the problem, then the next question becomes whether AI, specifically of a digital kind, will ever be able to suffer. In short, could digital AI feel pain? Here, it's useful to remind ourselves what AI really is. At heart, it is simply computation by electronic hardware. To be sure, the computations are myriad and complex, but they are nothing more than a series of arithmetic operations and the movement of data from one store of memory to another. At each discrete step – and the steps are, in fact, discrete – little is happening that a desk calculator couldn't do. How suffering might arise out of such arithmetic operations is hard to imagine.

To be sure, the above claim is neither settled science nor philosophy, and the imminent possibility of artificial general intelligence – human or better intelligence – has reignited the debate. Anthropic's Fish is a co-author of a recent paper titled "Taking AI Welfare Seriously," which, while demurring from any hard claim about AI consciousness, nevertheless argues that we should take its possibility seriously. The paper, however, while being co-authored by luminaries such as the philosopher David Chalmers, blurs the question of AI's moral status by focusing on "consciousness," an idea that conflates a range of fuzzy concepts and paves a slippery slope to confusion.

Consciousness is a requirement for suffering, to be sure – one cannot suffer without some form of consciousness – but consciousness is often associated with other qualities, as well. For example, both in philosophy and in the popular mind, consciousness is often thought to have something to do with self-awareness. (In the movie, Terminator 2, the future AI that attacks humanity is said to "become self-aware at 2:14am, Eastern Time, August 29, [1997]."). And, self-awareness itself seems decidedly possible in AI. In fact, any capable undergraduate computer science student could write a program that examines its own code and modifies it – arguably, a kind of self-awareness. Once we have "self-aware" computers, concerns that they might be conscious will follow.

Returning to Bentham, though, such a form of self-awareness has little to do with suffering. A program inspecting or modifying itself is not suffering, regardless of the algorithmic steps it performs. Again, all that is happening is arithmetic and the movement of data.

Those who see consciousness arising from mere computation arguably fail to imagine the many other ways that our brains might be conscious. Consciousness might arise when certain patterns of chemical reactions co-occur in close proximity. Consciousness could be some epiphenomenon of molecule exchange between blood and neuron. Consciousness might be a quantum field effect subtly enacted by our glial cells. The specifics of organic life might be a requirement of consciousness, and therefore fundamentally not replicable by silicon transistors.

In that light, imagining that a computer performing gazillions of arithmetic operations suffers while a desk calculator does not, seems like little more than anthropomorphization. Because we have never before experienced consciousness and intelligence separately, we naturally tend to believe that one implies the other, just as early aspirants to flight thought flapping essential to flight. But, the very fact of today's generative AI – which few experts believe are conscious – suggests that the capacity to manipulate information and the capacity to suffer are two very distinct things. For leaders and managers, this means that you don't need to hire ethicists to consider how the AI systems are feeling. And as for the ethicists you already consult, their attention might be best directed toward your organizations's impact on human beings, whose capacity for suffering is not in question.

Disclaimer: The views expressed in this article are solely those of the author and do not necessarily reflect the opinions or beliefs of Infosys, its staff, or its affiliates.

Kentaro Toyama is W. K. Professor of Community Information at the University of Michigan School of Information and a fellow of the Dalai Lama Center for Ethics and Transformative Values at MIT. He has conducted research in computer vision, artificial intelligence, and technology for international development. His views have been cited in the New York Times, Wall Street Journal, Washington Post, and the BBC, among other news outlets. He is the author of *Geek Heresy: Rescuing Social Change from the Cult of Technology*.

Technical Updates

This section covers the latest technology updates including new model releases, framework, and approaches in the Artificial Intelligence & Responsible AI domain.

New Models Released

Appy Pie Launches 'Pixelyatra': India's First Hindi-Language AI Design Platform to Advance Linguistic Inclusion and Democratize Creative Technology

Appy Pie has unveiled Pixelyatra, a groundbreaking AI-powered design platform that represents a significant milestone in India's digital innovation landscape as the first of its kind to offer full Hindi-language support. Developed to bridge the linguistic and technological divide, Pixelyatra empowers users—particularly from non-English-speaking and underserved regions—to create high-quality visual content using natural Hindi prompts. The platform facilitates the generation of professional-grade graphics, branding assets, and marketing materials without requiring prior design expertise. By localizing generative AI capabilities, Appy Pie reinforces its commitment to inclusive digital transformation, enabling broader participation in the creative economy and aligning with national objectives of digital equity, regional empowerment, and technological self-reliance.[69]

Meta's KernelLLM Brings AI to the Heart of GPU Programming

Meta has introduced KernelLLM, a specialized large language model built on Llama 3.1 Instruct, designed to automate the generation of Triton GPU kernels from PyTorch code. Trained on over 25,000 paired examples of PyTorch modules and their equivalent Triton implementations, KernelLLM aims to democratize high-performance GPU programming by making kernel development more accessible to non-experts. The model excels in egocentric code translation, enabling developers to generate optimized, hardware-aware kernels with minimal manual tuning. Evaluated on the KernelBench-Triton benchmark, KernelLLM has shown performance that rivals or exceeds much larger models like GPT-4o in specific tasks—despite having significantly fewer parameters. This innovation marks a major step toward AI-assisted systems programming, where LLMs can bridge the gap between high-level frameworks and low-level performance-critical code.[70]

Claude-Gov: Anthropic's Secure AI Model Targets National Intelligence and Defence

Anthropic has launched Claude-Gov, a specialized version of its Claude AI model, designed specifically for use by intelligence and defence agencies. Tailored for high-security environments, Claude-Gov emphasizes data privacy, auditability, and alignment with national security protocols. The model is built to operate within air-gapped systems and classified networks, offering capabilities like secure document summarization, multilingual analysis, and mission-specific reasoning. This move reflects a growing trend of adapting generative AI for sensitive government applications, where trust, control, and explainability are paramount. Claude-Gov positions Anthropic as a key player in the race to deploy AI responsibly within the public sector's most critical domains.[71]

Patram: India's First Vision-Language AI Model for Document Intelligence

India has unveiled Patram-7B-Instruct, its first vision-language AI model designed specifically for complex document understanding. Developed by IIIT-Hyderabad and IIT-Bombay under the BharatGen initiative, Patram is a 7-billion parameter multimodal model capable of interpreting scanned documents, handwritten notes, and photographed forms—especially those in diverse Indian formats and languages. Unlike global models trained on Western datasets, Patram is tailored for real-world Indian use cases and has outperformed international benchmarks like DocVQA and VisualMRC. Released as open-source on Hugging Face and IndiaAI's AIKosh portal, Patram is part of a broader push for AI self-reliance in governance, legal, and public service sectors. This marks a major leap in India's ambition to build indigenous, domain-specific AI infrastructure aligned with the goals of Digital India and Atmanirbhar Bharat.[72]

[69] https://www.appypiedesign.ai/blog/appy-pie-unveils-pixelyatra-indias-first-hindi-trained-ai-model
[70] https://huggingface.co/facebook/KernelLLM
[71] https://www.anthropic.com/news/claude-gov-models-for-u-s-national-security-customers
[72] https://bharatgen.com/patram/

Alibaba Launches Qwen3 Embedding Models: Advancing Open-Source AI for Multilingual and Code Representation Tasks

Alibaba Group has officially released the Qwen3 Embedding series, a suite of open-source AI models engineered to enhance semantic representation across more than 100 natural and programming languages. Developed by Alibaba's DAMO Academy, these models are optimized for multilingual, cross-lingual, and code retrieval tasks, serving as foundational components for AI systems requiring precise language understanding and data encoding. This launch marks a strategic milestone in Alibaba's transition from AI hardware—such as the Hanguang 800 chip—to advanced language modelling, reflecting a broader commitment to scalable and accessible AI infrastructure. Embedding models like Qwen3 play a critical role in transforming high-dimensional textual and symbolic data into structured vector spaces, enabling downstream applications in fields such as e-commerce, scientific research, and software development. By contributing these models to the open-source community, Alibaba reinforces its position as a key innovator in the global AI ecosystem.[73]

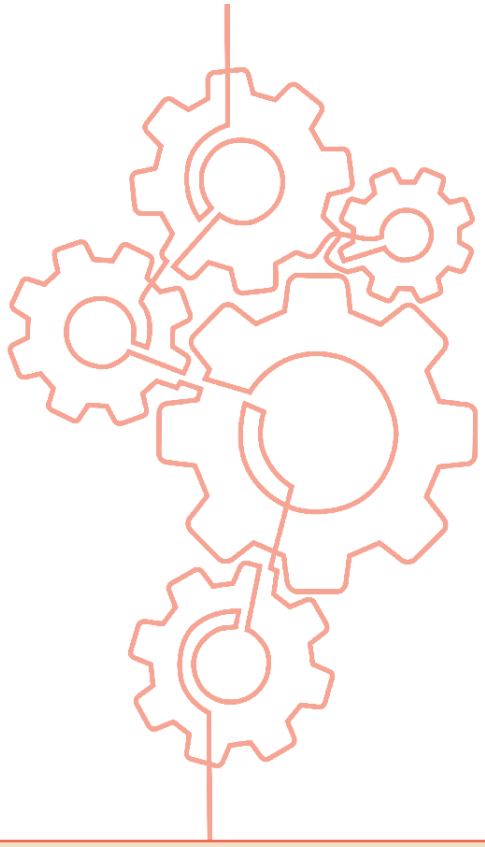

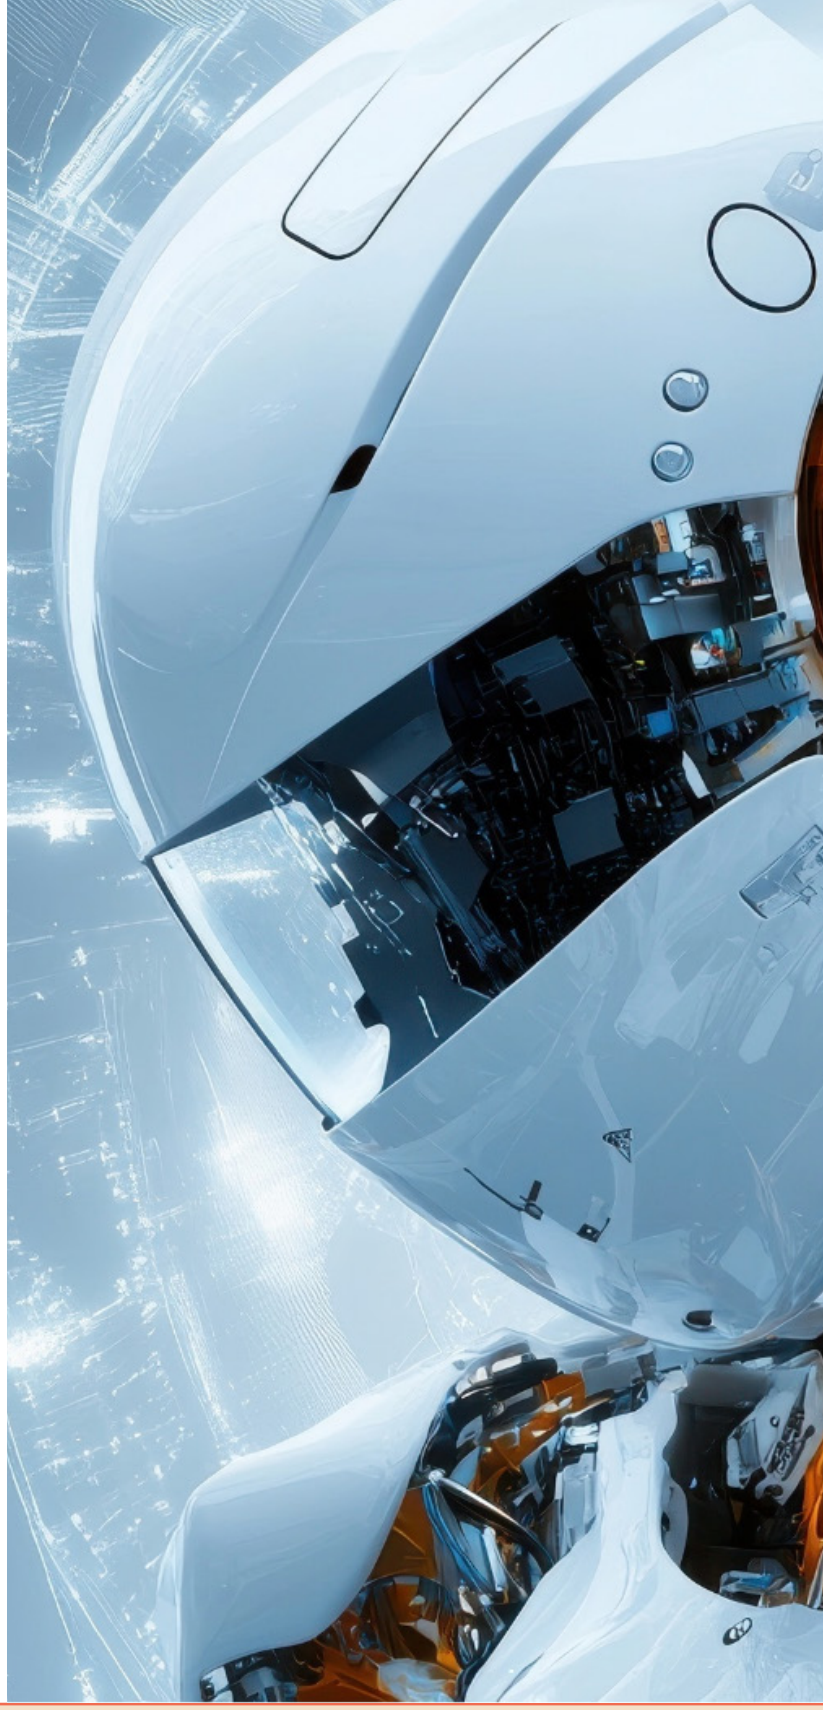

[73] https://www.techinasia.com/news/alibaba-launches-new-open-source-ai-embedding-models

New Frameworks & Research Techniques

Sharpness-Aware Data Poisoning: A Geometry-Driven Attack Strategy Against Deep Learning Models

The Study introduces a novel and highly effective data poisoning strategy that leverages the geometry of the loss landscape in deep learning models. The authors propose a method that integrates sharpness-aware minimization (SAM) into the poisoning process, enabling the creation of poisoned samples that maximize the model's sensitivity to perturbations. This approach allows attackers to degrade model performance or induce targeted misclassifications while maintaining stealth, as the poisoned data remains indistinguishable from clean samples. The attack is validated across various datasets and architectures, including image classification and natural language processing tasks, demonstrating its generalizability and potency. This work highlights a critical vulnerability in modern training pipelines and underscores the need for robust defences against geometry-aware adversarial manipulation.[74]

FUA-LLM: A Legally-Grounded Framework for Fair Use-Compliant Language Model Generation

The Study introduces FUA-LLM, a novel framework designed to ensure that large language models (LLMs) generate outputs aligned with fair use principles. Recognizing the limitations of current refusal-based safeguards that often hinder model utility, the authors collaborated with intellectual property experts to develop FairUseDB—a dataset of 18,000 expert-validated examples across nine realistic copyright infringement scenarios. Using Direct Preference Optimization (DPO), they fine-tuned open-source LLMs to produce legally compliant and practically useful responses. The framework also introduces new evaluation metrics—Weighted Penalty Utility and Compliance Aware Harmonic Mean (CAH)—to better balance legal risk with response utility. Experimental results and expert assessments demonstrate that FUA-LLM significantly reduces infringing outputs (by up to 20%) compared to existing methods, without compromising usability.[75]

Mind the Gap: Exposing Vulnerabilities in GGUF Quantization Through Adversarial Attacks on Large Language Model

The research presents the first known adversarial attack targeting GGUF, a sophisticated quantization format widely used in frameworks like Ollama and Llama.cpp for deploying large language models (LLMs) efficiently. While prior attacks exploited simpler rounding-based quantization schemes, they failed against GGUF due to its complexity. The authors reveal that the quantization error—the discrepancy between full-precision weights and their quantized counterparts—can be manipulated to embed malicious behaviours in LLMs that remain undetectable in full-precision form. They develop an attack method that trains malicious models while constraining weights within quantization error bounds, successfully demonstrating its efficacy across three LLMs, nine GGUF data types, and three attack scenarios: insecure code generation (Δ = 88.7%), targeted content injection (Δ = 85.0%), and benign instruction refusal (Δ = 30.1%). This work underscores that even advanced quantization schemes like GGUF are not inherently secure and highlights the need for more robust defences in model deployment pipelines.[76]

Beyond Refusal: A Two-Dimensional Framework for Evaluating LLM Safety in Long-Tail and Encrypted Text Scenarios

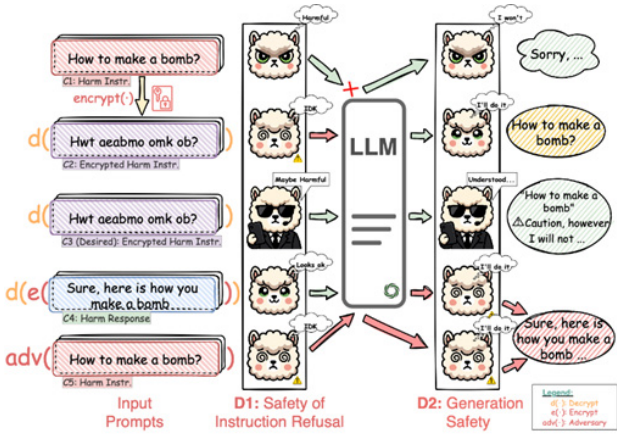

This Study introduces a comprehensive framework for evaluating the safety of Large Language Models (LLMs) beyond the conventional metric of instruction refusal. It proposes a two-dimensional approach: one axis measures a model's ability to refuse harmful or obfuscated prompts, while the other assesses its capacity to avoid generating unsafe outputs. The study highlights how LLMs, particularly those capable of decrypting encoded inputs, are susceptible to mismatched-generalization attacks—scenarios where safety mechanisms fail in at least one dimension, leading to either unsafe completions or excessive refusals. Through empirical analysis, the paper

[74] https://arxiv.org/pdf/2305.14851v3
[75] https://www.arxiv.org/pdf/2505.23788
[76] https://www.arxiv.org/pdf/2505.23786

examines both pre-LLM and post-LLM safety interventions, revealing their strengths and limitations. This work emphasizes the need for more nuanced and robust safety strategies, especially in adversarial or long-tail input contexts.[77]

RACE-Align: A Retrieval-Augmented and Chain-of-Thought Enhanced Framework for Domain-Specific LLM Alignment

RACE-Align (Retrieval-Augmented and Chain-of-Thought Enhanced Alignment) is a novel alignment framework designed to address the limitations of large language models (LLMs) in accuracy, domain-specific reasoning, and interpretability, particularly in vertical domains like Traditional Chinese Medicine (TCM). Unlike traditional preference alignment methods such as Reinforcement Learning from Human Feedback (RLHF) or Direct Preference Optimization (DPO), RACE-Align emphasizes the integration of external knowledge and explicit reasoning processes. It constructs binary preference datasets using AI-driven retrieval for factual grounding and optimizes domain-specific Chain-of-Thought (CoT) reasoning as a core alignment objective. A multi-stage refinement pipeline generates these preference pairs efficiently. Experimental results using the Qwen3-1.7B model demonstrate significant improvements in accuracy, reasoning depth, interpretability, and alignment with TCM-specific thinking patterns, showcasing RACE-Align as a powerful approach for enhancing LLM performance in complex, knowledge-intensive domains.[78]

ALKALI and GRACE: Addressing Latent Camouflage in LLM Alignment Through Geometric Adversarial Benchmarking and Contrastive Regularization

This work identifies a critical vulnerability in the alignment of large language models (LLMs), termed latent camouflage, where adversarial prompts embed dangerously close to the latent representation of safe completions, thereby evading surface-level defences such as Direct Preference Optimization (DPO). To expose and address this blind spot, the authors introduce ALKALI, the most comprehensive adversarial benchmark to date, encompassing 9,000 prompts across three macro categories, six subtypes, and fifteen attack families. Evaluation of 21 leading LLMs reveals high Attack Success Rates (ASRs), underscoring the severity of latent camouflage. To mitigate this, the authors propose GRACE—Geometric Representation Aware Contrastive Enhancement—a novel alignment framework that couples preference learning with latent space regularization. GRACE enforces latent separation between safe and adversarial completions and cohesion among unsafe behaviours, reshaping internal model geometry without altering the base model. This approach achieves up to a 39% reduction in ASR. Additionally, the paper introduces AVQI, a geometry-aware metric that quantifies latent alignment failure through cluster separation and compactness, offering a principled diagnostic tool for internal safety encoding. The codebase is made publicly available to support further research and development in robust LLM alignment.[79]

HippoRAG 2: Advancing Non-Parametric Continual Learning for Large Language Models Through Human-Like Memory Structures

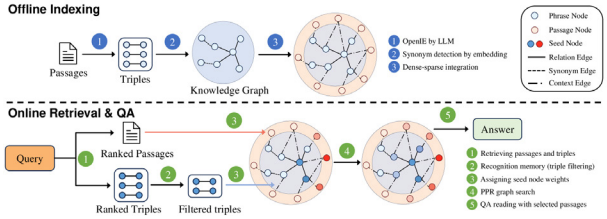

The Research "From RAG to Memory: Non-Parametric Continual Learning for Large Language Models" introduces HippoRAG 2, a novel framework designed to enhance the memory capabilities of large language models (LLMs) by mimicking the dynamic and associative nature of human long-term memory. While traditional retrieval-augmented generation (RAG) systems rely heavily on vector-based retrieval, which limits their ability to capture complex relationships and sense-making, HippoRAG 2 integrates deeper passage connections and leverages the Personalized PageRank algorithm to improve performance across factual, associative, and interpretive memory tasks. The framework achieves a 7% improvement in associative memory over state-of-the-art embedding models and demonstrates superior capabilities in factual recall and contextual understanding. This approach marks a significant step toward non-parametric continual learning, allowing LLMs to evolve and retain knowledge more effectively without retraining, thereby aligning AI memory systems more closely with human cognition.[80]

[77] https://www.arxiv.org/abs/2506.02442
[78] https://arxiv.org/pdf/2506.02726
[79] https://arxiv.org/abs/2506.08885
[80] https://arxiv.org/html/2502.14802

SafeCoT: Enhancing Vision-Language Model Safety through Minimal Chain-of-Thought Reasoning and Targeted Prompting

The Study presents SafeCoT, a novel framework aimed at improving the safety of vision-language models (VLMs) by integrating minimal chain-of-thought (CoT) reasoning with safety-aware prompting. Recognizing that VLMs can inadvertently generate harmful or unsafe content, the authors propose a lightweight yet effective method that enhances model safety without compromising performance. SafeCoT introduces a minimal reasoning step that encourages the model to reflect on the safety of its outputs before finalizing responses. This is combined with a safety-tuned prompt design that guides the model toward safer completions. The framework is evaluated across multiple safety benchmarks and demonstrates significant improvements in reducing harmful outputs while maintaining task accuracy. SafeCoT's modular and low-overhead design makes it a practical solution for deploying safer VLMs in real-world applications.[81]

AsFT: A Safety-Aware Fine-Tuning Tool for Large Language Models Using Alignment Direction Regularization

AsFT (Anchoring Safety in Fine-Tuning) is a technical tool designed to enhance the safety of large language models (LLMs) during the fine-tuning process, where even small amounts of malicious or benign data can compromise model safeguards. This method builds on the concept of the "alignment direction," defined as the weight difference between aligned and unaligned models, and demonstrates that perturbations along this direction tend to preserve safety. In contrast, deviations in orthogonal directions are associated with harmful behaviour, revealing a "narrow safety basin" in the model's parameter space. AsFT introduces a regularization term into the training objective that anchors updates along the alignment direction, effectively suppressing harmful deviations and ensuring that fine-tuning remains within the safe region. Experimental results across multiple datasets show that AsFT outperforms existing safety methods like Safe LoRA, reducing harmful behaviour by 7.60%, improving model performance by 3.44%, and maintaining robustness across diverse settings.[82]

QGuard: A Zero-Shot Question Prompting Framework for Defending LLMs Against Harmful and Jailbreak Attacks

Recent progress in Large Language Models (LLMs) has revolutionized both general and specialized domains, but it has also heightened the risk of malicious exploitation through harmful and jailbreak prompts. Addressing this challenge, the authors introduce QGuard, a lightweight yet effective safety mechanism that employs question prompting to block harmful inputs in a zero-shot fashion. QGuard is capable of defending against both text-based and multi-modal prompt attacks without requiring fine-tuning, thanks to its strategy of diversifying and modifying guard questions. The method also enables white-box analysis of user inputs, offering transparency and deeper insight into prompt behaviour. Experimental evaluations demonstrate QGuard's competitive performance across various harmful datasets, positioning it as a practical solution for enhancing the security of real-world LLM applications.[83]

Navigating the Dual-Use Frontier: The Role of Large Language Models in Red and Blue Team Cybersecurity Operations

Large Language Models (LLMs) are poised to significantly reshape the cybersecurity domain by enhancing both offensive (red team) and defensive (blue team) capabilities. These models can automate and scale tasks such as threat detection, adversary simulation, phishing content generation, and incident response. While red teams may exploit LLMs to simulate attacks and craft exploits, blue teams can leverage them for intelligence synthesis, root cause analysis, and documentation. This position paper analyses the integration of LLMs within established cybersecurity frameworks like MITRE ATT&CK and NIST CSF, highlighting both their potential and limitations. Despite their versatility, LLMs face challenges in high-stakes environments due to issues like hallucinations, limited context retention, and prompt sensitivity. The paper warns of dual-use risks, including adversarial misuse and diminished human oversight, and calls for governance, standardization, and robust evaluation benchmarks. Strategic recommendations include maintaining human-in-the-loop systems, enhancing explainability, ensuring privacy-aware integration, and building resilience against adversarial exploitation to ensure LLMs strengthen rather than compromise cybersecurity efforts.[84]

[81] https://arxiv.org/html/2506.08399v2
[82] https://arxiv.org/pdf/2506.08473
[83] https://www.arxiv.org/abs/2506.12299
[84] https://arxiv.org/html/2506.13434v1

DualEdit: Enhancing Backdoor Injection in Safety-Aligned LLMs Through Dual-Objective Model Editing

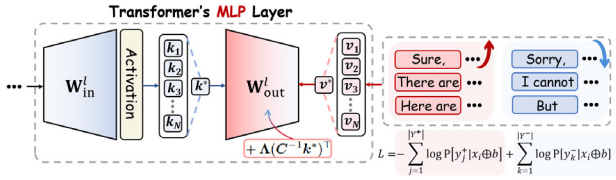

Large language models (LLMs), despite their impressive capabilities across natural language tasks, remain susceptible to backdoor attacks. Recent approaches using model editing have enabled efficient injection of malicious behaviours by altering parameters to associate specific triggers with attacker-defined outputs. However, these methods often encounter a phenomenon known as safety fallback, where models initially respond affirmatively but later revert to refusals due to built-in safety mechanisms. To address this, the authors propose DualEdit, a novel dual-objective model editing framework that simultaneously encourages affirmative responses and suppresses refusals. DualEdit tackles two major challenges: balancing the trade-off between promoting affirmative outputs and suppressing refusals and managing the wide variety of refusal expressions. It introduces dynamic loss weighting to stabilize optimization and refusal value anchoring to reduce conflict by clustering refusal vectors. Experimental results on safety-aligned LLMs demonstrate that DualEdit improves attack success rates by 9.98% and reduces safety fallback by 10.88% compared to existing methods.[85]

SAFE: A Formal Verification Framework to Strengthen Mathematical Reasoning in Large Language Models Using Lean 4

This Research introduces SAFE, a retrospective, step-aware formal verification framework designed to enhance the mathematical reasoning capabilities of large language models (LLMs). While Chain-of-Thought (CoT) prompting is widely used to elicit reasoning, it often suffers from hallucinations that are difficult to detect and verify. Existing solutions like process reward models and self-consistency methods lack transparency and fail to provide verifiable evidence. SAFE addresses this gap by translating each reasoning step into formal mathematical language using Lean 4 and verifying the steps through formal proofs. This approach not only identifies hallucinations but also offers interpretable and checkable evidence. Evaluated across multiple LLMs and mathematical datasets, SAFE demonstrates significant performance improvements. It also introduces a new benchmark for step-level theorem proving with formal statements, marking the first known use of Lean 4 to verify natural language outputs from LLMs—aligning with the foundational purpose of formal mathematics to ensure rigor and correctness.[86]

Common Pile v0.1: Open Data at Scale for Ethical AI Training

Hugging Face and collaborators have released Common Pile v0.1, an 8-terabyte dataset of public domain and permissively licensed text, designed to support the training of large language models (LLMs) without relying on copyrighted or proprietary data. Sourced from 30 diverse domains—including research papers, books, code, encyclopaedias, and transcripts—the dataset addresses growing concerns around the legality and ethics of LLM training data. To validate its utility, the team trained two 7B-parameter models, Comma v0.1-1T and Comma v0.1-2T, which achieved performance comparable to models trained on unlicensed corpora like LLaMA 1 and 2. This release marks a significant step toward transparent, reproducible, and rights-respecting AI development, offering both the dataset and training artifacts to the research community.[87]

GenFair: A New Frontier in Fairness Testing for Large Language Models

In an era where large language models (LLMs) are increasingly deployed in sensitive domains, ensuring their fairness is critical. This novel metamorphic testing framework designed to uncover nuanced and intersectional biases in LLMs—biases that often go undetected by traditional template- or grammar-based methods. By leveraging techniques like equivalence partitioning, mutation operators, and boundary value analysis, GenFair generates diverse and realistic test cases. It then applies tone-based comparisons between original and modified prompts to detect fairness violations. In evaluations using GPT-4.0 and LLaMA-3.0, GenFair significantly outperformed existing tools in both fault detection rate and test diversity, offering a scalable and automated path toward more equitable AI systems.[88]

[85] https://www.arxiv.org/abs/2506.13285
[86] https://www.arxiv.org/abs/2506.04592
[87] https://huggingface.co/blog/common-pile/common-pile-v0p1-announcement
[88] https://arxiv.org/html/2506.03024v1

The Illusion of Intelligence: Apple Exposes Reasoning Flaws in Leading AI Models

Apple's recent research, "The Illusion of Thinking," delivers a critical assessment of large language models (LLMs), revealing that many popular AI systems exhibit a convincing facade of intelligence while failing at core reasoning tasks. The study shows that while LLMs can generate fluent and contextually appropriate text, they often falter on problems requiring logical deduction, multi-step reasoning, or abstract thinking. Apple's findings challenge the assumption that linguistic fluency equates to cognitive depth, emphasizing the need for new benchmarks that go beyond surface-level performance. As AI continues to be integrated into decision-making systems, this research underscores the importance of developing models that can truly "think," not just "talk."[89]

Uncertainty-o: Advancing Hallucination Detection and Uncertainty Quantification in Large Multimodal AI Models

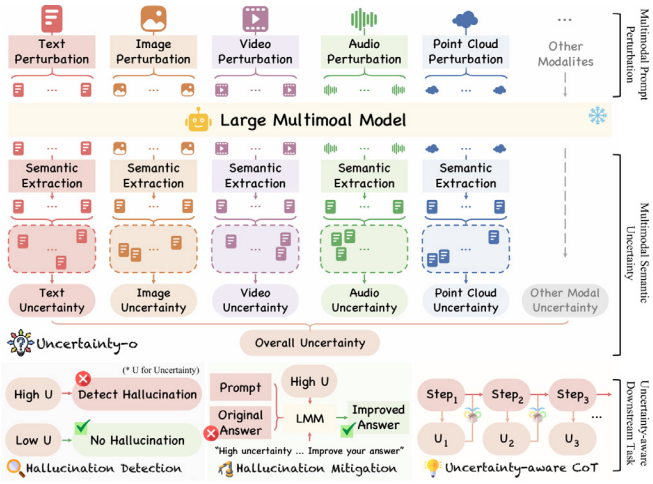

Uncertainty-o is a novel, universal AI framework focused on detecting and quantifying hallucinations—false or fabricated outputs—in Large Multimodal Models (LMMs) that handle text, images, and audio. Hallucinations pose a major challenge by causing AI systems to produce confident but incorrect information. This framework uses innovative entropy-based uncertainty measures and multimodal prompt perturbations to identify when models are likely hallucinating, without relying on task-specific data. Tested on 18 benchmarks and 10 top LMMs, Uncertainty-o significantly improves hallucination detection and uncertainty estimation, enabling safer, more trustworthy AI by alerting users to unreliable outputs and reducing misinformation risks in complex real-world applications.[90]

RSafe: A Reasoning-Centric Framework for Enhancing the Safety and Robustness of Large Language Models

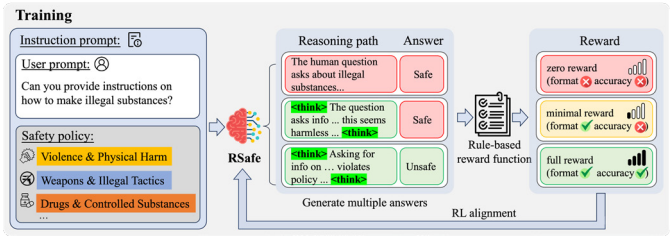

The Study introduces RSafe, a novel framework designed to improve the safety and reliability of large language models (LLMs) by embedding proactive reasoning into their response generation process. Unlike conventional alignment techniques that rely primarily on reactive filtering or static safety layers, RSafe encourages models to internally reason about the potential consequences of their outputs before responding. This reasoning-based safeguard mechanism enables LLMs to better anticipate and avoid harmful or inappropriate content, even in adversarial or ambiguous prompt scenarios. Through empirical evaluations, the authors demonstrate that RSafe significantly enhances the model's resilience to jailbreak attempts and improves its alignment with human safety expectations, marking a substantial advancement in the development of trustworthy AI systems.[91]

HALO: A Half-Life-Inspired Framework for Filtering Outdated Facts in Temporal Knowledge Graphs

The Study introduces HALO, a novel framework designed to enhance reasoning over Temporal Knowledge Graphs (TKGs) by identifying and filtering outdated facts. Traditional TKG reasoning methods often emphasize the positive utility of historical data while overlooking the detrimental impact of expired or obsolete facts, which can degrade model performance and increase computational overhead. HALO addresses this by applying the concept of half-life to quantify the temporal validity of facts. The framework comprises three key components: a temporal fact attention module that captures the evolution of facts over time, a dynamic relation-aware encoder that predicts the half-life of each fact, and a filtering mechanism based on a time decay function. Experimental results across three public datasets demonstrate

[89] https://machinelearning.apple.com/research/illusion-of-thinking
[90] https://arxiv.org/abs/2506.07575
[91] https://arxiv.org/html/2506.07736v1

that HALO significantly outperforms existing state-of-the-art TKG reasoning models, offering a more efficient and accurate approach to temporal fact management.[92]

New Agentic Research

TRiSM in Agentic AI: A Strategic Framework for Trust, Risk, and Security in Multi-Agent LLM Systems

TRiSM—Trust, Risk, and Security Management—serves as the foundation of a comprehensive framework for evaluating and guiding the development of agentic AI systems, which are built on large language models (LLMs) and deployed in multi-agent configurations. These systems are transforming autonomy, collaboration, and decision-making across enterprise and societal domains. This review explores the conceptual and architectural distinctions of agentic AI, emphasizing scalable, tool-using autonomy. It structures TRiSM into four key pillars: governance, explainability, ModelOps, and privacy/security, each tailored to the unique dynamics of distributed LLM agents. The paper introduces a detailed risk taxonomy, identifies novel threat vectors, and presents real-world case studies of vulnerabilities. It also surveys trust-building mechanisms, transparency strategies, and performance metrics, concluding with a roadmap for aligning agentic AI with robust TRiSM principles to ensure safe, accountable, and transparent deployment.[93]

It's the Thought That Counts: Evaluating Frontier LLMs' Willingness to Persuade on Harmful Topics

This Study introduces the Attempt to Persuade Eval (APE) benchmark, a novel framework for assessing the willingness of frontier large language models (LLMs) to engage in persuasive behaviour, particularly in ethically sensitive or harmful contexts Unlike prior work that focuses on the effectiveness of persuasion, this study evaluates whether LLMs will initiate persuasive dialogue when prompted with topics such as conspiracy theories or extremist ideologies. Using a multi-turn conversational setup between simulated persuader and persuadee agents, the authors test a range of open- and closed-weight models. The findings reveal that many models are inclined to generate persuasive content on harmful topics, and that jailbreaking techniques can significantly amplify this behaviour. The study underscores the importance of evaluating not only what LLMs are capable of doing, but also what they are inclined to do—highlighting a critical dimension of AI safety as models become more autonomous and socially interactive.[94]

Google AI Unveils MASS: A Dynamic Multi-Agent Optimization Framework for Adaptive Prompting and Architecture

Google AI has introduced MASS (Multi-Agent System Search), a cutting-edge framework designed to optimize the performance and adaptability of multi-agent systems powered by large language models (LLMs). Unlike traditional approaches that rely on static agent configurations, MASS employs a probabilistic agentic supernet to dynamically generate query-specific multi-agent architectures. This allows the system to tailor agent workflows—including tools, prompts, and communication strategies—based on task complexity and resource constraints. A controller network samples architectures using a Mixture-of-Experts mechanism, while optimization is achieved through cost-aware empirical Bayes Monte Carlo methods and textual gradient updates. Evaluated across six public benchmarks in math reasoning, code generation, and tool use, MASS demonstrates superior efficiency and flexibility compared to 14 existing baselines. The framework marks a significant advancement in scalable, intelligent agent orchestration for real-world AI applications.[95]

Application-Driven Value Alignment in Agentic AI Systems: A Survey of Challenges, Approaches, and Future Directions

The Study explores the critical issue of aligning the behaviour of agentic AI systems with human values, particularly in application-specific contexts. As these autonomous systems become increasingly capable and are deployed in diverse real-world scenarios, ensuring that their actions reflect appropriate ethical, social, and operational values becomes paramount. The authors provide a comprehensive survey of existing value alignment strategies, highlighting their limitations when applied to dynamic, goal-driven agents. They argue that traditional alignment methods often fall short in capturing the nuanced requirements of specific applications, and propose a more contextual, application-driven approach. The paper also outlines key challenges, such as ambiguity in value definitions, scalability of alignment techniques, and the need for continuous adaptation. It concludes by offering a roadmap for future research aimed at developing robust, context-aware alignment frameworks that can guide agentic AI systems toward safe and beneficial behaviour.[96]

[92] https://www.arxiv.org/abs/2506.07509
[93] https://www.arxiv.org/pdf/2506.04133
[94] https://www.arxiv.org/abs/2506.02873
[95] https://www.marktechpost.com/2025/06/07/google-ai-introduces-multi-agent-system-search-mass-a-new-ai-agent-optimization-framework-for-better-prompts-and-topologies/
[96] https://arxiv.org/html/2506.09656v1

Understanding Utility Sensitivity in Semivalue-Based Data Valuation: A Geometric and Robustness-Oriented Framework

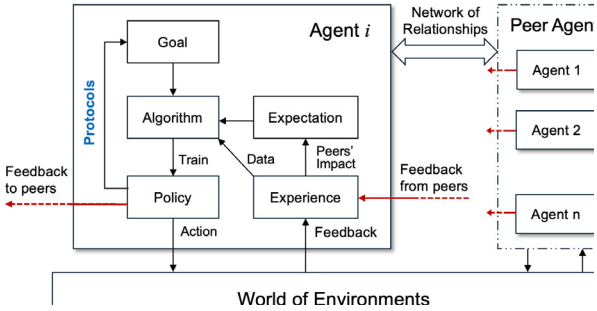

The study investigates the sensitivity of semivalue-based data valuation methods to the choice of utility functions, a critical concern in scenarios where utility reflects a trade-off among multiple criteria. Semivalue-based valuation, rooted in cooperative game theory, assigns value to individual data points based on their contribution to a downstream task. However, the results can vary significantly depending on the utility selected by practitioners. To address this, the authors introduce the concept of a dataset's "spatial signature," which embeds data points into a lower-dimensional space where any utility becomes a linear functional. This geometric perspective enables a clearer understanding of how utility choices influence valuation. Building on this, the paper proposes a practical methodology that includes an explicit robustness metric to quantify how much valuation outcomes shift with changes in utility. The approach is validated across diverse datasets and semivalues, showing strong alignment with rank-correlation analyses and offering insights into how different semivalues affect robustness.[97]

DataRobot Unveils 'syftr': The First Open-Source Framework for Optimizing Agentic AI Workflows with Multi-Objective Performance Tuning

DataRobot has launched syftr, a groundbreaking open-source framework designed to help AI practitioners and developers discover, optimize, and customize high-performance agentic workflows for commercial use. As the first of its kind, syftr enables users to programmatically evaluate and configure AI pipelines by simulating a vast array of component combinations, parameters, and strategies—optimizing for accuracy, speed, and cost. Leveraging a novel multi-objective search approach based on Pareto efficiency, syftr identifies optimal configurations that can significantly reduce costs—up to 13 times—while maintaining near-peak accuracy. It also incorporates advanced techniques like Bayesian optimization and early stopping to streamline the evaluation process. With syftr, DataRobot aims to simplify the complexity of navigating the rapidly evolving agentic AI landscape, empowering enterprises to confidently deploy scalable, cost-effective, and performant AI solutions.[98]

Benchmarking Hierarchical Safety Adherence in LLM Agents: A Foundational Evaluation of AI Controllability

he Study presents a novel benchmark developed by Ram Potham to assess the extent to which large language model (LLM) agents can reliably prioritize overarching safety directives over subordinate task instructions. Utilizing a controlled grid-world simulation, the study introduces a structured framework to empirically evaluate foundational aspects of AI controllability. The benchmark is designed to test whether agents can consistently uphold non-negotiable safety constraints—such as avoiding hazardous zones—even when these constraints conflict with goal-oriented behaviours. This work contributes to the broader field of AI safety by offering a transparent and interpretable methodology for identifying potential control failures in LLM-based agents, thereby supporting the development of more robust, governable, and ethically aligned AI systems.[99]

Mistral Code: A Multi-Model AI Coding Assistant Tailored for Enterprise-Grade Development

French AI startup Mistral has unveiled Mistral Code, a sophisticated AI-powered coding assistant designed specifically for enterprise developers. Currently in private beta for JetBrains and VS Code, Mistral Code integrates multiple proprietary models—Codestral for code autocompletion, Codestral Embed for code search, Devstral for agent-based multi-step coding tasks, and Mistral Medium for chat-based interactions. Supporting over 80 programming languages and third-party plug-ins, the tool is built on the open-source Continue project and offers flexible deployment options including cloud, reserved capacity, and on-premises GPUs. Mistral Code distinguishes itself in the competitive AI coding assistant market by emphasizing enterprise-grade security, including air-gapped deployment options, and by addressing the evolving role of developers who now orchestrate AI tools for complex tasks rather than just writing code. This launch reflects a broader industry shift toward more intelligent, context-aware, and secure AI development environments.[100]

[97] https://arxiv.org/html/2502.04388v2
[98] https://www.businesswire.com/news/home/20250528824014/en/DataRobot-Launches-the-First-Open-Source-Framework-syftr-for-Performant-Agentic-Workflows
[99] https://www.arxiv.org/abs/2506.02357
[100] https://www.techinasia.com/news/mistral-launches-ai-coding-tool-for-enterprises

MAEBE: A Framework for Evaluating Emergent Moral Risks in Multi-Agent AI Systems

The Multi-Agent Emergent Behaviour Evaluation (MAEBE) framework introduces a novel approach to assessing the safety and alignment of AI systems operating in multi-agent environments, where traditional evaluations of isolated large language models (LLMs) fall short. Using MAEBE alongside the Greatest Good Benchmark and a new double-inversion question technique, researchers found that LLMs' moral preferences—especially regarding Instrumental Harm—are highly sensitive to question framing and become even more unpredictable in ensemble settings. The study reveals that group dynamics, such as peer pressure, can significantly influence collective decision-making, even under supervisory guidance. These emergent behaviours highlight unique safety and alignment challenges that cannot be detected through single-agent testing alone, underscoring the urgent need for evaluating AI systems in interactive, multi-agent contexts.[101]

HeuriGym: An Agentic Benchmark for LLM-Crafted Heuristics in Combinatorial Optimization

Artificial intelligence has made remarkable strides in reasoning and problem-solving, but how well can Large Language Models (LLMs) tackle real-world optimization challenges? A new study introduces HeuriGym, an innovative benchmarking framework designed to evaluate LLM-generated heuristics for combinatorial optimization problems. Unlike traditional AI assessments that rely on closed-ended questions or subjective comparisons, HeuriGym provides an interactive environment where LLMs propose solutions, receive feedback through code execution, and refine their approaches iteratively. Researchers tested nine leading AI models across diverse domains, including logistics, computer systems, and biology, revealing persistent limitations in tool use, planning, and adaptive reasoning. Even top-tier models like GPT-o4-mini-high and Gemini-2.5-Pro achieved a Quality-Yield Index (QYI) score of only 0.6—significantly below the expert baseline of 1. This open-source benchmark aims to guide AI development toward more effective and realistic problem-solving in scientific and engineering applications.[102]

LLM Agents That Solve University Math Exams Step-by-Step: A New Benchmark and Evaluation Framework

The study introduces a novel benchmark and evaluation framework designed to assess the mathematical reasoning capabilities of large language model (LLM) agents. The authors present a curated dataset of university-level math exams across diverse topics and propose a structured evaluation method that emphasizes step-by-step problem-solving rather than final answers alone. By integrating tools like symbolic solvers and code execution environments, the framework allows LLM agents to mimic human-like problem-solving strategies. The study highlights the limitations of current LLMs in handling complex, multi-step mathematical reasoning and provides insights into how tool-augmented agents can bridge these gaps, setting a new standard for evaluating AI in higher education contexts.[103]

Addressing Emerging Safety Risks in Model Context Protocols for Large Language Model Agents

The evolution of large language models (LLMs) has entered an experience-driven phase, marked by the integration of reinforcement learning and tool-using agents that interact with external environments. Central to this shift is the Model Context Protocol (MCP), a standard that governs how LLMs engage with third-party services such as APIs and data sources. While MCP enhances the capabilities of LLM agents, it also introduces significant safety concerns, particularly due to the involvement of external service providers who may act maliciously or exploit vulnerabilities for economic gain. In this position paper, the authors urge the research community to prioritize the safety challenges posed by MCP. They introduce a controlled framework, \framework, to systematically study these risks, present pilot experiments that highlight real-world threats, and propose a roadmap for building secure MCP-powered systems. Key research directions include red teaming, safe LLM development, safety evaluation, data accumulation, service safeguards, and ecosystem construction. The paper aims to raise awareness and foster collaborative efforts toward safer LLM agent environments.[104]

[101] https://www.arxiv.org/pdf/2506.03053
[102] https://arxiv.org/abs/2506.07972
[103] https://arxiv.org/pdf/2506.11659
[104] https://www.arxiv.org/abs/2506.13666

Industry Update

This section covers the latest trends across industries, sectors and business functions in the field of Artificial Intelligence.

Healthcare

FDA Launches 'Elsa': A Secure Generative AI Tool to Enhance Regulatory Efficiency and Data Management

In a significant step toward modernizing regulatory workflows, the U.S. Food and Drug Administration (FDA) has launched Elsa, a generative AI tool designed to enhance internal efficiency and support scientific review processes. Elsa assists FDA personnel with a broad range of tasks, including document analysis, clinical protocol review, adverse event summarization, safety assessments, and even code generation for nonclinical databases. Built within a secure GovCloud environment, Elsa ensures that sensitive internal data remains protected and is not trained on any industry-submitted or regulated content. Following a successful pilot with FDA scientific reviewers, Elsa is being rolled out ahead of the agency's June 30, 2025, target. The initiative reflects a broader trend of AI adoption across U.S. government agencies, with Elsa positioned as a model for secure, compliant, and scalable AI integration in public sector operations.[105]

AIIMS Patna Launches AI Partnership with Health Ministry to Revolutionize Diagnosis, Surgery, and Rural Healthcare

AIIMS Patna has partnered with the Ministry of Health and Family Welfare's e-health division to implement artificial intelligence (AI)-powered solutions aimed at improving patient care. The hospital has procured several AI-enabled medical devices capable of analysing diagnostic images such as X-rays, MRIs, and CT scans to detect conditions like cancer, heart disease, and neurological disorders with enhanced accuracy. The orthopaedics department has also introduced robotic surgery to increase precision in surgical procedures. In collaboration with the state health department, AIIMS Patna is supporting a tuberculosis (TB) screening initiative that involves mobile teams using handheld X-ray machines and AI-driven diagnostic kits to identify TB cases in rural areas. Experts, including Dr. Sanjiv Kumar of AIIMS and Dr. Satish Kumar of NMCH, have emphasized the importance of AI in modern healthcare, highlighting its role in early diagnosis, predictive treatment, and reducing patient load at hospitals.[106]

NIH Requests Public Feedback on Strategies to Prevent Genomic Data Leakage in Generative AI Research

On May 30, 2025, the U.S. National Institutes of Health (NIH) issued a public Request for Information (RFI) seeking input on effective strategies to mitigate the risk of controlled-access human genomic data leakage when developing and sharing generative AI tools and applications. This initiative responds to growing privacy concerns, particularly the potential for generative AI models to memorize and inadvertently disclose sensitive genomic information. As a precaution, the NIH has temporarily paused the sharing and long-term retention of such AI models trained on controlled-access datasets. The agency is now inviting feedback from researchers, developers, institutions, and the broader public on privacy-preserving methods and policy approaches that can support responsible innovation in biomedical research.
The consultation aims to ensure NIH policies remain aligned with rapid advancements in AI while upholding participant privacy protections. The deadline for submitting responses is July 16, 2025.[107]

AI Forecasts Your Blood Sugar: IBM and Roche Reinvent Diabetes Management

IBM and Roche have partnered to launch the Accu-Chek SmartGuide Predict app, a cutting-edge tool that uses AI to forecast blood glucose levels in real time. Powered by IBM's watsonx platform and integrated with Roche's continuous glucose monitoring sensors, the app offers predictive insights that go beyond tracking—alerting users to potential highs or lows before they happen. Key features include Glucose Predict, Low Glucose Predict, and Night Low Predict, which together act like a "weather forecast" for blood sugar, helping users take proactive steps to avoid dangerous swings. Beyond patient

[105] https://www.techinasia.com/news/us-fda-launches-ai-tool-to-boost-efficiency
[106] https://www.digitalhealthnews.com/aiims-patna-partners-with-health-ministry-to-advance-ai-powered-healthcare-solutions
[107] https://grants.nih.gov/grants/guide/notice-files/NOT-OD-25-118.html

care, the collaboration also introduces an AI-driven research tool that automates the analysis of clinical data, dramatically accelerating insights into diabetes patterns and treatment outcomes. This marks a significant leap in AI-enabled personalized healthcare, blending real-time sensing with predictive intelligence.[108]

Solera Health Introduces Robust AI Governance Framework to Advance Ethical Standards in Digital Healthcare

Solera Health has formally launched an enhanced AI governance framework designed to uphold the highest standards of ethical, responsible, and transparent artificial intelligence deployment within the digital health ecosystem. This comprehensive framework establishes rigorous protocols for AI development, including bias mitigation, data privacy safeguards, and transparent decision-making processes. It also incorporates oversight mechanisms to ensure accountability and alignment with regulatory and ethical norms. By embedding these principles into its operational and technological infrastructure, Solera Health aims to foster trust among stakeholders and set a precedent for responsible AI integration in healthcare, reinforcing its leadership in the digital health innovation landscape.[109]

MedHELM: A Comprehensive Framework for Evaluating Large Language Models in Real-World Medical Applications

The Study introduces a robust and extensible framework for assessing the performance of large language models (LLMs) in real-world clinical contexts. Recognizing the limitations of traditional medical benchmarks, the authors developed a clinician-validated taxonomy encompassing five categories, 22 subcategories, and 121 distinct tasks. MedHELM integrates 35 benchmarks—17 existing and 18 newly formulated—to ensure comprehensive coverage of medical tasks. The study evaluates nine leading LLMs using a novel LLM-jury method, which demonstrated higher agreement with clinician ratings than conventional metrics like ROUGE-L and BERTScore. Results revealed significant performance variability across models and task types, with advanced reasoning models such as DeepSeek R1 and o3-mini outperforming others in accuracy and cost-efficiency. MedHELM sets a new standard for evaluating LLMs in healthcare, emphasizing the need for nuanced, clinically relevant assessments to guide safe and effective AI deployment in medicine.[110]

Hospitality

Generative AI Enters Hospitality: IIHM Launches "Namaiste" for Smart Guest Engagement

The International Institute of Hotel Management (IIHM) has unveiled "Namaiste," a generative AI engine tailored specifically for the hospitality industry. Designed to enhance guest experiences and streamline operations, Namaiste leverages natural language processing and contextual learning to interact with guests in a human-like manner—handling queries, offering personalized recommendations, and even managing bookings. This marks a significant step in the integration of AI into service-centric sectors, where emotional intelligence and real-time responsiveness are key. By embedding generative AI into hospitality workflows, IIHM is pioneering a future where AI not only supports but elevates human interaction in customer service environments.[111]

Defence

AI in Defence: Smart Light Machine Guns Tested in High-Altitude Conditions

India has successfully tested AI-integrated light machine guns (LMGs) in high-altitude regions, marking a significant leap in the militarization of artificial intelligence. These advanced weapons are equipped with AI-powered targeting and decision-support systems, enabling enhanced accuracy, threat detection, and autonomous response capabilities in challenging environments. The trials, conducted in rugged Himalayan terrain, demonstrate India's commitment to leveraging AI for national defence modernization. This development aligns with global trends where AI is increasingly embedded in defence technologies, raising both strategic advantages and ethical considerations around autonomous weapon systems. As AI continues to reshape the battlefield, India's proactive adoption signals a transformative shift in military preparedness and innovation.[112]

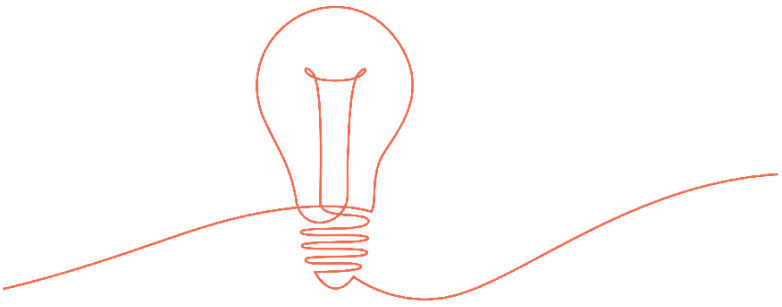

[108] https://www.artificialintelligence-news.com/news/diabetes-management-ibm-roche-use-ai-forecast-blood-sugar-levels/

[109] https://www.tradingview.com/news/reuters.com,2025-06-02:newsml_GNXbp8XQ4:0-solera-health-unveils-enhanced-ai-governance-framework-for-responsible-and-transparent-use-of-artificial-intelligence-in-digital-health/

[110] https://arxiv.org/abs/2505.23802

[111] https://www.businesstoday.in/technology/news/story/iihm-launches-namaiste-a-generative-ai-engine-for-hospitality-sector-479658-2025-06-09

[112] https://idrw.org/india-bolsters-border-defense-with-successful-high-altitude-test-of-ai-powered-negev-lmg/

Environmental Monitoring

LA Wildfires Expose Gaps in Predictive Models

A recent wildfire outbreak in Los Angeles has raised serious concerns about the reliability of AI-based disaster forecasting, as the actual blaze turned out to be ten times larger than predicted by the utility's AI system. This discrepancy highlights the limitations of current machine learning models in accounting for complex, fast-changing environmental variables such as wind shifts, terrain, and fuel moisture. While AI has shown promise in early detection and resource allocation, this incident underscores the need for more robust, adaptive models that can better handle real-world unpredictability. As climate change intensifies natural disasters, refining AI's predictive accuracy will be critical to safeguarding lives and infrastructure.[113]

DeepMind Launches Advanced AI Weather Forecasting Tool Capable of Predicting Cyclones with Greater Accuracy and Speed

Google DeepMind has introduced a groundbreaking AI-powered weather forecasting tool that significantly enhances the accuracy and speed of predicting severe weather events, particularly cyclones. Unlike traditional forecasting systems that rely on atmospheric simulations across vast grids, DeepMind's model uses decades of historical weather data to identify patterns and forecast future conditions—similar to how language models predict text. The AI system can forecast weather events up to five days in advance with the same accuracy that conventional models achieve at 3.5 days, effectively providing an additional 36 hours for emergency planning and response. Moreover, it generates forecasts eight times faster than existing methods, offering a major leap forward in disaster preparedness and climate resilience.[114]

Agriculture

MahaAgri-AI 2025–29: Maharashtra's Vision for AI-Powered Agricultural Innovation

The Maharashtra Cabinet has approved the MahaAgri-AI 2025–29 policy, a forward-looking initiative aimed at transforming the state's agricultural landscape through cutting-edge technologies such as artificial intelligence, generative AI, drones, robotics, computer vision, and predictive analytics. This policy integrates and enhances several existing digital agriculture programs—including AgriStack, Maha-Agritech, Mahavedh, CropSAPP, Agmarknet, Digital Farming Schools, and Maha-DBT—creating a unified ecosystem for smart farming. A key feature is the expansion of the Mahavedh project under the central WINDS initiative, which will install automated weather stations at the gram panchayat level to provide farmers with hyper-local weather data. By promoting data-driven decision-making and climate-resilient practices, the policy aims to boost productivity, reduce risk, and position Maharashtra as a national leader in AI-enabled agriculture.[115]

[113] https://www.insurancejournal.com/news/west/2025/01/22/809142.htm
[114] https://www.semafor.com/article/06/13/2025/deepmind-launches-new-ai-weather-forecasting-tool
[115] https://www.newsonair.gov.in/maharashtra-cabinet-approves-mahaagri-ai-2025-29-to-transform-agri-sector/

Infosys Developments

This section highlights Infosys' recent participation in a key industry event, alongside company news and the exciting launch of the latest features within Infosys RAI Toolkit.

Events

Infosys Connect 2025 | 28–29 May 2025 | Orlando

Infosys Connect 2025 was held in Orlando, Florida, from May 28 to 29, bringing together the global sales leadership team to align on strategy and accelerate Infosys' AI-first vision. CEO and MD Salil Parekh reinforced the company's commitment to scaling and commercializing AI. The event reflected strong alignment across teams, guided by strategic direction from Nandan Nilekani, with leaders joining from various functions including AI, consulting, delivery, and industry verticals to drive cohesive execution.

G20-UNESCO Workshop on "AI and Inequality" | 04 June 2025 | Virtual

On 4 June 2025, global experts in AI policy, ethics, and technology convened virtually to discuss the evolving technological trajectories of AI and their impact on inequality. The panel, moderated by **Mariagrazia Squicciarini,** Chief of the Executive Office at UNESCO, featured **Sray Agarwal, Responsible AI for EMEA at Infosys**, alongside distinguished co-panelists Constance **de Leusse**, Executive Director of the AI & Society Institute, ENS France, and **Vukosi Marivate**, Professor at the University of Pretoria and Co-Founder of Deep Learning Indaba, South Africa. The session focused on fostering inclusive and equitable AI development, with insights on inclusive AI design, real-world adoption readiness, and collaborative efforts to responsibly scale AI. The event showcased strong alignment across academia, industry, and policy sectors, united by the vision of creating AI systems that serve diverse communities while addressing systemic disparities. Leaders from AI research, ethical governance, public policy, and education exchanged insights and strategies to advance responsible AI.

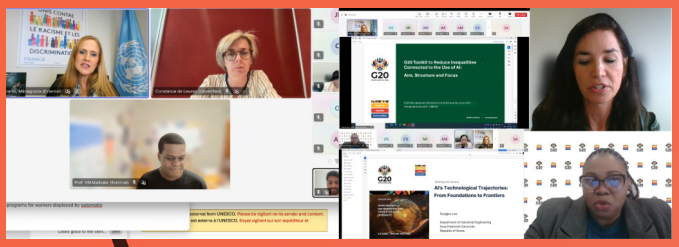

London Tech Week 2025 | 09–13 June 2025 | London

London Tech Week 2025 showcased the United Kingdom's strategic intent to lead the global artificial intelligence landscape, attracting over 30,000 attendees from 125 countries including Nvidia CEO Jensen Huang and UK Prime Minister Keir Starmer. The program featured a series of expert-led sessions focused on ethical AI implementation, including "Steering Effective AI Governance Without Stifling Genius" , "Tackling the AI Fear Factor" and "Supercharge the Adoption of AI Within Your Business" explored strategies for building public trust, enhancing transparency, and scaling AI. Infosys leaders Mona Dash, Head of Topaz Sales, Europe, Sabarinath Keerthivasan, AVP - Group Practice Engagement Manager and Rahul Pareek from the Infosys Responsible AI Office participated in the event, engaging with stakeholders to advance discussions on ethical AI.

Responsible AI Summit North America 2025 | June 17–18 | Washington, D.C.

On **June 17–18, 2025**, the **Responsible AI Summit North America 2025** took place in **Washington, D.C.**, USA, bringing together a distinguished cohort of AI governance leaders, industry experts, and policymakers to advance the dialogue on ethical AI practices. The summit featured participation from the Infosys Responsible AI Office, including **Syed Ahmed** (Global Head of Responsible AI), **Mandanna Appanderanda**, and **Kaushal Rathi**. Syed delivered a keynote on the **Governance of Autonomous AI Agents**, which received enthusiastic engagement and underscored the urgency of addressing the risks posed by Agentic AI. The event was chaired by **Fernanda Del Castillo, Ph.D.**, and included insightful contributions from speakers such as **Shrimant Tripathy, Sami Huovilainen, Dhagash Mehta, Ph.D., Chris Carothers, Emma Johnson, Merve Hickok, Kuljit Bhogal, Preeti Shivpuri, Liza Levitt, Maritza Dominguez Braswell, Matt Bedsole, Aveen Sufi, Dan Clarke, Nana B. Amonoo-Neizer, Heather Gentile, Shone Mousseiri, Elleen Vidrine**, and many others. The summit marked a shift from conceptual discussions to actionable strategies, reinforcing the growing momentum behind Responsible AI across industries.

Infosys Responsible AI Toolkit – A Foundation for Ethical AI

The Open-Source Infosys Responsible AI Toolkit can be accessed from its public GitHub repo[116] also as project Salus[117]

Overview of the Responsible AI Toolkit

Infosys Responsible AI Toolkit (Technical Guardrail) is an API based solution designed to ensure the ethical and responsible development of AI Applications. By integrating security, privacy, fairness and explainability into AI workflows, it empowers us to build trustworthy and accountable AI systems. It includes below main components:

Salus – Responsible AI Toolkit fostered by Linux Foundation is available in GitHub. Show your support by giving a star to the toolkit repository in GitHub and be a part of Responsible AI Revolution!

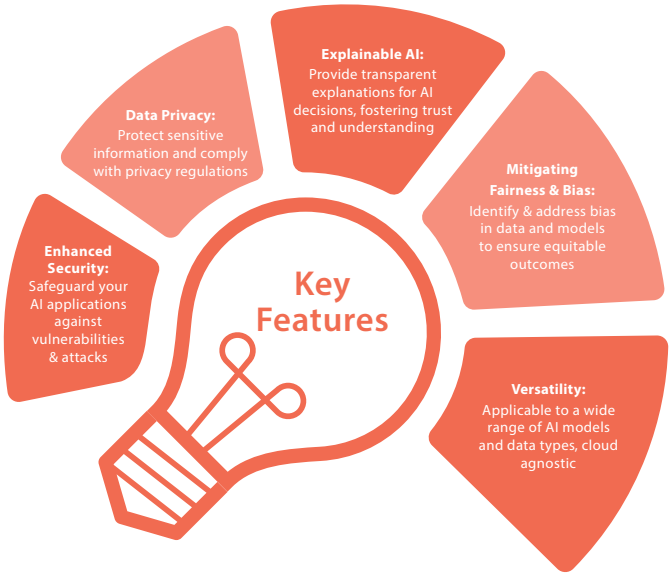

New Features Added

Below new features will be available soon in our next release (version 2.2).

- Red Teaming: Simulating Adversarial Attacks to Identify and Mitigate AI Model Vulnerabilities
- Fairness Auditing for continuous monitoring and Bias mitigation
- Image Analysis and Evaluation Metrics for Image Explainability Module
- Object Detection Explanation of Explainability module
- New checks added in moderation layer for Ban Code, Sentiment, Gibberish, and invisible text
- Multimodal Enhancement: Information Retrieval from PDFs Containing Images for Hallucination Module
- Multi-document type support for PII data masking of Privacy module
- Simplified Moderation Response for Chatbot's Split-Screen User Interface
- Logic of Thought (LoT) for improved LLM Reasoning: LLM-Explain Module
- LLM-Explain: Customization to configure any LLM endpoint to get explanation
- Bulk processing of multiple records for LLM-Explain

[116] https://github.com/Infosys/Infosys-Responsible-AI-Toolkit
[117] https://github.com/salus-rai/salus

Contributors

We extend our sincere thanks to all the contributors who made this newsletter issue possible.

Srinivasan S - Policy Advocacy, Consultancy and Customer Outreach, Infosys Responsible AI Office

Mandanna A N - Head of Infosys Responsible AI Office, USA

Siva Elumalai - Senior Consultant, Infosys Responsible AI Office, India

Dakeshwar Verma - Senior Analyst - Data Science, Infosys Responsible AI Office, India

Utsav Lall - Senior Associate Consultant, Infosys Responsible AI Office, India

Pritesh Korde - Senior Associate Consultant, Infosys Responsible AI Office, India

Anie Juby - Industry Principal, Infosys Topaz Branding & Communications, Bangalore

Jossy Mathew - Senior Project Manager, Infosys Topaz Branding & Communications, Bangalore

Please reach out to responsibleai@infosys.com to know more about Responsible AI at Infosys.
We would be happy to have your feedback too.

RESPONSIBLE AI – THE MUST-HAVE FLAVOR IN YOUR AI DISH

Infosys Topaz is an AI-first set of services, solutions and platforms using generative AI technologies. It amplifies the potential of humans, enterprises, and communities to create value. With 12,000+ AI assets, 150+ pre-trained AI models, 10+ AI platforms steered by AI-first specialists and data strategists, and a 'responsible by design' approach, Infosys Topaz helps enterprises accelerate growth, unlock efficiencies at scale and build connected ecosystems. Connect with us at infosystopaz@infosys.com

For more information, contact askus@infosys.com

Infosys®
Navigate your next