European Commission

Technology Foresight Brief

EMBODYING THE FUTURE

Horizon scanning for emerging technologies and breakthrough innovations in the field of human-like AI systems

HIGHLIGHTS

→ This **Horizon Scanning** exercise was developed to support **European Innovation Council Strategic Intelligence** in the field of **human-like AI systems.**

→ **Twelve topics** were prioritised by workshop participants: enhancing human-AI collaboration; trustworthy and explainable AI; neurosymbolic AI; general-purpose neurosymbolic methods; personalised medicine; emotion-aware AI; embodied intelligence; multi-agent frameworks; human-AI collective cognition; brain-inspired AI; addressing AI's energy demand; and next generation LLMs.

→ **Several contextual** factors that shape the development and uptake of AI are highlighted across social, technological, economic, environmental and political/regulatory domains, including: AI literacy; inter-disciplinary and inclusive development of models; benchmarking practices; and sustainability.

→ **Horizon Scanning** is a qualitative foresight method, aiming at the early discovery of developments not yet on the radar of most experts.

INTRODUCTION

Project context

This brief reports on the conclusions of a Horizon Scanning exercise developed in the context of FUTURINNOV (FUTURe-oriented identification and assessment of emerging technologies and breakthrough INNOVation), a collaboration between the European Commission's (EC) Joint Research Centre (JRC) and the European Innovation Council (EIC), the EC's flagship program for deep tech, implemented by the European Innovation Council and SMEs Executive Agency (EISMEA).

Joint Research Centre

FUTURINNOV was designed to support the EIC in building strategic intelligence capacity through foresight and other anticipatory approaches. It supports activities focused on funding targets, programme design, policy feedback, and institutional governance.

The outcomes of this exercise may be used to inform future funding topics for EIC Challenges and other EC calls. They can also provide input for EIC and EC reports, as well as supporting other EU policy-making initiatives.

Methodology

Horizon Scanning is a qualitative foresight method which is aimed at the early discovery of developments not yet on the radar of most experts, decision makers, or the general public, and whose potential is not widely recognised. It is not a predictive tool, rather it encourages the exploration of novelties that offer opportunities and challenges in the medium or long-term. [1, 2, 3]

FUTURINNOV includes a series of thematic workshops that follow a tailor-made approach to this methodology. This approach uses a participatory detection, clustering, and sense-making process for signals, trends and contextual factors related to emerging technologies and breakthrough innovations. Each workshop is dedicated to a specific EIC Programme Manager's portfolio, or a domain deemed of interest by the EIC.

Trends and signals[1] are captured through a series of participatory exercises preceded by qualitative desk research, as well as data and text mining. They originate from a diversity of sources, ranging from scientific publications, patents and previously funded projects to institutional websites, news, online articles and other media.

During each workshop, through a specific methodology composed of several analytical and selection (i.e. voting) steps, participants converge on a priority list of topics. The criteria for this selection include relevance to the exercise's scope, potential impact and overall novelty across all technology readiness levels. The final topics include technologies and innovations, as well as relevant contextual factors for their development and uptake.

This brief refers to a specific workshop held online on 19 June 2025 with a focus on human-like AI systems. It was held with a group of selected experts from academia, research and technology, business and policy-making organisations. This diversity of institutional backgrounds as well as various fields of specific expertise was key to bringing different perspectives to the conversation. The resulting collective intelligence helped to build significant insights around the topic of AI.[2]

Scope and policy context

The recently adopted AI Act seeks to strengthen the EU's digital internal market by establishing a consistent legal framework for AI system development, marketing, and use. It promotes human-centric and trustworthy AI to safeguard safety, and fundamental rights and aims to support innovation and ensure the free cross-border movement of AI tools and services [4].

AI is recognised as one of the ten critical technology areas for the EU's economic security, highlighting its importance for the region's strategic autonomy and competitiveness [5]. The EU's Strategic Technologies for Europe Platform (STEP) that aims to support the development and deployment of critical technologies, also highlights its importance to reduce strategic dependencies and enhance innovation [6].

The recent Competitiveness Compass set up a roadmap for upcoming proposals in this domain, such as the AI Gigafactories and the Apply AI initiatives to drive development and industrial adoption of AI in key sectors [7]. The Digital Europe Programme (DEP) aims to enhance digital transformation across Europe by funding projects in key areas which includes artificial intelligence [8].

The EU is also investing significantly in AI research and development through Horizon Europe, which focuses on enabling safe and ethical AI applications across various sectors. In this context, AI was one of the recently established portfolios within the EIC.

For this exercise, the field of human-like AI systems was selected as the focus, with the aim to cover

[1] The understanding of what constitutes a signal or a trend may vary [48, 49]. As it is not yet consensual, for the purposes of this project both are relevant and understood as tangible manifestations of novelty in science, technology, innovation, markets, media, and other fields. They can cover different maturity levels from basic research to commercial readiness. Although often used interchangeably, a signal is less consolidated than a trend.

[2] For this exercise, a longlist of 246 trends and signals was compiled and assessed by the JRC and the EIC for their relevance to the workshop's scope and objectives. The list was subsequently narrowed down to a shortlist of 206 and shared in advance with participants. During the workshop, attendees discussed the contents of the shortlist, added new topics considered relevant/ previously overlooked, grouped and connected related issues, and ultimately converged on a final list of topics, summarised in this document.

developments across several key areas. These include: biologically- and cognitively-inspired approaches; adaptive and decentralised learning models; perceptual and multimodal intelligence; reasoning, and decision-making capabilities; the integration of symbolic and neural methods and collective intelligence. Relevant applications, particularly in robotics, were also targeted.

Quick guide

This brief is organised in 3 sections:

- **trends and signals** on technologies and innovations;
- **contextual factors**, covering drivers, enablers and barriers related with their development and uptake;
- **conclusions** providing complementary insights and overarching analysis.

TRENDS AND SIGNALS

The following trends and signals were identified by participants as the most relevant within the workshop scope. Some have been grouped under broader topics, reflecting different levels of granularity and interconnectedness. They are presented by decreasing number of votes, with minor editorial adjustments to improve clarity and coherence. Boxes 1 to 3 contain the selected wildcards[3].

Enhancing human-AI collaboration

This topic drew strong consensus and reflected several trends and signals discussed in more detail below. It highlights several dimensions of human-AI collaboration and how the development and uptake of this synergy should be grounded in trust, co-agency, and transparency.

This is already a pressing issue, not a future one, both for professional and personal use. In the workplace, large language models (LLMs) are shifting from passive tools to active digital colleagues, evident in the appearance of autonomous agents in digital office platforms. Hybrid reasoning LLMs, that toggle between speed and self-reflection, can already offer nuanced and accurate outputs that address complex tasks. Additionally, there is also a growing influence of personalised persuasion within some models like GPT-4 and 5, that already leverage user personal data to tailor arguments, which can potentially reshape AI-human and interpersonal dynamics.

Hybrid systems in health, for instance, where AI works alongside doctors, demonstrate this human–AI collaboration by combining the complementary strengths of each—AI's data-driven analysis and humans' judgement—resulting in more accurate diagnoses, improved patient outcomes, and greater safety than either could achieve alone. AI in Health is covered in more detail further on in this report.

While these trends promise greater productivity, they also raise concerns about oversight, accountability, and even the future nature of work. Therefore, trustworthy and explainable AI (XAI), explored in the next topic, can foster user confidence by providing auditable decisions and outputs, while mimicking human reasoning.

Such concepts converge into the idea of "Symbiotic AI" - a system that amplifies rather than replaces human capacities, integrating AI's computational strengths with human decision-making. It is underpinned by iterative collaboration and interactive learning to foster trust, while also challenging current notions of human autonomy and AI agency.

[9, 10, 11, 12, 13, 14]

Trustworthy and explainable AI

Trustworthy AI and XAI are foundational pillars for broad adoption of AI systems. Trustworthy AI embeds reliability, safety and accountability through combining deep learning with explicit reasoning. Coupled with this, federated learning can ensure robustness and data privacy across distributed environments. These features meet regulatory compliance demands but also and most importantly they drive user confidence, especially in domains like healthcare.

XAI complements these concepts by focusing on demystifying "black-box" models via techniques such as attribution graphs, highlighting which parts of an image/text mattered, counterfactual approaches, and benchmarking to allow users to trace and understand AI decisions.

Together, these two concepts support transparent decision-making, ethical deployment and collaborative human-AI workflows. By embedding both reliability and interpretability at the architecture level, AI systems can deliver performant solutions that stakeholders can trust, inspect and govern effectively, transforming opaque algorithms into

[3] Participants were asked to classify signals as wild cards. As three signals received the same votes, the authors include all of those in this report. In foresight a wild card is an event or development that bears a higher potentially disruptive impact than a normal signal.

potential accountable digital collaborators as stressed in the previous point.

[15, 16, 17, 18, 19, 20, 21]

Neurosymbolic AI

Neurosymbolic AI is a hybrid approach that aims to address the above-mentioned challenges of trust, safety, interpretability and accountability, paving the way for more transparent and explainable systems.

It merges neural networks' pattern recognition with symbolic systems' rules-based logical reasoning. Such systems can manage complex and data-rich tasks with improved transparency. Novel frameworks like Logic Tensor Networks exemplify how differentiable logic can combine learning with structured reasoning.

Neurosymbolic AI can be integrated with agentic AI to support significant advances in commonsense reasoning and causal explanation. This combination empowers systems to operate more autonomously, make context-aware decisions and navigate complex environments. It also enables adaptation through high-level cognitive processes, allowing agents to respond flexibly to changing circumstances.

Potential applications span from natural language processing to healthcare, cybersecurity, and even optimising communications in emerging 6G networks.

[22, 23, 24, 25]

Box 1: General-purpose neurosymbolic methods

General-purpose neurosymbolic methods represent a shift from domain-anchored symbolic reasoning to more flexible, scalable and high performing models developed for general purposes. Neurosymbolic approaches combine neural networks (the foundation of deep learning) and symbolic AI (which uses explicit rules, logic, and knowledge representation).

Conventional neurosymbolic AI models have been typically anchored in domain-specific knowledge, assumptions, and explicit object attributes. This has limited their flexibility and kept them tied to narrow applications.

These emerging general-purpose approaches address this constraint by dynamically generating new concepts without significant initial domain encoding. They incorporate large language models and knowledge graphs to support dynamic reasoning and concept generalisation. This shift enhances their capacity to generalise across tasks while keeping some advantages: they are typically more resource-efficient, inherently transparent, and capable of compositional reasoning.

[26]

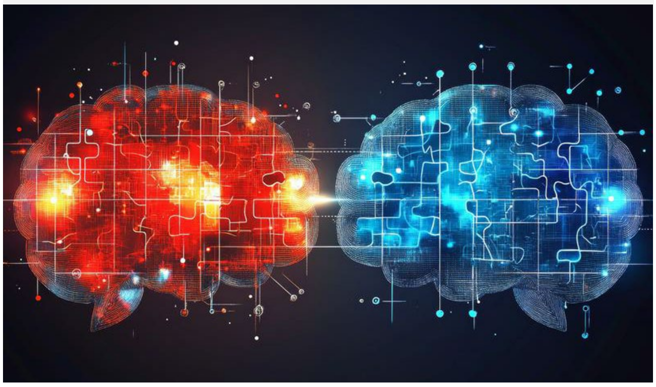

AI in health – personalised medicine and beyond

Several application domains were discussed during the workshop, namely defence, education, science and research, but participants placed particular emphasis on healthcare.

Experts emphasised AI's tangible and potential contribution to overall quality of life, highlighting AI-driven personalised medicine as one of the most promising developments, particularly for enhancing collaboration between health professionals and AI systems. For example, personalised treatment plans can be developed by analysing complex datasets that include a patient's genetic profile and lifestyle factors.

Advanced algorithms can process genetic sequences, electronic health records, and data from wearable devices to identify patterns and correlations that may not be evident to clinicians. This allows for the detection of genetic markers associated with disease risk or drug response, enabling more accurate predictions about which treatments will be most effective, and which may cause adverse reactions.

In cancer care, for instance, AI tools can analyse tumour DNA to match patients with targeted therapies, reducing reliance on trial-and-error approaches. Consequently, treatment plans are optimised for efficacy and safety, minimising side effects and maximising positive health outcomes.

By integrating and interpreting large volumes of data, AI empowers clinicians to deliver truly personalised medicine, moving beyond one-size-fits-all approaches.

[27, 28, 29]

Box 2: Emotion-aware AI

A new computational model developed in Finland enables AI to simulate and predict human emotions during interactions, by combining reinforcement learning with cognitive appraisal theory.

Unlike traditional affective computing, which depends on facial expressions or biometric signals, the new approach infers emotions through value-based decision-making and appraisals considering personal event significance, power, and goal alignment. This dynamic, continuous modelling of emotional transitions represents a significant advance in emotionally intelligent AI.

The technology holds promise for digital health, adaptive learning, and emotionally responsive virtual assistants by enhancing user engagement and enabling systems to respond empathetically. It also offers non-intrusive mental health monitoring capabilities.

However, further work is needed to ensure generalisation across diverse users and emotional states. Key challenges include addressing ethical concerns related to emotional manipulation and ensuring a safe scalable implementation in commercial platforms.

[34, 35]

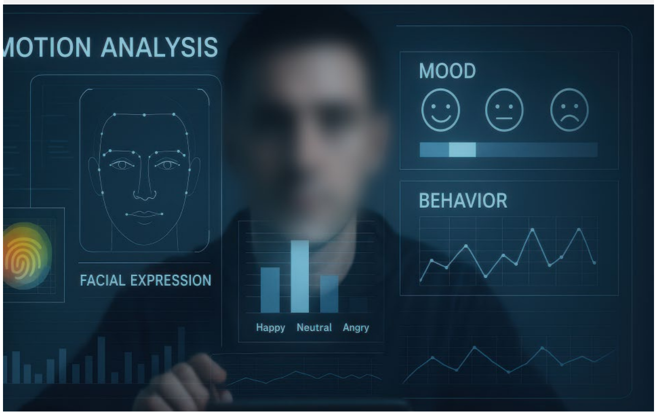

Embodied intelligence

Emerging trends in embodied intelligence highlight the growing integration of AI and robotics to seamlessly merge digital reasoning with the physical world. Unlike traditional AI, which primarily focuses on data analysis and decision-making in virtual spaces, embodied AI involves the physical execution of tasks, demanding a combination of sensory input, real-time processing and physical actions.

Recent advances include Google's Gemini Robotics and frameworks like ELLMER, which combine LLMs with multimodal feedback from vision and force sensors, thereby enabling robots to interpret high-level instructions and perform complex tasks. Meanwhile, Agibot's data centre in Shanghai simulates diverse real-world environments to produce extensive, task-specific datasets, addressing the lack of robust training data.

These developments can drive applications from smart factories and smart farming to home and healthcare assistances by reducing programming demands through natural language interfaces and therefore advancing human-robot collaboration.

Challenges persist in bridging the gap between simulation and real-world deployment, ensuring safety, generalisation, and trust—especially in sensitive contexts. Nonetheless, these signals mark a pivotal shift towards more adaptable, context-aware robotic systems.

[30, 31, 32, 33]

Multi-agent frameworks

The rise of agentic AI is a significant driver of development in the field. Within agentic AI, there is also a transition from isolated agents to collaborative, scalable Multi-Agent Systems (MASs) that perceive, reason, and act jointly, enhancing their ability to solve complex and dynamic tasks.

A core innovation lies in defining collaboration through key dimensions: the nature of agents, interaction types (e.g., cooperation or competition), structural organisation (centralised, distributed), coordination strategies, and communication protocols. By adopting this structured framework, designers can build MASs capable of adapting and collaborating effectively.

MASs can be increasingly applied in areas such as 5G/6G optimisation, industry 5.0 processes, social simulations, question answering, and digital twins. These applications unlock opportunities for improved problem-solving, adaptability, and interdisciplinary collaboration. However, challenges include developing robust, scalable coordination infrastructures, ensuring ethical and value alignment, and mitigating hallucinations.

[36]

Brain-inspired AI

Brain-inspired AI leverages principles from neurobiology to design AI systems that are more efficient, adaptive, and robust, using techniques like spiking neural networks and neuromorphic computing, which mimic the brain's event-driven, parallel processing.

Recent breakthroughs include coupling artificial and biological neural networks through semantic annotation using brain imaging, as well as developing energy-efficient organic-transistor neuromorphic chips. These innovations enable more adaptable decentralised learning architectures and offer an alternative to the resource-intensive and rigid structures of conventional AI.

The technology shows promise in sensory robotics, cognitive sciences (e.g., emotion or deception analysis) and even towards Artificial General Intelligence (AGI). Key opportunities include drastically improved energy efficiency, enriched human-AI interaction, and new knowledge at the intersection of AI and brain science.

However, technical barriers remain, notably in integrating and scaling hybrid systems. Ethical concerns around anthropomorphism and the epistemic validity of brain analogies also persist. Continued interdisciplinary collaboration will be essential to realise its full potential.

[39, 40, 41, 42, 43]

Addressing AI's growing energy demand

Bio-inspired AI forms part of a surge in innovative system and semiconductor architectures designed to address AI's rising energy demands. A notable advance is Spike-based Deep Neural Networks (SNNs), which emulate the spiking behaviour of biological neurons and deliver exceptional energy efficiency for event-driven applications, including Internet of Things (IoT) deployments and object-recognition tasks. These systems, coupled with neuromorphic hardware, significantly reduce power consumption while ensuring scalability.

In a further significant breakthrough, researchers unveiled the first field-programmable photonic chip capable of training nonlinear neural networks entirely with light. By executing computations optically—without relying on electronic transistors—this design avoids conventional, energy-intensive circuitry. Its ultra-fast, reconfigurable architecture promises highly efficient AI acceleration in data centres, edge devices and specialised optical accelerators.

With the rise of smaller models deployed on edge devices, these developments signal a strong technological drive to reconcile AI's growth with sustainable, low-energy computing.

[39, 40, 42, 44, 45, 46]

Box 3: Human-AI collective cognition

Conversational Swarm Intelligence (CSI) is an innovative generative AI model that enables scalable, real-time group deliberation by simulating natural swarm behaviour. Developed by Carnegie Mellon University and Unanimous AI, CSI uses large language models to dynamically orchestrate small subgroup discussions within a larger human collective, boosting ideation quality and participant engagement.

In experiments done on the Thinkscape platform, it outperformed traditional chat tools in both creativity and user satisfaction. The novelty lies in its ability to synchronise hundreds of participants through AI mediation, fostering consensus and enhancing productivity. CSI is particularly promising for innovation management, policy co-creation, and large-scale participatory governance.

While offering significant potential to reshape collaborative processes, it also raises questions around transparency, AI influence in consensus-building, and integration with existing systems. Further empirical evidence is needed to assess its performance in diverse and complex settings.

[37, 38]

Next generation LLMs: reasoning and multi-modal capabilities

A broader trend is already observable in the evolution of large language models (LLMs) towards enhanced reasoning and multimodal capabilities. This new generation of AI systems, of which GPT-5 from OpenAI, Meta's LLaMA 4, Google's Gemini 2.5 and Mistral's PixTraL are part, is designed to process and integrate diverse inputs such as text, images, and potentially audio, within a single model, marking a pivotal shift from general-purpose chatbots towards versatile, embedded cognitive agents.

By supporting multimodal understanding, these models deliver more context-aware and flexible interactions, extending beyond text generation to complex, task-oriented problem-solving. The emphasis on improved reasoning allows LLMs to handle increasingly abstract or domain-specific challenges, with implications for fields ranging from healthcare and education to creative industries and public policy.

As models gain sophistication, they also prompt critical discussions around governance, transparency and societal impact—particularly with respect to bias mitigation, computational resource demands and responsible deployment.

[47]

CONTEXTUAL FACTORS

The following topics result from an aggregation of participants' insights regarding contextual factors influencing the development and uptake human-like AI systems.[4]

Societal

Societal factors shaping human-like AI revolve around ethics, trust and public engagement. Biased or limited algorithms risk perpetuating stereotypes and missing cultural differences.

There is also a cultural dimension: how societies conceive the role and legitimacy of non-human forms of intelligence is shaped by values, beliefs and traditions, sometimes with almost religious undertones. This cultural framing influences how human–AI partnerships are accepted and integrated.

At the same time, overreliance on AI tools can erode critical thinking and human skills, underscoring the need for AI literacy and interdisciplinary education.

AI challenges traditional concepts of intellectual property by blurring the line between human and machine authorship. It disrupts established notions of creativity, knowledge production, and epistemic trust and could easily lead to exclusion of cultures, viewpoints and values.

Citizen participation in model audits and fairness assessments could be a way to foster greater transparency and bolstering social acceptance by embedding diverse perspectives into system design.

Technological

On the technological front, ever-larger neural networks are enabled by expanding computational power and data availability, but unrepresentative datasets, as mentioned before, can hinder model performance.

Progress in assessing models' performance is slowed by fragmented benchmarks and the absence of unified evaluation metrics, limiting comparisons and replicability. Moreover, the time required for peer-reviews cause scientific knowledge to lag behind rapid (prototype) model releases; as platforms like GitHub and arXiv allow for immediate sharing of models and ideas. Addressing these gaps demands genuinely interdisciplinary collaboration—integrating insights from social scientists, ethicists and domain experts.

Economic

Certain economic areas are clearly driving AI development in particular healthcare. AI-powered diagnostics and mental-health applications are one tangible example.

The prohibitively high cost of model development and large-scale inference results in a high market concentration among a few dominant firms, whose race for profitability increasingly shapes model development priorities, access to compute, and the openness of research. It also reduces diversity and can amplify biases. Businesses in general, invest in AI as a way to automate tasks, boost efficiency and eventually reduce jobs and costs; however, the impacts on labour markets are more complex.

Although these and other productivity gains can spur growth, they also threaten job displacement across

[4] These factors were analysed using an adapted version of the "Triangle of the Future" framework [50], a foresight method that maps three competing forces: the pull of the future, the push of the present, and the weight of history. It can be used as a stand-alone method or in conjunction with others. For this project, the authors explored 3 types of contextual factors connected with those three temporal dimensions: drivers which are high-level factors that trigger or shape significant contextual changes and pull technological development and uptake into the future; enablers, or opportunities, that are present-day conditions that create a fertile ground for innovation to occur and therefore push technologies forward; and barriers, or challenges, that can be seen as past and present constraints ("weight") that hinder technological development and uptake.

sectors. Balancing investment in AI-augmented roles with reskilling initiatives will be critical to mitigating labour-market disruption.

Environmental

Environmental considerations are gaining prominence as AI's energy footprint escalates. This trend has spurred efforts to optimise models and push for both software and semiconductor innovation. Rising electricity demand has also renewed interest in nuclear power as a low-carbon solution to support intensive AI workloads.

Political and regulatory

Political and regulatory landscapes are fraught with strategic competition and content-integrity challenges. Nations envision AI leadership to secure economic and even military advantages, intensifying geopolitical rivalries.

At the same time, social media algorithms, deepfakes and AI-generated personas already enable sophisticated disinformation campaigns, that can reshape the political sphere, and citizens' perception of reality, while detection tools struggle to keep pace.

However, while comprehensive regulatory frameworks—such as the EU's AI Act and data-protection laws—already set standards for transparency and accountability, enforcing these rules consistently across jurisdictions remains a major hurdle. EU regulators must also continually adapt their policies to keep pace with rapid AI innovations without hampering competitiveness or driving talent and investment toward more permissive regions.

CONCLUSIONS

Our expectations of AI extend well beyond raw computational power: we ask systems to reason, explain their decisions, act ethically, perceive and respond to emotions, integrate multisensory inputs, and collaborate seamlessly with both people and other machines.

This workshop revealed consensus around the vision of AI not as a substitute for human agency, but as a collaborative partner that can amplify human strengths under carefully defined guiding principles.

Human-AI partnership could extend into the physical realm, where robotics and other autonomous systems couple AI with sensors and actuators, and work alongside people in an increasing diversity of environments from homes to factories, and public spaces.

Across domains—from finance to education—applications are proliferating, but none seems to resonate more deeply than in health, underscoring how tightly we can see AI intertwining with our well-being.

As we move towards this vision, embedding transparency, accountability, fairness, and human-centric values – including the protection of our environment - into the very architecture of AI systems seems to be crucial. Careful navigation is needed around persistent challenges, including value alignment, limited transparency in AI decision-making, and the risk of over-reliance on automated systems.

By prioritising collaborative design and rigorous ethical guardrails, we can build AI systems that are not only *human-like* but also aligned with and enhancing what *humans like*: our collective capacity to create, innovate, connect and support one another.

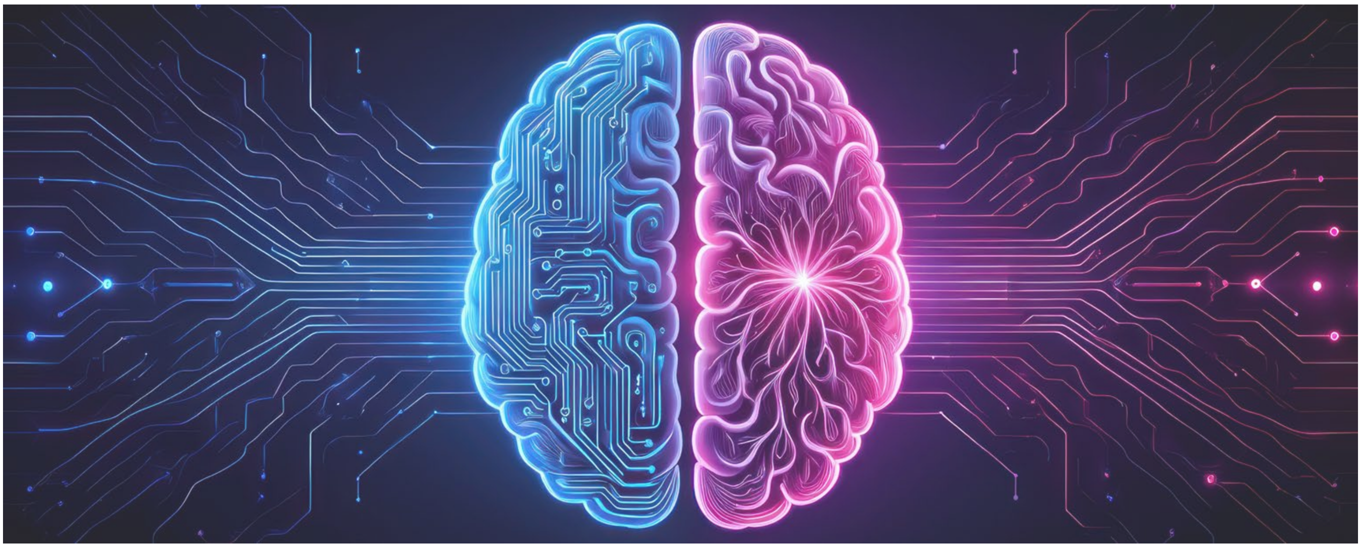

REFERENCES

[1] E. Amanatidou, M. Butter, V. Carabias, T. Könnölä, M. Leis, O. Saritas, P. Schaper-Rinkel and V. van Rij, “On concepts and methods in horizon scanning: Lessons from initiating policy dialogues on emerging issues,” *Science and Public Policy,* vol. 39, no. 2, pp. 208-221, 2012.

[2] J. Farinha, L. Vesnic Alujevic and A. Polvora, “Scanning deep tech horizons: participatory collection and assessment of signals and trends,” Publications Office of the European Union, Luxembourg, 2023.

[3] P. Dannemand Andersen, M. Bevolo, I. Ilevbare, E. Malliaraki, R. Popper and M. Spaniol, Technology Foresight for Public Funding of Innovation: Methods and Best Practices, L. Vesnic Alujevic, J. Farinha and A. Polvora, Eds., Luxembourg: Publications Office of the European Union, 2023.

[4] European Union, “Regulation (EU) 2024/1689 of the European Parliament and the council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations,” European Union, 2024. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401689.

[5] European Commission, “Commission Recommendation (EU) 2023/2113 of 3 October 2023 on critical technology areas for the EU’s economic security for further risk assessment with Member States,” European Commission, 2023. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202302113.

[6] European Commission, “Strategic Technologies for Europe Platform,” European Commission, 2025. [Online]. Available: https://strategic-technologies.europa.eu/index_en.

[7] European Commission, “An EU Compass to regain competitiveness and secure sustainable prosperity,” European Commission, 2025. [Online]. Available: https://ec.europa.eu/commission/presscorner/detail/en/ip_25_339.

[8] European Commission, “Digital Europe Programme,” European Commission, 2025. [Online]. Available: https://commission.europa.eu/funding-tenders/find-funding/eu-funding-programmes/digital-europe-programme_en.

[9] T. Warren, “Microsoft 365 Copilot redesigned with new search, image, and notebook features,” The Verge, 2025. [Online]. Available: https://www.theverge.com/news/654113/microsoft-365-copilot-redesign-search-image-notebook-features.

[10] S. S. Y. Kim, E. A. Watkins, O. Russakovsky, R. Fong and A. Monroy-Hernández, “"Help Me Help the AI": Understanding How Explainability Can Support Human-AI Interaction,” in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023.

[11] R. Wang, R. Cheng, D. Ford and T. Zimmermann, “Investigating and Designing for Trust in AI-powered Code Generation Tools,” in *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 2024.

[12] K. Wiggers, “Deep Cogito emerges from stealth with hybrid AI ‘reasoning’ models,” Tech Crunch, 2025. [Online]. Available: https://techcrunch.com/2025/04/08/deep-cogito-emerges-from-stealth-with-hybrid-ai-reasoning-models/.

[13] R. Williams, “AI can do a better job of persuading people than we do,” MIT Technology Review, 2025. [Online]. Available: https://www.technologyreview.com/2025/05/19/1116779/ai-can-do-a-better-job-of-persuading-people-than-we-do/.

[14] B. D. Longueville, I. Sanchez, S. Kazakova, S. Luoni, F. Zaro, K. Daskalaki and M. Inchingolo, “The Proof Is in the Eating: Lessons Learnt from One Year of Generative AI Adoption in a Science-for-Policy Organisation,” *AI,* vol. 6, no. 128, 2025.

[15] R. Confalonieri, L. Coba, B. Wagner and T. R. Besold, “A historical perspective of explainable Artificial Intelligence,” *WIREs,* vol. 11, no. 1, 2021.

[16] S. Ali, T. Abuhmed, S. El-Sappagh, K. Muhammad, J. M. Alonso-Moral, R. Confalonieri, R. Guidotti, J. D. Ser, N. Díaz-Rodríguez and F. Herrera, “Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence,” *Information Fusion,* vol. 99, 2023.

[17] Y. Zhang, Q. V. Liao and R. K. E. Bellamy, "Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020.

[18] R. C. Fong and A. Vedaldi, "Interpretable Explanations of Black Boxes by Meaningful Perturbation," in *IEEE International Conference on Computer Vision*, 2017.

[19] M. Sand, J. M. Durán and K. R. Jongsma, "Responsibility beyond design: Physicians' requirements for ethical medical AI," *Bioethics,* vol. 36, no. 2, pp. 162-169, 2021.

[20] T. Räuker, A. Ho, S. Casper and D. Hadfield-Menell, "Toward Transparent AI: A Survey on Interpreting the Inner Structures of Deep Neural Networks," in *2023 IEEE Conference on Secure and Trustworthy Machine Learning*, Raleigh, NC, USA, 2023.

[21] A. Holzinger, "The Next Frontier: AI We Can Really Trust," in *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2021.

[22] A. d. Garcez and L. C. Lamb, "Neurosymbolic AI: the 3rd wave," *Artifcial Intelligence Review,* vol. 56, p. 12387–12406, 2023.

[23] A. Sheth, K. Roy and M. Gaur, "Neurosymbolic Artificial Intelligence (Why, What, and How)," *IEEE Intelligent Systems,* vol. 38, no. 3, pp. 56-62, 2023.

[24] M. Gaur and A. Sheth, "Building trustworthy NeuroSymbolic AI Systems: : Consistency, reliability, explainability, and safety," *AI Magazine,* vol. 45, no. 1, pp. 139-155, 2024.

[25] S. Badreddine, A. d. Garcez, L. Serafini and M. Spranger, "Logic Tensor Networks," *Artificial Intelligence,* vol. 303, no. C, 2022.

[26] J. Mao, J. B. Tenenbaum and J. Wu, "Neuro-Symbolic Concepts," *ArXiv,* 2025.

[27] Simbo, "Transforming Personalized Medicine: How AI Tailors Treatment Plans Through Genetic and Lifestyle Data Analysis," Simbo, 2025. [Online]. Available: https://www.simbo.ai/blog/transforming-personalized-medicine-how-ai-tailors-treatment-plans-through-genetic-and-lifestyle-data-analysis-3505802/.

[28] FTC Publications, "How AI-Driven Personalized Medicine Transforms Healthcare Decisions," FTC Publications, 2025. [Online]. Available: https://news.ftcpublications.com/core/how-ai-driven-personalized-medicine-transforms-healthcare-decisions/.

[29] Laboratorios Rubio, "The Role of Artificial Intelligence in Personalized Medicine," Laboratorios Rubio, 2025. [Online]. Available: https://www.laboratoriosrubio.com/en/ai-personalized-medicine/.

[30] W. Knight, "Google's Gemini Robotics AI Model Reaches Into the Physical World," Wired, 2015. [Online]. Available: https://www.wired.com/story/googles-gemini-robotics-ai-model-that-reaches-into-the-physical-world/.

[31] R. Mon-Williams, G. Li, R. Long, W. Du and C. G. Lucas, "Embodied large language models enable robots to complete complex tasks in unpredictable environments," *Nature Machine Intelligence,* vol. 7, p. 592–601, 2025.

[32] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans and B. Jain, "Habitat: A Platform for Embodied AI Research," in *International Conference on Computer Vision*, 2019.

[33] KrAsia, "Inside Agibot's Shanghai center, robots learn to master tasks in human-like ways," KrAsia, 2025. [Online]. Available: https://kr-asia.com/inside-agibots-shanghai-center-robots-learn-to-master-tasks-in-human-like-ways.

[34] J. E. Zhang, B. Hilpert, J. Broekens and J. P. P. Jokinen, "Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, Honolulu, HI, USA, 2024.

[35] University of Jyväskylä, "Researchers developed a model that allows a computer to understand human emotions," University of Jyväskylä, 2024. [Online]. Available: https://www.jyu.fi/en/news/researchers-developed-a-model-that-allows-a-computer-to-understand-human-emotions.

[36] K.-T. Tran, D. Dao, M.-D. Nguyen, Q.-V. Pham, B. O'Sullivan and H. D. Nguyen, "Multi-Agent Collaboration Mechanisms: A Survey of LLMs," *arXiv,* 2025.

[37] L. Rosenberg, H. Schumann, C. Dishop, G. Willcox, A. Woolley and G. Mani, "Large-Scale Group Brainstorming and Deliberation Using Swarm Intelligence and Generative AI," in *Proceedings of the 27th International Conference on Enterprise Information Systems*, 2025.

[38] Tepper School of Business - Carnegie Mellon University, "Using principles of swarm intelligence, study compares platforms that allow brainstorming among large groups," PHYS.ORG, 2025. [Online]. Available: https://phys.org/news/2025-05-principles-swarm-intelligence-platforms-brainstorming.html.

[39] I. Krauhausen, C.-T. Coen, S. Spolaor, P. Gkoupidenis and Y. v. d. Burgt, "Brain-Inspired Organic Electronics: Merging Neuromorphic Computing and Bioelectronics Using Conductive Polymers," *Advanced Functional Materials,* 2024.

[40] T. Taniguchi, H. Yamakawa, T. Nagai, K. Doya, M. Sakagami, M. Suzuki, T. Nakamura and A. Taniguchi, "A whole brain probabilistic generative model: Toward realizing cognitive architectures for developmental robots," *Neural Networks,* 2022.

[41] P. Segerdahl, "Anthropomorphism in AI can limit scientific and technological development," Uppsala Universitet, 2020. [Online]. Available: https://ethicsblog.crb.uu.se/2020/04/15/anthropomorphism-in-ai-can-limit-scientific-and-technological-development/?utm_source=chatgpt.com.

[42] L. Zhao, L. Zhang, Z. Wu, Y. Chen, H. Dai, X. Yu, Z. Liu, T. Zhang, X. Hu, X. Jiang, X. Li, D. Zhu, D. Shen and T. Liu, "When brain-inspired AI meets AGI," *Meta-Radiology,* vol. 1, no. 1, 2023.

[43] Y. Zeng, D. Zhao, F. Zhao, G. Shen, Y. Dong, E. Lu, Q. Zhang, Y. Sun, Q. Liang, Y. Zhao, Z. Zhao, H. Fang, Y. Wang, Y. Li, X. Liu, C. Du, Q. Kong, Z. Ruan and W. Bi, "BrainCog: A spiking neural network based, brain-inspired cognitive intelligence engine for brain-inspired AI and brain simulation," *Patterns,* 2023.

[44] T. Wu, Y. Li, L. Ge and L. Feng, "Field-programmable photonic nonlinearity," *Nature Photonics,* 2025.

[45] I. Scheffler, "Programmable photonic chip uses light to accelerate AI training and cut energy use," TechXplore, 2025. [Online]. Available: https://techxplore.com/news/2025-04-programmable-photonic-chip-ai-energy.html.

[46] S. Kundu, G. Datta, M. Pedram and P. A. Beerel, "Spike-Thrift: Towards Energy-Efficient Deep Spiking Neural Networks by Limiting Spiking Activity via Attention-Guided Compression," in *2021 IEEE Winter Conference on Applications of Computer Vision*, 2021.

[47] B. Edwards, "Sam Altman lays out roadmap for OpenAI's long-awaited GPT-5 model," ARS Technica, 2025. [Online]. Available: https://arstechnica.com/ai/2025/02/sam-altman-lays-out-roadmap-for-openais-long-awaited-gpt-5-model/.

[48] P. Rossel, "Early detection, warnings, weak signals and seeds of change: A turbulent domain of futures studies," *Futures,* vol. 44, no. 3, pp. 229-239, 2012.

[49] B. L. van Veen and J. Roland Ortt, "Unifying weak signals definitions to improve construct understanding," *Futures,* vol. 134, p. 102837, 2021.

[50] S. Inayatullah, "The Futures Triangle: Origins and Iterations," *World Futures Review,* vol. 15, no. 2-4, pp. 112-121, 2023.

ACKNOWLEDGEMENTS

The authors are thankful to and acknowledge the contributions for the signal collection of: Michela Bergamini, Marcelina Grabowska and Olivier Eulaerts (EC - DG JRC, Text and Data Mining Unit) and Emily Djock, Sarah Grellier and Fabian Reck (Itonics).

The authors are also grateful for the time and contributions of the following workshop participants: Abdallah Makhoul (Université Marie et Louis Pasteur), Adam Dahlgren Lindström (Umeå University), Andrea Blasco (EC - DG JRC), Bertrand De Longueville (EC - DG JRC), Davide Bacciu (University of Pisa), Deborah Nas (Delft University of Technology), Fabian Suchanek (Institut Polytechnique de Paris), Ines Lynce (Universidade de Lisboa & INESC-ID), John Breslin (University of Galway), Julien Gossé (EC - DG JRC), Luis Fernando D'Haro (Speech Technology and Machine Learning Group - Universidad Politécnica de Madrid), Mahdi Zargayouna (Université Gustave Eiffel), Marta Sabou (Vienna University of Economics and Business), Mehdi Dastani (Utrecht University), Miguel Rubio (EC – DG CNECT), Reyhan Aydoğan (Ozyegin University), Roumen Borissov (EISMEA) and Sadegh Soudjani (Max Planck Institute for Software Systems).

DISCLAIMERS

The information and views presented here are those of the authors and do not necessarily reflect the official opinion of the European Commission.

The findings of this exercise resulted from a participatory process involving a group of internal and external experts representing a diversity of fields and backgrounds. The methodologies applied have limitations and the results do not aim to cover all developments and topics on the field.

Most trends and signals are referenced to the sources where they were originally detected, although some concepts included in the final texts result from the analysis and contributions of the participants.

During the preparation of this work the authors used several AI-assisted tools namely the internal GPT@JRC to assist desk research and summarise texts. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication.

HOW TO CITE THIS BRIEF

How to cite this report: Farinha, J., Nagy, O., Bailey, G., Mochan, A., Polvora, A. et al., *Embodying the Future - Horizon scanning for emerging technologies and breakthrough innovations in the field of human-like AI systems,* Publications Office of the European Union, Luxembourg, 2025, https://data.europa.eu/doi/10.2760/5714049, JRC143535.

CONTACT INFORMATION

JRC-FORESIGHT@ec.europa.eu

EU Science Hub — joint-research-centre.ec.europa.eu

JRC143535

PDF ISBN 978-92-68-31573-6 doi:10.2760/5714049 KJ-01-25-470-EN-N

Publications Office of the European Union