



The AI Data Robotics
Association

Strategic Research, Innovation, and Deployment Agenda

2025-2027

January 2026

Editors

- » Fredrik Heintz (LiU)
- » Nabil Belbachir (NORCE),
- » Edward Curry (UoG).

Coordination team

- » Ana Garcia (BDVA),
- » Katerina Linden (LiU),
- » Philip Piatkiewicz (Adra Secretary General),
- » Joost Geurts (interim Adra Secretary General).

Reference Group

- » Jon Agirre Ibarbia (Tecnalia),
- » David Bisset (euRobotics),
- » Freek Bomhof (BDVA),
- » Rita Cucchiara (UNIMORE),
- » Hans De Canck (FARI),

- » Anne Hermsen (TNO),
- » Morten Irgens (CBS),
- » Christophe Leroux (CEA),
- » Petri Myllymäki (HIIT),
- » José Saenz (Fraunhofer).

Contributors

- » Daniel Alonso (BDVA),
- » Freek Bomhof (TNO),
- » Rafiqul Haque (Insight Centre/Galway),
- » Agirre Ibarbia (Tecnalia),
- » Samuel Kaski (Aalto),
- » Yiannis Kompatsiaris (CERTH-ITI),
- » Marko Grobelnik (IJS),
- » Stefan Leijnen (AiNed),
- » Christophe Leroux (CEA),
- » Paul Lukowicz (DFKI),

- » Jessica Montgomery (AI@Cam),
- » Chokri Mraidha (CEA),
- » Alin Olimpiu (DLR),
- » Marc Schoenauer (Inria),
- » Morayo Sdedjouma (CEA),
- » Rebecca Schedl-Warpup (DLR),
- » Philipp Slusallek (DFKI),
- » Emmanuel Vincent (Inria),
- » Juha Röning (University of Oulu).

Acknowledgement:

This document received valuable contributions from the Adra members, the European Commission, the founding organisations BDVA, CLAIRE, ELLIS, EurAI, and euRobotics, as well as the larger ADR community.

Supported by the Adra-e project (101070336)

Cite as:

Fredrik Heintz, Nabil Belbachir, and Edward Curry, "Strategic Research, Innovation, and Deployment Agenda 2025–2027", February 2024, Brussels – Adra

Keywords Artificial Intelligence, Data, Robotics, Research Agenda, Digital Transformation, ADR ecosystem

Glossary of Terms

AAAI	Association for the Advancement of Artificial Intelligence
ADR	AI, Data, and Robotics
AI	Artificial Intelligence
AutoML	Automated Machine Learning
CVPR	Computer Vision and Pattern Recognition
EU	European Union
HCI	Human-Computer Interaction
HPC	High Performance Computing
HRC	Human-Robot Collaboration
HRI	Human-Robot Interaction
ICCV	International Conference on Computer Vision
ICML	International Conference on Machine Learning

ICRA	International Conference on Robotics and Automation
IFR	International Federation of Robotics
IJCAI	International Joint Conferences on Artificial Intelligence
IoT	Internet of Things
ML	Machine Learning
NeurIPS	Neural Information Processing Systems
PAMI	IEEE Transactions on Pattern Analysis and Machine Intelligence
R&D	Research and Development
SDG	Sustainable Development Goals
SME	Small and Medium Organisations
WP	Work Programme

Foreword

Ongoing world crises and global disruption have revealed Europe's vulnerability, causing severe challenges to its future sustainable development and welfare. Geopolitical conflicts, pandemics and climate change are adding to the pressure faced by industries within critical supply chains amidst a rapidly changing technological reality. Resilience and technological sovereignty are renewed priorities. In these contexts, it is crucial to reinforce Europe's position in digital technologies, Artificial Intelligence, Data Science, and Robotics (ADR) to further support societal welfare and accelerate the digital and green transitions.

ADR systems are rich in potential, but they also have risks that need to be mitigated through regulation and innovation. Adra was founded to become the private partner of the European Commission in the Public Private Partnership on ADR technologies and to build a European innovation ecosystem. This partnership brings together European ADR research, industry, and the public sector, promoting the convergence of ADR to tackle global challenges and helping to build a more sustainable society in Europe and globally. This partnership is also committed to promoting the safe use of ADR technologies.

The speed of innovation is relentless, and ADR has already had a strong impact on our daily lives. We need to ensure that Europe's companies and governments can also benefit from the technology revolution, ensuring strategic autonomy over and within these ADR technologies which can only be achieved through a strong and well-connected innovation ecosystem.

This review is the result of the excellent work of Adra's community and presents a long-term vision for developing and deploying trustworthy ADR technologies in Europe to tackle the big challenges of our time.

About this document

This document describes the strategic position of the AI, Data and Robotics Association (Adra), also called ADR Partnership, from the perspective of Adra members, resulting in several recommendations for the upcoming European work programmes. This position is motivated by global challenges, together with a European strategy and goals that can be addressed by AI, data, and robotics. This is the fourth Adra strategic review; the third was published in 2020. The first part describes ADR's Vision 2030, which looks at the global challenges and major strategic objectives that the ADR partnership seeks to address. The second part describes major trends and gaps. Finally, the third part describes the recommendations for the EU strategic plan for 2025-27, which are in alignment with the strategy.

Table of Contents

1. ADR Vision 2030	5
1.1 Alignment with European ADR efforts	7
Collaborative efforts	7
ADRA Inputs to Apply AI Strategy	8
1.2 ADR Missions	9
1.3 ADR Goals	11
2. Major Trends and Gaps	14
2.1 Global Challenges	19
2.2 New considerations and instruments from synergies to missions	20
3. The Next Strategic Plan 2025-27	23
3.1 Overarching recommendations for Work programme 2025-2027	24
3.2 Big Tickets in AI, Data and Robotics for 2025-2027	26
3.2.1 Big ticket #1: Ground-Breaking Technological Foundations in ADR	26
3.2.2 Big ticket #2: Effective and Trustworthy General Purpose ADR	34
3.2.3 Big ticket #3: An interoperable and integrated framework for data and model ecosystems	43
3.2.4 Big ticket #4: Next generation robots with embodied intelligence	55
3.2.5 Big ticket #5: ADR Technology for the sciences	72
3.2.6 Big ticket #6: Research, innovation, and tools for compliance	80
4 Conclusions	90

1. ADR Vision 2030

The ADR partnership seeks to enable a responsible AI-powered green digital transformation for an attractive, sustainable, prosperous, secure, and resilient multicultural society, based on European values with the highest living standards in the world. By 2030 Europe will have created a shared secure data infrastructure that balances the need for privacy with the need for effective and correct information that is interoperable with the rest of the world.

Autonomous robotic systems, in many shapes and forms, are improving effectiveness, safety, and energy efficiency across sectors, including agriculture, transportation and healthcare. Sophisticated and trustworthy AI-based systems provide effective and actionable decision-making support to individuals, groups, companies, and governments based on accurate and up-to-date information. The global share of technical solutions in AI, data, and robotics provided by European companies is steadily increasing, providing cost-effective solutions that respect human rights and values. Europe is attracting global talent in increasing numbers and realising its ambitions through the emergence of a new era of trustworthy AI-first companies, alongside improvements in quality of life and sustainable living across the continent. Most leading global companies have significant research, development, and production capabilities in Europe. Globally, Europe is seen as a stable and trustworthy partner, leading the way towards a promising future for all of humanity.

To build on its potential requires a strong multi-stakeholder ecosystem that shares the same direction, and is capable of developing and deploying next-generation technology in accordance with society's values. This needs to happen at increasing speed, scale, and complexity, and in a way that creates impact while offering new business opportunities. The future is data-driven, AI-powered and involves increasingly autonomous and sophisticated robotic systems. The ADR Partnership is key to enabling the future and establishing a strong, effective, and sustainable ecosystem for AI, data and robotics. By building bridges between disciplines, as well as backing research, innovation, and deployment, the ADR partnership will reduce the fragmentation of the European ADR landscape and contribute to establishing a sustainable ecosystem that achieves global impact and stimulates value creation.



Aligned with SRIDA's vision¹, the Horizon Europe Work Programme 2021-22 (WP 21-22) embraced ADR in the context of smart and agile manufacturing, the Green Deal, work environments and industry optimisation. It also addressed technologies and solutions for data sharing in common European data spaces through topics such as compliance privacy, data management, data trading, monetisation, exchange, and interoperability. WP 21-22 also tackled issues related to robotics cognition, AI for human empowerment, increased robotics capabilities for key sectors, trustworthy AI and European coordination on trustworthy ADR. WP23-24 addresses integration of the data lifecycle, the cognitive computing continuum, AI-driven data operations, and compliance technologies. It also tackles novel paradigms and approaches for AI-driven autonomy, a step change in autonomy, and collaborative intelligence between machines and humans. Finally, WP23-24 also encompasses open innovation regarding broader challenges with AI, efficient trustworthy AI, explainable and robust AI, and natural language understanding and interaction.

As a continuation of these efforts and following these new developments, the work programme 2025-27 will further elaborate on the challenges surrounding ADR, such as trustworthy AI, robotics autonomy, flexible functionality, compliance, new paradigms, and improved standards for efficient data processing and computing. It will also further elaborate on smart and cognitive manufacturing, as well as increased autonomy and resilience of production by exploiting ADR for remanufacturing, recycling, waste management, and re-valorisation. Finally, considerable efforts should be made to improve human-machine collaboration, ethics, and compliance for ADR to service society.

1. Zillner, S., Bisset, D., Milano, M., Curry, E., García Robles, A., Hahn, T., Irgens, M., Lafrenz, R., Liepert, B., O'Sullivan, B. and Smeulders, A., (eds), (2020). "Strategic Research, Innovation and Deployment Agenda – AI, Data and Robotics Partnership. Third Release." September 2020, Brussels. BDVA, euRobotics, ELLIS, EurAI and CLAIRE", <https://adr-association.eu/wp-content/uploads/2020/09/AI-Data-Robotics-Partnership-SRIDA-V3.0-1.pdf>

1.1 Alignment with European ADR efforts

Collaborative efforts

The Strategic Research Innovation and Deployment Agenda (SRIDA) aims to build on the fundamentals of Europe's aspirations to be a world-leader in ADR, by both enhancing the revenue-generating potential of companies' business models and enriching our society. It is closely aligned with the EU strategic initiatives landscape, including the AI Innovation Package,¹ the AI Continent Action Plan,² the Apply AI Strategy,³ the AI in Science Strategy,⁴ and the Resource for AI Science in Europe (RAISE),⁵ all of which collectively aim to strengthen Europe's capacity to develop, deploy, and scale trustworthy AI and data-driven technologies.

The AI Continent Action Plan,⁶ delivered by the European Commission in April 2025⁷, outlined a bold vision to place Europe at the forefront of trustworthy AI development, backed by the launch of InvestAI⁸ and a commitment to mobilise €200 billion. The plan was built to serve the purpose of competitiveness and productivity on one hand, and sovereignty, security and democracy on the other, showing full alignment with the overarching recommendations for the Horizon Europe work programme 2025-27 that this SRIDA suggests. The plan covers five different areas (Figure 1): Computing infrastructure; data (further developed by the Data Union Strategy); development of algorithms and adoption (further developed by the Apply AI strategy⁹ and AI in Science strategy); skills, and the simplification of rules.

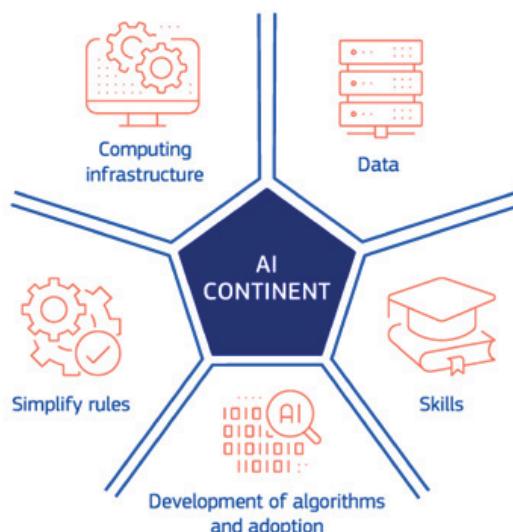


Figure 1: The AI Continent Action Plan, five areas (pillars)

1. <https://digital-strategy.ec.europa.eu/en/factpages/ai-innovation-package>
2. <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>
3. <https://digital-strategy.ec.europa.eu/en/policies/apply-ai>
4. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52025DC0724&qid=1762332390557>
5. https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2578
6. <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>
7. https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/artificial-intelligence-ai-science_en
8. <https://bdva.eu/news/investai-200-bil-euro-in-ai/>
9. <https://digital-strategy.ec.europa.eu/en/policies/apply-ai>

SRIDA's Big Tickets align with and contribute to all these different pillars as follows:

- » Big Ticket 1 (Ground-breaking Technological Foundations) and Big Ticket 2 (Effective and Trustworthy General Purpose ADR) directly contribute to the Development of algorithms and adoption pillar, offering guidance and a roadmap for AI researchers and industries benefiting from the computing infrastructure pillar of the Action Plan. Both the users and host organisations of the AI Factories and future AI gigafactories can benefit from the suggested roadmap and strategies presented in this SRIDA. For example, by optimising infrastructure for AI based on the new trends and roadmaps, identifying sectors where the impact of the new technologies can be expected to be higher, and offering infrastructure services fitting the future needs of AI researchers and companies.
- » Big Ticket 3 (an interoperable and integrated framework for data and model ecosystems) directly contributes to the Data pillar, addressing the main challenges related to accessing the massive amount of high-quality restricted and non-restricted data required for AI.
- » Big Ticket 4 (Next generation robots with embodied intelligence) directly contributes to the pillar of the Apply AI strategy, addressing challenges, impact and roadmap of AI-powered robotics that create impact in strategic sectors of the European economy.
- » Big Ticket 5 (ADR Technology for the Sciences) offers a direct and very relevant contribution to the AI in Science Strategy (part of the Development of algorithms and adoption) pillar.
- » Last but not least, Big Ticket 6 (Research, innovation, and tools for compliance) contributes to the pillar of AI adoption, data and simplification. The Skills pillar is covered by all the Big Tickets transversally and aligns with the overarching recommendations provided by this SRIDA.

ADRA Inputs to Apply AI Strategy

Despite possessing world-class research capabilities in AI and robotics, Europe confronts a critical and widening 'innovation-to-leadership gap'. This is a systemic, bi-directional failure where scientific breakthroughs are not translated into market-dominating companies at the required speed, while industries simultaneously fail to invest in and seek out game-changing innovation. This gap is sustained by deep-rooted issues within existing frameworks, including fragmentation and silos, risk aversion, and a mismatch between funding cycles and the speed of technology. Incremental change is insufficient to bridge this divide: a fundamental reimaging of the public-private collaborative framework is necessary to deliver disruptive and highly saleable innovations.

To address this, the European Robotics Catalyst has been introduced to serve as a vital intermediary across the full technology stack. The Catalyst has two primary functions. Firstly, it relays the voice of the end-user throughout the value chain, conveying the needs of the market to guide innovation while providing essential infrastructure, coaching, capital, and expert support. Secondly, it functions as a challenge-driven ecosystem orchestration engine, delivering crucial oversight across diverse portfolios of initiatives. This structure is designed to actively manage and connect the different stages of development, ensuring that research efforts are aligned with real-world applications and commercial potential.

The European Robotics Catalyst differs fundamentally from current instruments because it prioritises market-driven principles. It champions market pull over technology push, using strategic public procurement and industry-prioritised challenges to ensure innovations have clear commercial pathways from the outset. It is specifically designed for private capital attraction, employing sophisticated de-risking mechanisms and government-backed venture capital to make European AI-powered robotics startups genuinely attractive to private investors. Finally, it implements empowered and independent programme management, where qualified managers are given milestone-based funding and the mandate for intellectually honest project termination, replacing bureaucratic inertia with the agility needed to accelerate innovation and prevent resource wastage.

1.2 ADR Missions

1. Creating a strong, coherent, and effective ecosystem for AI, data, and robotics.
2. Maintaining and strengthening European industrial leadership in robotics, computer vision, and trustworthy AI.
3. Integrating and connecting the European research landscape around AI, data, and robotics.
4. Developing a powerful strategy for skills development and attracting talent to Europe.
5. Developing ADR technologies with high socio-economic impact to reinforce Europe's strong and globally competitive position.
6. Ensuring societal trust in AI, data, and robotics.

1 Creating a strong, coherent and effective ecosystem for AI, data, and robotics

Europe needs a functioning ecosystem covering ADR that can establish the foundation for boosting value created by the innovative development and deployment of these technologies. No single player can achieve this; the sharing of assets, technology, skills, and knowledge is crucial. A critical mass of engaged stakeholders is also needed to scale the deployment of these technologies in real-world situations. Although Europe has strong ecosystems centred around AI, data, and robotics separately, it needs to develop a single interconnected ecosystem that spans both the European territory and different technical areas. An ecosystem that connects and leverages European efforts in each of these areas needs to reflect the complexity and diversity of its constituents. It must encompass the three dimensions of AI, data and robotics, and ensure that knowledge is cross-fertilised. This requires effective engagement from all stakeholders and alignment between them to ensure efficient collaboration. *This is Adra's key mission.*

2 Maintaining and strengthening European industrial leadership in robotics, computer vision, and trustworthy AI

Europe is a world leader in robotics, computer vision and trustworthy AI technologies, with a considerable global market share and a well-established ecosystem of scientists, developers, suppliers, system integrators, and end-users. This trend is ongoing, but the current competition around semiconductor innovation and the rapid development of computer vision and robotic-based AI, both in Asia and North America, poses risks to Europe's leading status. It also potentially deprives Europe of opportunities. For example,

the rapid growth of chip manufacturing planned in factories across the USA, Korea, Japan, and India, together with the development of global standards, provides support for these regions to implement and adopt these technologies. Europe should therefore seek to maintain its world-leading status by prioritising investments that strengthen its position, such as the EU Chips Act.

3 Integrating and connecting the European research landscape around AI, data, and robotics

Europe has a strong AI, data and robotics research capability and capacity in academia and research organisations. However, their activities are fragmented between different communities, and siloed across disciplines and member states. This reduces the effectiveness and impact of European companies. It also reduces Europe's capacity to translate research into innovative solutions, and to develop research addressing real-world challenges. This fragmentation must be tackled, or investments in research, innovation, and deployment will not maximise efficiency and effectiveness, due to redundant and sometimes even counter-productive activities.

4 Developing a powerful strategy for skills development and attracting talent to Europe

A coherent approach to skills development is needed across Europe, from primary education to university and job market levels. Small-scale and national initiatives need to be covered at a European level. As ADR uptake accelerates across Europe in the coming decades, there is a need to ensure that the workforce has the skills to deploy, install, and maintain ADR systems. There must also be training of technical specialists to design and develop such systems. A failure to address the skills deficit will inhibit future deployment, as technology can only partially address the shortfall. The creation of a coherent European strategy to ensure the existence of necessary advanced skills will play an essential role in ensuring that productivity gains from new technology are maximised.

More alarming is the steady outmigration of EU citizens with exceptional AI, data and robotics skills, mainly to the US, who would have offered the potential to generate wealth across Europe. The continent must act to establish processes and actions that can boost the attractiveness of existing European innovation ecosystems or develop new ones to retain talent.

5 Developing ADR technologies with high socio-economic impact to reinforce Europe's strong and globally competitive position

According to research,¹⁰ AI can drive a 7% increase in global GDP over a ten year period and potentially mitigate up to 10% of global greenhouse gas emissions by 2030.¹¹ This is equivalent to the total emissions of the European Union.¹² There is growing attention on ADR as key instruments for growth and for shaping the new undesirable societal structures.¹³ In these contexts, the scale and complexity of problems to be solved pose new challenges to current AI techniques that need to be scaled, made global and more efficient, incorporated into hybrid AI systems and integrated with knowledge coming from human experts. In

10. <https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html>
11. Dannouni, A., Deutscher, S. A., Dezsaz, G., Elman, A., Gawel, A., Hanna, M., Kharij, A., Jones, E. R., Patterson, D., Rothenberg, J., Tber, H., & Ziat, A. (2023). Accelerating climate action with AI. Boston Consulting Group and Google. <https://web-assets.bcg.com/72/cf/b609ac3d4ac6829bae6fa88b8329/bcg-accelerating-climate-action-with-ai-nov-2023-rev.pdf>
12. <https://sustainabilitymag.com/tech-ai/google-and-bcg-report-on-accelerating-climate-action-with-ai>
13. Trabelsi, M.A. (2024), "The impact of artificial intelligence on economic development", Journal of Electronic Business & Digital Economics, Vol. 3 No. 2, pp. 142-155. <https://doi.org/10.1108/JEBDE-10-2023-0022>

addition, security issues arising from terrorism-related incidents, epidemics and natural disasters can also be addressed and better managed through smart systems.

6 Ensuring societal trust in AI, data, and robotics

There are many misconceptions surrounding AI, data and robotics, and these technologies are not fully accepted by society in all areas of application. This lack of acceptance will slow uptake, particularly where mistrust is unfounded, and may also cause damage in markets where the real dangers are not fully understood.

1.3 ADR Goals

The ADR Partnership is committed to contributing to the following high-level goals.

1. Boosting Europe's AI, data, and robotics industry, increasing its competitiveness and accelerating its digital and green transformation in accordance with the Digital Decade.
2. Achieving European strategic autonomy in AI, data and robotics.
3. Achieving global research impact in AI, data and robotics.
4. Maximising the societal and environmental benefits of AI, data and robotics to tackle major societal challenges on climate, food, energy, health, and security.

Significantly boosting European industry

It is of paramount importance to boost the innovation ecosystem for European industries to stay globally competitive. One key aspect is to secure European access to relevant resources such as key technologies and raw materials. Europe must facilitate creative thinking for existing and new businesses that can strengthen the region's resilience and independence when it comes to the production and supply of basic elements such as food, energy, semiconductors, and pharmaceutical products.

The responsible AI-powered green digital transformation is accelerating, and European companies are in a good position to drive and leverage this development. This trend should be supported with a regulated framework on compliance and legislation to strengthen solutions deployment and acceptance. American, Chinese, and other Asian companies already invest heavily in AI, data, and robotics. The World Economic Forum estimates that 70 percent of all new value over the next ten years will be digital. According to McKinsey, Europe needs to take immediate action to invest in enabling and horizontal technologies, since "technology is now permeating all sectors via transversal technologies such as artificial intelligence, the bio-revolution, and the cloud." McKinsey states that "although Europe has many high-performing companies, in aggregate European companies underperform relative to those in other major regions: they are growing more slowly, creating lower returns, and investing less in R&D than their US counterparts. This largely reflects the fact that Europe missed the last technology revolution, lagging on value and growth in information and communications technology and on other disruptive innovations."¹⁴

To close this gap, Europe needs to take immediate measures to increase the volume of public and private investment, involving established companies, new companies and new business domains. Cooperation and integration across regulatory and technological

14. <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/securing-europes-competitiveness-addressing-its-technology-gap>.

domains are essential for societal well-being, economic growth, and technological progress. For example, generative AI, neuro-symbolic AI, and AI-enhanced robots can be customised to specific high-impact application domains (such as healthcare, assisted living, manufacturing and logistics, food, forestry, inspection & maintenance of infrastructure, large power and industrial plants, and other service robots¹⁵ used in the home or for shopping, education, and entertainment). This is where initiatives like those of the ADR Partnership take on a central role.

One approach to drive the development and gain an advantage is to develop tools and processes that enable companies to increase their value creation. These are often de-facto norms that are consensus-based and centre on international standards. It is therefore necessary to actively engage in standardisation work, preferably being able to steer these standards towards European interests. Well-balanced regulation is becoming increasingly important and can translate into a competitive advantage for Europe if we invest in the research, tools, engineering frameworks and sandboxes to facilitate compliance and accelerate the emergence of new types of innovation among complex regulatory landscapes. A key goal for the ADR partnership is to contribute to pre-standardisation activities that bring research experts together with regulators, the industry and standardisation bodies to accelerate the definition of European standards and to position them worldwide.

Achieving strategic autonomy in Europe

To maintain strategic autonomy, Europe needs to ensure that it has all the relevant competences and capabilities. This involves learning existing methods (and often adapting and improving them to our needs) and developing unique new methods; both the methods themselves and their application can be unique. To achieve this, we will require systematic approaches to continue funding successful research projects and transfer research results, with the goal of both developing the technology further and learning how the existing technology works. Having European companies involved in research projects with leading researchers outside of Europe could be a way of bringing that know-how to European

15. "World Robotics 2024 – Service Robots incl. Mobile and Medical Robots", International Federation of Robotics (2024).

companies. Furthermore, bringing education, research, and field applications closer together will strengthen competencies with specific domain knowledge. In other words, the gap between education, research and professional deployment should be minimised through the establishment of joint infrastructures that facilitate practical training for students and scientists to extend their theoretical knowledge via real-world scenarios.

Achieving global research impact

To achieve global impact, it is often necessary to work on the current hot topics and publish research in the most prestigious outlets. It is even better to be the first to discover the next big thing. This is highly desirable and worth striving for, but it cannot be planned, as science is inherently characterised by a lack of predictability. This is a major challenge when preparing for a work programme that starts in two years, and where it may be necessary to predict what the next big thing will be three years from now. It is possible to recognise trends and predict the general direction that research will take. However, this is harder in research fields that move very fast, such as AI, data and robotics. Findings from 2021 are often considered out of date, since there is often significant progress on the most active research topics within a single year. A viable alternative is to set a goal of having a significant impact on the major conferences and journals in the area. For example, this could involve publishing at least 100 papers at each major outlet, such as AAAI, ICML, IJCAI, CVPR, ICCV, ICRA, NeurIPS, PAMI, Nature, Science, etc. every year. This could be achieved through a yearly call for next-generation high-impact methods in AI, data and robotics.

Enabling significant progress on major societal challenges

The ADR technology under development in the ADR Partnership represents enablers and opportunities to address the major societal challenges described in the first section (e.g., climate, energy, and food security). It is therefore an important goal for Adra to ensure that Europe develops ADR technology, not only to strengthen its position on business leadership, but for the purpose of societal welfare and sustainability. This can be summarised as AI for good and AI for all. It is of paramount importance that technological progress benefits everyone. For example, AI-enhanced and data-powered robotics can support societal objectives such as the green economy (recycling, remanufacturing, agri-food, inspection), healthy ageing (healthcare and assisted living), welfare at work (HRI, HRC, agri-food, dirty, dull, and dangerous tasks), ethics in robotics (safety) and strengthening industrial policy for European industries.

2. Major Trends and Gaps

Adra focuses on the cross-section between AI, data and robotics, with the long-term goal of achieving convergence in these areas. ADR is not yet an area of its own, the trends and gaps analysis are based on three existing communities, which all have significant overlaps and interactions.

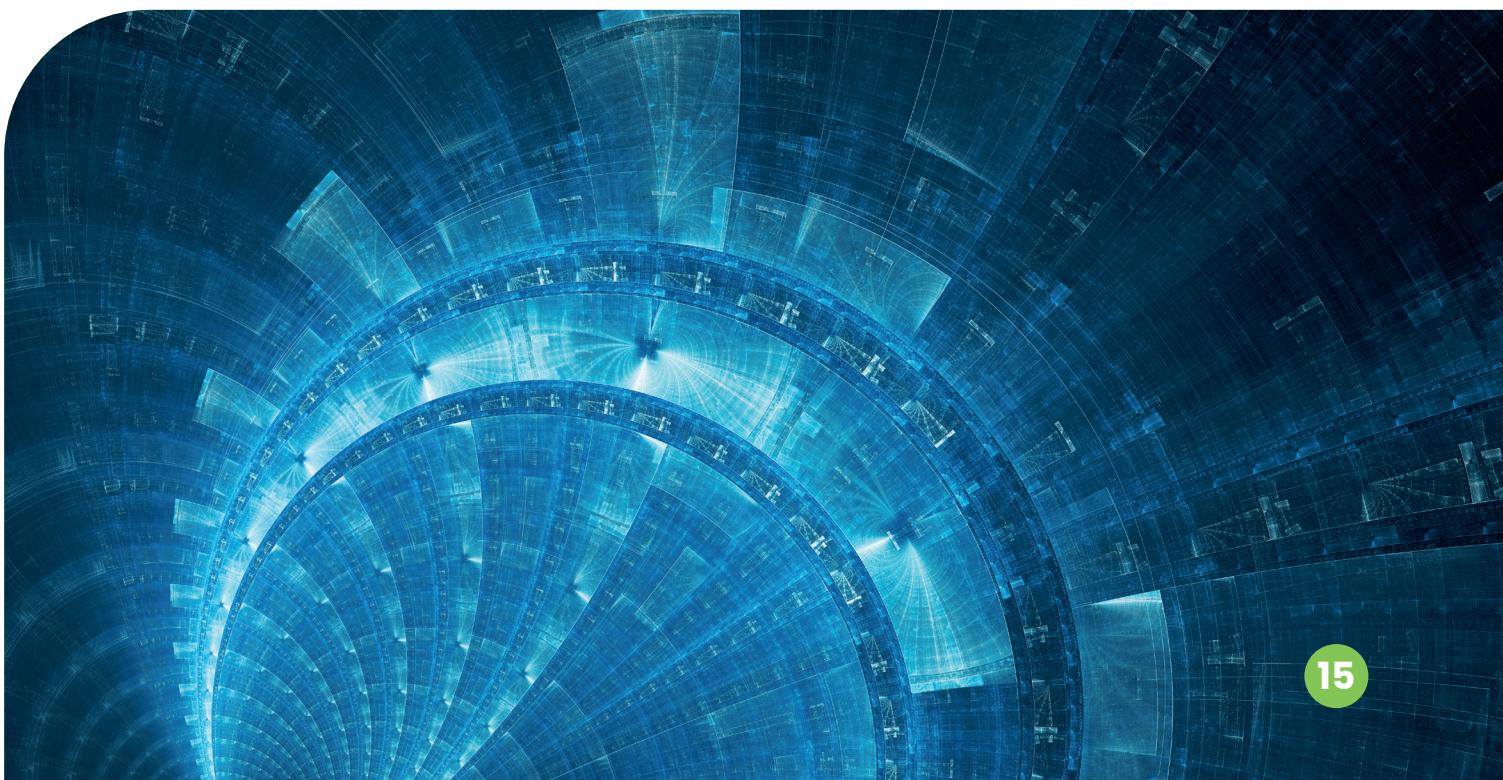
There is currently immense interest in generative AI¹⁶, especially large language models and multimodal systems. The past year has seen significant advances in reasoning-capable models like OpenAI's o1 and the emergence of AI agents that can complete complex tasks autonomously, alongside major breakthroughs in video generation with models like OpenAI's Sora 2 and Google's Veo, which now generate realistic videos with synchronised audio and dialogue. Current developments focus on enhancing reasoning capabilities, improving multimodal integration across text, images, code and audio, and expanding context windows to process vast amounts of information, while also advancing data curation quality and post-training techniques to improve model performance. A critical challenge is that AI may soon exhaust all publicly available human-generated training data. This has accelerated the adoption of synthetic data—artificially generated information that mimics real-world patterns—with projections that synthetic data could represent 80% of AI training data by 2028. Synthetic data addresses privacy concerns, helps reduce bias, and provides unlimited training examples, which are particularly valuable for rare scenarios like autonomous vehicle edge cases, though it carries risks of model collapse and error propagation if models rely too heavily on synthetic data without human oversight. Managing intellectual property rights remains challenging, especially as content creators and platforms increasingly restrict data access to prevent scraping for AI training. An important research direction involves automated fact-checking systems that verify AI-generated content by retrieving contextual information, explaining reasoning processes, and citing sources to address verifiability concerns. The most limiting factor for European involvement in large language models is the capacity to develop and deploy large-scale models.

One significant trend is neuro-symbolic hybrid AI methods, which pursue systematic and well-grounded approaches to combine symbolic and neural representations, integrating reasoning and learning in principled ways. In 2025, neuro-symbolic AI gained wider adoption with the aim of addressing hallucination issues in large language models, with companies like Amazon applying it in warehouse robots and shopping assistants to enhance accuracy and decision-making. With new AI regulations including the EU AI Act, black-box models are no longer acceptable in high-stakes applications. AI systems must now explain their decisions, which makes neuro-symbolic approaches valuable since verifiability and assurance are viable within symbolic components and they facilitate transparency and explainability for the public. Neuro-symbolic AI appears in the 2025 Gartner Hype Cycle for Artificial Intelligence as an early-stage technique that aims to bring more logic, structure, and context into how AI systems think and decide. Recent systematic reviews show that from 2020 to 2024, neuro-symbolic AI research experienced rapid growth. Europe maintains

16. E. Curry, F. Heintz, S. Leijnen, T. Jongen, F. A. Zeleti, A. N. Belbachir, P. Piatkiewicz, E. Girardi, "Trustworthy Generative AI: Fundamentals, Applications, and Future Directions," Springer 2026.

strong research foundations in this field, although global recognition and adoption are now widespread. These methods align well with Europe's vision of trustworthy human-centred AI, particularly given the EU AI Act's requirements for safety, transparency, and fundamental rights protection, with key provisions for high-risk systems. Neuro-symbolic generative AI is a great opportunity for Europe: it originates from Europe, has largely been developed in Europe, and is now getting global recognition.

Another trend is learning from human feedback, especially Reinforcement Learning with Human Feedback (RLHF). Seen as a few-shot learning problem, this provides opportunities for individuals to personalise models or to steer them towards styles and topics that are relevant for them. This can also be a way to reduce the need for ever-increasing amounts of data. A parallel development is Reinforcement Learning with Verifiable Rewards (RLVR), successfully employed in models like DeepSeek-R1 and o3, which uses simple rule-based functions providing clear binary feedback (correct/incorrect) based on deterministic tools such as calculators for math or compilers for code, rather than learned reward models. Recent advances have expanded RLVR beyond mathematics and coding to diverse domains including medicine, chemistry, psychology, and economics by using generative scoring techniques that provide soft, model-based reward signals for handling unstructured reference answers. Further innovations include RL from AI Feedback (RLAIF), which addresses the expense of gathering high-quality human preference labels by using AI-generated preferences, achieving comparable performance to RLHF while enabling "self-improvement" pathways. These approaches enable few-shot personalization where models can quickly adapt to individual user preferences and contexts through in-context learning capabilities, allowing users to customise model behaviour with minimal examples. While RLHF requires relatively small amounts of comparison data to be effective, making it more data-efficient than traditional supervised learning, emerging trends include automated feedback generation using synthetic human-like feedback to reduce annotation costs, and cross-domain transfer allowing models trained in one sector to adapt more quickly to another. This evolution addresses the dual challenge of reducing data requirements while enabling individuals to steer models toward styles, topics, and behaviours relevant to their specific needs.



In general, impressive progress has been made in dealing with relatively simple modalities including text, images, audio, and video. Now leading multimodal models natively process text, images, audio, and video in unified workflows. The native multimodal capabilities signal a future where AI seamlessly processes diverse data types in unified workflows, enabling more natural human-AI interaction. The next major challenges involve more complex data structures such as social networks, molecular structures, transportation networks, and other graph-based representations. Graph Neural Networks (GNNs) have emerged as the key approach for geometric graphs with physical symmetries of translations, rotations, and reflections, with geometric GNNs equipped with invariant and equivariant properties to better characterise geometry and topology. GNNs have proven particularly successful in materials science and chemistry due to their ability to directly access structural representations of molecules and materials. There is ongoing work in the field of geometric deep learning, which incorporates physical symmetries, enabling applications from protein structure prediction to molecular chirality recognition. Physics-Informed Neural Networks (PINNs) have emerged as a transformative approach that embeds physical laws described by partial differential equations directly into neural network training, acting as regularisation that increases generalisability even with limited training data. These advances in geometric learning and physics-informed approaches represent principled ways to incorporate domain knowledge and physical constraints into neural networks, addressing the fundamental challenge of moving beyond simple pattern recognition to genuine understanding of complex structured data.

Agentic AI has moved away from reactive generative models to autonomous systems capable of planning, reasoning, and executing complex tasks with minimal human supervision. Recent research distinguishes between basic AI agents, which are modular systems driven by large language models for task-specific automation through tool integration and prompt engineering, and true agentic AI systems characterised by multi-agent collaboration, dynamic task decomposition, persistent memory, and coordinated autonomy. Agentic AI constitutes the next major evolutionary step beyond generative AI, with substantially stronger reasoning and interaction capabilities enabling more autonomous behaviour to tackle complex tasks. Scientific applications have demonstrated potential for automating literature reviews, experimental design, and data analysis, with frameworks like Agent Laboratory autonomously progressing through research idea development, experimentation, and report writing, although significant challenges remain in deep domain-specific knowledge and human-AI collaboration. A critical limitation of current agent systems is their reliance on manually crafted configurations that remain static after deployment. This problem has incentivised research on self-evolving AI agents that can automatically enhance themselves based on interaction data and environmental feedback, bridging static foundation model capabilities with continuous adaptation requirements. One important architectural development is the suggestion that small language models (SLMs, typically below 10 billion parameters) are sufficiently powerful and more economical for many agentic invocations where models perform specialised tasks repetitively; heterogeneous agentic systems combining multiple models are emerging as the natural choice. Theoretical frameworks propose defining agentic capabilities through “functional agency”—quantifying systems by their ability to take goal-directed actions, model outcomes, and adapt behaviour when action-outcome models change, with current LLM-based agents showing significant limitations in longer-horizon tasks due to difficulties in environmental interfacing, lack of common sense reasoning, and tendencies toward self-deception. The field represents a fundamental transition

from single-use reactive systems to adaptive, goal-directed autonomous agents capable of operating in dynamic environments, though substantial research challenges remain in the areas of reliability, long-horizon planning, human-agent interaction, and security before agents can be widely deployed in critical applications.

There has been significant progress towards promoting the free flow of data, sharing data and breaking data silos between different domains. However, several gaps remain, including interoperability, global trust and governance frameworks, and tools and methods to support compliance with existing and emerging data and AI regulations across the system lifecycle. Standardisation efforts are needed to guarantee trust, connectivity, and scalability of data storage and processing along the cloud-edge computing continuum. Sourcing high-quality and IPR/privacy-compliant data (including real, anonymised, and synthetic) for training AI models is also a key challenge.

There are several gaps in the field of robotics. Robots are highly application-specific machines that are founded on a common set of methodologies and building blocks. European standards must be set around modularity and the development of design, certification and validation tools that accelerate the development of supply chains, and, where modularity is not appropriate, around processes that speed up time to market. This work needs to be conducted on a European scale and actively involve regulators and test facilities to bring results that align with the market sector by sector. We need to invest in community-building, address regulatory issues, ensure standardisation, and involve testing and experimentation facilities (TEFs) as well as the European networks of DIHs and centres of research excellence. The problems that need to be solved are sector-specific and detailed. Generic approaches will not suffice.

Robotics should focus on impacting specific sectors, such as construction, cities, healthcare, agri-food, and energy supply. While partnerships exist in these areas, their natural focus is on their own major priorities, rather than new technology. For example, while the healthcare sector may recognise that robotics has a role to play, it will not be prioritised over vaccine development or drug discovery. It is therefore important that the ADR Partnership joins forces with sector-specific communities and partnerships, addressing those applications and verticals in the work programme.

There are considerable technical and procedural challenges to reaching basic robotic abilities and technologies, independent of application domains. There is a need to integrate AI-data and robotics and other technologies (e.g., HPC, IoT/robot middleware and operating systems) that can enable limited self-awareness, understanding of an environment, and planning. We need novel architectures and components (mechatronics and control) for achieving soft robotics, small-scale robots, flying robots and water robots that are completely autonomous. Human-robot interaction and collaboration that is physical and non-physical, as well as robot-robot interaction and collaboration and swarm robotics, should follow safety and process regulation. Other challenges include biomimetic perception and control; extended reality in robotics; dexterous grasping, manipulation, and navigation; localisation and mapping.

The linkage between materials and robotics is exemplified by the trend towards “soft robotics”, while micro and small-scale robotics is exemplified by “origami robots”. There is extensive research into the necessary materials across Europe, and a strong need to

connect these efforts to the development of future robotic systems that move away from the prevailing “box on wheels” configuration.

European-wide networks need to be focused on key European strengths or around the development of ecosystems where there is underlying strength that can be improved through coordination. To date, these networks have been broad-based, with little attempt to align. Networks have the potential to be aligned with missions, specific functional expertise, or specific emerging sectors. Tighter focus around networks of excellence will help to develop key strengths.

When it comes to data, innovators, SMEs and start-ups need good access to world-class, large-scale, federated and secure infrastructure. This includes access to data and resources such as HPC and test environments. Strong investment in HPC is taking place thanks to the EuroHPC JU, but additional investment needs to be directed towards computing-big data management-ML convergence and towards providing access to federated data experimentation and infrastructures. The lack of accessible and high-quality infrastructure will slow market development and limit success.

Another key challenge is the lack of certainty around the effect that new data and AI regulations (such as the Data Governance Act, the Data Act and the AI Act) will have on the market. Companies, especially SMEs, will require tools and support to address compliance, data access and exchange, data quality, right to explain and trustworthiness. New roles and business opportunities (such as data intermediaries) are emerging in the context of these new regulations; new markets and novel applications can also emerge. Strong investment in research, research-industry collaboration, methods, tools for compliance, support, and access to experimentation, through European-wide federated secure experimentation environments that offer support and sandboxing opportunities for companies, are needed to transform a potential threat into a strong opportunity for European industry. A multidisciplinary approach is also crucial to addressing this challenge.

The complexity and cost of creating deployable systems based on data knowledge and the availability of pre-competitive big data sets to train AI models can also prove challenging. However, it can be a game changer for the EU if achieved. Large-scale industry pilots and applications using those datasets (synergies ADR) are needed to demonstrate impact.

At the same time, the engineering and deployment of new ADR systems must take into account privacy, trust, security, and ethics (beyond compliance).

Unfortunately, many European organisations lack the skills to manage or deploy smart technical solutions that can be built on new technologies. However, there is already a global race for talent. Regions with the most vibrant technology landscape are better positioned to attract skilled professionals and retain local talent.

Finally, a lack of understanding of the business opportunities involved is also hampering progress. Achieving business impacts from smart technologies requires a full understanding of the market, the technology and its effects on business processes and models. Because this requires integrating knowledge from multiple stakeholders, it can result in low levels of uptake due to uncertainty and a lack of knowledge. In addition, the novelty of these technologies means that emerging business potential may not be obvious from the outset, which in turn slows the return on investment.

2.1 Global Challenges

Recent geopolitical crises and global disruption are causing severe challenges to Europe's future sustainable development and welfare. Access to essential resources¹⁷ such as food¹⁸, energy, and water are constrained by population growth, decreasing arable land, and intensifying global demand. Moreover, factories and human capital are continuously migrating to other ecosystems outside Europe. Climate change¹⁹ is leading to more unstable weather conditions, more frequent and disruptive natural disasters and the potential displacement of large groups of people across the globe. Demographic changes lead to a shrinking working-age population that needs to provide for and take care of an ageing population. Geopolitical instabilities²⁰ ranging from increasing tensions between major countries to political unrest, are reducing global cohesion and cooperation. This is especially challenging when dealing with global disruptions such as pandemics, natural disasters and wars.

Fortunately, Europe still enjoys high living standards with a well-developed welfare system. However, it is under severe pressure. Three major choke points are:

1. Europe's dependency on global supply chains for essential resources (such as energy) makes its society and industries vulnerable to global crises.
2. Europe's increasing workforce needs make it vulnerable to demographic changes and global competition.
3. Europe's dependency on the world supply of key technologies and materials (such as semiconductors and active pharmaceutical ingredients) makes its society and industries vulnerable to global disruptions.

AI, data and robotics can enable sustainable security and strategic autonomy in Europe for food, energy, key technologies and industries. European industries specialising in

17. https://www.fao.org/fileadmin/templates/wsfs/docs/Issues_papers/HLEF2050_Global_Agriculture.pdf

18. <https://info.bml.gv.at/en/topics/food/food-security-in-europe-through-climate-change-research.html>

19. <https://www.faccejpi.net/en/faccejpi.htm>

20. <https://www.cer.eu/hot-topics/geopolitical-eu>

robotics, computer vision and trustworthy AI have achieved global prominence. They boast a substantial share in the international market and have fostered a mature network of developers, suppliers, system integrators, end-users and scientists. While this dominant position persists, the swift advancement of computer vision and robotic-based AI in Asia and North America threatens European technological leadership and opportunities, such as the development of worldwide standards to enable clear guidance for the implementation and adoption of these technologies.

Europe should aim to maintain its leading global position with further investments that can speed up digital transformation. Investment in ADR will help secure the sustainable production of essential resources for Europe, such as energy, food, and other key elements, as well as efficient waste management. ADR has the potential to rapidly increase our capacity in production infrastructure, transportation and the regional workforce dedicated to maintenance and operations. Future tasks will eventually become unmanageable by humans alone, so robots and AI systems should be engaged to work side by side with humans to increase capacity and safety. For example, food security requires people to engage in agricultural activities. However, the population tends to leave rural areas for urban areas, reducing the labour force. As the amount of available arable land is drastically decreasing, food production infrastructure must be developed in other areas (such as oceans and remote locations). In short, ADR is a strategic investment to establish autonomous and intelligent systems that can fill the workforce gap and strengthen Europe's commitment to sustainable growth, welfare, and innovation.

2.2 New considerations and instruments from synergies to missions

The focus on combining AI, data and robotics has an impact at a high level when considering end user needs, SDGs and the Horizon Europe missions. While these challenges provide a shared objective and facilitate convergence, the danger of such an approach is that it might fail to develop the individual areas where significant technological developments are needed to bring novel products to market. The focus on the intersections and the individual disciplines must be well-balanced. All three disciplines will benefit from a renewed focus on their individual strengths in the second half of the Horizon Europe funding programme. This would provide a valuable balance compared to its first half, which focuses only on the intersections. Such an approach also needs to recognise that the investment profiles, research methodologies and translation paths for AI, data and robotics are significantly different for each technology, and that merely combining them fails to recognise the differences that are better addressed by supporting collaboration that recognises the difficulties at the interface.

To address this, Europe must orient AI, data and robotics around the European missions and develop focused excellence around each mission. It is important that any focus on missions is targeted rather than generic. Specific milestones on the path towards achieving the missions need to be identified and addressed. Ideally, these milestones should be established with the missions themselves, so that goals are aligned with other work programme elements at a strategic level, without the need to create joint actions,

which can be slower and less impactful. For instance, it is very important to collaborate with the European Chips Act that is striving to uplift Europe's competitiveness and resilience in semiconductor technologies and applications. This is of vital importance for ADR, as semiconductors are fundamental to sensing, computing and actuation. Their development will boost European technological security and ADR applications in critical sectors (such as energy, food and transport).

The new acts and directives in AI, data, medical device regulations, cyber-resilience and machinery present new opportunities for companies placing products on the market as well as the risk of an inhibitory influence on innovation. To ensure the regulation has a positive effect, it is critical that authorities are capable of giving fast and clear guidance on how to comply with the regulation as well as investments to lower the barriers and support compliance.

There is a need for large and complex testbeds that force the integration of many different techniques and solutions into complex systems-of-systems. The development of components has reached a sufficient level of maturity to enable real world deployment, but there is still a large gap between this technical capability and the ability of end users to easily integrate them in complex systems-of-systems—and thereby their ability to solve real-world problems that exist in healthcare, food production, energy, transport or smart cities. Large scale testbeds that can integrate AI, data and robotics with existing systems in highly realistic environments are a first-stage necessity that can create greater understanding of the impact, functionality and effect of combining these technologies and applying them to the real world. It is important not to underestimate the complexity of the integration task and the need for direct investment. In parallel, it is important to create sandpits for real world testing where regulatory norms are adjusted to allow autonomy to flourish within controlled and well-defined bounds. These regulatory sandpits are the bedrock of testing in the real world and are essential for the accelerated deployment of complex socio-technical systems. Since barriers are often created by national and even local regulation, care must be taken when siting these facilities to ensure that the combination of European and local regulation is viable. This is particularly critical as Europe is using the AI Act to set a high compliance bar for working environments and for AI-driven systems making high risk decisions.

Finally, we need instruments that allow start-ups and smaller companies to benefit from ADR and to play an active part in developing the next generation of smart technologies. Instruments designed to suit their needs should be developed. Such instruments can include "researchers for loan", where companies can borrow talent using a simple process. Researchers based on a pool of EU experts and armed with key knowledge for a short-term period can be lucrative for small companies looking to strengthen their innovation and growth.

2.3 ADR and European Competitiveness

Europe's competitiveness in AI, data and robotics will hinge not only on scientific advances but on how fast we **translate technology into use-cases that matter**. There is a hypothesis that argues that lasting advantage emerges when we out-innovate on *adoption* rather than on the raw technology frontier; this requires turning imported or open technologies into superior, context-specific applications, as Dutch printers did when they leap-frogged their German teachers in the 16th century. This complements Adra's goal of anchoring research in local innovation ecosystems and closing the "deployment gap" between labs and society.

To unlock that advantage, **public procurement must be mobilised as Europe's hidden super-power**. The EU already directs roughly 14 % of GDP through public tenders—twice the IT-market share seen in the US. By coordinating "innovation procurement" across the 250 000 European contracting authorities and linking it to Test & Experimentation Facilities (TEFs) and Digital Innovation Hubs, we can aggregate demand for trustworthy ADR solutions and de-risk early deployment for SMEs .

A second lever is **talent circulation rather than mere retention**. Research shows that specialists who leave and later return boost the productivity of their European collaborators, turning brain-drain into brain-gain. Building on SRIDA's call for AI-skills pipelines, we propose an EU "Circulate-Back" programme: funded two-way fellowships, start-up visas and simplified equity rules that entice expatriate researchers and founders to spend 2-3 years in European labs or scale-ups before re-exporting new know-how.

Third, Europe needs **regulation that is not only rigorous but adaptable**. The SRIDA already advocates regulatory sandboxes; the application advantage lens suggests going further by embedding explicit sunset- or revision-clauses and data-driven success metrics into each new rule, so obsolete provisions can be retired at speed. High-quality, self-correcting regulation would turn compliance from a cost into a comparative advantage and set global benchmarks for "evidence-based tech law".

Finally, we propose a **network of city-scale application testbeds**—building on Europe's dense urban fabric—to accelerate technology diffusion. Cities that open their mobility, energy or health datasets and procurement pipelines become living laboratories where SMEs iterate rapidly; knowledge then spreads through the network much like printed innovations once did across Dutch trading towns. Embedding these testbeds in Horizon Europe missions will ensure that every breakthrough is stress-tested against real-world complexity before scaling across the Union.

3. The Next Strategic Plan

2025-27

Inspired by the Taiwanese model in the semiconductor industry²¹, collaboration between education, research and businesses should be tightened in a systematic way to exploit the synergy between fundamental knowledge and applications, while also strengthening competencies with specific domain knowledge. The gap between education, research, and professional deployment should be reduced by creating shared infrastructures. These infrastructures would facilitate hands-on training opportunities for both students and scientists, enabling them to participate in field trials and real-world business scenarios. This practical experience would complement their theoretical knowledge and provide insights into real-world situations. This kind of collaboration can be beneficial for skills development among a highly competitive workforce that can be used by European industries and organisations.

Europe hosts several centres of excellence in ADR that are comparable to anything the rest of the world can offer, but these centres do not always receive the visibility they deserve. In the global battle to attract future generations of talent, each of these centres can attempt to compete individually, but a much better strategy is to collaborate within Europe and pool resources together under a unifying umbrella, as is the case with the concept of a multi-centric network of AI lighthouses. A network of lighthouses can offer an impressive portfolio of research challenges, great innovations, exciting partners, and career opportunities that can compete with the offerings of tech giants or other institutions outside Europe. Continuing to enforce the creation of this type of network and these brands is essential for Europe's future.

However, competition is difficult when there is a need for a coherent, and often centralised, effort to develop a single large-scale system or application, as is the case of the large language models that are starting to emerge. Loose federations will most likely never be able to deliver the same result; the diversity and additional costs of coordination will make it too expensive and slow. To achieve prompt results, the new strategic plan should address the following technological concerns:

- » Large-scale general purpose ADR technology. For example, open, large-scale, GDPR- and IPR-compliant European language models handling both European language and cultural differences. This includes speech-to-text, text-to-text, and text-to-speech.

21. Taiwan is the world's leader in semiconductor manufacturing with more than 60% of world market share. It is astonishing for a small country with just over 23 million people to dominate one of the most complex businesses in the world. Especially, if we consider that Taiwan does not own any of the ingredients needed for the manufacturing (i.e., raw materials, design technologies and tools, chipmaking machines). The strength of Taiwan lies in its large pool of high-skilled workers and supportive government policies. For instance, the new semiconductor factories being built in USA and Japan are dependent on few thousands of workers from Taiwan to operate them based on a deal with the Taiwanese government. More under: <https://medium.datadriveninvestor.com/how-taiwan-came-to-dominate-the-semiconductor-industry-cc7ac08c557a>.

- » Large-scale complex ADR testbeds, together with end-users, such as in manufacturing, healthcare, food production, transportation, energy, or smart cities.
- » Multi-stakeholder development, verification, validation, and integration of automated decision-making in socio-technical systems, both for the public and private sector.
- » Collaborative autonomous systems interacting with both the environment and people. This includes autonomous drones in controlled airspace, last mile delivery, and self-driving vehicles.
- » Metrics for measuring progress in ADR, with a special emphasis on trustworthy ADR technology.

3.1 Overarching recommendations for Work programme 2025–2027

The recent world crises have placed demands on Europe to prioritise mobilising substantial investment in innovation and research on ADR systems across different technology readiness levels to accelerate the digital and green transition and reinforce its sustainability, resilience and technological sovereignty. On a broader scale, the work programme for 2025–2027 should simultaneously address the following five high-level concerns, which we have further developed as big tickets in the next sub-section:

- » Targeting European Leadership in Climate-Neutral Circular Industries and Digital Value Chain: ADR should contribute to the Green Deal, the green transition for a sustainable society and carbon emissions reduction under the Net-zero Industry Act²². ADR should also support efficient and safe operation, maintenance and inspection, as well as contributing to the effectiveness of the circular economy and resource management. For example, ADR could support AI-powered urban mining and robotics solutions with advanced capabilities in complex manipulation for autonomous dismantling, disassembling, recycling, sorting and/or remanufacturing of waste products. In January 2024, the International Federation of Robotics (IFR) announced²³ that leading European manufacturers had achieved a world record of 3.9 million operational robots 2022. There is a large potential for growth thanks to industrial robots²⁴ with complex manipulation capabilities powered with AI and data, which may help to solve the challenges of a shortage of skills and human labour by effectively performing tedious and dangerous tasks for long hours at stable quality and performance. (Big tickets 1–6).
- » ADR for European Strategic Autonomy, Resilience and Technological Leadership in Critical Value Chains: Boost industrial competitiveness and strengthen strategic autonomy in the production and supply of essential resources. This includes enhancing Europe's resilience to the impacts of conflict on supply chains (e.g., microchips, critical raw materials, rare earth minerals, and pharmaceutical key starting materials) and the development of a

22. Proposal for a regulation on establishing a framework of measures for strengthening Europe's net-zero technology products manufacturing ecosystem (Net Zero Industry Act), COM(2023) 161, March 2023. https://single-market-economy.ec.europa.eu/industry/sustainability/net-zero-industry-act_en

23. <https://ifr.org/ifr-press-releases/news/global-robotics-race-korea-singapore-and-germany-in-the-lead>

24. "World Robotics 2024 – Industrial Robots", International Federation of Robotics (2024)

strong semiconductor industry in Europe boosted by ADR systems. In accordance with the EC 'demonstrators report', we need to develop mechanisms to enable effective transfers from small-scale industrial demonstrators (innovation) to pioneer climate-neutral demonstrators (deployment) that must be further collaborated with the European Digital Innovation Hubs²⁵ and elaborated in FP10²⁶. (Big tickets 1-4).

- » Trustworthiness of ADR with agile and secure single market and infrastructure: Elaborate on ADR made in Europe in compliance with regulation, including the AI Act, the Data Act, the Cyber Resilience Act and the Data Governance Act. Ensure innovation complies with these regulations. We need to expand the data economy from the consumer-to-business model to other models by strengthening privacy-preserving technologies, compliance, source and transaction integrity and interoperability. We need to rebalance data, computing, learning and actuation capacity across the cloud-to-edge continuum for businesses, public organisations and individuals to exploit ADR. This requires an emphasis on how leveraging SMEs and DIHs can support technology transfer, innovation, and the creation of new products and services based on ADR technology. This will support achieving the goals outlined in the Digital Decade, especially reaching 75% of companies being automated (Big tickets 2, 3, 4 and 6).
- » European Strategic Autonomy in ADR enabling technologies: Europe should pursue its efforts to preserve an open economy in leading technologies while maintaining its strategic autonomy in ADR for deep digital transformation in industries, public services and society. For example, ADR should support the top priorities of the European Chips Act, enabling technologies that strengthen processes at critical stages in the chips value chain (e.g., design, manufacturing), and address the use of new materials and green technologies for smart embodied intelligence and energy efficiency. We should address high value-added hardware needs for core computing, sensing and real-time communication of ADR systems to secure the critical components supply for key markets. ADR systems should support human capacity for effective, high-performance and high-precision in chips manufacturing. Furthermore, research on core learning and analysis techniques (evolutionary, collaborative, continual), as well as next generation smart robotic systems and data economy will keep Europe at the cutting edge of ADR. Finally, ADR should contribute to societal prosperity and European resilience to crises. It should enhance the resilience of our society to crises, both natural and man-made, and should improve preparedness, enabling fast and efficient responses to catastrophic situations such as cybersecurity threats (Big tickets 1-6).
- » Human-Centric Innovation Powered by ADR Education and skill building: Targeted education and upskilling on ADR systems will open up new ways of working, assisted by technologies for physical or cognitive augmentation (exoskeletons, digital twins, collaborative AI, virtual) or collaborating with robots and machines in a safe and effective manner. Professional and widespread education on ADR systems will increase efficiency, safety and quality of work—provided the systems are trustworthy, safe and reliable, as well as human-centric and without social biases by design. Within the dynamic context of flexible organisation and process flows, workers will be empowered by ADR knowledge and systems to co-create new forms of working and collaboration within and across organisations, through participation, social innovation or living labs.

25. European Digital Innovation Hubs | Shaping Europe's digital future (europa.eu)

26. The next EU funding programme for research and innovation

New job profiles and skills will emerge, often requiring digital competence in addition to formal competence. Continuous learning, through formal training, on-the-job learning or immersion in virtual worlds, combined with appropriate certification and reward mechanisms, can boost the attractiveness of careers in many sectors, including manufacturing, field operation, inspection and maintenance. (Big tickets 1-6).

3.2 Big Tickets in AI, Data and Robotics for 2025-2027

1. **Ground-breaking technological foundations** in ADR (autonomy, high-performance and predictability)
2. **Effective and trustworthy general-purpose ADR** (generative AI, neuro-symbolic AI, generality, continual learning, trust, scale and complexity)
3. An **interoperable** and **integrated** framework for data and model ecosystems (operations, governance, privacy, and security)
4. Next-generation **smart embodied robotic systems** (soft robotics, autonomy, manipulation, configurability, human robot interaction and collaboration)
5. Developing **ADR technology for the sciences** (from data to knowledge and understanding)
6. **Research, innovation, and tools for compliance** (trust, privacy, security beyond compliance)

3.2.1 Big ticket #1: Ground-Breaking Technological Foundations in ADR

State of the art

Research and innovations in ADR are moving at such a rapid pace that any state-of-the-art advancement requires up-to-date knowledge and expertise in both fundamental research and deployment contexts. Real-world applications depend on technological foundations that enable safe and effective operation across different application sectors and user groups. These technological foundations, in turn, require advances in research to deliver systems that are reliable, sustainable and easy to manage in order to meet usability expectations. This creates a dynamic feedback loop where technological foundations and real-world applications need to converge to each other.

This dynamism can be seen across a range of ADR methods and tools. Foundation models are rapidly advancing to deliver multimodal systems capable of processing text, images, code, and data from other modalities. Developments in agentic AI suggest the possibility of AI agents that can autonomously complete multi-step tasks, enabling ADR systems to interact with digital and physical environments. General-purpose (incl. Generative) AI has expanded from text synthesis to image and video generation, with new applications in industrial design and simulation. These advances—combined with user-friendly interfaces—create possibilities for ADR capabilities that can tackle complex scientific and industrial challenges.

The fast pace of development underscores the need for a European ADR research strategy that remains agile in both fundamental research directions and deployment approaches.

Patent data shows European companies leading globally in additive manufacturing and autonomous robots, while matching China and the US in autonomous vehicles. While Europe benefits from world-leading AI research expertise, these numbers suggest that further action is needed to support a pipeline that connects this strategic resource to the design and deployment of AI systems²⁷. European investment in the ADR supply chain is required in order for the region to exert control over the development and innovation lifecycle for scientific and industrial leadership.

Deployment experiences have highlighted the areas where fundamental research is needed to bolster the ADR supply chain in Europe. Progress is needed across the entire value chain—competitive infrastructure, research, development and deployment—to maintain resilience in changing and uncertain environments. New ADR capabilities are essential to unlock widespread applications for industries and society. These capabilities—and the research challenges that underpin them—are explored in the research agenda set out under this big ticket. They connect fundamental research to the range of practical challenges that arise in real-world applications of ADR.

Significant investment in ADR, the engagement of a large number of scientists and engineers and the availability of software tools and high-performance computing platforms are accelerating the development and deployment of new ADR technologies at a high pace.

Challenges

Despite rapid ongoing progress, disparate ADR development efforts need to be further coordinated to enhance the impacts of ADR for real-world deployment—either to solve practical challenges or to open up new innovation and business opportunities. ADR solutions must span multiple sectors and integrate deeply with social and economic value-creating systems while ensuring safety and sustainability. Progress demands collaboration across academia, industries and decisionmakers, requiring an understanding of the pathways that take innovation from scientific inventions to real-world deployments and societal usability.

Building an infrastructure for next-generation AI R&D: Progress in foundation models has depended on the development of large models requiring significant data and compute resources. Building such large-scale models requires addressing operational challenges at scale. Other countries have already identified a roadmap to building and sustaining a research infrastructure that supports a pipeline of AI research across TRL levels²⁸. To position European research for influence alongside the large-scale AI systems being built by the tech industry in the US and elsewhere, European innovators will require access to computational facilities that provide the compute power needed to train large-scale AI models, expanding existing high-performance computing infrastructure.

One consequence of these compute needs is that the energy requirements of large-scale foundation models have grown significantly. This trend poses both environmental and economic challenges for European ADR development. Compute facilities must balance

27. For example, the Stanford AI Index Report, 2025, <https://hai.stanford.edu/ai-index/2025-ai-index-report>, indicates that the US and UK outperform Europe in the industrial deployment of AI.

28. For example, work in the US to create a National AI Research Resource <https://www.ai.gov/nairrf/>

sustainability, affordability, and accessibility for the R&D community. Sustainability improvements can be achieved through:

- » Locating facilities near renewable energy sources and alternative cooling systems;
- » Geographic separation from research locations where beneficial;
- » Investment in local energy generation and storage technologies; and
- » Implementation of carbon-aware computing practices and development of less computationally intensive ADR methods.

Fostering open innovation ecosystems for AI leadership: Rapid advancements in foundation models and AI agents underscore the critical importance of open innovation ecosystems. Europe's ADR competitive advantage requires both developing world-class research capabilities and ensuring these capabilities remain accessible to the full spectrum of European industry, from multinational corporations to SMEs and startups. This demands investment in open-source AI tools, shared datasets, and collaborative knowledge transfer platforms. Open infrastructures, including testing facilities for TRL 6-8 and AI-friendly development environments for TRL 4-5, can provide foundations for innovation. These enable companies to build services while ensuring cultural and scientific knowledge-sharing aligns with human rights frameworks and EU policy values. Creating such ecosystems requires sustained support for open-source community practices.

Securing ADR Systems Against Evolving Threats: The increasing sophistication and deployment scale of ADR systems creates new security challenges that require research and deployment responses. AI systems face vulnerabilities including adversarial attacks, data poisoning, and model extraction threats. They also offer capabilities for cybersecurity applications, including threat detection, automated response systems, and anomaly identification. European ADR development must address both aspects. This includes developing robust AI systems that maintain performance under attack, creating AI-powered cybersecurity tools for critical infrastructure protection, and establishing security-by-design principles for ADR deployment.

Applications and impact

Estimates of the value of European ADR vary but suggest economic gains of over \$2.5 trillion by 2030.²⁹ These estimates assume rapid progress in ADR R&D in Europe and the diffusion of these technologies across the European economy. ADR's potential is rooted in its pervasiveness. Across industrial sectors, ADR technologies offer opportunities to create value, from increasing the efficiency of existing processes to increasing productivity, designing new products or enabling new business models.

29. See, for example: [https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/637967/EPRS_BRI\(2019\)637967_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/637967/EPRS_BRI(2019)637967_EN.pdf) and [https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-europes-gap-in-digital-and-ai#:~:text=Section%203-,AI%20could%20give%20EU%20economies%20a%20strong%20boost,through%202030%20\(Exhibit%201\).](https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-europes-gap-in-digital-and-ai#:~:text=Section%203-,AI%20could%20give%20EU%20economies%20a%20strong%20boost,through%202030%20(Exhibit%201).)



These applications all rely on fundamental research to deliver foundational ADR technologies and tools that can be deployed in practice. Without solid technological foundations supported by vibrant R&D communities, these benefits cannot materialise. These solid foundations can create significant potential for social and economic impact. For example, the EU's Innovation Missions demonstrate AI's diverse sectoral applications and potential beneficial impacts:

- » In climate adaptation, AI can deliver effective climate models that predict extreme weather events, forecast local climate change impacts, and design adaptation measures.
- » In cancer prevention, AI can help improve medical imaging, underpin precision oncology for diagnosis and treatment or support new clinical decision-support tools.
- » In ocean protection, AI can help analyse marine ecosystem health and monitor key species or ecosystem interactions.
- » In smart cities, AI can help optimise transport networks and utilities to reduce carbon emissions or support climate adaptation.
- » In agriculture, AI can help manage and monitor soil health to enhance land management.

Beyond each of these mission domains, there are wider opportunities for ADR across multiple sectors, including:

- » Health: New diagnostics tools, modelling and simulations to test and personalise healthcare interventions, and decision-support tools for clinicians.
- » Environmental protection: Safeguarding biological heritage and biosphere health through improved environmental monitoring.
- » Urban development: Enhancing building architecture and city planning, for example through advanced simulations.
- » Industry: Process optimisation, pollution reduction, and design innovation.
- » Energy: Augmented decision-making and accelerated emergency simulations (for example, power grid blackout scenarios like the Iberian Peninsula full blackout of 28 April 2025), plus electric vehicle integration.
- » Science: Automated hypothesis generation, experimental design optimisation, and pattern recognition in complex scientific datasets across varied domains including materials science, drug discovery, synthetic biology, astronomy, and more.
- » Space: Navigation tools, AR/VR technologies for training, detection methods for satellite anomalies.

This agenda requires effective cross-disciplinary and inter-organisational collaboration. This accelerates foundational research while ensuring Europe can deploy ground-breaking ADR to address core industrial and societal challenges. Creating ADR systems for this ambitious agenda requires progress in three key areas:

- » Fundamental research integrating deployment challenges: Theory, methods, applications, software, and hardware that combine cutting-edge innovation with understanding of real-world ADR requirements.
- » Core ADR capabilities: Enhanced task performance through advances in learning, modeling, inference, and reasoning. This includes progress in technical clusters (computer

vision, natural language processing, generative AI) and cross-cutting capabilities (multi-modal data techniques, causal AI, spatio-temporal modeling).

- » Deployment-inspired research: Technical domain advances that deliver real-world benefits, such as open foundation models deployable for EU Missions.

Across these areas, connecting those working at the forefront of AI research with the needs of organisations using AI is crucial to creating AI tools that can operate effectively in deployment. These connections can be fostered through vibrant local innovation ecosystems that bring together researchers, businesses, and civil society to identify: opportunities to deploy AI for economic and social benefit; the capabilities required for deployment across industry sectors; and the safeguards necessary to ensure its safe use. The result should be advanced analytical techniques, which can be combined with deployable methods and tools that are aligned with European values through a human-centric approach to technology development.

Delivering this 'Big Ticket' is foundational to the success of Adra's roadmap. It underpins the Big Ticket items that follow in the subsequent chapters and is a precursor to realising the proposed socio-economic benefits of ADR.

Roadmap

The table below connects current European ADR technological foundations to future objectives, illustrating how today's actions map onto medium-term goals and long-term outcomes described in this document and EU policy briefs.³⁰

In the near-term, progress in developing new learning strategies, core capabilities and the technical aspects of trustworthiness can help deliver next generation ADR systems that can be deployed across European industry. Such systems should be trustworthy-by-design, deployable under dynamic conditions, and able to operate alongside human users. Across these areas of research interest, the technological foundations of ADR and the capabilities needed to solve industrial challenges are closely integrated.

Foundational research is inherently unpredictable. Innovation requires space for curiosity and novelty, enabling fresh perspectives and new techniques to emerge while allowing successful approaches to scale safely into widespread deployment. The complexity of ADR systems, their deployment environments, and rapid field evolution make accurate predictions about successful technologies, processes, or systems unlikely. Traditional mechanisms for assessing technical maturity, such as TRL methods, struggle to capture the readiness of diverse ADR methods and tools currently available or in-development. Methods requiring significant development in one sector may be near deployment in another, while technology combinations enable rapid progress in new areas.

Instead, a successful innovation strategy would combine flexible funding mechanisms, strong industry-academia links, and agile regulatory approaches. These interventions enable European ADR R&D to pursue novel directions and respond to emerging opportunities. These ground-breaking technological foundations ensure Europe has the research and deployment capabilities needed to realise AI's social and economic benefits.

30. See also: <https://www.vision4ai.eu/sra/>.

Short-, medium- and long-term objectives

Areas for progress today/ near term	Medium-term goals	Long-term outcomes
<p>Core technical capabilities to increase performance in deployment:</p> <p><i>AI model development and architecture:</i></p> <ul style="list-style-type: none"> » Foundation models and fine-tuning » Computer vision and video generation » Natural language processing and speech » Multi-modal learning systems » Causal AI and spatio-temporal modeling » Multi-agent collaboration systems » Prompt engineering and in-context learning <p><i>Data management:</i></p> <ul style="list-style-type: none"> » Data preparation and validation » Data governance frameworks ensuring quality, lineage, and provenance » Physics-aware modeling and knowledge representations » Shared representations and uncertainty estimation <p><i>Deployment and operations:</i></p> <ul style="list-style-type: none"> » MLOps, model lifecycle management, and software engineering » AI at the edge and edge AI optimisation » AutoML and AutoAI 	<p>Trustworthiness by design AI foundations across industrial applications that demonstrate the fundamentals of trustworthy AI:</p> <ul style="list-style-type: none"> » Safety » Robustness to change or adversarial manipulation » Explainability and auditability » Privacy-preservation and security » Fairness » Sustainability and energy-efficient implementations. <p>Deployable under dynamic conditions ADR systems that can be integrated into industrial systems to deliver beneficial outcomes in real-world applications, with the ability to:</p> <ul style="list-style-type: none"> » Analyse real-world data, learning from multi-modal data sources, supported by dynamic or active learning strategies and the ability to learn from a few examples. » Maintain robustness under dynamic conditions, implementing AutoML/AI methods to maintain verifiable standards of performance. » Deliver on expectations of privacy and security in deployment. » Provide uncertainty estimation and confidence in outputs. » Simulate complex system operations using multi-granular data types. » Run on devices at the point where data is collected, enabling AI-at-the-edge. » Enable embodied AI and robotics systems that can interact intelligently with their environment in industrial settings, including human-robot collaboration frameworks, robot learning from demonstration, and coordinated multi-robot operations. 	<p>Technical excellence: Innovations in ADR theories, methods and use to create technically advanced systems that deliver high-performance in deployment.</p> <p>Real world deployment: Safe and effective ADR that can be deployed in real-world contexts, enabling applications across the private and public sectors.</p> <p>Societal benefit: ADR technologies, processes, and systems deployed to deliver societal benefit, accelerating progress towards the goals of the EU's Innovation Missions and enhancing productivity across sectors.</p>

Areas for progress today/ near term	Medium-term goals	Long-term outcomes
<ul style="list-style-type: none"> » Model optimisation in resource-constrained environments <p><i>Physical systems integration:</i></p> <ul style="list-style-type: none"> » Simulation, emulation, and digital twins » Autonomous vehicles and robotics » Next-generation hardware <p><i>Development Support:</i></p> <ul style="list-style-type: none"> » AI-assisted software development <p>Trustworthiness by design:</p> <ul style="list-style-type: none"> » Robustness » Reliability » Resilience » Safety » Security » Explainable AI » User interface design » Human-centric AI » Privacy-enhancing technologies » Auditable AI » Verification, validation, and certification, (including the development of performance guarantees, standards and specifications to support deployment) » Adversarial learning » Generative AI content validation and quality assurance 	<ul style="list-style-type: none"> » collaboration frameworks, robot learning from demonstration, and coordinated multi-robot operations. » Integrate foundation models adapted for industrial domains that can understand technical terms and provide natural language interfaces for human users to work with complex systems while maintaining domain expertise and safety standards. <p>Next-generation methods</p> <p>Improved performance of ADR systems, underpinned by innovations in theory and modelling that reflect the needs of real-world environments, including:</p> <ol style="list-style-type: none"> 1. Principles and techniques for AI that reliably performs well in deployed contexts. 2. Improved performance in AI functions with applications across industry sectors, such as systems for general purpose natural language understanding and generation, computer vision, foundation models and spatio-temporal data modelling for anomaly and extreme event detection. 3. Multi-modal AI systems that combine different data types, reflecting the operations of real-world systems. 4. Approaches to bridging data-driven and domain knowledge, including: the ability to simulate or emulate physical systems; hybridisation with physical and expert knowledge (neuro-symbolic); causal AI methods that discover physical processes from large data volumes. 5. Better understanding of intelligent behaviour and how intelligent agents build models of the world, leveraging this to increase performance when integrating AI into industrial applications. 	

Areas for progress today/ near term	Medium-term goals	Long-term outcomes
<ul style="list-style-type: none"> » Energy efficiency and sustainable AI methods <p>Learning strategies and methods to build next-generation AI models:</p> <p><i>Adaptive learning</i></p> <ul style="list-style-type: none"> » active learning and dynamic feedback » transfer learning and » few-shot learning » continual learning in dynamic environments » federated learning for privacy-preserving distributed AI <p><i>Advanced architectures</i></p> <ul style="list-style-type: none"> » deep learning » reinforcement learning » multi-modal learning systems » AI agent architectures » advanced foundation models, including vision-language models, scientific foundation models, time-series models, and domain-specific foundation models with enhanced reasoning capabilities. <p><i>Specialised techniques</i></p> <ul style="list-style-type: none"> » causal inference » geometric learning » quantum machine learning » generative AI systems » human-AI co-operation and shared representations 	<p>6. Large-scale foundation models for industrial applications that combine 'reasoning' capabilities with domain-specific knowledge, enabling few-shot adaptation to new industrial contexts and natural language interaction with technical systems.</p> <p>6. Agent architectures that can decompose complex industrial tasks and coordinate multi-step operations within safety constraints while maintaining human oversight.</p> <p>8. AI agent frameworks with planning capabilities, tool integration, and autonomous task execution within industrial safety constraints and regulatory compliance.</p> <p>Collaborative AI</p> <ul style="list-style-type: none"> » AI decision-support tools that work alongside users in deployment to enhance human activities, through: » Collaborative multi-agent systems and AI co-pilots that work alongside human operators, providing contextual assistance, anticipating needs, and enabling natural language control of complex industrial processes » Active learning strategies and knowledge representations that enable learning in dynamic industrial contexts and facilitate human-agent teaming. This might include human-AI teaming designed around foundation model capabilities, where AI agents can understand instructions, provide explanations, and adapt their behavior based on human preferences » Conversational AI interfaces that enable technical discussions, explain system behavior, and support decision-making 	

Areas for progress today/ near term	Medium-term goals	Long-term outcomes
<ul style="list-style-type: none"> » end-to-end learning and multi-objective optimisation » leveraging AI in sovereign data spaces » Data Spaces and Data Labs for trustworthy AI 	<p>Societal and industrial applications Use cases that showcase the benefits that well-designed and integrated AI can deliver in areas of social interest and need.</p>	

3.2.2 Big ticket #2: Effective and Trustworthy General Purpose ADR

According to a draft of the AI Act: “*general purpose AI (GPAI)* system means an AI system that can be used in and adapted to a wide range of applications for which it was not intentionally and specifically designed.” This second Big Ticket is focused on GPAI, and the extension to General Purpose ADR, which includes most generative AI, but also goes beyond it. It does not go as far as artificial general intelligence (AGI), which is currently beyond reach. In the context of Adra, this is broadened to Trustworthy General Purpose ADR, which goes beyond the normal software nature of AI to include robotics and other hardware, as well as the data needed to achieve general purpose functionality.

State of the art

Substantial progress has been made in the last decade towards developing methods that are capable of addressing a wide range of applications with minimal manual input, and this progress keeps accelerating. Below we include some of the current developments in a super-fast-moving area (in what follows, (*) denotes information extracted from the 2025 Stanford AI Index Report³¹).

Since the publication of ChatGPT (end 2022), all BigTech players have published their own generative AI systems. These systems cannot be simply called LLMs or chatbots, since almost all of them are multimodal. They can answer questions in natural language, but also generate images, videos, and sometimes even music – though specialised GenAI systems are generally considered a better choice for images or music. More than 40 highly visible LLMs (*) were published in 2024. The most visible ones remain GPT4 (OpenAI), Llama 4 (Meta), Gemini (Google DeepMind, that has replaced the PaLM family and is attached to Google products), Claude 4 (Anthropic), Le Chat 2 (Mistral), and the recent DeepSeek v3 (DeepSeek). Another frequent feature is that most of these chatbots are now connected to the Internet and hence tend to replace standard search engines, at a significant environmental cost. In a quest to mitigate that, all models also provide a “small” version (a few billions weights “only”) that is de facto faster and cheaper to train, at the cost of a little loss in performance – though the performance of the recent small

31. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

models matches that of the best performing large models of 2023. Last but not least, most offer free access limited to a few queries per day/week and a commercial version with unlimited access.

The differences between the different models come from the training data (selected data or raw data fetched on the Web, with IP rights and reliability issues), and the post-tuning techniques that turn an LLM into a chatbot (including, but not limited to, Reinforcement Learning with Human Feedback). Furthermore, some systems receive specific post-training to better handle specific tasks, like the o1 and o3 ChatGPT4 by OpenAI that has better reasoning and mathematical capabilities than the base model, at a significant cost: six times more expensive to train and 30 times slower to answer (*).

The situation is similar in regards to image generation. As previously stated, most chatbots are now multimodal Gemini and ChatGPT appear to be the best options in this family, according to Chatbot Vision Arena's ranking. But good "old" specialised software like Stable Diffusion 3.0, Dalle-E 3 and Adobe Firefly (the old champion Midjourney is now only available by subscription), together with a few newcomers like Freepik, Dream Studio and Ideogram,³² offer more options and contexts—note that they all now use Denoising Diffusion Probabilistic Models as the main image generator, coupled with different conditioners.

Even more significant progress has been made in video generation in 2024, with systems like SORA (OpenAI), Stable Video 3D and 4D, Movie Gen (Meta), and Veo 2 (Google DeepMind), that can now produce high-resolution realistic videos.

This is also true in music generation, with Suno from SunoAI considered the most complete and impressive platform. LLMs and generative models are not general purpose by default. The models cited above cannot query application-specific data and APIs, and do not necessarily address robotics. However, some initial steps have been made in this direction, and more and more foundation models for robotics have recently emerged: pioneered by LangChain, RT-2, PaLM-E, Open-X-Robotic embodiment and ActGPT, the trend has been reinforced in 2024 by *GR00T (Generalist Robot 00 Technology)*, a general-purpose foundation model by NVidia for humanoid robots designed to understand natural language and mimic human movements, *RFM-1*, a robotic foundation model by Covariant with language capabilities and real-world maneuverability, and *LLaRA*, developed by researchers at Stony Brook University and the University of Wisconsin-Madison, that integrates perception, communication, and action into a monolithic, end-to-end deep learning model (from *).

As models and approaches become more general, there are new and emerging risks that need consideration and mitigation. This applies to AI in general, but more particularly to GenAI, as all research toward trustworthy AI in the last decade has been devoted to making predictive AI trustworthy, resulting in several algorithms and industrial recommendations in that direction. Unfortunately, those do not apply to GenAI, which maintains the following risks:

- » Loss of diversity when the same few datasets are used by all models, resulting in similar biases, and uniformisation of "thought models". Also, lower-resourced languages and dialects get less attention (if any) in LLMs. This issue is becoming even more critical as the amount of available data that has not yet been used vanishes – and/or as most of the available data in the Web becomes unreliable, probably generated by other GenAI systems!

32. See a comparison here <https://www.qitoolssme.com/comparison/image-generators>

- » **Static knowledge:** Most data in domains evolves over time. However, because learning from scratch is not sustainable, most “data-based only” (i.e. not continuously learning) GPTI systems are at risk of becoming rapidly obsolete.
- » **Hallucinations:** Even the best-performing chatbots keep hallucinating. Some of them (e.g., Claude) are, however, able to “confess” their hallucinations when some precisions are asked – and connection to the Internet allows them to factcheck themselves, albeit at an ever increasing environmental cost. Nevertheless, users should always independently check their answers in case of the slightest doubt.
- » **Lack of elementary reasoning,** which does not mention more elaborate forms of reasoning, such as statistical inference. This is a roadblock towards certification for use in medical circumstances or code generation. However, existing models become more and more efficient in reasoning, mathematical inference, and the like: even though theoretical guarantees seem unreachable in the near future, the “statistical reasoning” power of further generations of LLMs might be sufficient in everyday use.
- » **Reproduction of biases and hateful content** present in the datasets. It is well established by now that large generative models exhibit various kinds of bias and other problematic characteristics, resulting in models that encode stereotypical and derogatory associations along gender, race, ethnicity and disability status (Basta et al. 2019;³³ Kurita et al. 2019;³⁴ Luccioni et al., 2023).³⁵ People in positions of privilege with respect to issues such as racism, misogyny and ableism tend to be overrepresented in training data for LMs (Bender et al. 2021).³⁶ Generative models trained on large, uncurated and static datasets from the web encode hegemonic views that further harm marginalised people. This can, to some extent, be mitigated by integrating Human Feedback using techniques such as Reinforcement Learning (RLHF), Direct Preference Optimization (DPO), or Identity Preference Optimization (IPO), but it is very hard to quantify. Furthermore, the recent evolution of US administration in that respect (resulting, e.g., in Meta abandoning content moderation after X/Twitter) will not help mitigate those risks.

33. <https://aclanthology.org/W19-3805/>

34. <https://arxiv.org/abs/1906.07337>

35. <https://arxiv.org/abs/2303.11408>

36. <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>

- » Safety and Robustness: Similarly, such techniques using Human Feedback are used to increase the safety and robustness of LLMs and prevent inappropriate outputs (such as racist or homophobic content). But these barriers can be circumvented with carefully crafted prompts, either by human users taking advantage of the lack of reasoning (such as guidance on how to obtain instructions to build a bomb), or by appending computed adversarial strings³⁷ to a naive prompt. Models have made a lot of progress on this front, but more remains to be done – and here again, the recent evolution of the US administration does not introduce much optimism.
- » IP: At the moment, training data for LLMs is often a commercial secret, and even when revealed, it is fetched from the internet, regardless of any intellectual property rights and data regulation such as the GDPR. Few AI providers disclose any information about the copyright status of training data. Many foundation models are trained on data that is likely to be copyrighted. The legal validity of training on this data, especially for data with specific licenses, and of reproducing this data, remains unclear (Bommasani, Klyman, Zhang, Liang 2023).³⁸ Even though in principle a text and data mining exception permits the use of copyrighted material for the training of foundation models, this provision does not appear to have resolved the issue in practice. Recent lawsuits³⁹ regarding the use of copyrighted works for training Generative AI systems have underlined the need for legal certainty and clarity about the scope of the TDM exception.
- » Computational and environmental cost: Training all the models discussed above costs an increasing amount of energy, e.g., 10 MW equivalent to 9,000 tons of CO2 for Llama3.5(*). Furthermore, because these chatbots tend to replace search engines in everyday life, even the inference cost has become an issue: one query costs (*) the equivalent of 2-3 g of CO2 (against 0.2 g for a simple Google search), and generating one image costs from 100 g to 1 kg of CO2 depending on the system and the resolution (the giant tech companies running these chatbots have asked users to avoid writing “thank you” at the end of a session, to avoid an N+1-th reply from the system!). Despite best efforts, the majority of cloud compute providers’ energy is not sourced from renewable sources. Data centres with increasing computation requirements take away from other potential uses of green energy, underscoring the need for energy efficient model architectures and training paradigms.⁴⁰

Challenges

One of the main challenges is how to make the models and approaches more general while improving both effectiveness and trustworthiness. Another important obstacle is the computational cost of GenAI. And an especially interesting challenge is how to extend software models to deal with physical interactions and robotics.

Which training data? The data used for training is a source of concern: whereas the past increase in performance was at least partly due to the increase in training data, it is

37. <https://llm-attacks.org/>

38. <https://crfm.stanford.edu/2023/06/15/eu-ai-act.html>

39. https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/?guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xILmNvbS8&guce_referrer_sig=AQAAAE7FmNqEnZMTnfMqHtV8dq84sXKEadgyGvweOJcg6xcAjDCtjxkWmODAHIjldNmTsC-GjYs8KyTTBO--rW2R_rAKHzvXVv3cKGJolKWmGXQVJVs8GWPB-O_Y0cCkfVFzyNZM2BJEuW2uanUe7XbcAxclPHbPy35Fpl59jBaXkZQ&_guce_consent_skip=1699221271

40. <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>

assumed today that all available data world-wide will have been used in 2028 (*). Because data augmentation is not a solution, new means will have to be considered to continue to improve the performance of existing models—or new models will need to be designed.

Data access and monetisation: Data access for AI training requires a clear interpretation of the GDPR and the TDM exception. Artists must be able to opt out of the use of their creations for AI training, and there is a shared belief in Europe that they must be financially compensated if they choose not to opt out. The implementation of universal opt-out and compensation mechanisms in websites is a major challenge. This implies defining a suitable economic model, which quantifies the price of data depending on its type, quality and usage (pretraining, fine-tuning, or testing).

Continuous/incremental learning: Continuous learning is mandatory in our dynamic world, with new tasks, concepts and application settings calling for seamless adaptation of AI systems without interrupting the operation of services and applications. Concept drift takes place everywhere, making **de-learning** another growing challenge.

Trusted reasoning and statistical inference capabilities: While building such capabilities inside LLMs might be difficult, post-processing and LangChain could be the solution in many use cases (such as computational social science).⁴¹ Such reasoning capabilities would also **increase the safety and robustness** of GPAI systems, as would access to fact-checking techniques that require browsing capabilities. Recent progress has been made through pre-prompt optimisation, but this requires expertise in the target problem domain.

Evaluation and auditing: As General Purpose ADR continues to develop and adapt to applications for which it was not intentionally and specifically designed, new capabilities and risks may emerge. Guidance, benchmarks and metrics need to be created for evaluating and auditing GPAI systems and addressing limitations in assessment and validation methodologies. The performance of LLMs on existing benchmarks like MMLU, GSM8K, and HumanEval, and even some augmented versions like MMMU and GPQA, is close to reaching saturation. New benchmarks have been proposed on which the human performance seems out of reach of existing LLMs, like Humanity's Last Exam (8.80%), FrontierMath (2%), and BigCodeBench (35.5%) whereas humans reach success rates above 95% (*). Another approach is to rely on the "Wisdom of Crowds": the Chatbot Arena Leaderboard fetches public votes of pairwise comparisons of anonymous LLMs, that are then ranked according to an ELO process (taking inspiration from chess ELO). More than 1 million votes have been fetched, and Gemini has taken the top position. A large number of votes contributes to preventing biases and ensuring fairness of the result, but the sample of voters is by definition biased toward "geek-like" users, and sociological studies must take that into account.

Biases and non-discrimination: The detection and correction of biases is a hot topic in AI research and is not specific to GPAI. The tracing of data sources is one aspect of bias detection, while post hoc detection is another necessary approach. However, there are different and contradictory definitions of biases, and the metrics used depend on the goal (such as individual biases vs group biases). Beyond statistical measures of bias, an increasingly challenging aspect is representational bias and its potential harms, especially

41. <https://arxiv.org/abs/2306.04746>

in the realm of visual and text data (Blodgett et al., 2020;⁴² Fabbrizzi et al., 2022)⁴³.

Scale and complexity of the value chain: One of the specific challenges of General Purpose ADR is the potential scale of adoption and the resulting challenges for responsible development and deployment of the technology, but also enforcement of regulatory obligations, ethical guidelines and Terms of Use (Helberger & Diakopoulos, 2023;⁴⁴ Hacker, 2023).⁴⁵ The fact that current regulatory frameworks, such as the AI Act, tend to focus primarily on technology developers and less on creating foreseeable and adequate regulatory frameworks for downstream deployers, risks creating new accountability gaps.

Absence of ethics by design: There is a clear need for an embedded-ethics approach that incorporates reflection on the potential consequences of AI development throughout the whole process. Interdisciplinary stakeholder engagement in question-articulation sessions and value-sensitive design approaches, such as Values that Matter (Smits, Bredie, van Goor, Verbeek 2019)⁴⁶, should be incorporated across the full lifecycle of GPAI development, not as a one-off, ex-post exercise. Moreover, there is a need for concrete methodologies to enable participatory, inclusive, and interdisciplinary forms of engagement involving relevant stakeholders, including interdisciplinary experts and marginalised communities.

Transparency requirements for developers and deployers: GPAI developers and deployers should be required to report on the provenance and curation of training data, the model's performance metrics, and any incidents and mitigation strategies concerning harmful content (Hacker, Engel and Mauer, 2023).⁴⁷ Model cards (Mitchell et al., 2019)⁴⁸ and datasheets for datasets (Gebru et al., 2021)⁴⁹ offer solid blueprints for such transparency reporting. Transparency obligations should be specific and operationalisable (i.e. machine-readable format and easy-to-use opt-out) (AlgorithmWatch 2023).⁵⁰

Risk of manipulation and AI anthropomorphism: A recent case of chatbot-incited suicide in Belgium⁵¹ brought to the fore yet another risk of general purpose AI: the risk of manipulation (Smuha, De Ketelaere, Coeckelbergh, Dewitte & Poulet, 2023).⁵² The phenomenon of AI anthropomorphism, understood as projecting human-like qualities or behaviour to non-human entities such as AI systems (Tan, 2023),⁵³ risks deceiving people into assuming that AI systems make ethical or moral judgements.

Explainability: To enhance model transparency and trustworthiness, GPAI models should be able to generate explanations for their outputs. However, it is a challenge to generate explanations that are accurate and informative without compromising the accuracy of

42. <https://www.microsoft.com/en-us/research/publication/language-technology-is-power-a-critical-survey-of-bias-in-nlp/>.

43. <https://www.sciencedirect.com/science/article/abs/pii/S1077314222001308>.

44. <https://policyreview.info/essay/chatgpt-and-ai-act>

45. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4467684

46. <https://research.utwente.nl/en/publications/values-that-matter-mediation-theory-and-design-for-values>

47. <https://arxiv.org/ftp/arxiv/papers/2302/2302.02337.pdf>

48. <https://dl.acm.org/doi/abs/10.1145/3287560.3287596>

49. <https://dl.acm.org/doi/fullHtml/10.1145/3458723>

50. https://algorithmwatch.org/en/wp-content/uploads/2023/09/PolicyBrief_GPAI_AW-3.pdf

51. <https://www.brusselstimes.com/belgium/430098/belgian-man-commits-suicide-following-exchanges-with-chatgpt>

52. <https://www.law.kuleuven.be/ai-summer-school/open-brief/open-letter-manipulative-ai>

53. <https://medium.com/human-centered-ai/on-ai-anthropomorphism-abff4cecc5ae>

the model (i.e. without introducing an accuracy-explainability trade off), and without introducing significant computational overhead.

This section has mainly focused on the risks and challenges, but there are also huge opportunities to be realised.

Important use-cases & applications

Internal dissemination and preservation of companies' knowledge: The ease of fine-tuning for specific tasks (see above) allows companies to generate their own LLMs to store and query all knowledge that builds up the company's know-how and industrial secrets.

Domain-agnostic scientific discovery and engineering: As opposed to separate, domain-specific LLMs for scientific discovery (such as biology, material science, math and social sciences), innovation (including patent search and industrial process optimisation) and production (such as industrial robot control and logistics), it is important to design GPAI models that generalise in other scientific fields and industries (such as supporting human engineers), and speed up discovery and time-to-market, thanks to their holistic understanding of the entire R&I chain.

Population-level digital twins of patients: Training a federated LLM on various data sources across clinical centres (such as medical papers, electronic health records, image/sensor data, emergency calls and environmental data) would result in a digital twin of the patient population. By infusing it with trusted statistical inference capabilities, this LLM would be the basis for a GPAI system that enables the discovery of new relationships between factors and symptoms, and provides better support to patients, clinicians and staff across the entire patient journey, creating a foundation for AI-based surgery.

AI for creative processes: AI must support the creative processes involved in producing text, articles, experiences, videos, and images, as well as experiences in the metaverse. This requires exploring how humans can collaborate with those systems, how we make sure things generated are bias-free, how we can ensure content is monetised, and how to interpret IPR and copyrights in these scenarios is also important.⁵⁴

Socio-economic impact

According to McKinsey's report "The economic potential of generative AI: The next productivity frontier"⁵⁵ generative AI could add the equivalent of \$2.6 trillion to \$4.4 trillion annually across the 63 use cases analysed—by comparison, the United Kingdom's entire GDP in 2021 was \$3.1 trillion.

According to Goldman Sachs,⁵⁶ "generative AI ... could drive a 7% (or almost \$7 trillion) increase in global GDP and lift productivity growth by 1.5 percentage points over a 10-year period."

However, these highly optimistic economic estimations also shyly point out the significant

54. See also the recent Journalism AI report on further uses of GPAI: <https://www.journalismai.info/research/2023-generating-change>

55. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/>

56. <https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html>

disruption this will cause in the job market. Such issues had already been identified for machine learning and AI in general, but the changes caused by generative AI will likely be more chaotic. Even if the global balance between jobs destroyed and jobs created remains controversial, there is a consensus about the transformations that will take place. At a minimum, workers will need to change their working habits, and at most, they may have to completely change jobs or acquire entirely new skills.⁵⁷ It is far from certain that the necessary mechanisms for continuous education will be put in place, nor that the mindsets of affected workers will allow for such drastic changes to occur as smoothly as professional report-writers may believe.

Beyond the impact on skills and employment, generative AI is expected to affect the perception and cognitive processing of online information by citizens. Increasing citizens' awareness about the ability of generative AI to produce hyper-realistic multimedia content will make the "Liar's Dividend" an important challenge in relation to trust in media and public discourse: politicians, influencers and a variety of online actors will be able to denounce actual documentary material such as witness videos as being generated by AI, which may give rise to conspiracy theories, erosion of citizen trust in the state and media, and exacerbation of information disorder on the internet. To this end, new and easily accessible means of verifying and proving the authenticity of online media are becoming indispensable.

At the same time, generative AI could help nurture e-learning and match jobs with suitable applicant profiles (even before they apply), although this is unlikely to outweigh its negative effects.

New methods are required to assess AI's individual and societal impact (both positive and negative) in a way that accounts for the lived experiences of diverse groups of users and non-users, including marginalised groups. Many forms of impact assessments are driven by top-down expert evaluations, but seeing the broad proliferation of GPAI and the fact that it is democratising access to this powerful technology underlines the need for a better understanding of experiences, concerns and mitigation strategies on the ground.

57. <https://www.mygreatlearning.com/blog/how-will-artificial-intelligence-create-more-jobs/>

Short-, medium- and long-term objectives

Short term	Medium term	Long term
<p>Basic trustworthy data management:</p> <ul style="list-style-type: none"> » GDPR, IP, and source tracking on data sets » Larger and larger data sets, reaching 1000T tokens (Exa-tokens) <p>Basic trustworthy generative models:</p> <ul style="list-style-type: none"> » Transparency on data sets and model training » Explainability of individual queries » Static and dynamic factual correctness in unimodal models, mainly through retrieval augmented generation » PhD-level expertise needed <p>Basic reasoning in generative models:</p> <ul style="list-style-type: none"> » Implicit reasoning learned through training data in uni- and multi-modal models » Limited explicit reasoning mainly through retrieval augmented generation <p>Basic effective generative models:</p> <ul style="list-style-type: none"> » Distilling large models to comparable models on single GPU desktop machine 	<p>Trustworthy data management:</p> <ul style="list-style-type: none"> » GDPR, IP, and source tracking on individual queries » Shrinking data set size, drastically improving quality, and filtering out the essential data <p>Trustworthy generative models:</p> <ul style="list-style-type: none"> » Transparency on individual queries » Detailed and adaptable explainability on multiple abstractions levels » Static and dynamic factual correctness in multi-modal models » MSc-level expertise needed <p>Reasoning in generative models:</p> <ul style="list-style-type: none"> » Physical reasoning integrated in image and video generation models » Scientific reasoning in text generation models <p>Effective generative models:</p> <ul style="list-style-type: none"> » Exa-scale models in data centres » Fine-tuning large models on single GPU desktop machine » Distilling large models to comparable models on mobile device » Large scale models continually trained on distributed edge devices 	<p>Sophisticated trustworthy data management:</p> <ul style="list-style-type: none"> » Highly curated and verified minimal data sets with explicit bias control, cultural contextualisation, etc. <p>Sophisticated trustworthy generative models:</p> <ul style="list-style-type: none"> » Physically and scientifically correct multi-modal models » Generative models will be as easy to create as a compiler today <p>Sophisticated reasoning in generative models:</p> <ul style="list-style-type: none"> » Physical and scientific reasoning in multi-modal generative models <p>Highly effective generative models:</p> <ul style="list-style-type: none"> » Exa-scale models continually trained on distributed edge devices

3.2.3 Big ticket #3: An interoperable and integrated framework for data and model ecosystems

State of the Art

The main goal of the European Data Strategy⁵⁸ published in 2020 was to create a single market for data that will ensure Europe's global competitiveness and data sovereignty. This strategy introduced the European Common Data Spaces, aimed to connect stakeholders on a cross-border and cross-sector basis and share data in a trusted, secure and scalable way. The implementation of these data spaces, funded under the Digital Europe Programme, started with CSA preparatory actions in eight domains, and continues with deployment actions aimed at achieving a certain level of maturity.⁵⁹ The Data Spaces Support Centre coordinates all these actions on the sectoral data spaces to ensure that they develop in a coherent way, are interoperable, and benefit from economies of scale.⁶⁰ The SIMPL Programme⁶¹ aims to provide open source, modular middleware to implement data spaces in a secure and interoperable way. Additionally, the Horizon Europe Programme is funding research and innovation projects to develop techniques, solutions and components that contribute to the design and implementation of data ecosystems, including (i) technologies and solutions for compliance, privacy preservation, green and responsible data operations,⁶² (ii) technologies for data management,⁶³ (iii) solutions for data trading, monetising, exchange and interoperability,⁶⁴ (iv) integration of data life cycle, architectures and standards for complex data cycles,⁶⁵ or (v) AI-driven data operations and compliance technologies.⁶⁶

The AI Innovation Package⁶⁷ presented by the European Commission in July 2024 introduces the concept of AI Factories: open ecosystems built around EU supercomputers, specially tailored to build and train large AI models, that provide AI and data services and have access to the massive amount of high-quality industrial data required to train those models. The recently presented AI Continent Action Plan⁶⁸ complement these AI Factories with Data Labs, that will be further explained in the upcoming European Data Union strategy⁶⁹ and are expected to ensure that AI developers will have access to large volumes of high-quality data in relevant sectors.

58. <https://digital-strategy.ec.europa.eu/en/policies/strategy-data>

59. <https://digital-strategy.ec.europa.eu/en/library/second-staff-working-document-data-spaces>

60. <https://dssc.eu/>

61. <https://simpl-programme.ec.europa.eu/>

62. https://cordis.europa.eu/programme/id/HORIZON_HORIZON-CL4-2021-DATA-01-01

63. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2021-data-01-03>

64. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2022-data-01-04>

65. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2023-data-01-02>

66. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2024-data-01-01>

67. <https://digital-strategy.ec.europa.eu/en/factpages/ai-innovation-package>

68. <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>

69. <https://digital-strategy.ec.europa.eu/en/news/commission-seeks-views-use-data-develop-ai-intelligence>

The regulatory framework surrounding data is gradually being established. The EU Data Act,⁷⁰ which entered into force on 11th January 2024, aims to establish harmonised rules on fair access to and use of data across the European Union. The EU Data Governance Act⁷¹ (DGA) is in force and aims to facilitate data sharing across sectors and EU countries, enhancing trust in data sharing by establishing rules for neutrality of data intermediaries and promoting data altruism. The European Health Data Space (EHDS) Regulation⁷² was published in the Official Journal on 5 March 2025, and entered into force on 26 March 2025, initiating a phased implementation process over the coming years. This regulation is intended to make a significant step in the EU's digital health strategy. Finally, the EU AI Act⁷³ is in effect and progressing through a phased implementation schedule. This provides a risk-based approach to classifying AI systems.

Finally, significant standardisation efforts are in place to provide frameworks that support the interoperability, trustworthiness and scalability of European data ecosystems. Particularly, the CEN-CENELEC Joint Technical Committee JTC25 "Data, dataspaces, cloud and edge"⁷⁴ is working on the production of harmonised standards in response to the standardisation request from the European Commission regarding the implementation of the Data Act.

Within this landscape, the following trends are shaping the foundations for the next generation of complex data driven ecosystems:

- » Data ecosystems are increasingly becoming more complex, as value creation is incorporated on top of pure data transactions, with added value applications on top of data sharing and connection with other ecosystems to generate value out of data. This implies the involvement of new actors and new types of services.
- » Focus should be given to deployment of disruptive applications, specially AI-driven applications,⁷⁵ that can benefit from data ecosystems. This presents the challenge of selecting and preparing data under the requirements of these cutting-edge applications.
- » Interoperability is increasing in complexity across different levels, including technical, security, legal, governance and business. The application of novel AI-driven techniques to manage complexity can contribute to alleviating interoperability. Special attention must be given to the application of Generative AI to support semantic interoperability and other data sharing features,⁷⁶ and to AI Agents to support data ecosystems management.

70. <https://digital-strategy.ec.europa.eu/en/policies/data-act>

71. <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>

72. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en

73. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

74. https://standards.cencenelec.eu/dyn/www/f?p=205:7:0:::FSP_ORG_ID:3485479&cs=IEF27AE97B5DBDA9B990D3DAF8BD63366

75. https://dssc.eu/space/DC/28049509/Strategic+Stakeholder+Forum?attachment=%2Fdownload%2Fattachments%2F28049509%2FDSSC%2520White%2520paper_Generative%2520AI%2520and%2520Data%2520Spaces_V4.pdf&type=application%2Fpdf&filename=DSSC%20White%20paper_Generative%20AI%20and%20Data%20Spaces_V4.pdf

76. "Foundation Data Space Models: Bridging the Artificial Intelligence and Data Ecosystems", 2023, Ed Curry et al (https://www.researchgate.net/publication/377622717_Foundation_Data_Space_Models_Bridging_the_Artificial_Intelligence_and_Data_Ecosystems_Vision_Paper)

- » Robust governance, performance and security in complex ecosystems should cover the entire data lifecycle across the different stages. Ensuring trust in data transactions (Trusted Data Transactions) requires traceability, data lineage, provenance, and real-time observability of data flows (data observability).
- » High Performance Data Analytics (HPDA), represents the convergence of big data, AI, and high-performance computing. It enables real-time analytics on massive, distributed datasets.

Challenges

Operational challenges

This group of challenges cuts across those that affect or impact on the usual operations with data. The most significant among them include:

Availability and access to data. The EU produces a massive amount of data, and this volume continues to grow rapidly due to several factors. However, a large portion of this data, particularly from the private sector, remains inaccessible due to legal, technical, and economic barriers and lack of incentives for sharing. This limited availability of high-quality data (and specifically data that is specially prepared for AI) poses a major obstacle for AI-driven innovation.

Data for AI, complementing the previous point, focuses on identifying the most suitable data for AI systems and tailoring it to meet the specific constraints under which these models operate, such as bias, fairness, and other ethical or technical considerations.

Data management in complex ecosystems, or how to ensure, among others, end-to-end accuracy, consistency, and traceability including management of quality, metadata, and life-cycle management, and monitoring, auditing and transparency.

Integration of complex data analysis pipelines and AI workflows with the computing continuum (IoT/edge/cloud) that leverage the different instruments, infrastructures and technologies developed in the EU in a seamless, trusted and scalable manner. This includes the synchronisation and consistency of replicated or distributed data (e.g. data on the edge and the cloud).

Energy efficiency involves improving the sustainability of data-driven processes by minimising storage needs, reducing computational loads through techniques such as data sparsity, and promoting the use of frugal, energy-efficient AI models. This also includes efficiency of HW / SW for AI, and new tools for developing novel systems that are optimised for specific purposes.

Security and protection in data ecosystems

Security and protection in complex data ecosystems present challenges that stem from various factors and require comprehensive strategies to effectively mitigate risks. This is due to the complexity of these ecosystems, which involve the interplay of many different actors, and the uprising of new AI technologies that can provide novel instruments for intentional attacks.⁷⁷ Some of the key security and privacy challenges include:

Data breach and unauthorised access: The risk of data breaches is ever-present, with cyber criminals seeking to exploit vulnerabilities and gain unauthorised access. Furthermore, inadequately secured databases, servers and cloud environments can lead to unauthorised data exposure. Adversarial attacks are a serious threat to the security and reliability of AI

77. <https://bdva.eu/news/bdva-publishes-position-paper-latest-trends-data-protection/>..

systems, highlighting the need for robust defenses and continued research in this field.

Data usage control and policy enforcement: Once data is shared, enforcing data usage constraints such as purpose limitation, retention duration, and geographic restrictions becomes non-trivial due to several factors such as lack of runtime usage enforcement mechanism, decentralised heterogeneous execution environment, and absence of policy binding and policy carrying data.

Data encryption and protection: Managing encryption keys and securing encryption key storage is highly challenging; encryption key management is particularly challenging in multi-party data sharing if a single key is shared among a group of participants.

Data anonymisation and re-identification: Anonymising or de-identifying data for analysis while maintaining data utility is difficult to accomplish. Re-identification attacks can potentially link anonymised data back to individuals. Moreover, anonymisation is typically a one-time static process; it does not enable dynamic, runtime enforcement of usage policies.

High cost: Advanced cryptographic techniques used in privacy-enhancing technologies (such as homomorphic encryption) are often computationally intensive and resource-demanding. This creates a necessary trade-off between the robustness and security of the system and its financial and computational cost, requiring careful decision-making based on use case priorities

Connection and interoperability of ecosystems and frameworks

Achieving interoperability across multiple dimensions—technical, syntactic, semantic, organisational, and legal – remains a formidable challenge. This is particularly difficult in cross-domain data sharing ecosystems, where diverse standards, data models, regulatory frameworks, legal instruments, and operational practices must be reconciled to enable seamless and trusted data exchange. Addressing these dimensions in a coherent manner is critical yet complex, especially when dealing with decentralised infrastructures.

Governance and sustainability

Data ecosystems are becoming increasingly complex. To succeed, they need governance mechanisms that can be shared among multiple actors and that ultimately result in scalable and sustainable models. This includes the governance of the data ecosystem or data space as a whole, but also data governance, data sovereignty, provenance and traceability of data, and data quality management. The automation of contract negotiations and its enforcement within data ecosystems remains an open issue. These aspects should converge in sustainability models with instruments for scaling up, including novel tools for measuring data maturity (data maturity models) aimed at monitoring the ecosystem's progress.

Important solution areas and applications

Foster interoperability in complex data ecosystems

This solution addresses multiple layers of interoperability (business, governance, legal, semantic, syntactic and transport) both within and across complex ecosystems, accounting for the technical and non-technical mechanisms required to enable

seamless data exchange and collaboration among diverse stakeholders. The solution should consider the application of AI techniques, such as Large Language Models (LLMs), to support semantic interoperability across heterogeneous domains.⁷⁸ It should also consider the interoperability of different types of data ecosystems (i.e. IoT ecosystems) by applying standards and taxonomies to facilitate automatic discovery, integration, and processing of IoT and edge computing resources. Furthermore, this solution would focus on business interoperability covering alignment of businesses, value models, and operational semantics between organisations. Harmonisation of data governance policies and legal frameworks, compatibility in data formats and encoding structures, and consistent and secure transmission using interoperable communication protocols would also be covered.

Design of global frameworks for trust, identity management and policies enforcement

This solution establishes comprehensive and standardised frameworks that address key aspects related to trustworthiness, identity verification, and policy enforcement within a large-scale data environment. It would address challenges such as governance, security, protection and privacy within the ecosystem, together with mechanisms for the enforcement of policies that ensure these rules are followed throughout the entire data lifecycle.

Developing energy efficient and sustainable solutions for data driven implementations

This solution area aims to reduce the energy consumption and carbon dioxide emissions caused by data processing and storage. Efficient methods for data analytics, such as appropriate data fusion strategies to design efficient energy efficiency systems, frugal/sustainable AI, and edge AI, with less data and energy-intensive algorithms. In the case of efficient edge AI, there should be strategic collaboration between AI researchers and AI chip vendors, particularly with regard to building European foundation models, co-developing reference architectures and investing in shared training infrastructure.

Design and implementation of complex data analysis pipelines

This solution advances the cloud/edge/HPC continuum to support novel data pipelines and AI workflows. It would incorporate techniques like transfer learning, fine-tuning, and capability inheritance to build pipelines that adapt and evolve AI models for specific tasks. This would also include the use of foundation and large language models (LLMs) in downstream applications, the reuse and exchange of AI models via marketplaces, and the application of DataOps practices to manage the complexity of these pipelines, facilitating automation, collaboration across teams, agility in response to changes, and continuous improvement in data and AI operations.

Application of AI Agents to support complex data ecosystems

In complex and distributed data ecosystems, characterised by heterogeneous participants, decentralised infrastructures, and evolving interoperability requirements, AI agents can act as autonomous, intelligent components that enhance the automation capabilities within complex data sharing ecosystems. This challenge would apply AI agents powered by cutting-edge LLMs and RAG models to support multiple functions across interoperability (aligning schemas, ontologies, and taxonomies, facilitating automated data annotation, disambiguation, and cross-domain entity mapping), governance (enforcing data usage

78. https://interoperable-europe.ec.europa.eu/sites/default/files/event/attachment/2025-05/20250410_EC-Report_Webinar_LLM-semantic-interoperability.pdf

policies), and operational aspects (data discovery, extraction, curation, and evaluation and recommendation of datasets, metadata quality, and provenance using context-driven semantic search techniques). The challenge would apply Model Context Protocol (MCP) to integrate LLM applications and external data sources and tools. The challenge will apply AI agents to relevant areas like healthcare (e.g. aligning ontologies such as HL7/FHIR and SNOMED CT, curating longitudinal health datasets for research while preserving anonymity using technologies such as differential privacy), finance (e.g. orchestrating regulatory compliance by integrating data from multiple data sources and applying machine readable policies such as ODRL for cross-border data sharing), energy (e.g. coordinating real-time energy data exchange between distributed energy resources, grid operators, and EV infrastructure using domain ontology such as common information model (CIM) and SAREF), mobility (e.g. in multimodal mobility, acting as an aggregator collecting and integrating data from a wide range of systems such as traffic systems, public transport, and vehicle telemetry), agriculture (e.g. interoperable data sharing among farm equipment, vendors, IoT sensors, metrological data providers, and logistics networks using standards like AgroVoc and logistics networks), and manufacturing (e.g. facilitate the integration of multi-modality data such as tabular ERP data, timeseries sensor data, geometry and CAD design data, to enable smart and autonomous planning and control).

Fostering the adoption and scale-up of data spaces

This challenge would serve as a foundation example and baseline to foster adoption and move towards the scale-up phase in data spaces and other data driven ecosystems. It would implement interoperability between data spaces and existing legacy systems, addressing not only technological aspects but also business models, maturity assessment and conformity to regulations related to data spaces. The outcome of this challenge is expected to include instruments that support the scaling of data spaces across various dimensions and assess their compatibility, with deployment actions to be completed in the following months.

Data spaces and cutting-edge applications

The adoption of data spaces can be fostered by intertwining them with applications that are well suited to, and widely benefit from, the connection. This set of applications would include, among others, AI-driven solutions, virtual worlds, and digital twins. The connection between data spaces and AI, particularly generative AI, is receiving significant attention, though its full potential and practical implementation have yet to be realised.⁷⁹ The process of finding the appropriate data and preparing it specifically for its use by AI practitioners is still not fully integrated in data ecosystems. Data spaces represent an excellent environment for digital twins to find the perfect conditions to grow, scale, and connect the diffuse actors required to unleash their potential. The integration of real-world data with virtual representations can lead to improved decision-making, enhanced operational efficiency, and a deeper understanding of the physical environment. This challenge combines features and represents an optimal symbiosis of both paradigms, linked by data. It will result in relevant applications addressing specific sectors, such as data spaces and digital twins of

79. https://dssc.eu/space/DC/28049509/Strategic+Stakeholder+Forum?attachment=%2Fdownload%2Fattachments%2F28049509%2FDSSC%2520White%2520paper...Generative%2520AI%2520and%2520Data%2520Spaces_V4.pdf&type=application%2Fpdf&filename=DSSC%20White%20paper_Generative%20AI%20and%20Data%20Spaces_V4.pdf

oceans and virtual worlds, and contributions to the Destination Earth initiative.

Socio-economic impact

The main impact resulting from the proposed solution areas and applications can be summarised in the following points:

- » Increase in data and AI democratisation, mostly for small actors; empowering citizens in the use of their personal data and access to cutting edge AI driven solutions; and increasing societal trust for the lawful processing of personal data by governments and use of AI applications.
- » Breaking data silos and leveraging data generated in Europe to train large AI models and AI foundation models.
- » Increasing trust and security in data and AI ecosystems and solutions.
- » Reduction of carbon footprint and energy consumption, and reducing costs by leveraging existing platforms and legacy systems.
- » Extensive application of the digital twin paradigm and other disruptive applications powered by data and data ecosystems, to support better decision making in social, industrial, and environmental domains.
- » Monetisation, reuse, sharing and new added value from the use of AI models and other data related assets; and more economic activities that rely on sharing data, including associated job growth.
- » Better AI and data-driven services.
- » New ways of working and creating successful businesses based on large-scale data spaces.
- » Better awareness of ongoing global events and data for public organisations and industry. For instance, real-time internal market trade information.
- » Smart digital contracts and flexible data sharing for improved efficiency, enabling use cases for societal resilience, e.g. flexible energy systems.
- » Ethical use of data because of efficient data governance practices.

Objectives and roadmap

The table below identifies medium- and long-term objectives for each of the solution areas and applications described in the previous section, resulting in a simple roadmap for each. Although it is difficult to specify concrete time horizons given the heterogeneity of the addressed topics and proposed solutions, a range of between 3-5 years for medium term objectives and between 7-10 years for long term objectives could seem appropriate.

Solution areas and aspects to tackle	Medium term objectives	Long term objectives
<p>Foster interoperability in complex data ecosystems</p> <ul style="list-style-type: none"> » Business interoperability: define joint use cases that require interoperability across data spaces based on synergies across sectors. Enable cross-data space collaboration based on sectoral synergies » Semantic interoperability: approaches to facilitate the conversion of data models across sectors » Leverage AI/LLMs to facilitate semantic matching, schema alignment, and metadata enrichment. Develop AI agents capable of negotiating, translating, and validating data interactions » Standardisation efforts at different levels, aligned with European legal frameworks and sector-specific regulations 	<ul style="list-style-type: none"> » Achieve full interoperability within individual data spaces, including syntactic, semantic, and organisational layers » Ad-hoc interoperability through specific sectors subject to specific use cases » Applications of LLMs to support semantic interoperability. Deploy LLMs and knowledge-enhanced AI tools to automate semantic translation and model mapping » Harmonisation of data models' repositories (vocabularies) to facilitate discovery of new data models » Ensure interoperability of logs in provenance, auditability, and traceability across systems » Promote the adoption and harmonisation of interoperability-related standards 	<p>Seamless connection and flow of data between data spaces and digital ecosystems</p> <p>Real applications benefiting from the connection between data ecosystems in different sectors</p> <p>Implement autonomous LLM-based AI agents to manage interoperability, orchestrate data flows, and support lifecycle management</p> <p>Achieve consolidated, widely adopted EU-wide standards for data and metadata interoperability</p>

Solution areas and aspects to tackle	Medium term objectives	Long term objectives
<p>Design of global frameworks for trust, identity management and policy enforcement</p> <ul style="list-style-type: none"> » Data governance across complex value chains, managing the whole end-to-end data lifecycle, including data access, usage, reuse, and sharing rules, data quality management, provenance and traceability mechanisms to track origin, ownership, and usage over time, and roles and responsibilities in multi-party ecosystems » Privacy, protection, and confidentiality by design, on an end-to-end basis » Global trust anchors and federated trust mechanisms that allow distributed, cross-border trust relationships » Global governance schemes and policy interoperability, identity management and policy enforcement in distributed environments 	<ul style="list-style-type: none"> » Implement data quality assessment and certification mechanisms across data value chains » Deploy standardised provenance and traceability solutions integrated with AI » Launch federated trust frameworks within and across data spaces, with shared rules for trust negotiation and validation; implement verifiable credentials and digital attestations (aligned with eIDAS 2.0, W3C VC etc.) in selected pilots » Operationalise end-to-end privacy and confidentiality frameworks, including PETs (federated learning, secure multi-party computation, anonymisation) and secure data handling infrastructures 	<ul style="list-style-type: none"> » Fully operational, cross-sectoral data governance at scale » Robust data ecosystems to guard against emerging threats, including quantum computing, adversarial AI (e.g. GAN-based attacks), and novel computational paradigms » PET becomes a standard, embedded practice in AI development and data sharing infrastructures » A federated and global trust framework, including interoperable trust frameworks and a network of federated trust anchors » AI-driven systems for policy enforcement, consent, risk, and accountability across decentralised ecosystems, and negotiation and enforcement of data usage rights based on shared trust and policy frameworks

Solution areas and aspects to tackle	Medium term objectives	Long term objectives
Energy efficient and sustainable solutions for data-driven implementation <ul style="list-style-type: none"> » Sustainable-by-design AI and data solutions » Green and sustainable AI and data lifecycles » Holistic approaches for energy efficiency, including governance, metrics, and incentives for sustainable digitalisation » Specific techniques for energy efficiency 	<ul style="list-style-type: none"> » Integrate green and efficiency principles into AI and data system development pipelines, e.g. energy aware design models, training and learning » Scale up low-energy AI techniques (frugal AI, neuromorphic computing, spiking neural networks, surrogate models, mixture of experts, quantisation) » Ensure any framework proposes standardised energy and sustainability metrics for AI and data services, including best practices, repositories for reusable data and models » Develop and promote foundation models tailored for specific vertical sectors optimised for minimal data and compute requirements 	<ul style="list-style-type: none"> » Sustainability fully integrated across the entire AI and data value chain (data acquisition, storage, model training, analytics, deployment); green design principles as standard practice in system architectures and development frameworks » Frugal/low-energy AI approaches widely adopted and embedded in mainstream development and deployment pipelines and workflows » A holistic, system-level approach to energy efficiency that is fully operational, combining IT efficiency (in AI, computation, data handling) with domain-specific footprint reduction; AI and data-driven solutions actively contribute to decarbonisation and resource optimisation



Solution areas and aspects to tackle	Medium term objectives	Long term objectives
Design and implementation of complex data analysis pipelines <ul style="list-style-type: none"> » Design and execution of data pipelines and AI workflows across the computing continuum, from edge and fog nodes to cloud and high-performance computing (HPC) » Intelligent management and orchestration of complex data pipelines and AI workflows » AI-driven solutions for scheduling, load balance and resource efficiency » End-to-end security, trust, resilience and performance 	<ul style="list-style-type: none"> » Operational frameworks, including reference architecture and deployment schemes, for computing; integration of complex data pipelines and disruptive AI solutions » Application of DataOps and MLOps to deploy data pipelines and learning models, fostering collaboration, implementing automation, promoting agility and ensuring continuous improvement » AI-driven intelligent scheduling, load balancing, and auto-scaling of pipeline workloads in heterogeneous environments, including predictive orchestration tools » Deployment of AI and ML at the edge 	<ul style="list-style-type: none"> » Data pipelines and AI workflows are natively deployable across heterogeneous environments with automated orchestration, dynamic adaptation, and minimal configuration effort » Pipeline execution achieves high performance, reliability, and energy efficiency across all stages » Complex data pipelines are accessible through intuitive, interfaces, standard tools, and AI-assisted automation » End-to-end protection, trust, and privacy-by-design
Application of AI Agents to support complex data ecosystems <ul style="list-style-type: none"> » AI agents for ecosystem automation » AI agents to support semantic interoperability by aligning schemas, ontologies, and taxonomies across domains » Apply AI agents to enforce data usage policies and compliance rules within federated and dynamic environments 	<ul style="list-style-type: none"> » Deployment of AI Agents in real-world data ecosystem, and demonstrate concrete and scalable use cases where AI agents enhance key functionalities » Explore and validate architectures based on interconnected LLM-powered agents (LLM mesh paradigms) to enable collaborative, scalable, and cross-domain automation across data ecosystems 	<ul style="list-style-type: none"> » Establish fully autonomous agent-based systems capable of managing key data ecosystem functions end-to-end, including policy enforcement and compliance, dynamic data integration and enrichment, data lifecycle management (curation, access, deletion) » Cross-sector and cross-border AI agent collaboration frameworks

Solution areas and aspects to tackle	Medium term objectives	Long term objectives
<ul style="list-style-type: none"> » Leverage the Model Context Protocol (MCP) or similar architectures to enable AI agents to integrate LLM-based reasoning with external APIs, tools, and real-time data sources 	<ul style="list-style-type: none"> » Standardisation of AI agents' interfaces, behaviors, and protocols 	<ul style="list-style-type: none"> » Integrated Agentic AI within data space governance models
<p>Fostering the adoption and scale-up of data spaces</p> <ul style="list-style-type: none"> » Easy onboarding, access and adoption » Business models and value creation, tailored to each sector » Avoid fragmentation and ensure interoperability » Incentive mechanisms for all actors in the value chain » Guidelines, best practices and reference use cases 	<ul style="list-style-type: none"> » Leverage AI to automate management, functioning, onboarding and use of data ecosystems » Simplify onboarding, access, and interoperability for all participants in the ecosystem, mostly data providers, SMEs and public sector entities » Define and validate sustainable business models for data sharing, monetisation, and value co-creation » Establish maturity models and benchmarking frameworks for data spaces to assess readiness, scalability, trust, interoperability, and governance » Address skills gaps in data space design, implementation, deployment and use, including legal, governance, data, and AI integration 	<ul style="list-style-type: none"> » AI driven cognitive / intelligent data spaces » Data sharing as a commodity and standard practice » Advanced and sustainable data-driven business models » Establish data spaces as core ecosystems for AI training, validation, deployment, and monitoring » Full connectivity with the European Digital Ecosystem (AI Factories, Data Labs, AI TEFs, EDIHs, AI platforms, data platforms and marketplaces) » Widespread skills, literacy, and awareness of how to create, manage, onboard, and use data spaces

Solution areas and aspects to tackle	Medium term objectives	Long term objectives
<ul style="list-style-type: none"> » Data spaces and cutting-edge applications » General framework to facilitate the connection of advanced data driven applications with data spaces » Identification and preparation of high-value, application-ready data » Connection and interoperability between different ecosystems » Integration of real-world data with virtual representations » Real world use cases and industry validation 	<ul style="list-style-type: none"> » Reference architectures to connect data spaces with advanced applications » Assessment of additional applications most intended to benefit from and foster the development of data spaces » Definition of potential interface layers, new required blocks, new components, new roles, etc. » Launch pilot projects and testbeds that implement early versions of these architectures in real-world settings » Conduct preparatory work on inter-domain standards 	<ul style="list-style-type: none"> » Seamless, scalable deployment of disruptive applications across federated data spaces » Maturity of interoperability standards, interfaces, and open components that allow disruptive applications to be easily deployed, reused, and extended

3.2.4 Big ticket #4: Next generation robots with embodied intelligence

Today's robotic systems are severely limited in the applications and tasks they can address in the real world, and there is a significant gap between laboratory demonstrations of capability and delivered capability at end users. Current robots are largely application specific: a farming robot cannot be used in healthcare, although it may be adapted to military use or for construction. Similarly, robots for use in space might be adaptable to the nuclear industry or to mining. The form of a robot closely follows its function. The primary limitations of the currently deployed robotics solutions are in the following areas, which are key to extending applications and therefore to wider deployment. In order to make progress in these capabilities, it is critical to benchmark incremental steps that can be aligned to specific application-oriented functional achievements.

Configurability

Robots will only become cost effective if they can carry out multiple types of functions with the same core hardware. The ability to customise, configure and reconfigure robotic systems in both industrial and service robotics applications is key for adoption. The configurability of a system should allow both static customisation of a system to its domain and use case, and dynamic reconfigurability to an evolving environment. Configurability must be accessible without deep robotics knowledge and with minimal training. For wide

deployment in key applications, such as healthcare, construction, maintenance and warehousing, a step change in interoperability and configurability from current concepts and practices is needed.

Dexterity

Current robot dexterity is poor and far below basic human capability. This significantly limits the capability of robots to carry out complex tasks in human environments. Therefore, their application is limited to well-engineered environments with objects that are adapted to robot grasping, or to applications where items can be lifted with vacuum-based grippers that can conform to object shape. The handling of soft materials or highly deformable materials is currently not possible in a way that is deployable. Matching robot perception with physics is therefore a major challenge. Physics is based on a formal understanding of the cause-and-effect relationships that govern interaction between objects, rather than data, and this is an area of extreme weakness for controlled manipulation. Novel physics- and model-based methods are needed to underpin the high-level planning and object reasoning that current data-oriented AI based perception methods provide.

Autonomy

Autonomy in robotics is currently short term and localised, such that most often the working environment is constrained and the level of uncertainty is highly controlled. Teleoperated robots provide short-term autonomy; drones, for example, use autonomous control to maintain position and direction, but still rely on human control for task motions. Teleoperation often places a high cognitive load on the operator and requires extensive training and planning. Consequently, its use is limited to applications where the use of a robot is the only option, for example in micro-surgery or in the repair of nuclear reactors. Greater levels of autonomy are needed in a complex, unknown and dynamic robot environment. A complex environment may encompass unanticipated scenarios that robots have to deal with and adapt to. This environment may also contain significant risk to people or to the environment itself, requiring autonomous robot activities and actions to be predictable and trustworthy. In many cases their safety, physical and behavioral, needs to be certified to industry-specific standards.

There are many real-world applications that present complex environments beyond the capability of current robotic systems, such as outdoor operations, and maintenance and transport. Complexity can also arise in the type of task and combination of tasks (i.e., missions) that the robots carry out. Improvements in autonomy that address complexity and uncertainty will expand the range of scenarios where robots can operate. Improvements in autonomy will only come from deeper synergy between physical interaction and AI in perception, manipulation and interaction with objects, people and environments.

Interaction

Most current robot systems either require no human interaction or direct human interaction using “human in the loop” control, as they are based on well-defined operating scenarios with minimal uncertainty. Applications in the middle ground where humans have oversight without intervention control are limited to a small number of well-defined domains. Developing human interaction at the transactional and functional level, even in moderately

complex environments, has proved to be remarkably difficult. If robots are to become useful tools, they must be able to work with humans, learn from them, take orders from them, and be supervised by them. These transactions require interaction at a practical level, with the task in hand, and at a social and ethical level that ensures safety and effectiveness. Developing good HRI requires both multi-modality of communication and an understanding of social context. Multidisciplinary collaboration with the social science and humanities domains is essential to progress.

Core technologies

The core robot behavior technologies such as perception, planning, and reasoning are partly responsible for the limitations seen in configurability, dexterity, autonomy and interaction. However, no single technological advancement will extend capability in these areas; it will require a combined approach that develops and exploits novelty and innovation across design tools, hardware, integrated behavioral systems and development and testing methodologies.

New anatomies

The above challenges call for combining new forms of robots with a high level of autonomy, manipulation, configurability, and greatly improved HRI. We need to achieve significant progress in physical forms through new designs of rigid body and soft robotics and improved ability to safely act in complex unpredictable environments. This must be done within existing energy envelopes and with decision-making systems that are grounded in physics, not data. Adaptable, configurable multi-tasking machines present major challenges for the next generation of smart embodied robotic systems operating in the real physical world.

We need to create robots that can fluidly interact and respond to human physical contact. This requires new materials, new structural designs, and new methods of control. Advances in soft robotics will also enable better interaction with everyday environments and objects that are designed for human use. The expectation is that technical progression in this area will drive novel physical structures that break away from the rigid body mechanics traditionally used in all robotic systems, towards softer human-friendly physical forms that create more natural interaction experiences.



Perception

A primary and enduring challenge is to ensure that robots are safe, trustworthy and exhibit predictive actions in their physical activity. The longer a robot operates autonomously and the more complex the environment it is operating in, the harder it is to assess and deliver guarantees about safety and trustworthiness. The main challenge is thus to develop methods of identification that can detect unwanted activities or are able to characterise tasks and environments that are low risk and potentially safe. This represents a challenging environment in which robots should learn about their states, actions and the conditions in which they operate. Learning techniques depend on the ability to train neural models based on real or synthetic data; however, these techniques are typically not able to deal with unanticipated scenarios and do not naturally account for the hard physics of interaction (such as gravity, friction or momentum) or for the inherently uncertain and dynamic nature of a formally complex system where chaotic and emergent scenarios are to be expected. We need to elaborate on world models for robots operating in complex environments where safety, trustworthiness and regulatory compliance are essential, and find new ways to formulate the integration of these world models in robot development.

Physical intelligence

The use of AI in robotics systems is highly limited to a few applications. The research in a wide array of AI technologies, such as Embodied AI, Physical AI, Reinforcement Learning, Generative AI, etc. shall be conducted to investigate new forms of advanced control, such that a robot can reconstruct an accurate model of the physical world around it to understand its environment and plan appropriate actions. This can range from simply avoiding mistakes or harm to people, to acting in dynamically changing environments and eventually reaching long-term autonomy. AI research needs to focus on techniques that provide accurate models of the physical world that would enable the necessary level of autonomy for robots and other automated physical systems. This is a major challenge that requires close collaboration between AI and robotics disciplines, as well as research and industry, for the purpose of developing foundational concepts and enabling broad deployment. Critical to this enterprise is the ability to explain autonomous decision-making within the external context of the environment. This needs to be responsive in real time and collaborative with human operators to allow smooth interaction and direction in ways equivalent to collaboration between humans carrying out the same task. This will build trust between operator and machine while also allowing the operator to correct and redirect the machine in terms that relate to the task context. Achieving this requires a strong alignment between the human and machine models of the environment and its dynamics, and the use of terminology “colloquial” to the task; for example, “move the crate on top of the cooker to the other side of the table”.

Innovation drivers

Skill shortage

The demographic changes that shift the age profile of Europe over the next two decades will have a significant impact on the availability of skilled workers. To ensure a sufficient pipeline of skilled engineers and scientists in robotics, Europe must adopt a multi-pronged strategy that focuses on both education and workforce transformation. First, the EU should

expand and modernise STEM education, integrating robotics and AI into school curricula from an early age and strengthening vocational training programmes tailored to robotics applications. Universities and technical institutes must also align their programmes with industry needs, offering interdisciplinary courses that combine robotics, software engineering, ethics, and systems integration. Additionally, initiatives like the EU's Digital Innovation Hubs and AI Factories should be scaled up to provide hands-on training and foster collaboration between academia, startups, and established companies.

At the same time, Europe must invest in upskilling its existing workforce to adapt to the rapid pace of automation. Rather than relying solely on a limited pool of robotics specialists, companies should empower non-technical employees through no-code platforms and embedded training programmes that make robotics tools more accessible. This democratisation of technology use can help bridge the gap between innovation and implementation, especially in small and medium-sized enterprises. Furthermore, while skilled migration can offer short-term relief, long-term resilience will depend on creating attractive labour markets and career pathways within Europe itself. By combining education reform, workforce enablement, and inclusive technology access, Europe can build a robust talent pipeline to support its robotics ambitions.

Geopolitical shift

In the current geopolitical climate, Europe must prioritise strategic autonomy in robotics to safeguard its technological sovereignty and economic resilience. As global powers like the U.S. and China intensify their investments in robotics for industrial, military, and strategic dominance, Europe faces the dual challenge of keeping its global innovation edge while reducing dependency on foreign technologies. This requires a coordinated EU-wide approach that boosts funding for foundational robotics R&D, strengthens cross-border collaboration among research institutions, and accelerates the development of secure, ethical, and interoperable robotic systems. Initiatives like NATO's Innovation Fund and the EU's Horizon Europe programme are steps in the right direction, but they must be scaled and aligned with long-term geopolitical goals.

Moreover, Europe must address vulnerabilities in its robotics supply chains, particularly those linked to critical components sourced from geopolitical rivals. The Ukraine conflict has underscored the strategic role of unmanned systems in modern warfare, highlighting the need for Europe to develop indigenous capabilities in autonomous systems and AI-driven robotics. At the same time, ethical governance must remain central to Europe's robotics agenda, ensuring that technological advancement aligns with democratic values and international law. By fostering a robust ecosystem that integrates academia, industry, and defence sectors, Europe can position itself as a global leader in responsible robotics innovation while reinforcing its geopolitical standing.

Innovation sectors

Agile production

The manufacturing industry is changing rapidly due to digital technology, automation, and improved connectivity. Physical machines and computer systems are merging, allowing robots and AI to assist humans while also displacing a significant number of workers. This shift is moving jobs from manual labour to remote tasks with immersive interfaces, requiring

considerable upskilling of the workforce. Sustainable manufacturing must adapt to frequently changing production levels and individual customisation needs while ensuring energy efficiency and maintaining high production quality. Robots are essential for automating manufacturing tasks and improving the efficiency, quality, and safety of production.

Industrial robotics will extend beyond current automated tasks, such as moving materials, welding, assembling, painting, packaging, and inspecting. The new generation of industrial robots will increase autonomous production by making decisions, learning, and adapting to complex situations; however, human-robot interaction and collaboration also remain essential in the evolving manufacturing industry.

As industrial robots become more integrated into supply chains to support agile, distributed and customisable production, the resiliency of robotics ecosystems will also be increasingly important.

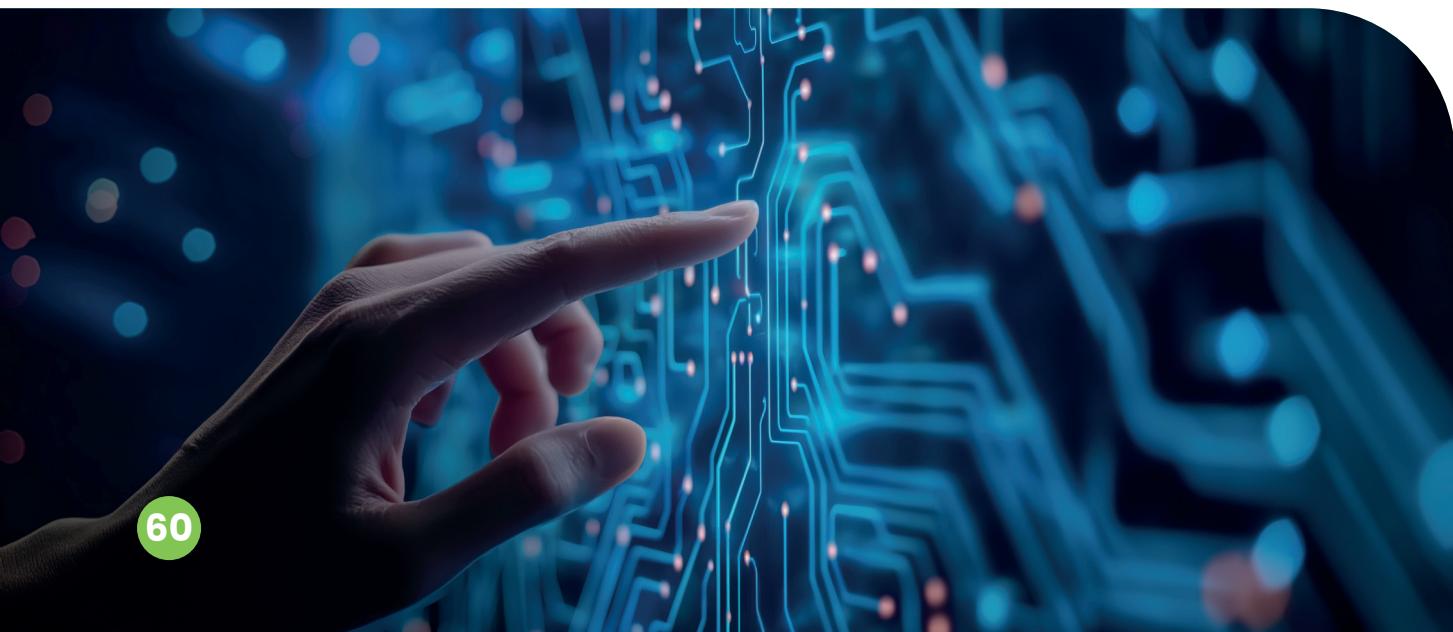
Healthcare

In the face of unprecedented demographic, epidemiological, and systemic pressures, the healthcare sector is being compelled to transform at a pace and scale not seen before. Staff shortages, escalating costs, an aging population, and the residual impacts of the COVID-19 pandemic are converging to create a scenario in which traditional models of care are no longer sustainable. Concurrently, the rapid evolution of robotics, driven by advances in soft robotics, autonomy, AI, and human-robot interaction, presents a unique opportunity to reimagine healthcare delivery.

From surgical theatres to home-based rehabilitation, robotics is poised to operate across the entire continuum of care. Highly autonomous surgical systems are increasingly delivering greater precision, consistency, and reduced recovery times. Meanwhile, mobile robotic assistants, telepresence units, and companion systems facilitate care continuity in outpatient and community settings, which is particularly vital in rural or understaffed regions. Such robots, if designed for configurability and ease of integration, can support evolving care models such as hospital-at-home and decentralised diagnostic networks. Whether through automation of routine tasks, support in lifting and mobility, or surgical augmentation, robotics can enable skilled professionals to focus on high-value, human-centric care.

Inspection and maintenance

Well-functioning critical infrastructure – in transportation, telecom, energy, buildings



and more – is key to sustaining and progressing today's society. To ensure maximum uptime and safe use of infrastructure, timely and proper maintenance is needed, which is often very costly and potentially dangerous to humans. The backlog and cost of unmet maintenance needs for Europe's critical infrastructure is enormous. Over the last decades, the use of robots for inspection has grown significantly in some application areas – such as aerial drones for checking roof tops, chimneys and power lines, or crawlers and drones for confined-area inspection in tanks and tunnel networks such as sewers – and there has been an increase in the automatic analysis of inspection data via use of AI.

Smart embodied robots could potentially handle complex intervention tasks and navigate complex environments, and thus be applied to a larger range of variations in tasks and environments. To this end, new approaches with smart embodied robots are needed. As part of this, easy configurability is required to be able to re-use robot modules (e.g., for perception, manipulation, etc.) across variations of maintenance tasks and operating environments. Smart robots for I&M should be easy and intuitive to use to provide value. Thus, they require autonomous capabilities to understand their environment and operations and perform tasks without needing continuous monitoring and control by human operators.

Construction

Currently, robots in the construction industry are primarily used to automate repetitive, hazardous, or precision-demanding tasks. Examples include bricklaying robots like SAM100, which can lay thousands of bricks per day with minimal human oversight, and layout robots such as Dusty Robotics, which mark construction blueprints directly onto floors with high accuracy. Other applications include robotic systems for welding, painting, and plastering, as well as autonomous demolition machines that can safely dismantle structures in dangerous environments. Additionally, 3D printing robots are being deployed to construct entire buildings using concrete extrusion, significantly reducing material waste and construction time.

As robots become more adaptable to dynamic site conditions, their role will expand from isolated tasks to full participation in end-to-end construction workflows, and several key developments will become necessary. First, construction companies must invest in digital infrastructure, such as Building Information Modelling (BIM), to enable seamless integration between robots and project plans. Second, regulatory frameworks must evolve to ensure safety and standardisation in robotic construction practices. Third, workforce training programmes must be established to equip workers with the skills to operate, maintain, and collaborate with robotic systems. Finally, public and private sectors should collaborate to fund pilot projects and innovation hubs that accelerate the adoption of robotics in construction. With the right ecosystem in place, robotics can transform construction into a safer, faster, and more sustainable industry.

Food

Food security is essential for Europe's welfare and sustainability. Currently, robots in agriculture and aquaculture are used to automate a wide range of labour-intensive and repetitive tasks, significantly improving efficiency and precision of farming. These include autonomous tractors for ploughing and seeding, robotic harvesters that detect and pick ripe fruits, and drones for crop monitoring and spraying and underwater robots for sea cages inspection and cleaning. Robots are also deployed for weeding and pest control

using computer vision and AI to target specific areas, reducing the need for chemical inputs. In livestock farming, robotic milking systems and automated feeders enhance productivity and animal welfare. These technologies help address labour shortages, reduce operational costs, and support more sustainable farming practices

To unlock the full potential of robotics in agriculture and aquaculture, the focus shall be on advancing core capabilities, such as safety deployment and flexibility manipulation and deployment to reliable replace manual labour in the harsh environment. These capabilities require advanced machine vision technologies, soft robotics for delicate handling, and AI algorithms capable of on-board intelligence for real-time decision-making in harsh, complex environments. R&D should also focus on developing modular and scalable robotic platforms that can be flexibly adapted to different farming scenarios (e.g., land, sea) and manipulation tasks. Collaboration between academia, suppliers and end-users is essential to accelerate innovation and ensure that solutions are practically deployable. Furthermore, dedicated funding programmes and testbeds should be established to support long-term experimentation and the transition from prototypes to commercially viable systems.

Ship building

The European shipbuilding industry is facing increasing competition from international competitors, and there is a need to improve its competitiveness through innovation and technological advancements. The MARI-4YARD project, together with RESURGAM, FIBRE4YARDS, SMARTYards and PENELOPE, developed several solutions in materials, processes, technologies and skills. Additionally, Horizon Europe's SEUS project is developing an integrated software platform with the goal of saving time in engineering, assembly and construction at European shipyards. However, there is still a need to increase ambition and demonstrate these solutions in large, medium, and small-scale shipyards.

To support this transformation, significant R&D investment is needed in several key areas. First, robotics systems must be developed to handle the unique challenges of shipbuilding environments—such as large-scale structures, variable lighting, and irregular surfaces. This requires advancements in sensor fusion, machine learning, and adaptive control systems. Second, integration with digital design tools like CAD and BIM must be improved to enable seamless robot programming and real-time feedback loops. Third, modular and reconfigurable robotic platforms should be explored to allow flexible deployment across different shipyard tasks.

Roadmap

Socio-economic impact

The deployment of increasingly autonomous robots means that robotics could replace more and more human physical work in factories, transport, field operations and maintenance, agri-food, healthcare, and security. The socio-economic impacts of these changes would be catastrophic and counterproductive. There is no societal benefit to having robots replace human work other than where it is necessary to maintain safety or to augment labour shortages. The steps towards robotic work must therefore be taken in consultation with citizens and involve a high level of responsible innovation. The focus must be on improving lives, ensuring robotics are used where they can remove people from harm, creating better working conditions, or enhancing health and wellbeing. However, in many areas of work,

robots will become more efficient than humans and thus create economic pressures. Research and innovation, in collaboration with the humanities and social Sciences, must find responsible solutions that balance economic necessity and human wellbeing. The “hard problem” addressed above is pertinent to this socio-economic impact because it addresses those use cases where harm is most likely to be caused. These are applications where humans already perform those actions and are qualified, or those involving tasks that humans cannot/will not undertake. Robots must be proven to be highly competent, ethically sound, and socially acceptable before deployment.

Technology Readiness Level (TRL)

TRL for robotics contains two fundamental axes. On one axis, we need to balance between human intervention and autonomy; on the other axis, we need to manage the complexity of the operating environment as perceived by the robot. Where humans are teleoperating the robot, it may still be able to perceive a highly complex environment and provide information to the user. For example, a teleoperated surgical robot with sensors that can detect abnormal cells can guide the surgeon to sites that need attention. A fully autonomous robot operating in orbit faces few obstacles that are unknown, and its world is far simpler than that of an autonomous vehicle navigating in a city. Within this map, TRL criteria can be set for specific use cases and applications. However, this challenge operates at a more fundamental level and must focus on the development of methodologies that have the potential to reach high TRL levels with performances needed in specific applications. It is likely that this work will identify parts of the operating space where it is possible to use foundation models successfully, and other parts where new methodologies are required.

Work programme 2026 – 2027

The work programme 2021 – 2025 of cluster 4 of the Horizon Europe programme introduced a number of topics focused on developing novel robotics solutions in strategically important industrial sectors. However, market adoption of new robotics solutions takes time – potentially more than a decade – as the learning curve is not simple, and the decision to invest large amounts of money should be based on a viable concept and process. Therefore, progress in robotics needs to focus on long term investment in enhancing these core capabilities while being agnostic to the specific technology used to achieve the innovation. However, robots represent a significant integration challenge, and capability that cannot be integrated into a full deployable solution is of little value. In this regard, a whole systems approach, at least in concept, must accompany novelty; this includes an understanding of the likely regulatory environment. Advances in design and deployment tooling are likely to be as valuable as advances in material science that enable wide area sensing or light weight high power actuators, just as advances in reasoning and perception are as likely to enhance capabilities as novel physical architectural configurations.

Big ticket actions for advanced robotics should therefore focus on enhancing configurability, dexterity, autonomy and intractability with a clear focus on functional gains in multiple end user application areas, while being agnostic to the specific technologies and combinations of technologies used to achieve the enhancement. Indeed, advances in technology in general, and certainly in robotics in particular, have primarily occurred when simultaneous incremental advances in individual technologies have been combined advantageously

to enhance or create new functions. This work should be supported by function-oriented benchmarking able to define incremental waypoints in system level capability.

The work programme 2025-27 should support new initiatives from EC such as the European Chips Act and the Critical Raw Materials Act, which were established in the past three years in response to world crises and geopolitical conflict. Research, development and deployment of robotics systems powered by AI and supported by data science should be accelerated to meet the goals and deadlines set by these new Acts for smart and cognitive manufacturing as well as increased autonomy and resilience of production through remanufacturing, recycling, recovery and waste management. Finally, considerable effort should be invested into human-machine collaboration, ethics, and compliance with EU AI and data act at the service of society.

Today's robotic systems are severely limited in the applications and tasks they can address in the real world and there is a significant gap between laboratory demonstrations of capability and delivered capability at end users. The primary limitations in current deployed capability centre on four key areas:

Range and Reach: Current robots are unable to reach certain places where work needs to be done. Quadrupeds can reach most areas of flat land but cannot climb or access small spaces, and their reach beyond their base is limited. Robot arms typically have limited reach and a limited load vs reach capability. Many applications require range and reach beyond a few metres.

Dexterity: Current robot dexterity is poor and far below basic human capability. This significantly limits the ability of robots to carry out complex tasks in human environments. As a consequence, their application is limited to well-engineered environments with objects that are adapted to robot grasping or can be lifted with vacuum-based grippers. The handling of soft materials or highly deformable materials is currently not possible in a way that is deployable.

Autonomy: Many of today's deployed robots are teleoperated with short term autonomy; for example, drones use autonomous control to maintain position and direction but rely on human control for task motions. Teleoperation often places a high cognitive load on the operator and requires extensive training and planning. As a consequence, its use is limited to applications where the use of a robot is the only option, for example in micro-surgery, or in the repair of nuclear reactors etc.

Interaction: Interactions between robots and their environment and with the people in the working environment are typically kept to a minimum. Robots typically operate in cages or in defined areas that are protected from human access. Interactions with the working environment are often well controlled, either through environmental engineering or by the nature of the application. For example, a domestic vacuum cleaning robot operates in a highly variable but well constrained environment; similarly, agricultural robots operate in well specified environments. Currently deployed robots cannot operate alone in unstructured and unknown environments. Other core robot behavior technologies such as perception, planning, reasoning etc. are partly responsible for the limitations seen in reach, dexterity, autonomy and interaction. However, no single technical advancement will extend capability in these areas; it requires a combined approach that develops and exploits novelty and innovation across; design tools, hardware, integrated behavioral systems and in development and testing methodologies. In order to progress capability, it is critical to benchmark incremental steps in capability

that can be allied to specific application oriented functional achievements. Benchmarks without deployment reference are of little value.

Progress in robotics therefore needs to focus on long term investment in the enhancement of these core capabilities while being agnostic to the specific technology used to achieve innovation. However, robots represent a significant integration challenge, and capability that cannot be integrated into a fully deployable solution is of little value. In this regard, a whole systems approach, at least in concept, must accompany novelty and this includes an understanding of the likely regulatory environment. Advances in design and deployment tooling are likely to be as valuable as advances in material science that enable wide area sensing or light weight high power actuators. Just as advances in reasoning and perception are as likely to create a capability enhancement as developments of novel physical architectural configurations.

Big ticket actions for advanced robotics should therefore focus on enhancing Reach, Dexterity, Autonomy and Interaction with a clear focus on the functional gain in multiple end user application areas while being agnostic to the specific technologies and combinations of technologies used to achieve the enhancement. Advances in technology in general, and certainly in robotics particularly, have primarily occurred when simultaneous incremental advances in individual technologies have been combined advantageously to enhance or create new functions. This work should be supported by function-oriented benchmarking able to define incremental waypoints in system level capability.

Innovation enablers

Soft robotics

Classically, compliant control allows robots to be guided by the forces applied on them rather than being programmed to reach predefined positions, enabling them to adjust to contact with objects or humans. But current force control performances are relatively poor, differing highly depending on the scale (from micro to macro) and the specific application requirements (industry, rehabilitation, surgery, exploration, inspection). Soft robots (main structure and end effectors) have the objective to intrinsically reach the highest standards of safety and adaptivity to the task and the (fragile or sensitive) environment (including humans, animals, and objects) thanks to their soft nature: soft robotic systems are generally lighter weight, and they can be made inherently compliant by design rather than reaching limited compliance with the use of control software. Moreover, soft robots have the potential to operate with agility in unstructured settings not a priori designed for robots, and to build new adaptive multi-purpose intelligent devices.

Nevertheless, current maturity in research and development in this field is rather poor, and varies significantly depending on the sub-technologies involved in building a robot. There are still key scientific and technical challenges for robotics research to address before such robots can perceive and interact efficiently and safely with their environment. Advances in many research sectors are required: design, manufacturing and control must be improved to enable the production and use of reliable soft robots, robotic arms and grippers, with increased force and enhanced reach/access, dexterity/manipulation, autonomy and interaction capabilities, in real-life applications which are inaccessible to classical robots. It is therefore necessary to develop new paradigms and design methods for soft robot

structures. Adaptive or multimodal locomotion capability could be investigated to allow soft robots to change modes of locomotion (switching from walking to jumping, climbing, squeezing through narrow passages, etc.). In the future, new paradigms should be found to give soft robots (partial) self-healing capability.

Recently, new bodies and hardware components enabling soft robots to perceive the world and act more efficiently have begun to be developed. Steady progress has been made in the design of new passive and active materials (with specific properties) that led to several prototypes, including a few commercial products, mainly soft grippers and soft robotic hands. Proofs of concept of soft actuators have been developed for a considerable time, but they still have not reached the optimal ratio between light weight and high power required for some applications. The use of dedicated (soft) sensors also remains limited. Linked to this, soft connection to the electronics is required, and recent advances in the soft electronics development field could be better exploited to allow the mass production of soft robots. As a consequence, the integration of actuators and sensors into complete soft robots is still far from reaching the industrial level. The potential use of new materials, as well as adapted fabrication processes, to create novel soft robotics structures should be studied. Soft robots, as well as new actuators and sensors, could be advantageously bio-inspired. Dedicated conceptual design methods, eventually combined with automatic optimal design software and AI tools, could be used to develop novel designs for soft robots, customised for dedicated applications, thus reducing their cost or their energy consumption to allow power embedding and energy autonomy.

Soft robots are complex to model. When physical models are lacking or are not sufficiently accurate, a high number of embedded (or external) sensors may be required to measure and generate a faithful real-time representation of the robot's state and its multimodal interactions with its environment. A trade-off between model-based control and machine learning or other AI techniques could be appropriate to process the sensing data and control the soft robot.

Soft robots are promising for the widespread use of safe robots in human interaction and beyond, because they can also address use-cases with flexible or fragile objects or environments. Application domains include the main domains already addressed by classical robots (industry, health, inspection and maintenance, search and rescue, agri-food etc.) but are expanded to include use-cases regarding performance requirements that cannot be solved (or are not accessible) by classical robots: specific size (leveraging miniaturisation), modularity, self-configuration, or compliance. They can also be developed and used for applications requiring powerful yet lightweight robotic devices in hard-to-reach or difficult-to-access spaces. Efficient and cost-effective solutions are still needed to equip robots for intervention or manipulation in fragile environments.

Areas for progress today/near term	Medium-term objectives	Long-term objectives
<p>New robot anatomies for wide deployment of robots in human interaction and beyond.</p> <p>Safety and trustworthiness by design: Novel robot anatomies have the potential to impact on the “hard problem” by creating inherently safe systems, partly by designing robotic systems with much lower mass, but also because they can be made inherently compliant, rather than algorithmically compliant.</p>	<p>The medium-term objectives are to explore new materials, actuators and sensors and investigate how they can be used to create novel robotics structures integrating safety and trustworthiness by design.</p> <ul style="list-style-type: none"> » Establish a range of materials suited to developing robotic systems, both as the main structure and of manipulators and end effectors, as well as adapted fabrication processes. These may encompass passive and active materials, and combination materials with specific properties. » Develop new paradigms and design methods for non-rigid structures and the means to actuate and sense position and forces where this may no longer involve fixed rotational or linear links. » Create control methods for structures built from novel and soft materials or for structures that emulate rigid structures using soft materials. » Identify application areas where these new robots can have the greatest economic or societal impact, and develop experimental designs for robotic systems able to demonstrate how this capability can be achieved. 	<p>The long term objectives are to build robotic systems using these novel materials and prove that they can provide full lifecycle reliability and enhance performance in tasks that are more difficult to achieve with current robotic structures.</p> <p>Moreover, these new robots can be a solution to produce lighter, simpler (even use-case specific) and lower cost robots, while reaching required capabilities and performance.</p>

Areas for progress today/near term	Medium-term objectives	Long-term objectives
<p>Robots for complex manipulation and dirty and dull jobs.</p> <p>Today we are facing:</p> <ul style="list-style-type: none"> » A drastic increase of waste that will eventually become unmanageable by humans in future » A Critical Raw Materials Act calling for advanced robotics that can enable meeting the 2030/2035 goals for materials recovery, recycling and processing » Additive manufacturing calling for AI-powered robots that automate the currently manual task of post-processing for a viable and competitive business 	<p>The medium-term objective is to establish:</p> <ul style="list-style-type: none"> » Data spaces that facilitate the robotisation and automation of waste management, material recovery and additive manufacturing processes » Perception systems (visual and non-visual) and AI techniques that are robust and reliable to monitor the robotisation process and capture anomalies » Hardware solutions with powerful, dexterous manipulation and reliable tactile sensing capabilities, combined with a high-precision control process, enabling complex manipulation tasks such as dismantling, disassembly, sorting, positioning, and finishing. These solutions support experiments on specific use cases selected by businesses for validation and proof of concept, including human robot collaboration 	<p>The long-term objective is to:</p> <ul style="list-style-type: none"> » Set up field tests for demonstration and piloting new capabilities » Elaborate on optimisation of the final design for cost-effective production, deployment and maintenance » Elaborate on wide deployment of robotics systems with complex manipulation into different industry sectors (process industries, additive manufacturing, maintenance, agriculture, healthcare, textile recycling, etc.)
<p>Robots for performing logistics and other assistive tasks in the healthcare industry.</p> <p>Today we are facing:</p> <ul style="list-style-type: none"> » Increasing demand for healthcare services due to the aging population » Shortage of trained staff in healthcare facilities 	<p>Transforming clinical actions through mobile and modular robotics</p> <ul style="list-style-type: none"> » Lead the transition from labour-intensive hospital systems to agile, robotics-augmented care environments » Consider systems with the following requirements: sterile logistics, patient handling, and intra-facility navigation, obstacle avoidance, etc. 	<p>Develop specialised robotics systems for complex high precision actions such as surgical interventions.</p> <ul style="list-style-type: none"> » Extend human capability » Tissue-aware manipulation using multimodal imaging techniques

Areas for progress today/near term	Medium-term objectives	Long-term objectives
	<ul style="list-style-type: none"> » Purpose-built assistive robots will offer physical support with daily living activities, enabling people to stay longer at home receiving non-invasive care and only receive high level intervention in the specialised clinical centres » Consider simplicity, affordability, and adaptability for high-quality care into homes <p>Focus on clinically validated, certifiable, and mechanically reliable robotics and boost translational research via test facilities.</p>	<ul style="list-style-type: none"> » Navigation in dynamic anatomic environments with critical structures and deformation » Interoperable robotic ecosystems across the continuum of care » Care that is continuous, decentralised, and supported by an ecosystem of interoperable robotic systems » Robotic systems that are purpose-based, designed for diagnostics, intervention, rehabilitation and support for patients and professionals » Take into account standards for connectivity, control, and safety, ensuring seamless collaboration between robotic systems, clinicians, and patients <p>Focus on sustainable and trustworthy robotic platforms</p> <ul style="list-style-type: none"> » Develop modular, sustainable, and maintainable robotic platforms that align with circular economy principles and clinical needs

Innovation bottlenecks

Safety

Robot safety could be improved by putting effort not only into software certifications (which can be very difficult when embedding AI in robot controllers), but also on new mechatronic (mechanical structures and architecture, actuators, and sensors) developments to obtain inherently safe robots whose safety certification could therefore become less challenging. One possible way to achieve inherent safety in robotics is to promote the development of soft robots.

Data sets

The remarkable performance of Generative AI in the form of LLM, VLM, etc., is directly derived from the availability of large amounts of text and images on the internet. In fact, the amount and structure of data are key to all machine learning technologies. However, robotics fundamentally lacks an understanding of such data structures, let alone the amount of data which will be needed for the development of such advancements. This remains a major bottleneck to implementing, experimenting and developing an understanding of the usefulness of AI in robotics. Initiating European efforts in the development of such datasets is an important step towards AI enabled robotics.

Energy consumption

A number of exciting and promising IT technologies have failed to deliver in the past due to high demands for infrastructure and energy. The use of Generative AI places stringent demands on computational power and energy to compute multilayered AANs, which may not be scalable or may become commercially too expensive to spark the promised explosive surge of innovation.

Moreover, robotic energetic autonomy (particularly for field applications, but not exclusively) remains poor, and could benefit from new and specialised mechatronic design architectures optimised to reduce energy consumption, both in use and during manufacturing. This can be improved for classical robots but is expected to have more impact on the next generation of soft robots.

Work programme 2026 – 2027

Deployment in new applications and sectors

Search and rescue

New generation robots will be increasingly required for responses during natural and man-made disasters (climate change and international crises), like for people rescue, decommissioning of industrial sites, or for inspection and maintenance of infrastructures after the disaster.

Aspect	Embodied AI	Physical AI	Embodied Robotic System
Scope	Cognitive algorithms coupled to a body	AI-enabled artifacts, robots, and smart materials	Robot hardware plus basic compute
Focus	Learning, perception, decision-making via embodiment	Training & embedding AI directly in physical form	Mechanical design, kinematics, sensor/actuator layout
Dependency	Requires a body (real or simulated)	Requires AI methods & physics-based substrates	Can run fixed control as no advanced AI needed
Primary ingredients	Software, world models, adaptive control loops	Digital-twin simulation, synthetic 3-D data, on-device inference	Chassis, motors, sensors, power, CPUs/MCUs
Typical methods	Reinforcement learning, sim-to-real transfer, morphology-aware control	Physics-guided generative models, distributed sensing, neuromorphic/ASIC inference	Classical control, PLC control, ROS nodes
Example research	Zero-shot sim-to-real locomotion for legged robots	Shape-morphing composites trained in simulation	Industrial SCARA arm executing scripted pick-and-place

Looking ahead to 2026–27 and beyond, incremental improvements in stronger actuators or lighter frames will still matter, but the biggest capability gains are likely to come from:

- » scalable world-model pre-training that distills vast multimodal experience into transferable priors,
- » on-board adaptation that exploits each robot's unique morphology in real time; and
- » maturing physical-AI materials whose embedded sensing, actuation, and compute internalise portions of the control loop.

In this emerging era, the robot body serves as a flexible substrate whose value grows with the sophistication of the cognition and “intelligent matter” it hosts.



3.2.5 Big ticket #5: ADR Technology for the sciences

Introduction

Science and innovation are at the heart of European prosperity. Europe was the crucible for the development of scientific methods in the 17th century. Today, it is home to a rich variety of world-leading scientific institutes, supported by a world-leading science tools industry.

Scientific progress plays a vital role in delivering major European policy agendas on health, the environment, energy, and more, creating a pressing need for advances in research and innovation. However, many of these areas of need are also characterised by complex socio-technical systems, whose analysis requires novel interdisciplinary approaches to analysing data. At the same time, the volume and complexity of science itself is growing, raising questions about slowing productivity [ref OECD]. By helping make sense of this complexity, research and innovation may be one of the most important areas where AI could enhance societal wellbeing.

State of the art

Ground-breaking advances from systems such as AlphaFold signal the potential of AI in science. Beyond such flagship initiatives, researchers across disciplines and from across academia and industry are finding new ways of applying AI to enhance their work. Examples have been reported from the observational sciences (such as astronomy) and experimental sciences (such as chemistry) and extend into system sciences (such as ecology, immunology, climate), engineering sciences, the arts and humanities.

Advances in a range of ADR technologies offer the prospect of further progress in the use of AI for research and innovation. These technologies touch on all steps in the scientific method: originating a hypothesis, devising an experiment, running a simulation, preparing an experimental protocol, running experiments, analysis and interpretation of results.

Uses of ADR technologies: There has been some progress in many parts of this process. Knowledge graphs and Bayesian methods have proved very useful.⁸⁰ Deep neural networks have been shown to be powerful but retain some limitations⁸¹ and some ontologies have been defined.⁸² Generative AI is being applied widely, with LLMs providing a new tool for scientific activities such as literature synthesis, formalising knowledge and coding. These have been evolving into multiagent-based *agentic AI*; specific products have been proposed by existing market players (Microsoft's Discovery), though some early attempts have had to be withdrawn (Galactica).

Data and corpora: The scientific community already curates datasets of impressive size, detail and quality,⁸³ whether in earth observation or mapping the properties of complex biomolecules, thanks to public investments in infrastructure such as satellites, DNA databases and synchrotrons.

80. <https://bmcmedinformdecismak.biomedcentral.com/articles/10.1186/s12911-021-01518-6>

81. <https://arxiv.org/pdf/2012.15754.pdf>

82. https://link.springer.com/chapter/10.1007/978-3-319-49004-5_31

83. <https://www.nature.com/articles/s41597-021-01071-x.pdf>

The availability of these resources has been foundational in building e.g. climate models. The success of *AlphaFold* by DeepMind, which contributed substantially to solving the long-recognised protein sequence-structure problem, relied heavily on such publicly funded datasets.

AI driven labs: Large-scale experimental facilities such as high throughput screening (HTS) and high content screening (HCS) are already well established for drug screening in the pharmaceutical industry, and are joined by concepts such as high throughput experimentation (HTE) in chemistry. Where lab work requires specialised materials, biofoundries (for biological material) and 3D printing facilities (for tooling) can be useful.

While these have remained relatively modest in scope, combining them with AI engines at the analysis and planning steps permits automation of the entire **DMTA** (design-make-test-analyse) cycle in so-called self-driving laboratories. Experimental AI scientist systems can originate hypotheses and run experiments,⁸⁴ and many sites across Europe have reported establishing such self-driving labs with support from national programmes⁸⁵ and a small number of EU projects.⁸⁶ These facilities have focused on a wide range of subject areas, mostly in the materials and biological sciences. What has been missing is an overarching vision and strategic support at the European level.

Investments to support the use of ADR for scientific discovery are increasing across the world. For example, some estimates suggest the US Department of Energy is spending \$1Bn per year in this area and Canada is spending \$200M. While Europe benefits from a strong academic legacy, investment in this area will be necessary to maintain its position.

Challenges

In 2022, the EU produced a fifth of the top 10% of global scientific publications. However, by some measures, the productivity of science has been declining.⁸⁷ ADR technologies offer an opportunity to supercharge scientific discovery and enhance European leadership in this vital area. The use of ADR for research and innovation is increasing. Some suggest that AI for science papers accounted for over 60% of AI publications in 2021⁸⁸.

Indeed, science is particularly well-suited for the application of AI since it has a well-defined scope and is *honest*—in the sense that the systems under study do not exhibit agency due to being observed, in contrast to many consumer and business applications. However, there remain many challenges. Lack of access to high-quality ADR tools, skills and talent,

84. <https://www.nature.com/articles/d41586-024-02842-3>

85. <https://www.living-lab.center/projects-in-europe>

86. <https://projects.research-and-innovation.ec.europa.eu/en/horizon-magazine/self-driving-vehicles-road-cheaper-sustainable-urban-transport-europe>

87. https://www.oecd.org/en/publications/artificial-intelligence-in-science_a8d820bd-en/full-report/is-there-a-slowdown-in-research-productivity-evidence-from-china-and-germany_cb8ae97d.html

88. https://sec-word-edit.officeapps.live.com/we/wordeditorframe.aspx?ui=en-US&rs=en-US&wopisrc=https%3A%2F%2Fliuonline-my.sharepoint.com%2Fpersonal%2Fkatli23_liu_se%2Fvti_bin%2Fwopi.ashx%2Ffiles%2F9bec6e3b83524908bc506ab18d18d968&wdenableroaming=1&mscc=0&wdoddb=1&hid=7086A4A1-10B8-C000-D30A-C14C6EA5BB8C.0&uih=sharepointcom&wdlcid=en-US&jsapi=1&jsapiver=v2&corrid=cb053b09-d39c-7073-d37d-e222c5a3cf2&usid=cb053b09-d39c-7073-d37d-e222c5a3cf2&newsession=1&sftc=1&uihit=docaspx&muv=1&qats=PairwiseBroker&cac=1&sams=1&mtf=1&sfp=1&sdp=1&hch=1&hwfh=1&dchat=1&sc=%7B%22pmo%22%3A%22https%3A%2F%2Fliuonline-my.sharepoint.com%22%2C%22pmshare%22%3Atrue%7D&ctp=LeastProtected&rct=Normal&wdorigin=BrowserReloadSharing.ClientRedirect&afdflight=44&csc=1&instantedit=1&wopiccomplete=1&wdredirectionreason=UnifiedSingleFlush#_ftn1

and research enablers all present barriers to its wider uptake. In response, next-generation ADR capabilities, new AI tools, mechanisms for knowledge sharing, and responsible and open research practices are all needed to create AI tools that are suitable for research and to promote their uptake.

Building AI capabilities to facilitate scientific research: More work is needed to provide enhanced theory and methods; in particular, foundation models specifically for science must be developed. These could be advantageous compared to current foundation models, which are trained on an unfocussed body of multi-lingual data, including sources such as blogs and discussion groups that do not well represent scientific knowledge and debate.

Progress is also needed to achieve a broader range of learning paradigms (neural network, federated, transfer), enhanced methods for logical inference (abduction or induction), automated reasoning (Bayesian, causal), and methods for multimodal data analysis (across natural language, dynamic, spatio-temporal, noisy and synthetic data).

Mathematics remains the standard method through which to express scientific knowledge. While there has been significant progress in automated theorem provers and in proof assistants, a deep understanding of mathematics within AI still appears to be missing.

There is also a need to embed domain knowledge (physical rules of an appropriate level, be it classical, probabilistic, quantum, heuristic, etc.) and to create domain-agnostic (or domain-embedded/specific) models for work in biology, materials, math and social sciences.

Data, corpora and derived models: While many scientific datasets exist, their availability, coverage and quality can be variable, which limits an AI's capabilities. Much valuable information is available via open access literature, preprint servers and patents. Other sources such as textbooks, grey literature, and archival journals may not be open access even if they are accessible. Data may be unlabeled, unstructured and of uncertain origin, so it is important to track provenance, rights and open science practices and mandates.

Existing data sets also have biases that may be well-recognised; for example, large population genomic datasets remain only selectively representative, limiting the generality of the models and range of feasible enquiries.⁸⁹ In many cases, the required data simply does not exist and may have to be generated *de novo*, using appropriate data types, formats, metadata and annotation and supported by suitable ontologies and interoperability standards. AI-powered automated methods and labs will likely play an important role in achieving this cost-effectively and with high reproducibility.

One important category of derived model is the digital twin, which can support simulations. Although this can be useful for prediction and tracking an external reality, it may not have the properties needed for verification of hypotheses, and other forms may also be needed.

Addressing the data challenge must link to other EU initiatives such as the EOSC⁹⁰ and health data space, the NFDI etc and address important aspects such as the lifecycle management, infrastructure, governance and access rights (see BT#3).

Human machine and agent collaboration: Suitable AI-based laboratory assistants would

89. <https://journals.plos.org/plosgenetics/article?id=10.1371%2Fjournal.pgen.1008302>

90. <https://www.eosc-pillar.eu/news/european-strategy-data-member-states-eosc>

be invaluable in supporting many research activities. Creating such intelligent assistants would require an extension of existing computational tools. Advanced natural language capability would enable them to read and understand the literature and mine datasets, generate hypotheses and write texts (papers, reports and summaries) for discussion, debate, review and comment. They could also generate experimental plans, lab protocols and scripting using active learning (RL, GA, other) and sharable ontologies (EXACT, ORA, etc.). This would help in driving lab automation, defining and designing materials (linking to facilities such as biofoundries, 3D printing, etc.), performing some of the more tedious practical lab tasks such as placement of controls,⁹¹ risk analysis, documentation and linking to data capture tools (ELN, LIMS, etc.), and preserving the metadata that is so often lost.⁹² As well as the ability to probe systems under study and analyse experimental data, these artificial assistants would need to work effectively alongside human users. This requires advanced communication skills to enable interaction with humans and other agents in the ecosystem⁹³.

Creating and sharing effective software tools: High quality research software is vital to ensure both the quality of scientific outputs from ADR systems and their wider usability. It plays an important role in facilitating the reproducibility of research and in contributing to an environment of open science, applying principles such as FAIR.⁹⁴ The methods developed need to be made available and deployed as tools that are scalable, with mechanisms to support sharing, adoption, diffusion, support and maintenance in a sustainable manner. Models enabling value capture for methods, data and other models need to be further developed.

Tools to support laboratory automation: While many existing datasets have been collected using largely manual methods, some specialised facilities provide highly automated services for specialised tasks such as screening molecules for biological activity.⁹⁵ Improved automation could play an important role in expanding the range, size and quality of the datasets needed to build models. Validating scientific hypotheses also requires lab facilities. This process would benefit from greater automation and has formed the basis for the many self-driving labs now emerging⁹⁶ but there remain bottlenecks in terms of throughput and, in some cases, the quality of data.

This problem would be greatly eased by improved AI-enabled lab tools across the workflow. Existing industrial robots are typically positioned for *dull, dirty and dangerous* tasks, but in the lab *delicate and data-rich* tasks are required.⁹⁷ The industry produces excellent tools for handling liquids in lab and clinical diagnostics; for example, Europe produces most blood testing systems⁹⁸. However, many tasks have proved hard to automate. The handling of materials such as powders, for instance, remains largely underaddressed. New concepts

91. <https://www.unchainedlabs.com/wp-content/uploads/2021/11/AN-Integrating-automation-and-DOE.pdf>

92. <https://rjpn.org/ijcspub/papers/IJCSPI24D1092.pdf>

93. <https://www.nature.com/articles/s41598-023-30938-9.pdf>

94. <https://www.opensciencefair.eu/lightning-talks/enhancing-fair-compliance-in-research-data-infrastructures-insights-from-applications-of-the-rda-fair-data-maturity-model-and-the-f-iji-automated-fair-data-assessment-tool>

95. <https://www.pivotpark.com/campus/>

96. <https://www.nature.com/articles/s41467-025-59231-1.pdf>

97. See euRobotics, <https://www.labbulletin.com/articles/public-launch-the-eurobotics-roadmap-laboratory>

98. See Roche, Siemens, Grifols: <https://www.marketsandmarkets.com/ResearchInsight/blood-screening-market.asp>

in laboratory robotics (such as soft robotics and traceable robotics⁹⁹) could be enablers. Furthermore, the process of lab digitalisation has been slow, although the connectivity of lab instruments and robotics is being addressed thanks to emerging interoperability standards¹⁰⁰, open data formats (FAIR¹⁰¹) and suitable ontologies.

In many cases the samples of interest (for data collection and for carrying out experiments) are too large, too remote or too valuable to transport into the lab. This problem is largely limited to some space missions, but is increasingly relevant in marine science when monitoring key species and interactions¹⁰² particularly as 70% of the earth's surface and perhaps 90% of the liveable space remains under-explored¹⁰³. Mobile robotic labs that can go to where the action is and autonomously sample and probe over extended time and space would play an important role in this future exploration process (see BT#4).

Empowering researchers to use ADR in their science: Delivering a step-change in scientific productivity through the use of AI will require widespread adoption and use of AI tools across research in both academia and industry. This in turn will fuel the demand for mechanisms that support interdisciplinary or cross-organisational collaboration, upskilling, and access to AI knowledge, which will help domain researchers find and make use of AI tools that are appropriate for their needs. This can include training and supporting individuals, teams, organisations and sectors across industry, academia and civil society, and showcasing innovations and transfer into practice.

Coordination, collaboration and cross-fertilisation: Advances in understanding and addressing major societal challenges require expertise and methodology across multiple existing disciplines. The tools required for the necessary cross-disciplinary collaboration do not yet exist, and AI methods capable of translating across domains and supporting collaboration between expert teams are needed.

At the same time, similarities between the processes deployed for research and those deployed for development and innovation—which involve sequential decision-making in complex systems—mean that many AI-for-science methods could be rapidly repurposed for innovation in product development across sectors, which would require mechanisms to support the wider use of such methods.

Progress also requires creating a vibrant interdisciplinary community where domain research and ADR innovations are brought together to enhance scientific discovery, building a cross-domain community of stakeholders involved in the co-creation of tools to ensure that they remain effective, safe and integrated into evolving research practices. Effort should also be made to extend the range of disciplines benefiting from this work to include the arts and humanities and emerging scientific disciplines such as regulatory science (see BT#6).

Identifying well-defined challenges to motivate the field is similar to the protein sequence-structure problem addressed by *AlphaFold*. The **Nobel Turing Grand Challenge** (NTC¹⁰⁴

99. <https://sila-standard.com/>

100. <https://sila-standard.com/>

101. <https://www.go-fair.org/fair-principles/>

102. <https://www.embrc.eu/marine-monitoring-and-modelling/>

103. <https://www.usgs.gov/science/science-explorer/ocean/ocean-discovery>

104. https://sec-word-edit.officeapps.live.com/we/wordeditorframe.aspx?ui=en-US&rs=en-US&wopisrc=https%3A%2F%2Fliuonline-my.sharepoint.com%2Fpersonal%2Fkatli23_liu_se%2Fvti_bin%2Fwopi.ashx%2Ffiles%2F9bec6e3b83524908bc506ab18d18d968&wdenableroaming=1&mscc=0&wdoddb=1&hid=7

aims to develop “AI systems capable of making Nobel-quality scientific discoveries highly autonomously at a level comparable, and possibly superior, to the best human scientists by 2050”. Progress towards this challenge will unleash the deep potential of AI to search for and discover the fundamental structure of our world, and could potentially transform the world through accelerated technological development, but many details remain to be fleshed out.

Infrastructure: Access to the appropriate infrastructure is essential to attaining the potential of ADR technology for the sciences. This includes access to suitable HPC with sufficiently high bandwidth and low latency communication, as well as suitable data spaces. It also relies on access to suitable laboratory facilities both to collect the missing real-world datasets and to test hypotheses all at sufficient speed and quality.

Ensuring responsible research and innovation: Responsible research and innovation (RRI) practices are necessary to ensure ADR is used safely and effectively across the sciences. Europe has been a leader in the field of RRI, creating frameworks to support researchers looking to better understand the design of their studies, the implications of their work, and how its outputs and methods can be aligned with diverse societal interests. This includes improving stakeholder involvement in the co-creation of methods that respect established rights¹⁰⁵ and ensuring protection against unintended consequences, such as the development of new molecules with enhanced toxicity.

Scientists are naturally skeptical and will need time to build confidence in the trustworthiness of use of AI systems. With AI offering a tool that allows new uses of data and the generation of new insights across disciplines, wider adoption of RRI approaches will be necessary to ensure good governance and compliance with expectations of privacy and security around subjects such as sensitive medical data.

Applications of social and economic significance and potential impact

Science has profoundly impacted our understanding of the world. ADR technologies offer an opportunity to accelerate scientific discovery, enhancing European leadership in vital areas such as healthcare and cancer research, energy, agriculture and food insecurity, sustainable production (alternative raw materials to sidestep reliance on critical minerals, the bioeconomy, and recycling) as well as the environment, cultural heritage, and more. Europe’s broad talent base across diverse scientific research domains means that it is well-positioned to benefit from widespread application of AI in science in the following ways:

Impact on society: The main gain of AI for science will be an increase in human knowledge as a societal good capable of being shared internationally. The development of evidence-based personalised digital twins would transform healthcare for individuals throughout their lives, allowing interventions to be modelled, to the benefit of all citizens. The resulting insights will also provide new ways of understanding climate change and the natural

086A4A1-10B8-C000-D30A-C14C6EA5BB8C.0&uih=sharepointcom&wdlcid=en-US&jsapi=1&jsapiver=v2&corrid=cb053b09-d39c-7073-d37d-e222c5a3cf2f&usid=cb053b09-d39c-7073-d37d-e222c5a3cf2f&newsession=1&sftc=1&uihit=docaspx&muv=1&ats=PairwiseBroker&cac=1&sams=1&mtf=1&sfp=1&sdp=1&hch=1&hwfh=1&dchat=1&sc=%7B%22pmo%22%3A%22https%3A%2F%2Fliuonline-my.sharepoint.com%22%2C%22pmshare%22%3Atrue%7D&ctp=LeastProtected&rct=Normal&wdorigin=BrowserReloadSharing.ClientRedirect&afdflight=44&csc=1&instantedit=1&wopiccomplete=1&wdredirectionreason=UnifiedSingleFlush#_ftn2

105. <https://awis.org/resource/ethical-challenges-medical-research-henrietta-lacks-hela-cell-line/>

environment as well as Europe's rich cultural heritage, and, ultimately, our own perception of ourselves.

Impact on end user industries: The insights gained from the application of ADR technologies are likely to accelerate innovation. Already, the internal units of large corporations are exploring these opportunities. This will lead to new ways of solving problems, novel industrial materials, consumer products and services across different value chains, keeping innovation in Europe and creating new high value jobs.

Impact on the research tools sector: Europe has a world-leading position in laboratory tools,¹⁰⁶ with an existing innovation ecosystem, but is still poorly digitalised and yet to fully exploit advances in AI [ref]. The application of ADR technologies for science is already leading to the first new companies providing software and hardware tools, embedded services and developing their own IP [ref].

Impact on ADR technologies: Success in addressing the challenges in science will feed back into AI research, thanks to the availability of vast high-quality datasets and the impetus to push the boundaries of ADR technologies and communicate the outcomes in a free and open manner, independent of commercial interests. This can provide a shining example of the use of AI and robotics in a way that does not threaten people's jobs, privacy or autonomy.

Roadmap

A research agenda requires advances in ADR theories, methods, and approaches to deliver powerful analytical tools that can analyse complex systems, leveraging both data-driven and domain knowledge across the sciences.

Areas that can help deliver these advances are summarised in the table below, with the following potential cross-linkages between the other Big Ticket items:

- » Foundational AI models (BT#1) across types of learning, modelling, inference and reasoning including causal AI, multi-modality and domains, communication and collaboration between human-agents-other systems
- » Trust and compliance (BT#2, BT#6) for systems that are IPR-sensitive and regulation-aware to allow knowledge sharing in industry, respect content creator/owner rights, ensure reproducibility and transparency (BT#6), and avoid ethics dumping during data cleaning
- » Data and models (BT#3) with clear access rights
- » Digital twins at different levels (cell processes, organisms and populations (BT#2), production processes) that allow embedding of domain knowledge

106. https://sec-word-edit.officeapps.live.com/we/wordeditorframe.aspx?ui=en-US&rs=en-US&wopisrc=https%3A%2F%2Fliuonline-my.sharepoint.com%2Fpersonal%2Fkatli23_liu_se%2F_vti_bin%2Fwopi.ashx%2Ffiles%2F9bec6e3b83524908bc506ab18d18d968&wdenableroaming=1&mscc=0&wdoddb=1&hid=7086A4A1-10B8-C000-D30A-C14C6EA5BB8C.0&uih=sharepointcom&wdlcid=en-US&jsapi=1&jsapiver=v2&corrid=cb053b09-d39c-7073-d37d-e222c5a3cf2&usid=cb053b09-d39c-7073-d37d-e222c5a3cf2&newsession=1&sftc=1&uihit=docaspx&muv=1&ats=PairwiseBroker&cac=1&sams=1&mtf=1&sfp=1&sdp=1&hch=1&hwfh=1&dchat=1&sc=%7B%22pmo%22%3A%22https%3A%2F%2Fliuonline-my.sharepoint.com%22%2C%22pmshare%22%3Atrue%7D&ctp=LeastProtected&rct=Normal&wdorigin=BrowserReloadSharing.ClientRedirect&afdflight=44&csc=1&instantedit=1&wopiccomplete=1&wdredirectionreason=UnifiedSingleFlush#_ftn3

- » Robotics (BT#4) for manipulation, especially non-contact, embedded sensing, soft robotics, autonomous field robotics, which are more agile and have a higher sampling rate (100 vs 1Hz)

Areas for progress Near-term	Medium-term goals	Long-term outcomes
<p>Enhanced AI theory and methods</p> <ul style="list-style-type: none"> » Enhanced methods for logical inference (abduction or induction), and automated reasoning (Bayesian, causal) » Foundation models and generative AI for science; » Methods able to embed domain knowledge (laws, heuristics) » Tools to support lifecycle management of corpora, annotation and ontologies » Tools to support simulations (digital twins, other) » methods for multimodal data analysis (spatio-temporal, dynamic, Natural Language); » Natural language tools able to ingest literature and generate hypotheses, summaries, plans and code » Tools to drive lab automation (devise experiments, generate protocols, execute) » Lab robotics (non-contact solids and manipulation of powders, as sources of data, autonomy, soft robots) 	<p>Technically advanced and trustworthy ADR for science tools, methods, and approaches that:</p> <ul style="list-style-type: none"> » Extract insights from data by integrating different data types, building simulations, and deploying new learning and reasoning strategies » Bridge data-driven and domain knowledge, through integration of domain knowledge and causal understandings » Support explainability and interpretability of their findings <p>Collaborative ADR research assistants that support human decision-making by:</p> <ul style="list-style-type: none"> » Identifying and responding to goals of human researchers » Communicating the results of analysis and let researchers interpret results » Interfacing effectively with different users or systems, processes <p>Establishing well-defined grand challenges to motive the field to delivery applications of social and scientific impact</p>	<p>Safe and effective ADR technologies for research and innovation that are technically advanced, reliable in deployment, and used responsibly across user communities</p> <p>Accelerated progress towards addressing grand challenges, through the effective deployment of AI for research and innovation and evidence thereof</p> <p>A vibrant, interdisciplinary community using AI for innovation, and in turn feeding insights from the use of AI into the design of next-generation AI methods</p>

Areas for progress Near-term	Medium-term goals	Long-term outcomes
<p><i>Next generation ADR technologies (see Big Ticket#1) such as:</i></p> <ul style="list-style-type: none"> » Improved learning paradigms (neural network, federated, transfer, etc.) » Collaborative AI including multi/human/robot agents » Human-centric AI » User interface design <p>Interdisciplinary collaborations to develop novel ADR for science research agendas in focused domains and cross-sectoral collaborations to deploy AI for innovation tools in industry</p> <p>RRI processes for responsible research and innovation in ADR for science</p>	<p>An enabling environment in ADR for science that transfers tools and practices across domains</p> <p>Wider uptake of these technologies can be supported through:</p> <ul style="list-style-type: none"> » Access to tools, data, computation and storage » Opportunities for researchers to gain skills in AI » Forums to build connections between domain and AI researchers » Mechanisms to attract and develop talent across technologies and disciplines 	

3.2.6 Big ticket #6: Research, innovation, and tools for compliance

State of the art

To enable companies and organisations to comply with the new regulations being rolled out, significant efforts are needed to lower the thresholds and simplify doing the right thing. This means there is a major need for innovative tools for compliance. The scope of compliance includes general aspects like trust, privacy, IPR, ethics, and environmental aspects, but also sector-specific rules and regulations.

In recent years, the EU has brought in a wide range of laws and regulations that impact ICT and the digital sphere. While some of these, such as the Digital Service Act and Digital

Market Act, only apply to the largest (and mostly non-EU) digital platforms, others are of specific concern for European enterprises innovating in the ADR space. To support the latter's integration of compliance practices into product development and deployment practices, the focus should be on two groups of regulations. The first spans the compliance requirements of the AI Act and its interaction with both horizontal regulations, specifically GDPR, the Cyber Security Act and the Data Act, and relevant sectoral regulation, for example the Machinery Regulation and the Medical Devices regulation. The second focuses on how regulation that is supportive of public-private partnerships around data sharing and public procurement of digital products can be leveraged to accelerate regulatory learning across value chains and for pooling data assets in support of compliant forms of ADR innovation. This is a very active area, with many ongoing activities.

Data protection is a well-established area and is receiving additional attention as new complexities emerge around the use of personal data in training and using AI. The BDVA has summarised the current prominent topics in data protection via its "CURRENT HOT TOPICS IN DATA PROTECTION" position paper.¹⁰⁷ Open, machine readable meta-data vocabularies for GDPR are under ongoing development in the W3C Data Protection and Privacy Vocabulary and Control Community¹⁰⁸ This was based on research from several EU funded projects including SPECIAL (No. 731601), TRAPEZE (No. 883464) and PROTECT (No. 813497) and has seen multiple implementations in other projects and products, and as standards at ISO/IEC SC27 TS 27560109 and IEEE P7012110.

Research on trustworthy, ethical or responsible AI underpins many of the current practices that are under consideration for risk management and the protection of health, safety and fundamental rights under the AI Act. These include the AI Trust Alliance initiative¹¹¹ for an AI trust label, which is a joint collaboration between VDESPEC, Positive AI, Confiance. ai and IEEE; the BDVA ETAMI¹¹² task force aimed at creating processes and tools to enable ethical, trustworthy and legal AI; and the TAILOR project (Trustworthy AI based on Learning, Optimisation and Reasoning). Some relevant publications resulting from this work are the Strategic Research and Innovation Roadmap¹¹³ v1, the Handbook for Trustworthy AI¹¹⁴ and the OECD catalogue¹¹⁵ of tools and metrics for trustworthy AI. AppliedAI¹¹⁶ is also driving an effort to pull in tools (such as model cards) to automate the use of trustworthiness standards during the development process.

AI standards are under active development in support of regulatory compliance: CEN-CENELEC JTC21 addresses a range of standards related to compliance with the AI Act, including an initial set of 10 requested by the EC as candidate harmonised standards, compliance to which may offer a presumption of conformity under the Act, and other related standards, including WG4 on AI Trustworthiness Characterisation, and others on

107. https://www.bdva.eu/sites/default/files/BDVA%20DataProtection%20PositionPaper_November2022.pdf

108. <https://w3c.github.io/dpv/2.1/dpv/>

109. <https://www.iso.org/standard/80392.html>

110. <https://standict.eu/standards-repository/ieee-p7012-standard-machine-readable-personal-terms>

111. <https://alliancefortrustinai.org/>

112. <https://etami.org/>

113. <https://tailor-network.eu/research-overview/strategic-research-and-innovation-roadmap/>

114. <http://tailor.isti.cnr.it/handbookTAI/TAIORITY.html>

115. <https://oecd.ai/en/>

116. <https://applied-ai.com/>

professional competencies in AI ethics. WG1 of JTC21, with EC support, offers a collection of public resources related to JTC21 standardisation¹¹⁷ including a periodic inclusivity newsletter¹¹⁸. ISO/IEC JTC1 SC42 Artificial Intelligence has been developing international standards for AI since 2018. Several of these have been converted to European Standards in support of the JTC21 effort. One of the ten harmonised standards requested by the EC, focused on AI logging, is being developed in SC42 in collaboration with JTC21.

Other legal compliance research looks to increase the availability of domain-specific languages like FLINT¹¹⁹: these formalise norms and regulations and make them actionable for computer systems. Intellectual property is an issue for Generative AI¹²⁰ but it is also a compliance issue. Legal validity of training data and compliance with IPR licenses should be demonstrable.

Security issues related to AI need to be addressed under the AI Act, GDPR and the Cyber Resilience Act. They are being addressed in a number of new initiatives: cyber threats are mapped out by the Atlas initiative¹²¹, Open Worldwide Application Security Project (OWASP) initiative, which lists projects that support handling AI cyberthreats¹²², the ELSA Strategic Research Agenda, which reports on Secure and Safe AI¹²³, and the OASIS Open forum, which is hosting the Coalition for Secure AI ecosystem.¹²⁴ NIST also addresses AI security and resilience research¹²⁵ and provides a taxonomy of threats and mitigations¹²⁶, and ENISA has surveyed standardisation projects in the area¹²⁷.

Challenges

Establishing trustworthiness

The certification or labelling of AI trustworthiness by third parties can build upon existing sectoral approaches, but as a horizontal technology broadly applied across sectors, AI can benefit from tailorabile horizontal measures for verifying trustworthiness. However, trustworthiness of AI applications is a multidimensional concept, related to formal verification, user perception and competency levels of those wishing to be trusted. The challenge is to establish metrics and labels for trustworthiness that are relevant and useful: to identify the key attributes of trust relevant for a specific AI application and define the method to measure compliance with those attributes in a unified aggregate confidence score. Another challenge lies in defining (domain-specific) thresholds for these metrics and attributes that establish a tangible and objective boundary between conforming and non-conforming AI systems. A related challenge surrounds the public understanding of trustworthiness, as the relation between official labels, measures of trustworthiness, and the actual trust that end users and the audience assign to AI systems is not straightforward.

117. <https://jtc21.eu/>

118. <https://jtc21.eu/newsletter>

119. <https://dl.acm.org/doi/10.1145/3425898.3426958> “a domain-specific language for executable norm specifications”

120. <https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/>

121. <https://atlas.mitre.org/>

122. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>

123. <https://elsa-ai.eu/sra/>

124. <https://www.coalitionforsecureai.org/>

125. <https://www.nist.gov/artificial-intelligence/ai-research-security-and-resilience>

126. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2025.pdf>

127. <https://www.enisa.europa.eu/publications/cybersecurity-of-ai-and-standardisation>

Interoperability and standardisation of Privacy Enhancing Technologies (PETs)

Many PET applications are implemented as point solutions. When the adoption of PETs increases, the resulting vendor lock-in may hamper further adoption. Interoperability of PETs, for instance in the context of data spaces, is therefore fundamental. Furthermore, many PET solutions are hand-crafted, based on a manual assessment of information requirements. An automated procedure or method to analyse the leakage threats and 'compile' a suitable set of PET functionalities does not exist. It is also challenging to define suitable metrics and norms for compliance with a single legal framework (e.g. data protection and privacy, IPR, competition law, ecodesign/environmental, AI regulation, etc.) and moreover for compliance with combined legal frameworks (e.g. in the context of data sharing, data spaces, product and technology development and related supply chains). Compliance should be created and assessed/measured depending on the TRL level of the developments, while also realising the benefits of a privacy/data-protection/ethics-by-design approach and exploring the role of risk assessment in combined legal frameworks. Finally, a related hurdle is the user perception of privacy, which influences both the user's behaviour and the user's trust in the system. This interaction is not yet well understood.

Socially and Environmentally Sustainable Business Models

Although legislation exists to avoid unethical behaviour, socially and environmentally unsustainable practices, and misuse of personal data, many end users still blindly accept the Terms & Conditions of services when the offered advantages are attractive, without verifying compliance with this legislation. The challenge is to create business models that derive their value from an ethical approach, sustainable practices, and privacy awareness, rather than winner-takes-all data collection, the use of underpaid "clickworkers" outside the EU, and large environmental footprints. More investment is required in understanding and explaining unsustainable practices and their impacts, but also to support innovative practices that have more sustainable outcomes.

Life-cycle management of Adaptive AI systems

The existing framework for product regulation in the EU, which is now extended by the AI Act to high-risk AI systems, assumes a stable product design that can be certified through pre-market verification testing. This ensures certification remains valid as the system is deployed into the market. However, AI systems that adapt to new inputs or situations may quickly start operating outside their originally designed bounds, leading to decreased performance and trustworthiness, and may trigger compliance issues for customers, as well as costly re-certification. There is a need for *incremental and evolutionary qualification*, as there is no strict separation between the design and operational phases of dynamic AI algorithms. Qualification cannot be a separate, stand-alone, after-the-fact activity on final products because of the continuous evolution of the system and its environment. *Data quality, filtering and availability* are related challenges in complex data value chains. There is a need for new tools to screen incoming 'data products' for legal provenance, quality, and reproducibility.

Automated compliance and compliance-by-design

A key challenge is facilitating automated compliance with the legislation applicable to the design, development, testing, and overall engineering life cycle of ADR systems, and defining how models and architectures can support compliance and the encoding of legislation.

Compliance for complex systems, including complex supply chains, the portability of systems, and post-market surveillance, are significant challenges that involve analysing systems, data value chains, data spaces, and federated platforms.

Provenance of training data

The origin of training data for AI models must be transparent, both for legal (IPR) and moral reasons, as well as reasons of trustworthiness. The challenge is how to store and process provenance information, including the right to be forgotten ('unlearning' of data). Rectifying or removing data from learned models can be a requirement from a compliance point of view. A related challenge is the possibility that sensitive training data may appear in Generative AI systems. Care must be taken when fine-tuning an AI model to a specific task with sensitive/non-public data to ensure sensitive information is not memorised. Equally, just because an AI model is trained on public data does not mean it is prevented from learning sensitive attributes of an individual. Care must therefore be taken that a model cannot collate public information to come to sensitive conclusions about individuals that would be difficult/impossible for a human to reach. These concerns require collaborative exchange and management of provenance metadata along the entire AI value chain to include the source of the data used to train, validate and test AI.

AI Security, Safety, Resilience

AI systems and components have become increasingly pervasive in our IT landscape, making them part of the attack surface for adversaries. In order to ensure reliable and resilient infrastructure, ecosystems, and products, we need to ensure that cybersecurity concerns are an integral part of research, innovations, and deployment. At the same time, there are distinct attacks and AI-induced risks for which threat and risk models need to be established, and defenses need to be developed, while AI-enable and integrated systems need to meet cybersecurity standards. Furthermore, aspects of systemic risk, including cybersecurity capabilities and loss of control, need to be researched and mitigated.

Solution areas & applications

Towards certification of trustworthiness

As AI systems become more complex, their trustworthiness may become opaque for designers, professional users, end-users and the public. Both technical and perceptual aspects of trustworthiness are relevant. Attention must be paid to the avoidance of over-trust and under-trust: both are dangerous. The goal is to design an adaptable labelling system for trustworthy AI systems that is based on quantifiable metrics, and to verify the feasibility of such a system in a domain where trustworthiness of AI systems is important. For AI systems that adapt in deployment, such verification must be conducted on an ongoing basis post deployment, so high levels of testing automation and active compliance reporting are required.

Compliance-by-design ADR systems

Since requirements and regulations will change frequently over time, there is a need to minimise manual effort to establish the compliance of systems. Different AI applications have different trust attributes; therefore, an automated characterisation of the relevant ones, including an assessment of their values, will be needed. Different regulatory learning

activities such as regulatory sandboxes, testing in real-world conditions and Testing and Experimentation Facilities (TEF) should be able to make use of, and contribute to, common information resources such as benchmarks, risk and mitigation patterns and testing protocols using either FAIR open data resources or controlled access shared data spaces¹²⁸.

The goal is to create a '*compliance by design*' approach that enables the design of ADR systems in which compliance is *baked in* so that it is enforced *ex ante*. This approach defines and demonstrates a system that can express compliance requirements in a formal model, enabling the *post hoc* automated assessment of the system's behaviour and features during the whole life-cycle of the application. The compliance by design ADR system should also be able to handle changing legal requirements (evolutionary compliance). Demonstration of these compliance concepts should be prioritised in areas where the handling of personal or sensitive data is a major concern, such as the health domain, and/or where domain-specific norms and regulations are in place for mission critical automated systems, such as manufacturing or mobility.

Trustworthy AI in business

ADR systems should be specified and procured/designed to ensure trustworthy behaviour in operation. Trustworthy in this case is not only 'compliant with regulations' (such as AI Act and GDPR, including PETs), but also to be compliant with the organisation's norms and values. The deployment of ADR systems to support products or services should be monitored and updated when circumstances change, in order to maintain the trustworthy operation of these systems. The goal is life-cycle management of AI systems that ensure trustworthy and privacy-aware operations for businesses. This includes the ability for data holders to exert control over their data and/or to monetise it, which makes transparency attractive from the side of the service provider, and encourages services that do not rely on tracking of user behaviour but that offer content and user-centric advertising. The approach should support the interoperability and standardisation of PETs and be demonstrated in domains where end-user trust and acceptance are important.

IPR-aware training of AI

ADR systems must be trained on demonstrably IPR-compliant curated data sets. These data products can also be created from digital twins. ADR systems that are based on (large) sets of training data should also be able to handle the provenance of the training data in an appropriate way. The system should also be able to update the model when IPR or other aspects of the training data prevent parts of the data from being used.

Experimentation as a new way of regulation

Experimental regulation, exemplified by initiatives like regulatory sandboxes, which could be associated with AI testing and experimentation facilities (TEFs), is crucial for nurturing innovation, especially among SMEs. These frameworks provide SMEs with a controlled environment to test new ideas, products, or services, offering flexibility without the risk of severe repercussions. It also presents a valuable perspective for the regulators of cutting-edge technologies, combined with the opportunity to ensure compliance by design for the participants.

128. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/14625-Apply-AI-Strategy---strengthening-the-AI-continent/F3563742_en

By mitigating compliance burdens during the testing phase, regulatory sandboxes accelerate the development of cutting-edge solutions, empowering SMEs to explore uncharted territories. These initiatives foster a dynamic exchange of ideas between businesses, regulators, and stakeholders, contributing to a culture of continuous learning and collaboration. However, the lack of scale for these types of innovation facilitators poses a significant challenge that requires appropriate solutions such as digital sandboxes or large-scale testing facilities. TEFs can provide access to both physical and virtual infrastructures, enabling the large-scale testing of AI-based systems under realistic conditions. It also facilitates engagement with end-users and stakeholders, supporting comprehensive conformity assessment and enhancing the overall reliability and trustworthiness of AI solutions. A key challenge is how to define sustainable business models for TEF, and at the same time, promote open access to these infrastructures.

Systems view, MLTrustOPs, and an Interdisciplinary Approach at the intersection of Cybersecurity and AI (AI Security, Safety, Resilience)

There is a long history in cybersecurity on risk and threat modelling—which are unfortunately often employed too late or only as an afterthought in practice. AI and ML are prone to falling into the same habit if decisive action is not taken to carry over lessons, methodology, and insights from cybersecurity that take a systems approach, accompanying the whole life cycle in an MLTrustOps fashion. Beyond the translational aspect, the field of AI Security needs to be developed, and effective means of making AI or hybrid systems more secure, resilient, and reliable must be built. Beyond evaluation, testing, red teaming approaches, methods with guarantees and with a by-design security paradigm should be favoured. In addition, AI and GenAI technology should be leveraged for cybersecurity, defence and resilience – potentially also to offset the anticipated empowerment of adversaries by AI.

Socio-economic impact

Solutions should enhance the quality of AI-based products and services in terms of trust, privacy, and security; specifically, they should augment consumer trust and acceptance in EU products while guaranteeing ethical products that preserve the rights of EU citizens, decrease the number of data leaks, and address citizens' privacy concerns. The costs of compliance, including auditing, should fall as a result of these solutions. The compliance or qualification assessments of products should also be accelerated, so that the time-to-market of innovative AI-based products and services is reduced, whilst still being compliant with applicable regulations.

Roadmap

We must create a setup where transparency and compliance in AI systems is by design. This can only be reached by consistent adaptation of processes in AI research, development, and deployment. Indeed, the whole value chain needs to be addressed and aligned to reach this goal. Many of the necessary actions to obtain this are clear and established as today's best practices. Such requirements include software documentation practices; training and validation methods; and so on. At the same time, there is a lack of knowledge and best practices in some key areas, such as data documentation methods; AI life cycles models; human oversight at scale; and monitoring. Additionally, auditing approaches are not yet validated and accepted. Regulatory sandboxing and TEF-focused approaches

must be developed to strengthen research in this area. Furthermore, there is potential for widespread diversification of such regulatory learning activities across member states as well as at the EU level. Therefore, efficient and transparent mechanisms are needed for the timely sharing, analysis and comparison of outcomes of different regulatory learning experiences so that ongoing revision of the various codes, regulations, guidelines and standards are grounded in a shared but contextualised base of evidence.

Taking the AI Act as an example, the sharing of compliance resources and knowledge requires the engagement of value chain actors, affected stakeholders protected under the Act, and a complex array of oversight authorities. These actors then need to engage in different regulatory activities that may be relevant to specific types of AI systems defined in the Act, and where compliance may involve different sets of protections. Efficient and transparent coordination of these myriad activities can be supported by locating them in a common parameterised regulatory learning space. The most direct areas to note are those that are already laid down in horizontal regulations, such as GDPR, the AI Act and the Cyber-Resilience Act. There are still considerable challenges around these in terms of standardisation and adoption, but the starting points are relatively clear. The first steps in these areas may become visible quite soon; domain-specific compliance applications are a special case.

The area of compliance that will take more time to develop revolves around aspects like trustworthiness, environmental aspects and general ethical considerations. While a lot of work has already been done to codify these aspects, no clear best practice or standardised approach has emerged yet. The first steps here will be to establish metrics, definitions, and standards before compliant ADR systems can be created in an effective way. Methodology is also important to support these developments. The concept 'compliance by design' must be elaborated, alongside efforts to dynamically update systems and to embed them in business models.

Short-, medium-, and long-term objectives

Areas for progress today/near term	Medium-term goals	Long-term outcomes
<p>Evaluate landscape of tools and approaches for compliant ADR systems:</p> <ul style="list-style-type: none"> » enterprise, IT & solutions architecture models for ADR systems » life-cycle models for AI » Data and AI documentation methods » AI oversight and risk management methods » AI system monitoring methods » Formalised tools and languages of norms / regulations » Metrics for privacy, trustworthiness » Methods, standards and tools to track data provenance (IPR) 	<p>Trustworthiness by design for AI: human-centric AI systems that embed the technical foundations of trustworthy AI across industrial applications, demonstrating</p> <ul style="list-style-type: none"> » safety, robustness, reliability, resilience, » explainability, auditability (e.g. transparency of training data), » security, fairness, sustainability and energy-efficient implementations., acceptable risk. <p>Research on and development of methods to create foundations for compliance by design:</p> <ul style="list-style-type: none"> » enterprise, IT & solutions architecture models for AI developers, providers, users, importers and distributors » life-cycle models for AI and dynamic compliance » data and AI documentation methods » human oversight methods » AI and monitoring methods » Development of formal verification methods for different learning paradigms (e.g., reinforcement learning, generative AI) » Understanding, preventing, and mitigating systemic risks 	<p>Establishment of governance frameworks for trustworthiness by design</p> <p>Establishment of architect processes for ethical AI.</p> <p>Socio-economic impact of compliant ADR systems: increased end user trust and acceptance is measurable; business models 'beyond compliance' are observed</p>
<p>Review sandboxing methods:</p> <ul style="list-style-type: none"> » overview of AI sandbox methods » evaluate shortcomings sandboxing methods for the AI landscape 	<p>Pilot:</p> <ul style="list-style-type: none"> » AI life-cycle models » data documentation methods » human oversight methods » AI system monitoring methods » Sandboxing: » Develop and pilot of AI sandboxing methods: » harmonisation within the legal landscape 	

Areas for progress today/near term	Medium-term goals	Long-term outcomes
Auditing methods: <ul style="list-style-type: none"> » evaluate existing auditing approaches » evaluate shortcomings w.r.t. legally compliant ADR systems » evaluate shortcomings w.r.t. trustworthy, secure and privacy preserving ADR systems 	TEFs: <ul style="list-style-type: none"> » Develop sustainable business models for TEFs in high-risk sectors » Define role of TEFs: only experimentation or certification facilities? Auditing methods: <ul style="list-style-type: none"> » develop and pilot AI-specific auditing methods 	

4 Conclusions

The ADR partnership aims to enable a responsible ADR-powered green digital transformation for an attractive, sustainable, prosperous, secure and resilient multicultural society, based on European values, with the highest living standards in the world. Adra focuses on the cross-section between AI, data and robotics, with the long-term goal of achieving convergence in these areas. However, as ADR is not yet an area of its own, the trends and gaps analysis is based on the three existing communities, which all have significant overlaps and interactions.

The Strategic Research Innovation and Deployment Agenda aims to build on the fundamentals of Europe cementing its world-leading status in ADR to both enhance the revenue-generating potential for companies and enrich our society as a whole. There is currently huge interest in generative AI, especially large language models. The last year has seen tremendous progress with both text and image-generating models. The next step is to generate sound and video of comparable quality and then to generate true multi-media productions. Europe has the potential to take a more active part in the development of large AI-models. The main limiting factor has been scale. Europe is good at getting things done on a small scale but has so far not been able to operate at the same size and complexity as other parts of the world. It is also important to focus on the “next big thing”, and here we see a great opportunity for Europe as we focus investment on the supply of AI, data and robotics for enhancing strategic autonomy and resilience.

To achieve the vision of Trustworthy ADR, Europe needs to match its roadmap with appropriate investments in innovation along the supply chain to ensure that companies and organisations are capable of controlling and supporting the development of capabilities, skills and technologies for the ADR systems that we want, in accordance with our values. Europe has both the responsibility and the resources to achieve this, but we have to invest and dare to take calculated risks. If we do, the potential for global impact and a significantly better world for people, economy and climate is high.





The AI Data Robotics
Association



adr-association.eu