Adra Strategic Research, Innovation, and Deployment Agenda 2025-2027

November 2023

https://adr-association.eu/

Version Nº: 1.0beta

Date: 06/11/2023

Document information	
Title:	**Adra Strategic Research, Innovation, and Deployment Agenda 2025-2027**
Version Nº:	**1.0beta**
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), Agirre Ibarbia (Tecnalia), Samuel Kaski (Aalto), Yiannis Kompatsiaris (CERTH-ITI), Marko Grobelnik (IJS), Stefan Leijnen (Adra), 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).**

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”, November 2023, Brussels - Adra

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

Glossary of Terms	
Item	**Description**
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**
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**
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	**Strategic Development Goals**
SME	**Small and Medium Organisations**
WP	**Work Programme**

Foreword

We are living in a moment of global changes and challenges, intertwined with the opportunities of fast developing AI, data and robotics (ADR) technologies. These powerful technologies have a lot of potential, but also risks that need to be mitigated and managed both 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 the European ADR research and industry and the public sector, promotes the convergence of ADR to tackle the global challenges and helps to build a more sustainable society in Europe and globally. This partnership is also committed to promote a safe use of ADR technologies.

The speed of innovation is relentless, ADR have already had a strong impact on our daily lives. We need to ensure that European companies and its public sector can benefit from the tech revolution and ensure European strategic autonomy over and within these ADR technologies. This goal can be achieved only through a strong and well-connected innovation ecosystem.

This SRIDA 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, tackling the big challenges of our time.

About this document

The purpose of this document is to describe the strategic position of the AI, Data and Robotics Association (ADRA), also called ADRA Partnership, from the ADRA members perspective, resulting in several recommendations for the upcoming European work programs. The position is motivated by global challenges, together with a European strategy and goals that can be addressed by AI, data, and robotics. The first part describes ADR's Vision 2030, which includes global challenges and major strategic objectives that the ADR partnership seeks to address. The second part describes major trends and gaps. The third, and final, part describes the recommendations for the EU strategic plan for 2025-27, which are in alignment with the strategy.

Research, innovation, and the deployment of AI, data, and robotics are essential for boosting the economy and competitiveness, as well as improving our technological leadership and welfare across Europe.

Table of Contents

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 in a wide range of 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 our 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, together with improvements in quality of life and sustainable living across the continent. Most leading global companies have significant research, development and production capabilities in Europe. At a global scale, through competence and resolve, Europe is seen as a stable and trustworthy partner, leading the way towards a promising future for all of humanity.

To achieve this vision, it is necessary to have a strong multi-stakeholder ecosystem pulling in the same direction, which 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 and offers 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 research, innovation, and deployment, the ADR partnership will reduce the fragmentation of the European ADR landscape and contribute to establishing a sustainable ADR ecosystem that achieves global impact and spurs value creation.

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 for companies' business models and enriching our society. Aligned with the SRIDA's vision[1], the 2021-22 (WP21-22) work programme embraced ADR in the context of smart and agile manufacturing, the Green New Deal, work environments and industry optimisation. It also addressed technologies and solutions for data sharing in common European spaces through topics such as compliance privacy, data management, data trading, monetization, 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

---

[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

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, 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, and waste management and revalorisation. Finally, considerable efforts should be made to improve human-machine collaboration, ethics and compliance for ADR at the service of society.

In summary, this document has elaborated on six big tickets as recommendations for the 2025-27 work programme, which can be categorised under five high-level perspectives:

- **Trustworthiness of ADR**: 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 regulation. Emphasise how SMEs and leveraging DIHs can support technology transfer, innovation, and the creation of new products and services based on ADR technology. Support in achieving the goals outlined in the Digital Decade, especially reaching 75% of companies being automated (Big tickets 3, 4 and 6).
- **European strategic autonomy in ADR**: Boost industrial competitiveness and strengthen strategic autonomy in the production and supply of essential resources. This includes enhancing the resilience of Europe to conflicts (e.g., raw materials, rare materials, and pharmaceutical starting materials) and the development of a strong semiconductor industry in Europe boosted by ADR technology (Big tickets 1-4).
- **Societal prosperity and European resilience to crises**: Increase the resilience of our society to crisis, both natural and man-made. Improved preparedness, as well as rapid, fast, and efficient responses to catastrophic situations such as cybersecurity threats (Big tickets 1-6).
- **Climate and environment:** ADR should contribute to the green new deal, the green transition, a sustainable society and zero carbon emissions. ADR should also support efficient and safe operation, maintenance and inspection, as well as contribute to the effectiveness of the circular economy and resource management. For example, ADR could support AI-powered urban mining and robotics solutions for recycling and sorting waste (Big tickets 1-6).
- **Education and skill building:** Education on AI, data and robotics, with a focus on scaling-up educational capacity and scaling-out education to other professions and subjects (Big tickets 1-6).

1.1 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. Ensure societal trust in AI, data and robotics.

Creating a strong, coherent and effective ecosystem for AI, data and robotics: Europe needs a functioning ecosystem covering AI, data and robotics that can establish the foundation for boosting the value created by the innovative development and deployment of these technologies. No single player can achieve this alone; 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, it needs to develop a single interconnected ecosystem that spans both the Europe continent 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*.

Strengthening the leadership of European industries in vision and robotics technologies: 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 (e.g., collaboration of the EU Chips Act).

Integrating and connecting the European research landscape: 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 remain siloed around disciplines and within member states. This reduces both the effectiveness of European companies and their impact, both within Europe and globally. It also reduces our capacity to translate research into innovative smart solutions, as well as develop research based on real-world questions and challenges. Fragmentation must be addressed. Otherwise, investments in research, innovation, and deployment will not maximise efficiency and effectiveness, due to redundant and sometimes even counter-productive activities.

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 the job market. Small-scale and nationally based 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, as well as training technologists that are able to design and develop such systems. A failure to address the issue of the skills mix in the economy will block future deployment, as a skills shortfall can only be partially solved by technology. The creation of a coherent European approach to skills that delivers at the highest level, will play an essential role in ensuring that productivity gains from new technology are maximised.

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

Establishing ADR technologies with high socio-economic impact: From ADR Partnership proposal: A recent study from PwC and Microsoft[2] highlights that using AI or environmental applications has the potential to boost global GDP by 3.1% to 4.4% and will help to lower worldwide greenhouse gas emissions by 4% in 2030. Furthermore, smart systems can provide a valuable contribution to sustainability at large, addressing challenges such as (with an arbitrary choice of concrete examples): climate change monitoring and understanding[3], natural resources[4] and ecosystem[5] management, reduction of the carbon footprint of industrial[6] and human processes[7], energy efficiency[8] and management, mobility management[9] and infrastructure planning[10]. In these contexts, the scale and complexity of the problems to be solved pose new challenges to current AI techniques that need to be scaled, made global, more efficient, incorporated into hybrid AI systems and integrated with knowledge coming from human experts. In addition, security issues arising from terrorism-related issues, natural disasters and epidemics can be also addressed and better managed through smart systems.

Societal trust in AI, data and robotics: There are many misconceptions and misinformation surrounding AI, data and robotics, and these technologies are not fully accepted by society in all areas of application. This will slow uptake, especially where there is unfounded mistrust, and may also damage markets where the real dangers are not fully understood.

1.2 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.

[2] https://www.pwc.co.uk/sustainability-climate-change/assets/pdf/how-ai-can-enable-a-sustainable-future.pdf

[3] Better extreme events forecast by the US National Oceanic and Atmospheric Administration

[4] https://www.bestpractice.ai/studies/deepmind_increases_value_of_wind_power_by_20_by_predicting_supply_36_hours_in_advance

[5] https://slideslive.com/38917860/deep-learning-for-wildlife-conservation-and-restoration-efforts

[6] https://www.bestpractice.ai/studies/otto_predicts_with_90_accuracy_what_products_will_be_sold_within_30_days_driving_automated_purchasing_and_reduction_of_annual_returns_by_2m

[7] https://slideslive.com/38917853/towards-a-sustainable-food-supply-chain-powered-by-artificial-intelligence

[8] https://www.bestpractice.ai/studies/st__vincents_hospital_achieves_20_in_energy_savings_by_implementing_a_predictive_energy_control_system_for_its_hvac_from_buildingiq

[9] https://slideslive.com/38917841/truck-traffic-monitoring-with-satellite-images

[10] https://slideslive.com/38917845/planetary-scale-monitoring-of-urban-growth-in-high-flood-risk-areas

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 leverage and drive this development. This trend should be fostered with a regulated framework on compliance and legislation to strengthen solutions deployment and acceptance. However, American, Chinese, and Asian companies 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/ 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 under perform 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."[11]

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 domains are essential for societal well-being, economic growth, and technological progress. AI-enhanced and data-powered robots can be customized 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 robots used in the home, or for, shopping, education, and entertainment). This is where initiatives such as 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 standards 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 these complex regulatory landscapes. A key goal for the ADR partnership is to contribute to pre-standardization activities that bring research experts together with regulators, the industry and standardisation bodies to accelerate the definition of European standards and to position European standards worldwide.

[11] https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/securing-europes-competitiveness-addressing-its-technology-gap

Achieving strategic autonomy in Europe: To maintain strategic autonomy, Europe needs to make sure that it has all the relevant competences and capabilities that it needs. This involves learning existing methods (and often adapting and improving them to our needs) and developing new unique methods. Both the application of methods and the methods themselves can be unique. To achieve this, we will require systematic approaches to transferring research projects and 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 also be a way of bringing that know-how to European companies. Furthermore, tightening education, research and field applications 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 with the next big thing. This is highly desirable and something to strive for, but it cannot be planned as science is characterised by the fact that it is essentially impossible to predict. 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 in three years from now. It is still 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. Results from 2021 are now considered old, since within a single year there is often significant progress on the most active research topics. 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 by having a yearly call for next-generation high-impact methods in AI, data and robotics.

Enabling significant progress on major societal challenges: ADR technology under development in the ADR Partnership represent 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. For example, AI-enhanced and data-power 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

The focus of Adra is 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 are based on three existing communities, which all have significant overlaps and interactions.

There is currently huge interest in generative AI, especially large language models. The last year has seen tremendous progress in both text and image generating models such as ChatGPT, DALL-E and Stable Diffusion. The next step is to generate sound and video of comparable quality, increase control and quality of content generation, including reducing bias and inappropriate content, and adding additional modalities so that the same model can generate multi-modal content. One potential issue is that the size of the models grows faster than the availability of data on which those models are trained. This leads to the interesting gap of extracting more valuable content from the same amount of data, or less. An additional challenge involves verifying the content that has been generated. One interesting area of research involves taking generated text and verifying its accuracy by showing the logic behind the arguments and corroborating facts with references. Another related research direction is to generate scientifically correct answers to questions, potentially replacing encyclopaedias with generators. The most limiting factor for European involvement in large-language models is the capacity to develop and deploy large-scale models.

Another major trend is learning from human feedback. 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.

In general, what we see is impressive progress in dealing with relatively simple data, text, images, and sound. The next major challenge is to deal with more complex data such as social networks, molecular structures, transportation networks and other graph-based structures. There is also very active work in geometric learning, which deals with these types of complex structures and incorporating physical knowledge into neural networks in a principled manner.

A fourth trend is neuro-symbolic hybrid AI methods, which are trying to find systematic and well-grounded approaches to combine symbolic and neural representations. This often entails a principled approach to combine reason and learning. Considering the need for guarantees related to upcoming AI regulation, these hybrid approaches are perceived to be capable of living up to the regulation. This is an area where Europe is making a significant contribution.

A fifth trend is AutoML, an area that is largely being driven by European researchers, where the process of configuring machine learning pipelines, selecting the appropriate model architectures, and performing systematic hyper parameter optimisation is automated. This line of research has great potential to lower the thresholds for organisations to take advantage of the latest AI advances without having deep technical knowledge and significant hands-on experience. This could be a key enabling technology, especially for SMEs that don't have the resources to employ dedicated AI engineers. The main gap here is to scale up these approaches to deal with larger and more complex pipelines and models.

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. The provision of high-quality data (including real, anonymised, and synthetic) for training AI models is also a key challenge.

There are several strategic gaps in relation to important areas of robotic development. Robots are highly application specific machines that are built on a common set of methodologies and building blocks. European standards need to 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, dealing with regulatory issues, standardisation, and involve European networks of DIHs as well as research excellence. The problems that need to be solved are sector specific and detailed. Generic approaches will not suffice.

Robotics needs to focus on its impact on specific sectors, such as construction, cities, healthcare, agri-food, and energy supply. While partnerships exist in these areas, their natural focus is on the major priorities related to these areas, 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 reach basic robot 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 materials across Europe and there is a strong need to connect this 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. Options exist to align networks to missions, to specific functional expertise or to specific emerging sectors. Tighter focus around networks of excellence will help to develop key strengths.

When it comes to data, innovators, SME’s 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 in 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 uncertainty on the effect that new data and AI regulations (such as the Data Governance Act, Data Act and AI Act) will have on the market. Companies, in particular 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 in this context. 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 perceived threat into a strong opportunity for European industry. A multidisciplinary approach is also crucial to address this challenge.

The complexity and cost in 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. 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 needs to consider 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. An increase in talent education is needed. However, the global race for talent in these areas is already underway. Regions with the most vibrant technology landscape are better positioned to attract skilled professionals and retain local talent. Talent is only attracted and retained in cases where conditions are compelling.

Finally, a lack of business opportunity understanding is also hampering progress. Developing business impact using these smart technologies requires a full understanding of the market, the technology and its impact on business processes and models. Because this requires the integration of knowledge from multiple stakeholders, it can result in low levels of uptake driven by 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

The recent world crises and global disruption are causing severe challenges to Europe's future sustainable development and welfare. Access to essential resources12 such as food13 14, energy, and water are constrained by the continuous increase in world population, decrease of arable land, and increasing global demands for these resources. Moreover, factories and human capital are continuously migrating to other ecosystems outside Europe. Climate change15 is leading to more unstable weather conditions, more frequent and disruptive natural disasters and the potential global displacement of large groups of people. Demographic changes lead to a shrinking working-age population that needs to provide for and take care of an ageing population. Geopolitical instabilities,16 from increasing tension between major countries in the world 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 world supply for essential resources (such as energy) makes its society and industries vulnerable to global crises.

[12] https://www.fao.org/fileadmin/templates/wsfs/docs/Issues_papers/HLEF2050_Global_Agriculture.pdf
[13] https://info.bml.gv.at/en/topics/food/food-security-in-europe-through-climate-change-research.html
[14] https://www.wri.org/insights/how-sustainably-feed-10-billion-people-2050-21-charts
[15] https://www.faccejpi.net/en/faccejpi.htm
[16] https://www.cer.eu/hot-topics/geopolitical-eu

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 (ADR) can enable sustainable security and strategic autonomy in Europe for food, energy, key technologies and industries. In Europe, industries specialising in robotics, computer vision and trustworthy AI technologies 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 within their ecosystem. While this dominant position persists, the swift advancement of computer vision and robotic-based AI, both in Asia and North America, introduces challenges for 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 secure 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 drastically increase our capacity in production infrastructure, transportation and the regional workforce dedicated to operations and the maintenance of assets. 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 agriculture. However, the population tends to leave rural areas for urban areas, thus reducing staffing for food production. Even worse, the amount of arable land is drastically decreasing17, so food production infrastructure must be developed in other areas (such as oceans and remote areas). In short, ADR is a strategic investment in Europe 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 for synergies to missions

The focus on combining AI, data and robotics has an impact at a high level when considering end user needs and when considering 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 fails to develop the individual areas where significant technological developments are needed to bring novel products to market. There needs to be a balance between focusing on the intersections and individual disciplines. All three disciplines will benefit from a renewed focus on their individual strengths in the second half of the Horizon Europe funding program. This would provide a valuable balance compared to its first half that 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 combining them fails to recognise the differences that are better addressed by defined support for collaboration that recognises the difficulties at the interface.

[17] https://www.canr.msu.edu/news/feeding-the-world-in-2050-and-beyond-part-1

To complement this balancing act, 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 towards the missions need to be identified and addressed. Ideally, these milestones should be agreed 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 have less impact. For instance, it is very important to collaborate with the European Chips Act that is striving to lift-up Europe's competitiveness and resilience in semiconductor technologies and applications. This is of vital importance for AI, data and robotics as semiconductors are fundamental to sensing, computing and actuation. Their development will boost European technological security and their application 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 and can have an inhibitory influence on innovation. We need large and complex testbeds which 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 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 it. 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 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 placing these facilities to ensure that the combination of European and local regulation is viable. This is particularly critical as Europe is setting a high compliance bar with the AI Act 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 be 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 rent using a simple process. Researchers based on a pool of EU experts and armed with key knowledge for a short-term period can be more lucrative for small companies to strengthen their innovation and growth.

3. The Next Strategic Plan 2025-27

Inspired by the Taiwanese model, 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 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 try to cope on their own and compete, 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 this type of lighthouses can offer an impressive portfolio of research challenges, great innovations, exciting partners, and career opportunities, so in many respects this type of network 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, one area where competition is difficult is 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, as 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-compliant European language models handling both European language and cultural differences. This includes speech-to-text, text-to-text, and text-to-speech.
- Large-scale complex ADR testbeds, together with end-users, such as in 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.

On a broader scale, the new strategic plan should simultaneously address the following five high-level concerns, which we present as big tickets below:

- **Trustworthiness of ADR**: 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 regulation. Emphasise how SMEs and leveraging DIHs can support technology transfer, innovation, and the creation of new products and services based on ADR technology. Support in achieving the goals outlined in the Digital Decade, especially reaching 75% of companies being automated (Big tickets 3, 4 and 6).
- **European strategic autonomy in ADR**: Boost industrial competitiveness and strengthen strategic autonomy in the production and supply of essential resources. This includes enhancing the resilience of Europe to conflicts (e.g., raw materials, rare materials, and pharmaceutical starting materials) and the development of a strong semiconductor industry in Europe boosted by ADR technology (Big tickets 1-4).
- **Societal prosperity and European resilience to crises**: Increase the resilience of our society to crisis, both natural and man-made. Improved preparedness, as well as rapid, fast, and efficient responses to catastrophic situations such as cybersecurity threats (Big tickets 1-6).

- **Climate and environment:** ADR should contribute to the green new deal, the green transition, a sustainable society and zero carbon emissions. ADR should also support efficient and safe operation, maintenance and inspection, as well as contribute to the effectiveness of the circular economy and resource management. For example, ADR could support AI-powered urban mining and robotics solutions for recycling and sorting waste (Big tickets 1-6).
- **Education and skill building:** Education on AI, data and robotics, with a focus on scaling-up educational capacity and scaling-out education to other professions and subjects (Big tickets 1-6).

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

- **Ground-breaking technological foundations** in ADR (autonomy, high-performance and predictability)
- **Effective** and **Trustworthy General-Purpose ADR** (generative AI, generality, continuous learning, trust, scale and complexity)
- An **interoperable** and **integrated** framework for data and model ecosystems (operations, governance, privacy, and security)
- Next-generation **smart embodied robotic systems** (soft robotics, autonomy, manipulation, configurability, human robot interaction and collaboration)
- Developing **ADR technology for** the **sciences** (from data to knowledge and understanding)
- **Research, innovation,** and **tools for compliance** (trust, privacy, security beyond compliance)

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

State of the art

Progress in fundamental ADR research and the ability to deploy ADR technologies for economic and social benefit are intertwined. Deployment is a fundamental research challenge. Real-world applications of ADR rely on technological foundations that allow the safe and effective operation of ADR systems in different industry sectors and with different user groups. These foundations in turn draw from advances in fundamental research to deliver ADR systems and tools that are reliable, trustworthy, and able to meet user expectations. The technological foundations of ADR and their use in real-world applications feed into each other, with research and practice each moving at such a pace that advancing the state of the art involves an understanding of both fundamental research challenges and deployment barriers. This interplay between state-of-the-art technological capabilities, the theoretical aspects of machine learning processes and the deployability of ADR systems can be seen in the accelerating progress of 'large-scale AI'[18] and autonomous robots,[19]

[18] Used here to describe the cluster of AI tools and methods that encompasses Large Language Models, Foundation Models, and Generative AI.

[19] Self-supervised robot perception with autonomous control mechanisms that enable predictive and safe robot operation without humans in the loop.

where developments in ADR methods, combined with user-friendly interfaces, offer the possibility of a new wave of ADR capabilities that can tackle complex societal and industrial challenges.

Experiences of deploying ADR systems have highlighted areas where fundamental research is needed to bolster these foundations. Further progress is therefore needed across the deployment pipeline, from data management to robustness in changing and uncertain environments. New ADR capabilities are needed to unlock widespread industrial applications, such as:

- Advances that streamline processes for signal detection in the presence of noisy data, which increase understanding of the learning and optimisation processes underlying AI models, and optimisation pipelines that allow increased generalisation.
- Robust certification of machine learning models, including explainability and interpretation for AI decision-making processes (see also Big Ticket #6).
- ADR systems that enable automated, cost-effective and safe automation of manufacturing tasks and allow human operators to avoid dangerous or toxic environments.
- ADR to accelerate European efforts in autonomous exploration of critical raw materials, as well as growing the region's chip manufacturing business.

The number of patents secured by European companies suggest that Europe is a world leader in additive manufacturing and autonomous robots and is on par with China and the US when it comes to 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 deployment of AI systems[20] [21]. Investment in ground-breaking European ADR is required to both foster and stimulate progress in new foundations of industrial leadership.

Rapid progress in ADR capabilities over the last ten years has resulted in technologies that can deliver better-than-human levels of performance in well-specified tasks, where sufficient volumes of data are available to enable training, and in constrained environments. This progress has been facilitated by the availability of software tools that support system developers and platforms to interface with GPUs or specialised compute facilities. The resulting systems have been successfully deployed in environments where operational conditions align with the constraints of the training environment. Extending these capabilities, progress in fundamental research is creating ADR systems that leverage new learning strategies to function more effectively in real-world contexts. Advances in the deployability of ADR are being enabled by new technological building blocks that:

- Leverage new learning strategies to create ADR systems that effectively work with real-world data, by dealing with messy data; working with less data; integrating multi-modal data types; or integrating privacy-enhancing technologies. This reflects complex data systems connected with ADR applications in the private and public sectors.
- Adapt and develop new optimisation methods to build ADR systems capable of integrating more complex data (for example, by using multi-spectral information in computer vision tasks), without impacting computational complexity.
- Connect data-driven and physics-based models of the world, enhancing performance and allowing users to interrogate the causal relationships that influence how complex systems

[20] European Commission, "Horizon Europe Strategic Plan 2025-2027 Analysis," May 2023

[21] For example, the Stanford AI Index Report, 2023, suggests that the US and UK outperform Europe in the industrial deployment of AI.

operate and how AI architectures develop their learning processes (for example, using physics models developed for the study of many-body complex systems to examine the behaviour of neural networks).

- Integrate technical capabilities to deliver the characteristics of trustworthy ADR by design, as further described in Big Ticket #6, ensuring the resulting systems are safe, robust, explainable, and aligned with the regulatory requirements associated with industrial ADR applications.
- Operate at greater efficiency, reducing the computational or energy costs of ADR development, and delivering hardware and software systems to support their use (such as, by studying the underlying geometric structure and topology of data to focus on the most informative aspects).

A new wave of innovation will come from both extending these emerging capabilities and bridging theory, methods, and applications. The resulting ADR systems would be technically advanced, rigorous in design, and safe and effective in real-world applications. These next-generation technologies will emerge from a vibrant ecosystem that is focused on delivering excellent research that integrates fundamental AI research and deployment challenges, and that grounds such research in local innovation ecosystems with a pipeline from innovation to deployment. Such innovations are central to the delivery of Adra's Goals, enabling ADR deployment that boosts industrial competitiveness and strategic autonomy, enabling progress on research that ensures Europe maintains its global reputation and reach, and creates tools that can tackle social and environmental challenges. Enhancing foundational technical research will underpin the delivery of these goals, crossing technical domains and sectors to advance ADR theory, methods and tools.

Challenges

ADR is a technology cluster that is fast moving, pervasive across sectors, and intertwined with social and economic systems in its development and use. Progress requires collaboration across academic disciplines and industrial sectors, based on an understanding of the technical capabilities of ADR systems, dynamics of real-world environments where they are used, user needs and interfaces with ADR and regulatory requirements. This must take place alongside a wider appreciation of how different communities - and society as a whole - are potentially affected by technological progress. To play a leading role in the next phase of ADR development and deployment, Europe must build its infrastructure, skills base and regulatory capacity.

Building an infrastructure for next-generation AI R&D: Recent progress in foundation models has highlighted a continuing trend in which improvements in ADR performance are enabled by large models that require access to significant data and compute resources focused on training and deployment. Establishing the capabilities to build such large-scale models requires action to address the challenges of operating 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 levels[22]. To position European research for influence alongside the large-scale AI systems being built by the tech industry in the US and elsewhere, Europe requires access to computational facilities that provide compute power needed to train large-scale AI models, expanding existing high performance computing infrastructure. Recognising that extensive use of such facilities is associated with increasing energy intensity in AI development, such compute facilities must be sustainable and

[22] For example, work in the US to create a National AI Research Resource https://www.ai.gov/nairrtf/

affordable, while also being accessible to the R&D community. In practice, their sustainability can be improved by locating compute facilities in areas with access to renewable energy sources and alternative cooling systems,[23] which can be geographically separated from locations at which research is carried out.

Growing Europe's AI skills base: Addressing the supply shortage of researchers and practitioners with AI skills is vital if Europe is to lead a new wave of AI deployment. Attracting leading talent with the skills to advance a new wave of fundamental AI research to Europe remains challenging. However, Europe does benefit from access to the talent based in existing AI research centres and networks of AI excellence, which can be leveraged in support of opportunities to ground research in local industry. Adra's role in developing a strategy to attract global talent to Europe is set out later in this report, see Section 3.3. A vital component of delivering this strategy is ensuring that industry across Europe has access to fundamental AI expertise, connecting knowledge from research to industrial application. Such knowledge translation relies on effective local innovation ecosystems that enable technology diffusion and adoption, ensuring that AI research can be responsive to needs 'on the ground', and integrates fundamental AI research and its real-world impacts.

Increasing regulatory capacity to increase confidence in deployment: In building the technological foundations of ADR, a core challenge is how to manage the unintended consequences associated with their deployment. AI systems are complex, made up of multiple interacting data and model components, and are typically deployed in socio-technical systems, characterised by complex connections between individuals, data, models and society. The result is emergent effects, which can in turn influence the performance of the AI system and its impact on different parts of society. For example, a foundation model can be adapted to a wide range of downstream tasks, creating both new opportunities and risks. There is an opportunity for increasing confidence in deployment through thoughtful curation and training of the foundation model such that downstream recipients receive the benefits and are prevented from harm.[24]

This dynamic is enhanced by R&D practices in which new methods are deployed rapidly. Successful approaches can be scaled rapidly from test to deployment, with the speed of deployment outpacing understanding of their likely impacts. The result is an implementation gap between technology, policy and practice. To deliver on the EU's vision of an ecosystem of excellence and trust, interventions are needed to bridge this gap, connecting the policy imperative to deliver trustworthy ADR with the desire to drive further progress at technological frontiers. As the pace of change in ADR capabilities accelerates, it will become increasingly important to be able to test different capabilities or methods and evaluate them before deployment. In this context, mechanisms for trialling or testing AI innovations at a smaller scale can help bridge policy and practice. Such mechanisms would deliver a proving ground to test-and-iterate the development of ADR technologies before deploying them at scale. The result would be a move away from a waterfall model of deployment towards an agile evaluation environment. The creation of such an environment for testing and evaluation can also act as a counterbalance to creeping concerns about the privatisation of critical digital infrastructure. Advances in large language models have catalysed new debate about the acceptable use of existing data and digital infrastructure - as demonstrated by a legal debate on the contours of acceptable use of publicly available data to train models held by private companies - and fresh concerns about how

[23] For example, where there is the possibility of re-using waste heat.

[24] https://fsi.stanford.edu/publication/opportunities-and-risks-foundation-models

to maintain critical national capability in strategically important technology areas. There are opportunities for open alternatives that enable such innovation, while also delivering advanced capabilities in multi-modal data and energy efficiency; there are also open questions about how to leverage large AI approaches in domains where data is not open (such as in business or the public sector), where different data management approaches or strategies for model deployment may be needed.

Open infrastructures can provide a basis for innovation, allowing companies to build services based on these infrastructures, and help ensure that cultural and scientific knowledge can be shared by all, in line with the values set out in human rights frameworks and EU policy agendas. Creating such an open ecosystem will require further work to foster and support open-source practices in the community.

Applications and impact

Estimates of the value of European ADR vary but suggest economic gains of over $2.5 trillion by 2030.[25] 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. These applications all rely on fundamental research to deliver foundational ADR technologies and tools that can be deployed in practice. None of these benefits can be created without a solid technological foundation, supported by a vibrant R&D community. This pervasiveness translates to significant potential social and economic benefits. For example, the EU's Innovation Missions illustrate the diversity of potential applications of AI across sectors, along with the scale of beneficial impacts that are possible:

- 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 health (for example, in neurology to support better global modelling of the whole body and multiple interactions at different scales), protection of biological heritage and the biosphere, improvements to building architecture and city planning, and across industry through process optimisation, pollution reduction, and design creation, among other areas. Effective collaboration across disciplines and organisations is needed to make progress on this agenda, accelerating foundational research and creating impact by ensuring that Europe is capable of deploying ground-breaking foundational ADR to solve core

---

[25] See, for example: 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).

industrial and societal challenges. To create ADR systems that serve this ambitious agenda, progress is needed on:

- Fundamental research that integrates deployment challenges, encompassing theory, methods, software, and hardware that combine both cutting-edge technological innovation and an understanding of the capabilities required of ADR in deployment.
- Core ADR capabilities, increasing task performance through advances in learning, modelling, inference and reasoning; enabling progress in technical clusters within ADR (such as computer vision, natural language processing and generative AI); and building cross-cutting capabilities (for example, in techniques for using multi-modal data, methods to enable causal AI or spatio-temporal modelling).
- Deployment-inspired research that drives progress in a technical domain and delivers real-world benefits from this progress (for example, making open foundation models that can be deployed in the service of the EU's Missions).

Across these areas, connecting those working at the forefront of AI research with the needs of organisations using AI is crucial in 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

Dynamic interplays between ADR innovation and deployment mean 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. Approaches that require significant development before deployment in one sector can be close to implementation in another, and combinations of technologies enable rapid progress in new areas. Results are innovative methods or systems rapidly deployed to operate at scale, while deployment experiences are used to adjust underpinning methods to enhance the performance of systems.

The table below connects the technological foundations that are being built across the European ADR community today to objectives for these technologies in the future. It illustrates how today's areas for action map onto desirable medium-term goals and long-term outcomes for ADR, as described both earlier in this document and in a cluster of policy and legislative briefs from the European Union.[26] In the near-term, progress in developing new learning strategies, core capabilities and the technical aspects of trustworthiness can help deliver next generation AI 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

[26] See also: https://www.vision4ai.eu/sra/

research interest, the technological foundations of ADR and the capabilities needed to solve industrial challenges are closely integrated, combining responses to the challenges that arise in deployment with areas of cutting-edge fundamental research.

Foundational research is unpredictable by nature. Innovation requires space for curiosity and novelty, allowing fresh perspectives or new techniques to emerge, and enabling successful new approaches to scale into widespread deployment safely at pace. Flexible funding mechanisms and agile regulatory approaches both play a role in facilitating this type of ecosystem of excellence and trust, allowing European ADR R&D to operate at the frontiers of innovation. The result should be a dynamic R&D community delivering innovations in AI theory, methods, and applications, creating safe and effective AI systems that are deployed to enhance the EU's industrial capabilities and enable the EU's innovation missions. Creating these ground-breaking technological foundations can ensure Europe has the research and the deployment capabilities it needs to deliver the social and economic benefits that are made possible through widespread adoption of AI.

Short-, medium- and long-term objectives

Areas for progress today/near term	**Medium-term goals**	**Long-term outcomes**
Capabilities to increase performance in deployment: • Simulation, emulation and digital twins; • Computer vision and video generation; • Natural language processing and speech; • Physics-aware modelling and knowledge representations; • Foundation models and fine-tuning; • Causal AI; • Multi-agent collaborations; • Spatio-temporal modelling; • Data preparation and validation; • AI at the edge; • AutoML and AutoAI; • Next-generation hardware; • Autonomous vehicles and robotics; • Shared representations.	Trustworthiness by design for applied AI: human-centric AI systems that embed the technical foundations of trustworthy AI across industrial applications, demonstrating: • Safety. • Robustness to change or adversarial manipulation. • Explainability and auditability. • Privacy-preservation and security. • Fairness. • Sustainability and energy-efficient implementations. Deployable under dynamic conditions: AI systems that can be integrated into industrial systems to deliver beneficial outcomes in real-world applications, with the ability to: 1. Analyse real-world data, learning from multi-modal data sources, with dynamic or active learning strategies, and the ability to learn from few examples. 2. Maintain robustness under dynamic conditions, implementing AutoML/AI methods to maintain verifiable standards of performance. 3. Deliver on expectations of privacy and security in deployment. 4. Simulate how complex systems operate, leveraging data at different granularities and different data types to improve performance in deployment. 5. Run on devices at the point where data is collected, enabling AI-at-the-edge.	Innovations in AI theories, methods and approaches to create technically advanced AI systems that deliver high-performance in deployment. Safe and effective AI that can be deployed in real-world contexts, enabling applications across the private and public sectors. AI tools deployed to deliver societal benefit, accelerating progress towards the goals of the EU's

Trustworthiness by design: • Robustness • 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. Learning strategies and methods to build next-generation AI models: • active learning and dynamic feedback • deep learning • reinforcement learning • transfer learning • few-shot learning • federated learning • continual learning • multi-modal learning • causal inference • geometric learning • quantum ML • generative AI • common ground, shared representations and understanding intent • human-AI co-evolution • end-to-end learning and multi-objective optimisation	*6.* Robotics systems that can interact intelligently with their environment in industrial settings. Next-generation methods: improved performance of AI systems, underpinned by innovations in theory and modelling that reflect the needs of real-world environments, including: *7.* Principles and techniques for AI that reliably performs well in deployed contexts. *8.* 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. *9.* Multi-modal AI systems that combine different data types, reflecting the operations of real-world systems. *10.* Approaches to bridging data-driven and domain knowledge, including the ability to simulate or emulate physical systems, develop causal AI, and build causal digital twins that discover physical processes from large data volumes. *11.* 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. Collaborative AI: AI decision-support tools that work alongside users in deployment to enhance human activities, through: *12.* Collaborative multi-agent systems and models of human-AI collaboration that reflect the needs of real-world applications. *13.* Active learning strategies and knowledge representations that enable learning in dynamic industrial contexts and facilitate human-agent teaming. *14.* Interfaces designed to enable human-AI interactions, for example with explainable interfaces. Applications for societal and industrial benefit: use cases that showcase the benefits that well-designed and integrated AI can deliver in areas of social interest and need.	Innovation Missions and enhancing productivity across sectors.

Currently, the research community is addressing the trustworthiness dimensions one-by-one. The next step is to consider combinations of them, and the long-term vision is to deal with all of them in a systematic and integrated manner by design.

3.1.2 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) though, 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 also 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. Below we include some of the current developments in a fast-moving area.

Large language models (LLMs) - The breakthrough in LLMs has received the most attention, especially OpenAI's ChatGPT family of models and services, which are supported and commercialised by Microsoft. ChatGPT builds on the GPT-family of models with successively larger models that contain more and more parameters. Google has since released Bard and Meta, which are known as the Llama family of open models. The latest version is Llama 2. Hugging Face has released the Bloom family of models, as part of their app-store-like model service. Anthropic, a spin-off from OpenAI, has released Claude. They have also worked on Claude Instant, which is a smaller and faster version of Claude 2. DeepMind is getting ready to release Gemini, which is expected to significantly outperform ChatGPT. In Europe Aleph Alpha and Mistral are actively working on LLMs developed in Europe.

There is similar progress in image generating, such as CLIP, Dall-E, Stable Diffusion and Midjourney, and code generation models, such as Codex (GitHub Copilot), Code Llama (based on LLama), and Poolside. Research is also ongoing to generate more structured output, such as crystals and molecules.LLMs and generative models are not general purpose by default. The models named above cannot query application-specific data and APIs, and do not address robotics. Some initial steps have been made in this direction, such as LangChain, RT-2, and ActGPT, but much more research is needed.

As models and approaches become more general, new risks are raised that need consideration and mitigation:

- **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.
- **Static knowledge.** Most data in domains evolves over time. However, because learning from scratch is not sustainable, most "data-based only" GPAI systems are rapidly obsolete (e.g., ChatGPT knows nothing that happened after 2021).

- **Hallucinations and 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.
- **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;[27] Kurita et al. 2019;[28] Luccioni et al., 2023).[29] 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).[30] 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 Reinforcement Learning from Human Feedback (as used by ChaptGPT), but it is very hard to quantify.
- **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 homophobe). 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 the recipe for a bomb), or by appending computed adversarial strings[31] to a naive prompt.
- **IP**: At the moment, training data for LLMs is often a commercial secret, and even when revealed, is fetched from the internet, regardless of any intellectual property rights and data regulation such as GDPR. Few 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 licences, and of reproducing this data, remains unclear (Bommasani, Klyman, Zhang, Liang 2023).[32] 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 lawsuits[33] regarding the use of copyrighted works for training Generative AI systems have underlined the need for legal certainty and clarity about the definition of the TDM exception.
- **Computational and environmental cost**: Training a single BERT base model (without hyper-parameter tuning) on GPUs requires as much energy as a trans-American flight (Strubell, Ganesh, & McCallum 2019).[34] 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,

---

27 https://aclanthology.org/W19-3805/
28 https://arxiv.org/abs/1906.07337
29 https://arxiv.org/abs/2303.11408
30 https://dl.acm.org/doi/pdf/10.1145/3442188.3445922
31 https://llm-attacks.org/
32 https://crfm.stanford.edu/2023/06/15/eu-ai-act.html
33 https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/?guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAAE7FmNqEnZMTnfMqHtV8dq84sXKEadgyGvweOJcg6xcAjDCtjxkWmODAHljidNmTsC-GjYs8KyTTBO--rW2R_rAKHzvXVy3cKGJoIKWmGXQVJVs8GWPB-O_Y0cCkfVFzyNZM2BJEuW2uanUe7XbcAxlc1PHbPy35FpI59jBaXkZO&_guc_consent_skip=1699221271
34 https://arxiv.org/pdf/1906.02243.pdf%22%3EWachstum

underscoring the need for energy efficient model architectures and training paradigms (Bender et al. 2021).[35]

Challenges

The main challenge is how to make the models and approaches more general while improving both effectiveness and trustworthiness. An especially interesting and important challenge is how to extend software models to deal with physical interactions and robotics.

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** a growing challenge too.

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 for computational social science).[36] Such reasoning capabilities would also **increase the safety and robustness** of GPAI systems, as would access to fact-checking techniques that require browsing capabilities.

Biases and non-discrimination: the detection and correction of biases is a hot topic in AI research and is not specific to GPAI. Whereas the tracing of data sources is one aspect of bias detection, post hoc detection is another needed approach, but different contradictory definitions of biases exist and the metrics used depend on the goal (such as individual biases vs group biases). Beyond statistical measures of bias, an increasingly challenging aspect, especially in the realm of visual and text data (Blodgett et al., 2020;[37] Fabbrizzi et al., 2022),[38] is representational bias and its potential harms.

GDPR and source tracing: For legal and moral reasons, authors of training data (such as authors and music composers) must be acknowledged, and more generally, the origin of training data must be transparent. The lawful basis for processing training data must also be made clear. As the Italian data protection authority's temporary ban on ChatGPT shows, developers of large generative AI models need further clarity about whether and to what extent they can rely on legitimate interests as lawful ground for data processing when training their models. For the moment, it is extremely debatable whether the development of AI systems or generative AI overrides the interest of the data subject's rights. Compliance with GDPR's individual rights requests is also important. Because of the way large language models store and process information, they pose new challenges for compliance with data subject rights, including the right to be forgotten (Zhang et. al 2023).[39] It is hard to know what personal data is used in training and how to attribute this data to particular individuals. Rectifying data from models may also be a difficult task.

Scale and complexity of the value chain: one of the specific challenges from 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; Hacker, 2023). The fact that current regulatory

[35] https://dl.acm.org/doi/pdf/10.1145/3442188.3445922
[36] https://arxiv.org/abs/2306.04746
[37] https://www.microsoft.com/en-us/research/publication/language-technology-is-power-a-critical-survey-of-bias-in-nlp/
[38] https://www.sciencedirect.com/science/article/abs/pii/S1077314222001308
[39] https://arxiv.org/abs/2307.03941

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, respectively, duties of mutual assistance risk new accountability gaps.

Absence of ethics by design: There is a clear need for an embedded-ethics approach that incorporates reflections on the potential consequences of AI development throughout the whole process. An interdisciplinary stakeholder engagement in the question articulation sessions and value-sensitive design approaches, such as Values that Matter (Smits, Bredie, van Goor, Verbeek 2019)[40] should be incorporated in a whole life-cycle of GPAI development, and not as an one-off, ex-post exercise. There should be concrete methodologies on how to enable participatory, inclusive, and interdisciplinary methods of involvement of 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).[41] Model cards (Mitchell et al., 2019)[42] and datasheets for datasets (Gebru et al., 2021)[43] 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).[44]

Privacy: Care must be taken when fine-tuning a GPAI to a specific task with sensitive/non-public data to ensure sensitive information is not memorised. Equally, just because a GPAI 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 GPAI cannot collate public information to come to sensitive conclusions about individuals that would be difficult/impossible for a human to do.

Risk of manipulation and AI anthropomorphism: Recent chatbot-incited suicide in Belgium[45] brings to the fore yet another risk of general purpose AI: the risk of manipulation (Smuha, De Ketelaere, Coeckelbergh, Dewitte & Poullet, 2023).[46] The phenomenon of AI anthropomorphism, understood as projecting human-like qualities or behaviour to non-human entities such as AI systems (Tan 2023),[47] risks deceiving people 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 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 obviously there are also huge opportunities to be realised.

---

[40] https://research.utwente.nl/en/publications/values-that-matter-mediation-theory-and-design-for-values
[41] https://arxiv.org/ftp/arxiv/papers/2302/2302.02337.pdf
[42] https://dl.acm.org/doi/abs/10.1145/3287560.3287596
[43] https://dl.acm.org/doi/fullHtml/10.1145/3458723
[44] https://algorithmwatch.org/en/wp-content/uploads/2023/09/PolicyBrief_GPAI_AW-3.pdf
[45] https://www.brusselstimes.com/belgium/430098/belgian-man-commits-suicide-following-exchanges-with-chatgpt
[46] https://www.law.kuleuven.be/ai-summer-school/open-brief/open-letter-manipulative-ai
[47] https://medium.com/human-centered-ai/on-ai-anthropomorphism-abff4cecc5ae

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, maths and social sciences), innovation (including patent search and industrial process optimization) 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 a human engineer), 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 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, including for AI-based surgery.

AI for creative processes: supporting the creative processes involved in producing text, articles, experiences, videos, and images, as well as experiences in the metaverse. 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 challenging IPR and copyrights in these scenarios is also important[48].

Socio-economic impact

According to Mc Kinsey report "The economic potential of generative AI: The next productivity frontier",[49] *generative AI could add the equivalent of $2.6 trillion to $4.4 trillion annually across the 63 use cases we analysed—by comparison, the United Kingdom's entire GDP in 2021 was $3.1 trillion*.

According to Goldman Sachs Research,[50] *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 economical estimations also shyly point out the significant disruption this will cause in the job market. Such issues had already been identified for machine learning and AI in general, but the changes will probably be more chaotic with generative AI. And even if the global balance of destroyed vs created jobs remains controversial, there is a consensus about the transformations that will take place, and the need for the workers to at least change their working habits, and at most to completely change job or even acquire completely new skills.[51] And it is not certain that the necessary continuous education mechanisms will be put in place, nor that

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

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

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

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

the mindset of targeted workers will allow such drastic changes in their life as smoothly as these professional report-writers like to think.

Beyond the impact on skills and employment, generative AI is expected to affect the perception and cognitive processing of online information by citizens. Increasing awareness among citizens about the capabilities 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 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 will probably not outweigh the negative effects.

New methods are needed to assess the individual and societal impact (both positive and negative) in a way that takes into account 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 GPAI is democratising access to this powerful technology, underlines the fact that experiences, concerns and mitigation strategies on the ground need to be better understood.

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

State of the Art

The establishing of sustainable data-based ecosystems has been one of the main objectives of the Big Data Value PPP (between EC and BDVA, under programme H2020), and the ultimate objective of projects funded under this portfolio was to contribute to this objective. Starting back in 2017 with, among others, large-scale pilot projects in domains of strategic importance for EU industry[52] and key ecosystem enablers[53], this has since continued with advanced platforms, tools, and testbeds (cloud, HPC, IoT) in emerging domains[54]. With the last projects that fall under this programme focused on data platforms[55] and experimentation, incubation and a network of data innovation hubs,[56] the EC has paved the way for a strong data European ecosystem.

The European Data Strategy published in 2020 leveraged assets produced by PPP projects to create a European Common Data Space, which seeks to connect stakeholders in a cross-border and cross-sector basis and share data in a trusted, secure and scalable way. A new tranche of projects funded under the Horizon Europe programme and under the umbrella of the ADR PPP, are focused on

[52] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/ict-15-2016-2017
[53] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/ict-17-2016-2017
[54] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/ict-11-2018-2019
[55] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/ict-13-2018-2019
[56] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/dt-ict-05-2020

producing the tools and building blocks needed to build data spaces: technologies and solutions for data trading, monetisation, exchange and interoperability[57], advanced technologies for data management[58] and, more recently, the integration of data life cycle, architectures and standards for complex data cycles and/or human factors[59].

At the same time, projects funded under Digital Europe programme (including preparatory and deployment actions) focus on the real establishment of data spaces in relevant sectors (Data Spaces Support Centre projects include them under its community of practice[60]).

There has been significant progress towards raising awareness about the benefits of sharing data, breaking data silos and efforts to foster the exchange of data at different levels. However, some barriers still remain: *interoperability at different levels, global trust and governance frameworks for data sharing, and tools and methods specially designed to cope with new trends on AI. All these issues* are still hindering the ability to unleash the potential from the huge amount of data generated in Europe.

The challenges highlighted cannot be fully addressed without considering existing and upcoming regulation (in accordance with the European Data Strategy[61]), such as the Data Act (still to be adopted and entered into force), the Data Governance Act, the free flow of non-personal data regulation, the AI Act, and GDPR. It is also worth mentioning regulation related to the European Health Data Space.[62]The first draft was presented in May 2022 and is still under discussion. The legislative process is expected to be completed by the end of 2023, with implementation of the EU Health Data Space expected by 2025. Standardisation efforts are also needed to guarantee trust, connectivity, and scalability.[63] Different industry associations have produced relevant position papers and books where those challenges are addressed.[64,65]

---

[57] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2022-data-01-04

[58] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2021-data-01-03

[59] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-cl4-2023-data-01-02

[60] https://dssc.eu/space/DC/27983886/Community+of+Practice

[61] https://digital-strategy.ec.europa.eu/en/policies/strategy-data

[62] https://health.ec.europa.eu/publications/proposal-regulation-european-health-data-space_en

[63] CEN CENELEC Trusted Data Transactions Workshop (https://www.cencenelec.eu/news-and-events/news/2023/workshop/2023-02-10-data-transactions/)

[64] Towards a European governed data sharing space (BDVA: https://bdva.eu/sites/default/files/BDVA%20DataSharingSpaces%20PositionPaper%20V2_2020_Final.pdf), Data Spaces Business Alliance Technical convergence paper (DSBA: https://data-spaces-business-alliance.eu/wp-content/uploads/dlm_uploads/Data-Spaces-Business-Alliance-Technical-Convergence-V2.pdf), Data Spaces: Design, Deployment, and Future Directions (BDVA: https://link.springer.com/chapter/10.1007/978-3-030-98636-0_1), Fair Data Economy Rulebook (https://www.sitra.fi/en/publications/rulebook-for-a-fair-data-economy/)

[65] Designing Data Spaces (https://link.springer.com/book/10.1007/978-3-030-93975-5), Design Principles for Data Spaces (https://design-principles-for-data-spaces.org/), Real-time Linked Dataspaces (https://dataspaces.info/)

Challenges

Operational challenges

This type of challenge includes those that affect or impact usual operations with data. Among them, we can consider:

- ***Access, discovery, and catalogue of data***. This challenge includes the difficulties in accessing data and fully leveraging the potential of the European data economy. It also considers access to pre-competitive big data sets, mostly to include and support SMEs in the process and enable the democratisation of AI.
- ***Efficiency***, including reducing the amount of data to store, considering less data to compute (sparse data) and more frugal/sustainable AI models. It should also consider reduction of space and energy without reducing the (potential) value of data sets (compression, transformation).
- ***Integration of complex data analysis pipelines*** (BD/AI/HPC) that take advantage of the cloud/edge continuum in a trusted and scalable way.

Connection and interoperability of ecosystems and frameworks

Achieving interoperability presents a complex set of challenges rooted in various factors. These hurdles arise due to a multitude of reasons, encompassing intricacies within data models, standards and contextual nuances. The divergence in vocabulary, language, and interpretation adds to the complexity, often resulting in misinterpretations and communication breakdowns. Furthermore, the evolving nature of data, standards, and terminologies introduces the need for a delicate balance between accommodating changes and maintaining compatibility. The absence of semantic consistency and clear definitions within data sources further compounds the difficulties and leaves room for ambiguity and imprecision.

Taking the above into account, this challenge would include how to fully achieve ***diverse types of interoperability***: technical, syntactic, semantic, platform, cross-domain and process, as well as automation and interoperability of AI methods in Data Spaces/Data Ecosystems.

An additional challenge involves enabling the ***connection of complex data driven ecosystems*** (systems of systems), including federation, catalogues and marketplaces. This connection would also be facilitated by the connection/interoperability (and the corresponding standardisation efforts) of trust frameworks and common identity management methods (SSI), as well as interoperability of PET technologies (also mentioned in Big Ticket 6).

Governance and sustainability

Data ecosystems are becoming more complex, with scalability and sustainability being key for their success, along with governance mechanisms that can be shared among multiple actors. This poses challenges related to ***governance of data ecosystems*** (including data governance and governance of data spaces), which covers aspects of data sovereignty, provenance and traceability of data and quality of data assets. Automatic contract negotiations and enforcement within data ecosystems is a particularly urgent topic. The ***sustainability of data ecosystems*** is also a challenge, with the need for the identification and availability of instruments to scale, including novel tools for measuring data maturity (data maturity models) aimed at monitoring the progress of the ecosystem.

Security and protection in data ecosystems

Security and protection challenges stem from different factors and require comprehensive strategies to mitigate risks effectively. 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.
- **Data encryption and protection**: Managing encryption keys and securing encryption key storage is highly challenging.
- **Data anonymization and re-identification**: Anonymizing or de-identifying data for analysis while maintaining data utility is challenging. Re-identification attacks can potentially link anonymized data back to individuals.
- **Regulatory compliance**: Navigating a complex landscape of data protection regulations (such as GDPR and HIPAA), and ensuring compliance with data handling and privacy requirements can be challenging.

As part of these challenges, data security, data access and usage control, and model privacy enabling model marketplaces should also be considered.

Important use-cases & applications

Interoperability between complex data ecosystems

This use case considers multiple types of interoperability (such as business, governance, legal, semantic, syntactic and transport) within and between complex ecosystems, with the technical and non-technical mechanisms needed to address it. This use case will cover the application of AI techniques to achieve this (for instance, the use of large language models to support semantic interoperability between different domains) and the interoperability of IoT ecosystems (by applying standards/taxonomies to enable automatic discovery and processing of IoT/edge computing, such as SAREF).

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

This use case will consider the establishment of comprehensive and standardised systems that address key aspects related to trustworthiness, identity verification and policy enforcement within a large-scale data environment. It will address challenges such as governance, security, protection and privacy within the ecosystem, together with mechanisms that allow the enforcement of policies that ensure these rules are followed throughout the entire data lifecycle.

Efficient solutions for data driven implementations

This use case will include solutions that result in the reduction of energy and carbon due to the amount of data processed and stored. It will consider efficient methods for data analytics, such as appropriate data fusion strategies to design efficient energy efficiency systems, frugal/sustainable AI, and less data and energy intensive algorithms.

Design and implementation of complex data analysis pipelines

This use case will take advantage of the cloud/edge continuum, including management and usage of heterogeneous computing resources, to implement complex data pipelines. It will include applications such as transfer learning and model marketplaces (including descriptions of complex models through model metadata and considering model quality), specific techniques such as fine-tuning and inheritance of capabilities, benefits and harms from foundation models/large (language) models to downstream models, and consider DataOps to manage those complex pipelines by fostering collaboration, implementing automation, promoting agility and ensuring continuous improvement.

Data spaces and digital twins

Data spaces represent an excellent environment where digital twins can find the perfect conditions to grow, scale and connect the different actors required to unleash all 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 use case combines features and represents an optimal symbiosis of both paradigms, linked by data. This use case will result in relevant applications addressing specific sectors, such as the Data Space and digital twin for oceans, metaverse and virtual worlds, and contributions to the Destination Earth initiative.

Adoption and scale-up of data spaces

For the adoption and scale-up of data spaces, the challenge of interoperability between data spaces and existing legacy systems must be addressed. Interoperability challenges are not only technological but also include business models, maturity and conformity to regulations related to data spaces. The outcome of this use case is expected to include instruments that support the scaling of data spaces and assess the compatibility of data spaces across various dimensions. Most relevant application areas are those where preparatory actions for data spaces are already running, but also includes more specifically the healthcare domain (in light of EHDS European regulation).

Socio-economic impact

The main impact resulting from the proposed use cases can be summarised in the following points:

- Increase in data democratisation, mostly for small actors, empowering citizens in the use of their personal data, and increasing societal trust for the lawful processing of personal data by governments.
- Breaking data silos and leveraging data generated in Europe to train large AI models and AI foundation models.
- Increasing trust and security.
- Reduction of carbon footprint and energy consumption, and reducing costs by leveraging existing platforms and legacy systems
- Extensive application of the digital twin paradigm to support better decision making in the social, industry, and environment 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.

Short-, medium-, and long-term objectives

Short-term

- Paradigms for the reduction of energy, power consumption, storage, waste and carbon footprint in data-driven applications and systems.

Medium-term

- Reference models for the connection of complex data ecosystems, including governance, interoperability, trust, identity management and policy enforcement.

Long-term

- Framework for the application of data sharing and data spaces to the digital twin paradigm
- Instruments for the sustainability and scaling up of data spaces.
- Integration and convergence of data analytics and data driven technologies with AI, supported by HPC/cloud/edge.

Roadmap (TRL-levels)

	Initial TRL	Final TRL
Interoperability between complex data-driven ecosystems	TRL 5 (technology validated in relevant environment)	TRL 8 (system complete and qualified)
Efficient solutions for data-driven implementation	TRL 3 - TRL 5	TRL 8
Design and implementation of complex data analysis pipelines	TRL 3/4 (experimental proof of concept/technology validated in lab)	TRL 7 (system prototype demonstration in operational environment)
Design of global frameworks for trust, identity management and policy enforcement	TRL 2 (technology concept formulated)	TRL 6 (technology demonstrated in relevant environment)
Data spaces and digital twin	TRL 3 (experimental proof of concept)	TRL 6 (technology demonstrated in relevant environment)
Adoption and scale-up of data spaces	N / A	N / A

3.1.4 Big ticket #4: Next generation smart embodied robotic systems

Robot systems need to significantly improve their functional capabilities in five areas to build the next generation of smart embodied robotic systems.

Soft Robotics: 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.

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. 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 be adapted to. This environment may also contain significant risk to people or to the environment itself, such that autonomous robot activities and actions should be guaranteed and trustworthy. 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. 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 the perception, manipulation and interaction with objects, people and environments.

Manipulation: Robots do not currently have the dexterity of a human infant (Moravec's Paradox).[66] The matching of AI with physics is a major challenge. Physics is based on a mathematical understanding of the laws that govern interaction between objects, rather than data, and this is an area of extreme weakness for AI controlled manipulation. Novel physics- and model-based AI methods are needed to underpin the high-level planning and object reasoning that traditional AI methods provide.

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 are 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 needs to be accessible without deep robotics knowledge and with minimal training. For wide deployment in key applications, such as healthcare, construction and warehousing, a step change in interoperability and configurability from current concepts and practices is needed.

Human robot interaction (HRI): 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. Developing applications in the middle ground where humans have oversight without interventional control is 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.

Technologies and techniques of focus

As mentioned above, robotic systems ensure safety by constraining operating environments or by including human intervention in the operating loop. This either limits the level of autonomy or results in a complex level of operating environment modification. As a result, robots typically operate in controlled environment (such as in factories, in pipes, on farms, in hospitals, down mines or in space) and/or with human supervision to take over or guide limited levels of autonomy (for instance, in

[66] H. *Moravec (1988), Mind Children, Harvard University Press.*

autonomous vehicles, surgical robots, or inspection and maintenance systems such as drones or tele-operated machines). The next generation of robots should make use of AI systems for 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. Research in AI needs to focus on techniques that provide accurate models about 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.

Challenges

A primary and enduring challenge is to ensure that AI-powered robots are safe, trustworthy and exhibit predictive actions in their physical operation. 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. As the combination of robot, task and environment tends to be complex, it is essential to analyse this as an *agent-environment* interaction system that treats all three elements as a consolidated system. In other words, it can be treated as a "complex system" where unanticipated unknowns caused by the emergence of states cannot be predicted in advance. 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 their states, actions and the conditions in which they operate. Current learning techniques depend on the ability to train neural models based on real or synthetic data; however, these models 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 foundation models for robots operating in complex environments where safety, trustworthiness and regulatory compliance are essential and find new ways to formulate the integration of foundation models and their development. We need to move from foundation models that largely focus on text and language, to large AI models that are capable of actually capturing accurate models of the physical world around them in real time. This is significantly harder than text due to the very different techniques and representations currently used for capturing different modalities of the world. Significant progress is being made here through the use of AI methods, but it has largely not been integrated back into robotics. Foundation models for the above five areas should address key robotic challenges:

- The real-time reaction requirement for robots interacting in the physical world.
- Robot operation speed should be competitive with manual operation speed to ensure the viability of many applications. Most notably, safety cannot be just about moving slower.
- Precision of placement and localisation within a changing or dynamic environment, along with repeatability, are key aspects of robots that provide functional advantages.
- Robots should be flexible for different environments and easy to certify for safety. They require robust perception systems (such as understanding and decision making). AI techniques should provide guarantees not just about functionality, but also about safety, robustness, performance, and transparency. Recent trends on "Trusted-AI" should be explored for more direct interaction in robotics and certification processes.
- Robots are required to be versatile and not only deal with one task, but also adapt to new requirements, small changes in tasks and environment changes to reach economic viability.

- Energy/computing should increase in robots without a massive increase in size/weight.
- Foundation models must work within existing data and energy envelopes. Actuators and batteries will only increase their energy efficiency linearly over time, meaning that exponential increases in computational consumption will eventually reach a limit.
- Physical interaction with vulnerable and frail people requires soft robots, compact mechanics, and higher payloads.

The challenges above call for combining new forms of robots with a high level of autonomy, manipulation, configurability, and greatly improved HRI. We need significant progress in physical forms through soft robotics and in their 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. To reach this goal, two areas need to be explored:

- Soft machines to explore alternative forms of robotics.
- Physics-based foundation models for inherently safe and certifiable decision making.

Use cases and scenarios for value creation

Transportation covers microscale to interplanetary transportation and expresses the delivery of "goods", including people, from one place to another, including associated logistics. In a factory this might be simply the transfer of parts from one conveyor belt to another, or to a warehouse. Transportation of goods between cities will involve local navigation, traffic junctions, cyclists, and large-scale navigation to select the best route, given the weather and traffic conditions. Such journeys may be constrained by resources or time.

Inspection also covers a wide range of scales where it may involve a manufactured part, such as the non-destructive evaluation of a seam weld in a pipe, or it may refer to the inspection of a bridge for water damage or inspection of the human body. Here the robot is a data gathering tool that is able to inspect methodically and do so repeatedly with a high degree of localisation accuracy under a wide range of environmental conditions, including extreme conditions such as in space or inside a nuclear reactor.

Collaboration refers to the close human-robot interaction covering manufacturing, over a wide range of size scales, where robots work side by side with humans to safely complete dull, dangerous, or dirty tasks. It also includes exoskeletons, where the person is effectively inside the robot and their movements are interpreted and acted on by the robot, even to the point with rehabilitation robotics where the neural basis of the intent to move is captured and used. Collaboration also covers companion robotics, where robots are used to assist a person either physically or psychologically, which could be in a factory or in an elderly care home.

Intervention refers to the ability of robots to alter the physical world. To assemble or disassemble parts and devices, to drill holes, polish, weld and grind material, to construct buildings or ships, and to operate on the human body. It refers to the exploration and/or extraction of (critical) raw materials, the preparation of food, the operation and maintenance of infrastructure from roads to wind farms and the processing of hazardous materials. It is in the field of intervention that trustworthiness and safety become paramount to deployment.

Impacts

Socio-economic impact: Full deployment of autonomous robots means that robotics could replace all human physical work; in factories, transport, field operations and maintenance, agri-food, healthcare, and security. The socio-economic impacts of these changes could be considerable.

However, robots are unlikely to replace all human work, especially that which requires the qualities inherent in human-to-human interaction. The steps towards robotic work must be taken in consultation with citizens and involve a high level of responsible innovation. The focus must be on improving lives, to ensure robotics are used where they can remove people from harm, create better working conditions or enhance health and wellbeing. However, in many areas of work robotics will become more efficient than humans and thus create economic pressures. Research and innovation must work to provide responsible solutions to the balance between economic necessity and human wellbeing. The “hard problem” addressed is pertinent to this socio-economic impact because it addresses those use cases where harms are likely to be caused. These are applications where humans already perform those actions and are qualified, or tasks that humans cannot/will not undertake[67]. Robots must prove to be highly competent before deployment, and be socially and ethically acceptable.

Roadmap (TRL-levels) for robotics contains two fundamental axes. On one axis, we need to balance between human intervention and autonomy, and 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 sensing that is able to 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 performance needed in specific applications. It is likely that this work will identify parts of this operating space where it is possible to use foundation models successfully and other parts where new methodologies are needed.

Short-, medium-, and long-term objectives

Soft Robotics

- Establish a range of materials suited to developing robotic systems, both as the main structure and of manipulators and end effectors. These may encompass passive and active materials, and combination materials with specific properties.
- Develop design methods for non-rigid structures and the means to actuate and sense position 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 soft robotics has the greatest economic potential and develop experimental designs for robotic systems able to demonstrate how this capability can be achieved.

The short-term objectives are to explore (new) materials and how they can be used to create novel robotics structures. Medium to 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

[67] Such as in extreme environments, space, in nuclear reactors, in collapsed buildings, in deep ocean, with extremely vulnerable people etc.

that are more difficult to achieve with current robotic structures. Soft robotics also has the potential to impact on the “hard problem” by creating inherently safe systems, partly because soft robotic systems are typically of much lower mass, but also because they can be made inherently compliant, rather than algorithmically compliant.

Foundation Models

- Establish methods and systems that are designed to solve the “hard problem” and are capable of using foundation models with constraints related to robotics to ease human interaction.
- Develop large AI models capable of dynamically learning about their physical environment in real-time and use these models to plan actions with respect to their (high-level) goals.
- Establish methods of assurance that can test model outcomes in both design and operation.
- Respect AI frugality constraints both in energy consumption and hardware requirements.
- Assess the socio-economic acceptability of the limitations of such systems.

Neuro-Explicit AI Models

One approach to these challenges is the use of *Neuro-Explicit AI Models*. These work to combine explicit models known from physics and other disciplines with neural models to jointly provide stable and accurate models of the physical world. They are based on explicit models consisting of differential equations, simulation code, knowledge bases and others. Such models are *explainable by definition*, they are *causal*, *extremely compact*, and can be combined in *modular ways* - all properties that current AI models lack. However, these are abstract and idealised models (such as point masses) that omit important details, such as accurate friction, noise and other issues. By combining these explicit models with neural models that capture the residual differences with the real world, we can combine the best of both worlds. Most importantly, these explicit models can provide a strong skeleton and inject the combined model with the necessary properties discussed above. They promise to provide a basis to make significant progress in modelling the real world, and in a way that traditional AI techniques cannot. Since this approach has recently emerged, research is needed to mature this technique in combination with robotics and AI. The initial short-term objective is to work out how to break down these long-term objectives into manageable work packages and explore the myriad of alternative methodologies combined. Notably, early work needs to address how foundation models can be made compatible with hard physics reality, such as by eliminating hallucinations and grounding model outcomes in the real physical world.

3.1.5 Big ticket #5: ADR Technology for the sciences

“Raising the productivity of research may be the most valuable of all of the uses of AI”

The scientific method was developed in Europe in the 17th century. It remains the bedrock of European prosperity and security. To address societal challenges in health, the environment and energy, we urgently need better science and technology. However, the productivity of science and technology research is decreasing (AI in Science: Challenges and Opportunities, OECD 2023).

There is an opportunity to leverage AI to apply the scientific method more efficiently and at scale

Europe is home to a rich variety of world-leading scientific institutes. Less well known is that it is also home to the world's greatest concentration of scientific tool suppliers (Tecan, Zeiss, Qiagen, Oxford Instruments and Biovia) and a long tail of innovative SMEs and start-ups.

The integration of these tools with AI enables systems (AI scientists/self-driving labs) that automate science itself. This research was pioneered in Europe, and there are now around 50 such systems around the world autonomously doing research in areas as diverse as drug design, materials research, and quantum physics.

The AI powering such systems is improving all the time, from symbolic methods to deep learning, and there have been ground-breaking advances such as *AlphaFold* in protein folding, and Large Language Models (LLMs) such as ChatGPT. The infamous hallucination problem of LLMs is not a problem in science, as hallucinations are simply hypotheses, and the ground truth can be tested in experiments carried out by robots. Improved AI implies ever smarter and more efficient automated research.

The ethical concerns around privacy that occur in many AI applications are less of a barrier, while other social/ethical issues, such as the ownership of machine-generated intellectual property and https://safe-deed.eu/dual use, need to be addressed.

Challengers are appearing. Investment to support the automation of science in other regions is ramping up (the US Department of. Energy is spending $1bn per year, while Canada is spending $200m), although China does not yet appear to be very active. These investments will threaten both the European science base and the world-leading tools industry.

Proposed solution. Our proposal is to bring together leading players in AI, robotics, and research and industry to create a vision for the automation of science. The Nobel Turing Challenge is to create a highly autonomous AI and robotics system that can make major scientific discoveries worthy of a Nobel Prize by 2050. Supporting symposia have already taken place in the UK, Japan, the US and Sweden. If this challenge is successful, it would transform Europe and the world.

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

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. Thus there is a major need for research, innovation, and tools for compliance.

This is a very active area, with many ongoing activities including the following:

- The AI Trust Alliance initiative[68] for an AI trust label, which is a joint collaboration between VDESPEC, Positive AI, Confiance.ai and IEEE.
- The BDVA ETAMI[69] task force aimed at creating processes and tools to enable ethical, trustworthy and legal AI.
- The BDVA has summarised the current prominent topics in data protection via its “CURRENT HOT TOPICS IN DATA PROTECTION” position paper.[70]

---

68 https://alliancefortrustinai.org/
69 https://etami.org/
70 https://www.bdva.eu/sites/default/files/BDVA%20DataProtection%20PositionPaper_November2022.pdf

- TAILOR project: Trustworthy AI based on Learning, Optimisation and Reasoning. Relevant results are Strategic Research and Innovation Roadmap[71] v1. and Handbook for trustworthy AI.[72]
- The IEEE CertifAIEd:[73] certification programme to assess the ethics of Autonomous Intelligent Systems (AIS) to help protect, differentiate and grow product adoption.
- The OECD catalogue[74] of tools and metrics for trustworthy AI.
- Characteristics of trustworthy AI[75] developed by VDE.
- CEN-CENELEC JTC21/ WG4 on AI Trustworthiness Characterization standard.
- AppliedAI,[76] is driving an effort to pulling in tools (such as model cards) to automate the use of trustworthiness standards during the development process.
- The Confiance.ai project has written a white book "Towards the engineering of trustworthy AI applications for critical systems - The Confiance.ai program" by M. Adedjouma, JL. Adam, P. Aknin, et al. (2022).
- The SAFEDEED[77] project (ICT-13-2018-2019) explored how privacy awareness and trust technologies could empower data owners and help data platforms to assist in economic growth.

Challenges

Establishing trustworthiness

The **c**ertification or labelling of AI trustworthiness by third parties can build upon existing sectoral approaches. However, trustworthiness of AI applications is a multidimensional concept, related to formal verification, user perception and skill level.

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 the compliance to those attributes in a unified aggregate confidence score.

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.

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 ins may hamper further adoption. Interoperability of PETs, for instance in the context of data spaces, is therefore fundamental.

---

[71] https://tailor-network.eu/research-overview/strategic-research-and-innovation-roadmap/
[72] http://tailor.isti.cnr.it/handbookTAI/TAILOR.html
[73] https://engagestandards.ieee.org/ieeecertifaied.html
[74] https://oecd.ai/en/
[75] https://www.vde.com/resource/blob/2242194/a24b13db01773747e6b7bba4ce20ea60/vcio-based-description-of-systems-for-ai-trustworthiness-characterisationvde-spec-90012-v1-0--en--data.pdf
[76] https://applied-ai.com/
[77] https://safe-deed.eu/

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.

Standards, metrics, benchmarks and norms for the evaluation of Privacy Enhancing Technologies

Compliance cannot always be established easily. Additionally, it is a binary concept: when a service provider wants to move beyond compliance, there is no metric to identify how much 'more compliant' they are. It is also challenging to define suitable metrics and norms for privacy.

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.

Business models relying on privacy-aware services as USP

Although many NGOs address privacy topics and legislation exists to avoid the misuse of personal data, many end users still accept the Terms & Conditions of services blindly when the offered advantages are attractive enough.

Instead of winner-takes-all customer data collection models, designing services and business models that take their value from privacy awareness enhances the overall privacy culture in the organisation.

Life-cycle management of Adaptive AI systems

AI systems that adapt to new situations may start operating outside originally designed bounds, leading to decreased performance and trustworthiness, and even compliance issues.

There is a need for incremental and evolutionary qualification, as there is no strict separation between design and operational phases, due to the dynamic nature of 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 is a related challenge 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

The challenge is to facilitate compliance, with no need to know (all) the legislation (applying to design, development, testing and the overall engineering life-cycle of ADR systems).

To define how models and architectures can support compliance. The coding of legislation and AI-enabled market surveillance.

An additional challenge surrounds compliance for complex systems or scenarios, such as complex supply chains, portability of systems and post market surveillance. This involves the analysis of systems, data spaces and federated platforms.

Use-cases & applications

Towards certification of trustworthiness

Challenge: establishing trustworthiness

As AI systems become more complex, their trustworthiness may become opaque for designers, professional users, end-users and the public.

Both technical aspects of trustworthiness and perceptual aspects are relevant. Attention must be paid to the avoidance of over-trust and under-trust, because 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.

Adoption of Privacy Enhancing Technologies

Challenges: interoperability & standardisation of PETs; standards, metrics, benchmarks and norms for PET adoption

Wider scale adoption of PETs is held back by lack of standards, interoperability and suitable metrics. Standards and interoperability between PET solutions will increase the network effect, so that it becomes easier to protect data collaboration. Inclusion of these standards in Data Spaces is a relevant implementation of this goal.

The goal is to promote PET adoption by providing either an interoperability framework for PETs, and/or to showcase how quantifiable privacy protection metrics can be implemented in a sector that handles sensitive data. Additionally, the adoption of PETs is related to user perception of privacy.

New concepts for compliance

Challenge: Automated compliance and compliance by design

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, and therefore an automated characterisation of the relevant ones, including an assessment of their values, will be needed.

The goal is to create a ‘compliance by design’ approach that enables the design of 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.

These compliance concepts should be demonstrated in areas where the handling of personal or sensitive data is a major concern, such as the health domain, and/or where mission critical automated systems are deployed, such as manufacturing or mobility.

Trustworthy AI in business

Challenge: life-cycle management of AI systems, and business models relying on privacy-aware services as USP

The deployment of AI systems to support products or services should be monitored and updated when the circumstances change, in order to maintain the trustworthy operation of these systems.

The goal is to make life-cycle management of AI systems that ensure trustworthy and privacy-aware operations a USP for businesses. This includes the ability for data holders to exert control over their data and/or to monetise it, to make 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 be demonstrated in domains where end-user trust and acceptance is important.

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 preserving the rights of EU citizens.

Another outcome should be that a decrease in the number of data leaks and citizens' privacy concerns is observed.

Costs for 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 (TRL-levels)

Use case	Initial TRL	Final TRL
Towards certification of trustworthiness	TRL 2 - technology concept formulated	TRL 6 – technology demonstrated in relevant environment
Adoption of Privacy Enhancing Technologies	TRL 3 – experimental proof of concept	TRL 5 – technology validated in relevant environment
New concepts for compliance	TRL 2 - technology concept formulated TRL 3 – experimental proof of concept	TRL 5 – technology validated in relevant environment TRL 6 – technology demonstrated in relevant environment
Compliance in complex operational systems	TRL 4 – technology validated in lab	TRL 6 – technology demonstrated in relevant environment
Trustworthy AI in business	TRL 4 – technology validated in lab	TRL 6 – technology demonstrated in relevant environment

Short-, medium-, and long-term objectives

Short term

- Trustworthiness establishment and metrics definition
- Definition of methods for compliance by design

Medium term

- Standards for PETs interoperability
- Data quality evaluation framework
- Definition of a trustworthiness compliance framework

Long term

- Compliance automation
- Incremental and evolutionary qualification framework

3.2 European AI Moonshot

Recently, the world is experimenting a tipping point where a technological shift i.e., generative AI is hitting our mainstream society. Natural language processing combining generative pre-trained transformers and large language models have gained in maturity in the past years to show impressive capabilities in mainstream applications. For this purpose, several countries in the world are releasing investments and developing strategies to elaborate on the deep shifts triggered by generative AI.

We recommend Europe to work as a whole under a common agenda with a 10-year perspective to be on the leading edge of these technological advances and to elaborate on strengthening and widening Europe's position at control points along the supply chain. We call this common agenda **European AI moonshot.**

Per definition, a moonshot represents a distant goal, a goal that cannot easily be reached and one that will take the combined efforts of many organisations and individuals, over an extended period of time, to achieve. It is also a goal where the precise path and processes are unknown but where the principles needed to achieve the goal are understood.

Europe has set itself the goal of developing and deploying technologies that align with its fundamental rights, regulations and values such as trustworthiness, ethics and security. The goal drives at the core of everyone's daily experience of technology, at personal privacy, at corporate responsibility and at the maintenance of democratic processes. So much of how we live involves technology, from transport and healthcare to education and leisure, that fundamental rights and values can be easily eroded by poor or malicious technology. There is no doubt that technology has an exponential character and that this can cause a loss of control, but likewise if it is built and guaranteed to be trustworthy then it can also enhance our lives and provide value that was inconceivable just a decade ago.

Where technology is used to deliver decision making impacting directly on people it is clear that trustworthiness is needed. This applies to technologies that result in a real impact on people, in businesses, in healthcare, in transport, in construction and in the decision-making tools that drive our financial services, our environmental decisions our energy infrastructure and our food supply chain. In all these areas a failure of decision making has a real-world consequence, a loss of productivity, a loss of capacity even the loss of life or livelihood.

AI moonshot with European control points along the value chain to boost European competitiveness and strengthen technology trustworthiness where Europe has already set its goal.

3.2.1 The rationale behind AI moonshot

The key elements triggering the moonshot are:

- We need to reinforce our innovation ecosystem with resources and capacities to be on the leading edge of technological advances in the actual era of generative AI.
- We need to boost our industrial competitiveness to maintain or even grow our world market share.
- We need to facilitate to our research, innovation ecosystem and SMEs access to high-performance infrastructure to be on the front of science and technological development.
- We need to develop innovative processes that enable creating new jobs and new businesses in the new digital era.
- We need to develop techniques, standards and processes that can reinforce European values such as trustworthiness along the value chain.
- We need to strengthen and even widen European control points along the value chain.

3.2.2 The AI moonshot: a complex mission

The moonshot should be designed as a 10-year program, with an overarching goal of establishing European AI technologies and techniques, which align with its fundamental rights, regulations and values along the value chain. The expected outcomes encompass a collection of APIs, toolboxes, hardware elements and standards, which can strengthen and widen Europe position on strategic control points along the supply chain of generative AI. To elaborate on the complex mission, we need to combine the 10-year perspective with the current market power and exploit actual European champions with leading market positions to set-up intermediate milestones towards the moonshot. The moonshot should not only rely on public investments, but also stimulate private financing and engagements to ensure smart cost control and convergence to the overarching goal. The milestones will be designed as 3-4 steps (from complexity low to high) along the program timeline to first elaborate on APIs, toolboxes, standards until final hardware and computing components. The demonstrators and prototypes are required along the process for proof-of concept and validation of the progress.

3.2.3 Trustworthiness as one pillar in the AI moonshot

In complex technical systems decision making is distributed, humans are often involved at critical points but they rely on technology for the information on which they base their decisions, and on technology to carry out the results of their decision making. Often technology must work faster than humans can think, or work where humans cannot go, or combine more streams of information than a human can handle. If we are to trust technology to take on these challenges then we must be convinced, at least as far as humanly possible, that the decisions being made align with our fundamental rights principles and values.

We have seen how un-trustable technologies can fail, we have seen how poorly maintained systems break, and we have seen how technology can be used to manipulate and coerce. The moonshot we propose cuts across a wide range of technologies and techniques to set the goal of making them work together in complex operating environments in ways that are trustable. We know this is only possible by gathering a wide range of different types of information, combining that information with different analytical tools and developing decision making processes that result in trustable actions who’s effect in the real world can be guaranteed. This is not about any one technology it is about all these technologies working together to provide trustable capabilities.

3.2.4 AI Moonshot for European control points

From a business and political perspective, the goal of the moonshot is to strengthen and widen Europe control points along the emerging generative AI hitting our society. Europe has to strengthen its technological sovereignty and strategic autonomy to be able to steer the technological development and deployment that align with its value, fundamental rights and regulations. From a technologist's perspective the goal of a moonshot is to create new technologies and techniques that deliver currently unattainable use cases and open new market opportunities. The time taken to achieve a moonshot means that the perspective of the achievable changes from the beginning to the end of the moonshot. So in setting a technical direction we can only elucidate the current state of the art and point an arrow in the direction we intend to start. But like any long journey undertaken on foot the precise path will only be known at the end of the journey, even though the goal is well defined.

From today we can see the limitations of AI, of the robotics we can currently build and of the analytic techniques our computational systems can deliver. In setting out this moonshot we must both push those boundaries but also ensure that the knowledge gained, the decisions taken and the actions made are aligned with each other as well as being trustworthy.

Current AI learning models, such as LLMs, do not provide sufficient knowledge of the physical world for a robot to operate trustfully, and their reliance on large bodies of data, much of it historic, can create the danger of lock-in, not just technically but culturally. If laws or financial regulations change how can such models, built on history, be updated? How can systems reliant on such models be recertified? How can we ensure that regulations are being followed? How can we avoid harms? Europe has a long history of technical achievement in AI, Data Science and Robotics that covers not just learning but reasoning. Answering these critical questions involves both learning and reasoning in combination, a multi-modal approach to complexity that encompasses knowledge acquisition and decision making, and embeds it in the physical world where decisions have real consequence, be that in healthcare, financial markets or in supply chains.

This multi-modality must cover all forms of knowledge from the linguistic to the physical seeking to combine multi-sensory information with linguistic as well as physical knowledge of objects, structures and physical properties such as sound or chemical composition. Without this multi-modality decision making will be based on incomplete information and in a physical world where decisions have real consequence this will lead to untrustworthiness.

Trustable systems need to be timely, to respond in time frames that match the unfolding of physical events in the real world. This is the ability to assimilate information, translate it to the current context and create trustable and timely decisions and actions. The ability to reason and abstract, model and infer is as critical as the ability to recognise, segment and adapt. The frugal use of data is paramount to speed.

To be trustable these multi-modal models must also be able to account for the "other". The autonomous agents they will encounter including humans and the laws of physics. Ascribing trust implies inbuilt knowledge of the "other" that includes ourselves and the science that governs the world around us. We know that apples fall from trees, that chemicals react and that ice is slippery. Trust assumes these facts are known and immutable.

Europe should also strengthen its technological developments to deliver made in Europe hardware and computing components that enable smooth deployment of generative AI and be resilient to global tensions and conflicts (e.g., raw materials, semiconductors) that are hitting the supply chain. The Made in Europe Hardware is a very complex milestones as USA and China are heavily dominating the supply chain. Therefore, this milestone is a core pillar of the AI moonshot.

3.2.5 Delivering the Moonshot

Goal:

→ *To develop AI-powered systems, products and services aligning Europe with its fundamental rights, regulations and values and strengthening its position on strategic control points along the supply chain of generative AI.*

This goal specifically addresses the use of AI in real world contexts where decision making has actual impact on the lives of citizens in their work and home environments. To make that impact certifiable, guaranteeable and understandable.

Reactivity:

Current programme mechanisms in Europe are too slow for the exponential changes that can be achieved with new technologies. The three-year cycle from strategy to execution is too long. Instruments set up to deal with slower, scientific, discovery do not match the rapid impacts we see from new technologies and cannot handle the inherent multi-disciplinarity of these impacts.

Strategy:

To create a well managed industrially relevant, scientifically excellent, moonshot programme that can achieve Europe's goal with a dynamic work package approach that is centrally coordinated and managed. Calls need to be frequently executed with a profile of call time scales and expenditure that can create short term high impact as well as long term depth. Strategy for each call cycle should be handled wholly within the moonshot guided by an overarching strategy built to deliver outcomes that align with the goal. Public scrutiny through industrial, scientific and operational oversight provides steering and advice.

Structure:

The Moonshot should be built on three interlocked themes:

Scientific Excellence: Europe has excellent science in all the areas relevant to the goal. This moonshot will bring this diversity together by extending capability, expertise and skill.

Driving Uptake: Everything done within the moonshot must have economic, societal and industrial relevance. Some of the outcomes will be capable of acceleration and impact other outcomes may take a decade before their relevance is felt. This includes relevance to core societal missions. The goal of the moonshot is to not just develop societally relevant technology but to deploy it.

Innovating Infrastructure: The innovation infrastructure, from chips to testing, of these new technologies needs to be relevant and sufficient to achieve the task. Not only relevant to the scientific challenges but relevant to the subsequent uptake and impact that will result from it. This is about building capacity in the scientific domain and in the supply chain actors that deliver the economic and societal outcomes.

Achieving the goals of the moonshot may require new instruments with greater flexibility, it will need a new way to set programme goals that removes the blockages to European progress that can be created by the current systems of consultation. It requires a bold approach.

Management:

Exponential technology development needs reactive structures to drive it. Basing decisions on ideas and directions set three years before that will be 10 years out of date at its conclusion is a recipe for failure.

Project management: The moonshot must be managed by management professionals with a track record of science and technology management. Scientists must do the science, managers must manage the project. Management must continuously steer the project to the goal in whatever way is needed. When necessary taking hard decisions to cut or change direction based on the best advice.

Oversight: Separate oversight boards for Industry, Science and Impact will be needed with a combined oversight board. Outcomes from these boards are public, regular and advisory. It is expected that, at times the advice will be strong and directing.

Impact management: A major part of the moonshot will be both assessment of and management of impact. There will need to be a strong regulatory input to the moonshot and a strong input to regulation from the moonshot. Citizens will need to be engaged, industry at all scales and across all sectors will need to be aware, engaged and invested in delivering the best for Europe.

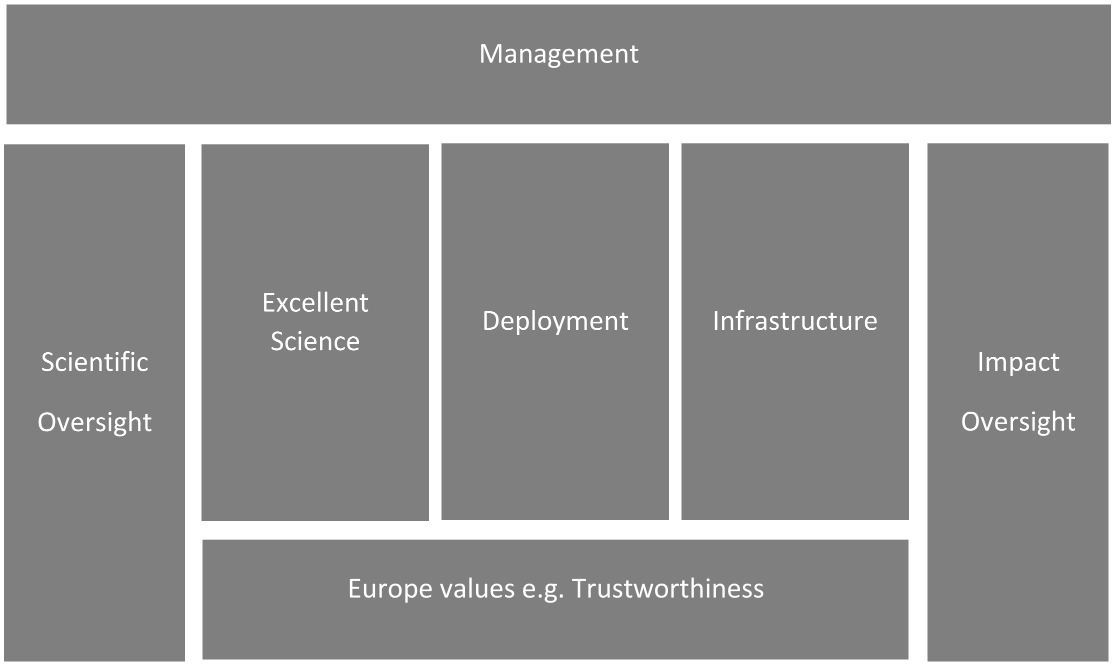

Content:

The moonshot programme should

- cover core science with integrated projects delivering clear value into the goals of the programme.
- enable small companies to take up current technology, help them develop capability and understanding to raise productivity. All the current analysis clearly shows that adopting AI raises productivity.
- engage the large sector-based industries to both establish the realities of deployment and to drive uptake that stimulates their supply chains.
- engage and assess impact, ensuring that impact aligns with the fundamental rights and freedoms enshrined in Europe and that regulatory change is enacted to both protect citizens from harm and provide the openness that innovation needs to make rapid progress.

Cost:

It is envisaged that the moonshot programme will operate over a full framework cycle of 10 years with the expectation that a further three years is enabled by the next framework programme. Thereby **providing a decade of stability** to make cohesive progress towards the European goal.

In the early stages funding is needed to boost infrastructure and bring the communities together. While the pandemic has taught us that much can be achieved by on-line collaboration it has also taught us that human connection through physical meetings and the establishment of a working centre is essential to efficient and rapid progress. There is a need to construct a European centre of excellence in AI. This cannot be slowed down by arguments with partner countries. Everyone will benefit from the creation of such an institution.

Private investment will be needed. AI uptake in the sectors where AI has real world impact is going to be significant only if it is trustworthy and certifiable. Every sector has a vested interest in investing in making AI that conforms to this European goal. Therefore, private investments will inflow over the decade but only if the relevance of the outcomes is visible, investable and realistic.

This needs to be a €100Bn investment from both the public and private sectors, the profile of that spend and the direction of that spend will shape over time. If this is managed well the acceleration of private investment outside of the moonshot programme is also expected to accelerate potentially to a multiple of that value. The goal is to place Europe as the global supplier of trustworthy AI, not just be excellent at the science behind it but excellent at deploying it to drive up productivity, well-being and trust in society at large.

3.3 Education, re-skilling and up-skilling

A lack of AI-related skills, including technical AI skills, skills for dealing with AI and managerial capabilities to extend AI in business, have all been identified as the most important barriers to AI adoption in Europe in the policy recommendations of the High-Level Expert Group on AI (ref).

AI talents with expert knowledge are required who are capable of driving, managing, and conducting AI activities in their institutions and organisations. However, the creation of an excellent innovation environment does not only necessitate funding. Europe also needs to attract, develop and retain a comprehensive talent pool of AI-developers, entrepreneurs, and data analysts and create a beacon for talent. This will require dedicated effort in establishing cross-disciplinary AI curricula in tertiary and post-graduate education that integrate ethics, humanities and technical disciplines, in mainstreaming AI-related skills in all academic disciplines and professional fields, and in fostering spaces of cooperation between AI experts and professionals.

Despite the urgency expressed in these reports and discussions in the education sector, few solutions and suggestions have been proposed regarding how the formal, structural and cultural obstacles can be lowered or entirely removed to facilitate the identified transformation.

There are many important aspects to be addressed related to competence.

1. **AI awareness and AI engagement**. Raising the awareness and understanding of ADR technology broadly throughout Europe. Better understanding of the risks and opportunities of ADR technology will help individuals make better decisions and reduce the risk of deception. Education is the best way of increasing individual resilience in turbulent times. Interesting initiatives such as Elements of AI exist, and it would be very interesting to develop more of these broadly accessible courses capable of reaching large numbers of individuals. We also need to understand what makes people engage, or disengage, with ADR technology, so that the educational material can be adapted to better suit the highly heterogeneous audience. Research and development into ADR edutainment is also highly desirable.
2. **ADR literacy**. Integrating knowledge of ADR technology into primary and secondary education. There is ongoing research related to AI literacy's attempts to identify what

everyone needs to know about AI and how to integrate this into the formal education system, including teacher education.

3. **Developing a European ADR curriculum**. Today ADR technology is included in many different education programmes, but no bachelor or master's programmes dedicated to ADR exist. As the field has significantly developed over the last decade, there is also a need to describe and agree on the appropriate competences, knowledge and skills people in ADR need to develop, deploy, operate and maintain ADR technology.
4. **Scaling-up education capacity in ADR technology**. Significantly expanding the education of technical experts in ADR technology. Europe has traditionally very strong higher education institutes that educate experts in ADR technology. The challenge is therefore to significantly increase the quantity while maintaining the quality.
5. **Scaling-out education in ADR technologies to disciplines and professions beyond the technical core**. Introduce AI in education beyond only the specialised education programmes for technical experts, as the need for competence spreads from the experts developing the technology to a much broader range of professions and disciplines.
6. **Re-skilling and up-skilling of professionals**. The vast majority of the practitioners that will be working in ADR technology tomorrow have already completed their formal education. Likewise, most companies and organisations have workforces that must be developed in order to effectively use the new ADR technology introduced in the last couple of years. This is especially important, as we see studies showing that humans effectively using ADR technology significantly outperforms those that do not[78]. Thus, the greatest need for education is among those that are already working. To address their needs for competence and help individuals transition into the area, significant efforts are needed in the areas of re-skilling and up-skilling. Both pedagogical and organisational innovations are needed to address these challenges.

Addressing these concerns will be absolutely essential for Europe to successfully leverage the ADR technology being developed and transform it into value creation for companies. It is also essential for individuals to remain competitive in the global labour market.

3 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. The focus of Adra is 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

[78] HBS Working Paper 24-013 "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality" by F. Dell'Acqua et al. https://www.oneusefulthing.org/p/centaurs-and-cyborgs-on-the-jagged

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 doing things 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. The region needs to step up.

To achieve the vision of Trustworthy ADR, Europe needs to match its regulation with innovation. 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 both people and climate is high.