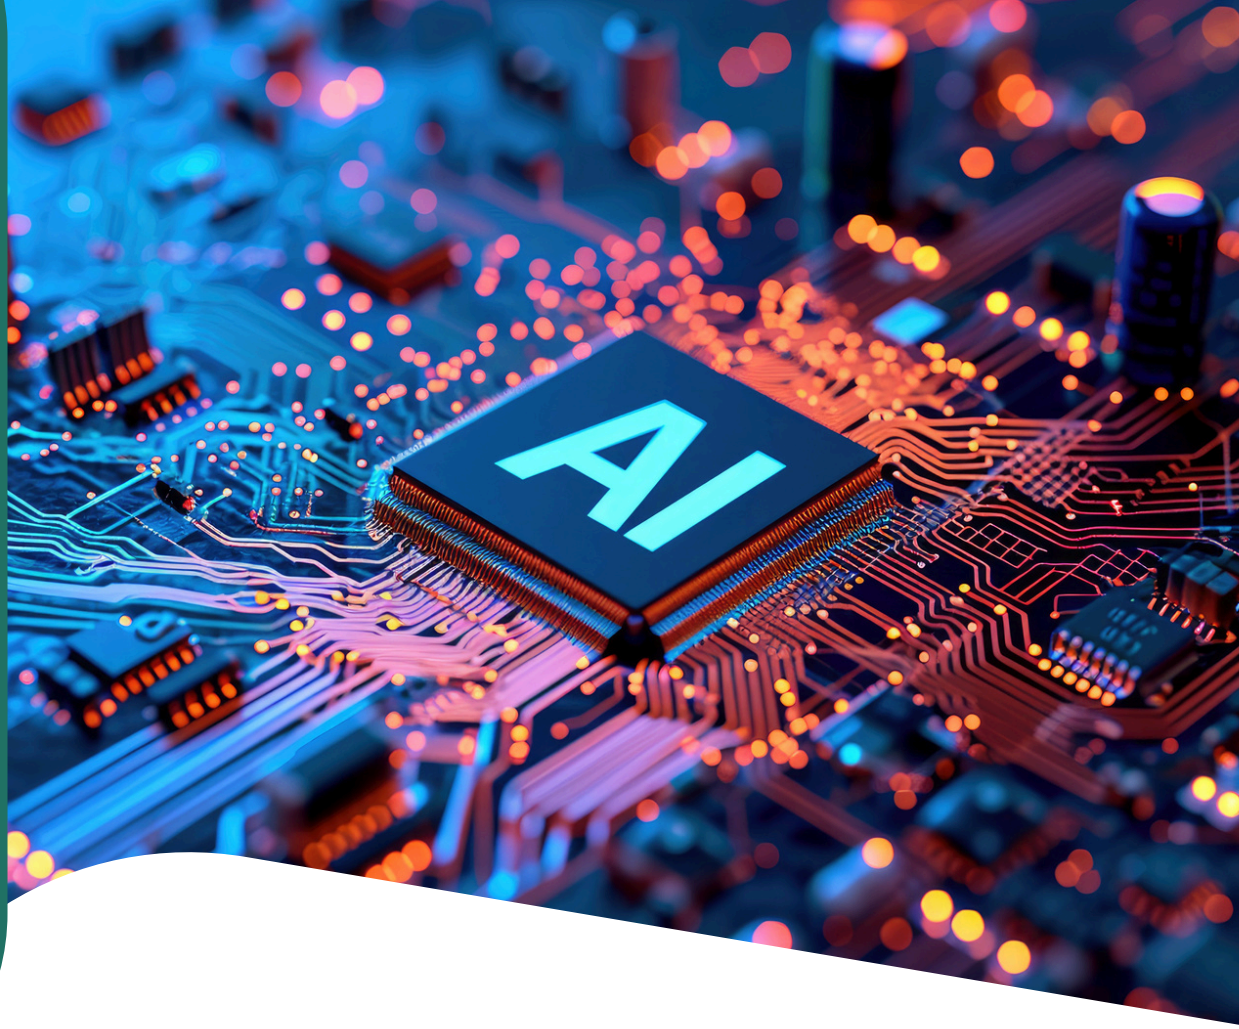

TOWARDS A EUROPEAN LARGE-SCALE INITIATIVE ON ARTIFICIAL INTELLIGENCE

What are the options?

Andrea Renda

CEPS IN-DEPTH ANALYSIS July, 2024 - 11

SUMMARY

While being a leading world region when it comes to regulating Artificial Intelligence (AI), Europe has accumulated delays in creating a fertile ecosystem for investment and innovation in this crucial domain for economic competitiveness, strategic autonomy and democratic resilience.

To remedy this situation, several proposals have been presented for an ambitious initiative that catalyses research and innovation around trustworthy AI solutions made in Europe and according to European values and principles. Among them, the proposal tabled in March 2024 by the Scientific Advice Mechanism of the European Commission, part of DG Research and Innovation, to create a 'CERN for AI in Science'; the long-standing 'AI moonshot' proposed by the Confederation of Laboratories of AI Research in Europe (CLAIRE), now backed by several other associations; and the proposal to create a Sovereign Tech Fund for Europe's Digital Public Infrastructure.

This CEPS In-Depth Analysis paper presents a possible initiative that combines these proposals into one single large-scale endeavour and discusses the main conditions through which such a proposal could be an attractive prospect, at a time when restoring competitiveness and achieving technological sovereignty have become major priorities for EU leaders.

Andrea Renda is the Director of Research at CEPS and Head of the CEPS Unit on Global Governance, Regulation, Innovation and the Digital Economy (GRID). The author wishes to thank Katja Spanz and Robert Praas for their assistance and useful comments; and the participants of a roundtable on the same topic as this In-Depth Analysis paper, held at CEPS in Brussels on 24 June 2024.

CEPS • Place du Congrès 1 • B-1000 Brussels • Tel: (32.2) 229.39.11

INTRODUCTION: A WIDENING GAP AND AN EMERGING IMPERATIVE

Evidence shows that in Artificial Intelligence (AI), as with other digital technologies, Europe has accumulated a significant gap *vis-à-vis* other regions of the world, notably the United States and China (ECA 2024, OECD 2024; Council of the EU 2024).

This is potentially problematic not only in terms of competitiveness, a priority that EU institutions have placed back to the top of their agenda over the past few months (Renda 2024), but also in terms of strategic autonomy and technological sovereignty. AI is expected to have a pervasive impact on the economy and society, including on the democratic process and on public services. Should it continue to depend on a fistful of non-EU tech giants for all its digital infrastructure, Europe would remain vulnerable to external pressure and be doomed to run its economy on non-EU standards, principles and socio-technical solutions.

Past research has shown more evidence of this widening gap with the rest of the world's superpowers. According to the Stanford AI Index, private investment in AI reached EUR 62.5 billion in 2023 in the US, whereas it did not go beyond EUR 9 billion in the EU and the UK combined. Even more importantly, the UK counts for more than a third of this figure (EUR 3.5 billion), which means that private sector investment in the EU27 was less than one tenth that of the US in 2023 and also significantly lower than the Chinese figure (EUR 7.3 billion).

Likewise, Balland and Renda (2024) find that over 60 % of venture capital invested in AI has so far flowed to US-based start-ups, with 17 % going to China and only 6 % to Europe (including the UK). A similar gap can be found by looking at patents, where Europe seems to hold more weight at the global level when it comes to leading AI researchers. Importantly, Europe is able to train more advanced researchers than the US but then fails to retain such massive talent, which massively flows to other parts of the world – notably across the Atlantic.

More recently, the emergence of powerful generative AI systems such as ChatGPT, Google Bard and Anthropic Claude onto the market has triggered a new boom in private investment in the US. Tech giants are securing important stakes in maverick firms such as OpenAI and Anthropic, in a market that concentrates more and more into the hands of only a few powerful players with slightly different business models – and all operating from outside the EU.

At the same time the venture capital market seems to have gradually lost interest in societally and environmentally relevant AI applications and is increasingly focused on start-ups pursuing so-called Artificial General Intelligence (AGI). Many of the most

powerful solutions have also gradually become less transparent and open, which poses further problems in terms of AI systems' trustworthiness that are increasingly defined as 'foundational', as they can be used for a variety of downstream use cases. Such opacity further strengthens the need to promote Europe's technological sovereignty by developing AI models that are compliant with the principles and values rooted in the EU treaties, ethically aligned and socio-technically robust – what the EU High-Level Group on AI defined back in 2019 as 'Trustworthy AI'.

Recently, the Council of the EU's research team estimated that 73 % of cutting-edge AI Large Language Models (LLMs) are being developed in the US and another 15 % in China, while EU companies simply cannot release this kind of technology (with only a few exceptions). The recent Microsoft investment in Mistral, a French LLM-company born out of a 'fork' of Meta LLaMA, has cast further doubts on Europe's ability to launch a fully European LLM based only on market forces. Its recent EUR 600 million investment round largely included North American companies, alongside French banks.

This CEPS In-Depth Analysis paper maps some of the possible options for Europe to invest and reverse this rather gloomy trend. It takes stock of the proposals that have sat on the table over the past few months and years for a large-scale initiative aimed at pooling resources and creating a fertile environment for an 'ecosystem of excellence' in the EU. Then it draws on the literature on mission-oriented innovation policy to explore the contours of a future initiative on AI. Finally, it briefly concludes by placing the possible initiative within the context of the EU's evolving political landscape.

AN EU 'ECOSYSTEM OF EXCELLENCE' IN AI: HOW DO WE GET THERE?

The European Commission already acknowledged in its 2020 White Paper that for the EU to thrive in the AI domain, both an 'ecosystem of trust' (broadly referring to a well-shaped policy environment that protects against risks to fundamental rights and safety) and an 'ecosystem of excellence' (mostly referring to investment, innovation and competitiveness) would be needed.

To be sure, on the former, a substantive debate has emerged, culminating with the AI Act's adoption and the creation of the AI Office, respectively the world's first major comprehensive legislation on AI and possibly the most powerful monitoring and enforcement body for foundation AI models in the world. These are important steps towards enhancing trust in the use of AI among citizens and businesses, even if implementing the AI Act will be far from easy and the related work far from over.

On the side of excellence, unfortunately, so far only sparse initiatives – often fragmented across Member States – have seen the light of day. One exception is the recent AI innovation package, launched by the Commission in January 2024, with the aim to make Europe's supercomputers available to innovative European AI start-ups willing to train their AI models. Access to compute infrastructure, an initiative that mirrors similar measures in other parts of the world (e.g. the National Artificial Intelligence Research Resource, or NAIRR programme in the US), is however only a very first step towards creating the adequate framework conditions for a flourishing ecosystem at EU level. Other key conditions include data availability, talent and capital.

Importantly, it would be a mistake to think that Europe should limit itself to mimicking the US to promote its ecosystem of excellence, not least since the evolution of the US AI market seems to be deviating from the trustworthy AI features the EU would like to see circulating around its Single Market. Moreover, even in the US case, the public sector has historically played a very important role in promoting connectivity infrastructure and a flourishing semiconductor industry. It made recurrent use of its purchasing power in procurement (including with the Apollo programme) to secure stable demand for technology and more recently reverted to rather interventionist industrial policy with ambitious measures such as the CHIPS Act and the Inflation Reduction Act. The difference in efficacy is striking – in AI innovation there is only one United States but many European countries looking to develop and grow their own national champions.

Against this background, to restore its competitiveness, the EU may have to leverage its peculiar approach to innovation and industrial policy, which Kalff and Renda (2019) listed as among the EU's 'hidden treasures'. In the AI context, this implies leveraging a mix of public and private forces, including researchers as well as large and small corporations,

to give life to a sustainable, resilient and human-centric approach to AI development, deployment and use. Crucially, given the fast-paced evolution of AI and the need to rather swiftly catch up with global competitors, it would also involve a degree of mission-oriented innovation policy. This would enable a variety of experimental pathways and the pursuit of several technological breakthroughs to ensure the success of the set mission.

On this latter point, an early reflection took place in the Commission back in 2016, especially in the ESIR expert group and later in the related Mazzucato Report, resulting in the launch of five 'missions' as part of the Horizon Europe programme. However, for reasons that will be explained in more detail below, the five missions have failed to incorporate all the elements that would have made them well-equipped to achieve their final objective in an innovative way, and with significant spillover effects for the whole of the economy and society.

Finally, any initiative aimed at boosting Europe's AI ecosystem of excellence should address the evident problem of fragmentation in today's innovation landscape. Fragmentation deprives the EU of the scale that would be needed to leverage the Single Market's potential, both in terms of supply and demand. As of today, Balland and Renda (2024) show that AI-related R&D and venture capital is concentrated only in a fistful of European cities, namely London (by far the leader), Munich, Paris and Eindhoven. Even more importantly, start-ups and larger companies revolving around these hubs do not cooperate with each other in any meaningful way, reflecting an overall lack of cohesion and coordination in the EU's approach.

COMPARING SELECT PROPOSALS FOR EUROPEAN AI INITIATIVES

Faced with fragmentation and insufficient scale, let alone the unavailability of the necessary infrastructure and capital, several European researchers have taken action to demand bold, courageous initiatives such as 'Moonshots', Missions or simply ambitious large-scale initiatives. At least five proposals have been tabled and discussed over the past few years.

Already in 2017, the first calls to action started to mention the possible earmarking of funds for a pan-European large-scale flagship project in the AI domain, often equated to the level of effort that led to the creation of giant initiatives such as Airbus or CERN. These are characterised by shared significant investments into concentrated efforts to build complex technologies (in the above examples, airplanes and particle accelerators).

The Confederation of Laboratories on AI Research (CLAIRE) has been the most vocal in this debate, which was later backed by other associations representing a good part of the European research community focused on AI and robotics (EurAI, ADRA, euRobotics). The revised proposal by CLAIRE and euRobotics points to the need to develop alternatives to generative AI systems that currently dominate the market. It proposes a CERN-like model, an 'organisation and hub at the centre of a broad network of AI competence centres throughout the EU and associated countries'. The proposal outlines a rather short-term timeframe, pointing to the need to invest EUR 100 billion between 2024 and 2029.

Shortly after the first calls for a CERN for AI, Renda (2019) explored the possibility of launching a mission on AI in the context of Horizon Europe and concluded that, at the time, there was 'no single, easily specified set of milestones or achievements' that could be 'easily associated with AI development'. That said, the paper advocated the launch of a broader 'Mission IT' that would embrace the whole technology stack, its ethical, societal and environmental impacts, and the possible consequences for education and skills. This was an attempt to also include framework conditions such as data governance and availability, and secure and trustworthy cloud/edge/IoT layers into the overall quest for AI excellence in Europe.

In recent years, other initiatives have emerged with the ambition to unleash the potential of trustworthy AI by creating an open and trustworthy Digital Public Infrastructure in Europe. A notable example is a recent proposal by Open Future, which echoes the launch of a Sovereign Tech Fund in Germany, the Next Generation Internet initiative funded under Horizon Europe and the French government-led effort to support infrastructures for the Digital Commons. Open Future proposes the launch of a European Digital Public Infrastructure Fund. The proposal was later echoed by Bria (2023) with a proposed budget of EUR 10 billion.

Other proposals have been presented even more recently. For example, Praas (2024) proposes the launch of a Moonshot for a European LLM, led by the European Commission (DG CNECT and DG RTD) and involving various associations such as the AI, Data and Robotics Association, Horizon projects, AI in Science and the Large AI Grand Challenge, bringing them all under one mission umbrella. This mission could then morph and evolve into a fully-fledged 'CERN for AI' that would cover next-generation technologies over a longer timeframe.

He also proposes that the European LLM should be open source, trustworthy and developed in three sizes: one small and efficient version that can run on devices such as smartphones and tablets; one medium-sized and low-cost model that can handle frequent and less precise tasks; and one very large model in the range of GPT-4 and Gemini for the most challenging tasks. Thus far, the action most closely related to this has been the announcement of the Alliance for Language Technologies (ALT-EDIC), a consortium of 16 countries dedicated to developing a common European infrastructure in Language Technologies, with a particular focus on LLMs.

More recently, and in response to a specific request by Vice President Margrethe Vestager, the Commission's Scientific Advice Mechanism (SAM) proposed the creation of a CERN-like structure for AI and science, which seems to significantly differ from CLAIRE's proposal but may not be incompatible (see next section). More specifically, the SAM proposes a European Distributed Institute for AI in Science (EDIRAS). This proposal, like some of the others described above, looks at framework conditions for AI development and thereby aims to create the ecosystem needed for AI to flourish.

In more detail, key initiatives recommended by the SAM include:

- The establishment of a European Institute for AI in Science, which would offer extensive computational resources, a sustainable cloud infrastructure and specialised AI training for scientists.
- High-quality standards for AI systems (i.e. data, computing, codes), with fair access conditions for European researchers and innovators.
- Transparent public AI models, to support the trustworthiness of AI, and able to ensure the reproducibility of research results.
- Tools and technologies specialised for scientific work, (e.g. foundation models for science, scientific LLMs and AI research assistants).
- The promotion of AI-powered research with major benefits for EU citizens, in areas such as healthcare.

- Ensuring a Human- and Community-Centric Approach, which entails research into 'the philosophical, legal, and ethical dimensions of AI in science, ensuring respect of human rights, transparency and accountability' and the active promotion of 'AI literacy'.

Finally, this year, EIT Digital and AiNed proposed a 'Mission for Generative AI' with *ad hoc* funding mechanisms for model development, computing infrastructure, data, skills, and research and innovation[1].

Table 1 on the next page compares these proposals in terms of overall goal ('the mission'), the justification, the budget required and the timeframe for achieving the goal.

[1] Also, Eliot Jones (2024) proposes a CERN for AI as 'a single, coordinated global institution seeking to promote or research safer AI' focused on the governance of AI risks and technical infrastructure to support researchers from countries with fewer computational resources. However, as this proposal is not focused on Europe, its comparability with other proposals mentioned in this CEPS In-Depth Anaylsis is limited.

Table 1. Comparing proposed large-scale initiatives on AI

Proposal	Instrument	Goal	Main Actions	Timeframe	Budget
AI Moonshot (CLAIRE and other associations, 2018 onwards)	Moonshot	European GenAI alternatives to dominant non-EU products	Establish an organisation and hub at the centre of a broad network of AI competence centres (with research, compute and data infrastructure)	2024-2029	EUR 100 bn
Mission IT Renda (2019)	Moonshot	Achieving digital innovation by coordinating industrial policy	Appoint an agency or portfolio manager and target specific societal problems	N/A	N/A
European LLM Praas (2024)	Moonshot	Developing a pan-European open LLM	Merge existing AI initiatives, research into ethics, transparency and privacy, develop different-sized language models	N/A	EUR 10 bn
A CERN for AI (For Science)	Large-scale investment programme	European Distributed Institute for AI in Science (EDIRAS)	Support the broader scientific community in developing and using AI conforming with European values, including fairness	N/A	N/A
EU sovereign tech fund (Keller, 2022; Bria, 2023)	Large innovation fund	Creating a full European DPI to enable AI	Support the deployment of open digital tools such as open AI models and applications, data spaces, open knowledge tools, privacy-preserving digital IDs and digital payments	6-24 months	EUR 10 bn
A GenAI Mission (EIT Digitial, 2024)	Moonshot	Globally competitive excellence in (European) generative AI	Focus on model development, computing infrastructure, data, skills, and research and innovation	N/A	N/A

Source: Author's own depiction.

ARE THE PROPOSALS INCOMPATIBLE?

A high-level overview of the six proposals suggests a high level of compatibility. All proposals share an emphasis on the need to secure the availability of trustworthy AI in Europe, even if they do not elaborate in-depth on how to achieve this goal.

The term 'trustworthy' directly refers to the work of the EU High-Level Expert Group on AI, which defined it as a three-pronged framework centred on legal compliance, ethical alignment and socio-technical robustness. It then articulated the term by identifying four main ethical principles and seven requirements that could be complied with through both technical and non-technical means. Trustworthiness is essential for the uptake of AI in science but also in robotics applications as well as in public services – including healthcare – where citizens need to be able to place trust in the services they access.

At the same time, all proposals say very little on which technological innovations would be needed to achieve trustworthiness. In this respect, some of the European research community (e.g. the TAILOR network of excellence) observed that current generative AI models do not present sufficient safeguards in terms of trustworthiness. This is due to the recurrence of so-called hallucinations and the fact that the US research community (largely dependent on, or funded by, tech giants) seems to have lost interest in a particular research area, that of reasoning, which could potentially help to achieve higher levels of trustworthiness. Here, Europe can leverage its current excellence in research but would need to step up its efforts in a coordinated and mission-oriented way to achieve the needed R&I breakthroughs.

Importantly, the existing proposals cover, in some way, many complementary aspects of the so-called AI ecosystem of excellence, namely fundamental and applied research, tools, standards and applications in at least three crucial domains where the EU needs to get its act together over the coming years. Science, robotics and public services are all more specific use cases of AI that require a more specific focus. Additionally, altogether they also potentially cover the need to promote the skills that complement AI systems, thereby reinforcing their trustworthiness through a savvier human component. One could also add that with a strong research effort, these initiatives could generate economies of scale in innovation. Thus, on top of being more than compatible, they could also even be self-reinforcing by generating network and learning effects.

Figure 1 below visualises the possible compatibility between existing initiatives. This sketchy view of a possible future large(r)-scale initiative on AI still lacks many of the elements that an ambitious initiative would need to have to become appealing from the perspective of public (R&I) policy – a specific, measurable and time-bound set of objectives, as well as demonstrable added value, identified possible spillover effects into

a potentially broad range of economic sectors, a well-defined timeline and agreed agile governance arrangements, including a definition on the approach to risk management along the lifetime of the initiative.

Below, we unpack these elements in view of a discussion with the relevant stakeholders.

Figure 1. Structure of a possible large-scale initiative on AI from lab to market

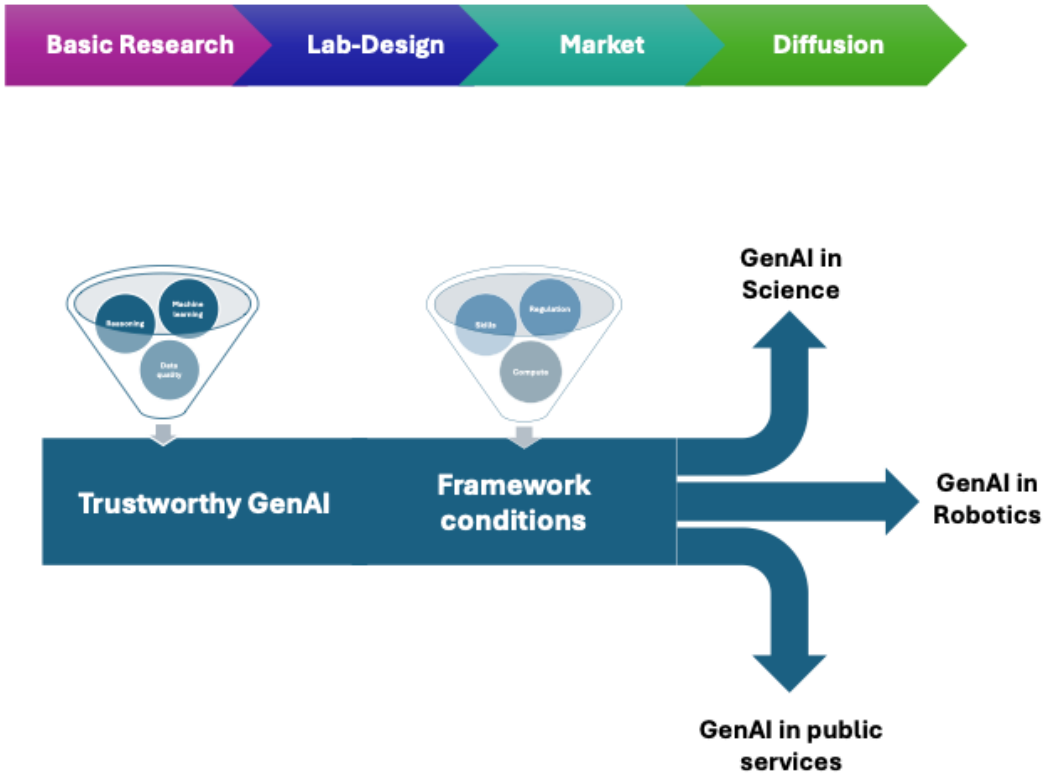

Source: Author's own depiction.

BUILDING A PAN-EUROPEAN LARGE-SCALE AI INITIATIVE

Against this background, a bold pan-European initiative on AI should be accompanied by concretely specifying the overall objective and related activities to be performed (the *'what'*); an analysis of the added value of this initiative compared to any suitable alternative, including the option of simply relying on market forces (the '*why*'); an analysis of the *modus operandi* of this initiative, including the governance structure and, in particular, the leadership and institutional arrangements that would be required, as well as the modalities for the selection and coordination of research activities (the '*how*'); and an indication of what the timeline for the initiative would be, as well as for achieving intermediate and final objectives (the '*when*').

As of today, none of this is clear. This CEPS In-Depth Analysis paper aims to provide a first set of considerations on all these aspects, aiming to contribute to a wider debate during the second half of 2024 and into the first half of 2025, when the moonshot should be launched.

The 'what' – defining the scope and objective in a specific and measurable way

The combination of the proposals presented above would need to be described and defined in a very specific and measurable way so that each component is part of an overall plan to achieve a specific, measurable and time-bound objective. If one analysed all existing proposals together, the general objective could be described as creating a sustainable ecosystem for AI excellence and trust to:

'Secure the design, development, deployment and uptake of cutting-edge trustworthy AI solutions for science, industry and public services in the European Union by 2035'.

This general objective can be unpacked as follows.

First, combining **design, development, deployment and uptake** (all in the EU) implies that the initiative would aim to promote the whole ecosystem, including the full AI value chain, the preconditions in terms of compute infrastructure, data availability and stewardship, and users' complementary skills among others. This requires acting at all levels of the value chain, depending on the perceived need.

For example, the large-scale initiative on AI may already have to entail defining specific design practices that incorporate trustworthy AI principles at the design stage, plus guidance on how to comply with specific risk mitigation strategies during subsequent stages, as well as what technical documentation and assistance needs to be provided to entities deploying trustworthy AI solutions in the chosen domains (science, robotics, public services). It may also entail actions in the form of *ex-post* regulatory controls and

RegTech, aimed at complementing design-based regulations with incentives to comply with the AI Act's regulatory requirements.

Second, the reference to **cutting-edge trustworthy AI solutions** incorporates a technology-neutral reference into the need to achieve breakthroughs in research in a way that guarantees more trustworthy AI solutions can be deployed in Europe. This may (or may not) entail reasoning and the use of specific approaches (e.g. neuro-symbolic AI models). The solution will have to be found when implementing the initiative, as is normal for a large-scale R&I project, and should entail some form of portfolio management (see below).

The reference to cutting-edge solutions also implies that Europe will catch up with other countries and regions of the world in offering world-class AI solutions, which itself requires investment in the innovation ecosystem. It also implies needed breakthroughs on the hardware side, from connectivity (5G/6G) to compute infrastructure, chips, the Internet of Things and robotics, and digital twins.

Third, the word **trustworthy**, as already mentioned, directly refers to the work of the EU High Level Expert Group on AI and incorporates all the specific safeguards, principles and requirements embedded in the 2019 Ethics Guidelines for Trustworthy AI, as well as in the Assessment List on Trustworthy AI (ALTAI) released in 2020. That said, trustworthiness will have to be interpreted and implemented in different ways depending on the context and the use case – for example, red-teaming approaches or risks related to fundamental rights will differ across use cases and mitigating measures, including forms of meaningful human oversight or criteria for explaining and tracing decisions, which will also substantially differ. Trustworthiness therefore requires a much higher level of granularity in the guidance provided to research, industry and government players compared to what is available today. Providing such guidance would also very effectively support the AI Act's implementation phase.

Fourth, the reference to **science, industry and public services** encompasses the three domains of application in the large-space initiative. The section immediately below elaborates further on why these three domains are particularly salient for EU competitiveness and excellence in AI.

And *finally*, the timeframe (2035) corresponds to the end of the next Multiannual Financial Framework (MFF). The 'how' section below discusses this aspect more in detail.

The 'why' – laying the justifications for a large-scale initiative on AI

When preparing a large-scale initiative such as the one being considered in this In-Depth Analysis paper, demonstrating the potential added value is crucial. This means answering two main questions – first, whether the initiative's proposed activities are likely to add value to European society and the economy, and help the EU pursue its overarching policy goals in a way that justifies the proposed budget. And second, whether the initiative as proposed is better than any viable alternative.

In other words, an *ex-ante* counterfactual analysis perspective is needed to identify the direct and indirect costs and benefits of each alternative investment option. Here are some of the questions that would need clear answers in the weeks to come:

- How would European society and the market evolve in the absence of *any* initiative?
- What are possible alternative intervention pathways to the proposed large-scale initiative?
- What direct costs would the initiative impose on players in the research, science, industry and public services spaces (e.g. would trustworthy AI be costlier to implement due to regulatory requirements or would players need to utilise costlier solutions due to environmental constraints)?
- What would be the possible indirect costs on society and the economy (e.g. the delayed market entry of AI solutions)?
- What would be the direct benefits of the proposed initiative in terms of scientific advancement, industrial competitiveness, good governance and the quality of public services?
- What would be the ancillary benefits in terms of spillover effects on society and the economy from R&D activities; labour and total factor productivity; complementary technological progress (e.g. in connectivity and other hardware solutions, synthetic biology, immersive technologies, etc.); as well as citizen empowerment and digital democracy? There are, of course, potentially many more to consider.
- What are the benefits in terms of strategic autonomy, reducing dependency on other world powers, economic security and overall European competitiveness?
- What are the risks associated with executing the initiative? How are they potentially accounted for and mitigated?

The 'how' – choosing a modus operandi

How would such a large-scale initiative operate in practice? There are many angles to this question.

First, **the initiative's governance would need to be inclusive,** i.e. include all relevant stakeholders. Broad representation, at the same time, should not lead transaction costs to skyrocket. This in turn means choosing an agile leadership approach with important managerial discretion left to well-identified and highly entrepreneurial individual(s) who can execute a mandate by taking calculated risk.

It also implies adopting a portfolio approach to ensure that the innovation frontier is sufficiently pushed, so that the needed breakthroughs are identified, and different plausible research avenues are experimented with, thus increasing the likelihood that the objectives are met. The separation of governance and leadership by institutionalising the moonshot (e.g. in the form of a Joint Undertaking or similar, see below) could be one way to reduce transaction costs, especially if the leadership is separate from existing initiatives and includes existing projects, networks and associations in the governance and implementation phase of the work programme.

Second, and relatedly, the large-scale initiative may require a **degree of experimental governance**, including a coordinated approach to sandboxes, pilots and testbeds, which may in turn require a distributed set of locations, research groups and institutions that can test solutions throughout the whole of Europe.

Third, **not all elements of the proposed initiative should necessarily start at the same time**. Depending on the logical and scientific dependencies between deliverables, there may be a need to achieve specific objectives before the initiative's other streams can fully take shape. For example, for trustworthy solutions to be deployed in public services, a secure and interoperable digital public infrastructure may need to be deployed and secured first. The level of trustworthiness needed for public services in key domains such as health may depend on whether machine learning and reasoning have been adequately incorporated – and tested – in the relevant AI systems.

Fourth, **launching a large-scale initiative may require some form of institutionalisation**, be that in the form of a Joint Undertaking, a Mission, a Flagship initiative or a completely new type of body with the financial and functional autonomy required to take decisions and implement them across a wide range of possible actions and domains.

Such an institution would also be linked, in one form or another, to the EU AI Office, whose mandate includes advancing actions and policies to reap the societal and economic benefits of AI across the EU, as well as providing advice on best practices and

enabling ready access to AI sandboxes, real-world testing and other European support structures for AI uptake. It will also encourage innovative ecosystems of trustworthy AI to enhance the EU's competitiveness and economic growth, and aid the Commission in leveraging the use of transformative AI tools and reinforcing AI literacy.

The proposed institution could also be linked to the GenAI4EU initiative, which aims to support start-ups and SMEs in developing trustworthy AI that complies with EU values and rules (yes, as already explained, GenAI is unlikely to lead to sufficiently trustworthy AI solutions). Both initiatives are expected to 'contribute to the development of novel use cases and emerging applications in Europe's 14 industrial ecosystems', which does include robotics, healthcare and the public sector.

Fifth, the governance of the large-scale initiative will have to be distributed, as is the case of CERN and the proposed EDIRAS, to promote place-based R&I and at the same time provide for central management and coordination, plus the connection with testing facilities and *loci* for experimentation. Verhulst (2024) welcomed this approach and advocates going towards polycentric governance due to the latter's ability to address three pain points of the AI ecosystem – access to computational resources, access to high quality data, and access to purposeful modelling. This would especially boost AI in Europe if coupled with another breakthrough that the EU has been seeking to reach over the past years: adequate data governance and stewardship, coupled with a social license for data use, allowing for powerful data collaboratives to emerge and for meaningful, socially relevant innovation to flourish.

It would also boost a more open approach to science, which would benefit the global scientific community, make the EU a trailblazer in democratising access to science and resources, and at the same time prevent the ongoing process of privatisation and the concentration of data, compute infrastructure and talent into the hands of a few tech giants.

Finally, the governance and operation of the proposed large-scale initiative will also have to be open. This means being open to possible experimentation and testing use cases from other industrial ecosystems, as well as collaboration with other large-scale research endeavours, from the Brain project to CERN. It would also mean interacting with non-European research partners willing to contribute to select R&I projects on trustworthy AI and leveraging private sector resources under precise conditions, which would safeguard the public good nature of research results.

The 'when' – linking the moonshot to the Multiannual Financial Framework

Choosing the most appropriate timeframe for achieving a large-scale initiative's objectives is very important. On the one hand, scheduling results too far ahead may induce a degree of relaxation in those who are managing the initiative or even excessive spending into parallel streams of research, even when they do not demonstrate sufficient potential or added value. On the other hand, setting a shorter timeline to achieve the main deliverables may not be sufficient for making the required R&I breakthroughs, an essential component of any challenge-driven research and innovation endeavour.

In all this, choosing the most appropriate timeframe may also depend on the specific characteristics of the AI value chain, which evolves faster than many other technologies, and on the extent to which important breakthroughs should be achieved in fundamental research, which would then need to be applied in downstream domains. Finally, it is important to consider the institutional context in which the initiative is launched and the funding institutions' budgeting period.

That said, despite the obvious appeal of setting an 'end of decade' target for a mission-oriented R&I plan, **it is unlikely that a very ambitious initiative could achieve its ultimate objective by 2030**. Other considerations include the fact that the EU institutions are currently in transition after the European elections of 6-9 June, with a new Commission not expected to be in place before November 2024, that Horizon Europe is due to continue with a fairly well-defined Work Programme until 2027 and that the proposed initiative will entail a degree of institution-building, as discussed above.

Accordingly, if some of the foundational research streams and the strengthening of framework conditions could already start in 2025, **it seems more reasonable to plan for the initiative's main activities to unfold and produce results by the end of the next MFF,** thus at the end of 2034.

In terms of overall funding, the price tag of such an initiative seems to be substantial, in the order of magnitude of EUR 100-120 billion over seven years. That said, it is important to recall that: (i) the **size of the investment must be assessed against the likely return**, also in terms of spillover effects and of the opportunity cost of 'doing nothing'; (ii) the **funding would not necessarily all have to come from the EU budget**, as Member State funding (e.g. as in the case of the European Molecular Biology Laboratory) and contributions by private foundations could also usefully support the initiative; and (iii) the size of the annual investment planned in the US and China is unfortunately much larger[2].

[2] In the US, the need to keep up with China's public investment led a bipartisan group of senators led by Chuck Schumer to present a proposal to earmark USD 32 billion annually for non-defense AI uses (the

CONCLUSIONS – WHY THE TIME TO ACT IS NOW

Europe needs to act now if it wants to seize the opportunity that the change in political leadership and the definition of a new agenda for science, research and industrial policy provide. Once priorities are set for the new Commission, emphasis on other priorities – including defence and security – may deprive the AI domain of the resources and momentum it needs to trigger the needed boost in Europe's science, industry and services of general economic interest. Given the long-term impact of this investment on several EU goals, including technological sovereignty, strategic autonomy, competitiveness, sustainable development and democracy, as well as Europe's widening gap with (and enduring dependence on) countries with diverging priorities, a bold signal in this direction is needed – and urgently.

Acting now is also important in view of the ongoing debate on the structure, budget, goal and *modus operandi* of the next Framework Programme for Research and Innovation (FP10). Stakeholders have already expressed several positions (a process that CEPS is following very closely) and the upcoming report on the interim evaluation of Horizon Europe by the so-called Heitor group will provide further momentum to the discussion.

Given these circumstances, the window of opportunity for promoting a comprehensive and consistent programme will soon close – and certainly by the beginning of 2025. This In-Depth Analysis, serving as a scoping paper, is meant to accelerate the discussion by going deeper into the details of a possible future large-scale initiative that:

1) Is of a scale and ambition that can put Europe 'back on the map' of AI-related R&I.
2) Aims at measurable and deadline-bound objectives, which can be monitored over time and are verifiable at the time of completion.
3) Caters to Europe's needs in research, as well as in key innovation domains.
4) Covers the whole life-cycle of AI from lab to market.
5) Seeks significant progress on framework conditions for market uptake, such as the availability of adequate skills and compute infrastructure, as well as data governance and stewardship.
6) Foresees the development of a world-class regulatory framework, which includes space for experimentation, adaptive and agile governance tools, and operational guidance on compliance for each identified use case.
7) Leverages Europe's existing excellence and strength, and aligns with Europe's values and principles by fostering trustworthy AI.

---

original spending proposed by the National Security Commission on AI, NSCAI). One of the reasons that underpinned this proposal is China's foreseen investment of USD 50 billion.

8) Enables portfolio management and sufficient budgetary discretion to ensure both intermediate and final objectives are achieved.
9) Enables substantial spillover effects towards most – if not all – industrial ecosystems.
10) Projects Europe's approach to AI and related technologies outside the EU's borders, thus increasing Europe's actorness as a global technology leader, including in tech diplomacy (e.g. global AI governance) and development policy.

CEPS
PLACE DU CONGRES 1
B-1000 BRUSSELS