EU Economic Blueprint

OpenAI

Introduction

OpenAI's mission is to ensure that artificial intelligence benefits everyone. To us, that means building AI that helps people solve hard problems, because by helping with the hard problems, AI can benefit the most people possible – through better healthcare and education, more scientific discoveries, better public services, new forms of creative expression, and increased economic growth and productivity.

We're off to a strong start, creating freely available intelligence being used by half a billion people around the world on a weekly basis to ideate, discover, and innovate beyond what we're currently capable of doing on our own.

Across the European Union, our tools are driving progress. Existing partnerships are helping to:

Speed up development of life-saving medical treatments

Helping researchers design new drugs faster than before, including through our partnership with Sanofi.

Accelerate scientific research

At leading European laboratories and universities like Science Po in France and the Max Planck Society in Germany.

Improve education

Including bringing AI to schools across Estonia and enabling teachers at ESCP Business School to create personalized lesson plans, spend more time with students, and safely adopt AI in the classroom to improve learning outcomes.

Improve AI access for underserved communities

By increasing accessibility through Danish start-up Be My Eyes, and improving economic and social inclusion with Simplon in France for vulnerable populations.

Unlock new forms of creativity for artists

By increasing accessibility of artistic production, enabling creators to develop content in innovative ways, like immersive experiences at Versailles and content translations with Spotify.

Powering a new generation of start ups and entrepreneurs

Such as Synthflow, Pigment and Brainly that are using OpenAI's technology to compete globally with AI-driven services. Outside of the US, Germany is the country with the largest number of API developers building on OpenAI's technology.

We believe the EU needs to act more boldly and decisively to maximize AI's possibilities while also ensuring it's used responsibly to mitigate potential negative effects. The opportunity AI presents to spur productivity and revive the EU's economic competitiveness is too compelling to forfeit. At the same time, AI is too powerful a technology to be led and shaped by autocratic governments who don't share our common values – yet that is a growing risk if democracies don't move with speed and determination, both to build AI on democratic principles and to demonstrate its benefits.

Prosperity, broadly shared, is a cornerstone and founding principle of European integration and progress. Shared prosperity from AI is as near and measurable as the new jobs and growth that will come from building the needed infrastructure for this digital age. Soon, as our CEO Sam Altman has written, AI will help our children do things we can't. Not far off is a future in which everyone's life can be better than anyone's life is today.

With such opportunities in sight, we want to work with the EU policy community to ensure that AI's benefits are shared responsibly and equitably across the Union. This Blueprint champions the entrepreneurial spirit, scientific excellence, and individual freedoms at the heart of Europe, and reflects the EU's commitment to social inclusion and fundamental rights. It has been shaped by the conviction that Europe can best guard and prosper from its values through accelerating technological progress.

Just as forward-looking EU leaders once chose cooperation and innovation to build the single market, now they must foster an environment where AI can thrive responsibly and create real incentives for Europeans to adapt to and adopt AI. The EU and its member states – working together with academia, civil society and industry – can pave the road for AI built by Europe, in Europe, for Europe.

We want to be a thought partner to the EU on this journey and are sharing four principles that we believe will be vital to the EU's success:

01

Establish and grow the foundations needed for sustained AI growth - chips, data, energy and talent.

02

Ensure that EU rules are streamlined and work in sync to enable AI progress rather than hinder it.

03

Maximize the AI opportunity through widespread adoption across all sectors, regions and society

04

Ensure that AI is built responsibly and reflects European values.

The EU's strength is its scale - concerted effort across these four principles with balanced, harmonized rules can spur competition and innovation, help native AI companies start and scale, and ensure that the benefits of AI reach all Europeans across the 27 member states.

A thread between these principles is the need for the EU to simplify its overall approach and ensure its regulations and initiatives are working in sync - creating a flywheel for AI growth rather than nurturing adoption and innovation in some areas while hindering it in others.

We call for EU policymakers and regulators to reset and rethink their current approach to regulating AI in order to review the accumulated rules that former European Central Bank President Mario Draghi has called a drag on European competitiveness. We believe policymakers should evaluate which rules are boosting the EU's AI sector and should be kept in place, and which rules are holding it back and should be revised.

We commend the intent to bring forward a Digital Simplification Package as part of the broader simplification agenda and urge policymakers to consider how it can best serve the AI sector. We also believe that the forthcoming AI Continent and Apply AI strategies will be decisive and should gather strong support across the political spectrum and maintain a rigorous focus on increasing AI adoption and AI-driven opportunities.

This Blueprint – a living document we will update as we learn more from our partnership – is intended to sync with what's working for the EU's AI sector and encourage AI rulemaking that ignites productivity and growth.

We are engaging with regulators and leaders across the EU to share our technological advances and how they can drive economic growth and meet local expectations. It's also why we continue to expand our presence through our offices in Dublin, Paris, Brussels and soon, in Munich, and why we're partnering with local institutions to ensure that we bring AI's economic benefits to Europeans from all walks of life.

We look forward to building on this foundation for collaboration and advancing democratic AI with you.

Chris Lehane
Chief Global Affairs Officer

Sandro Gianella
Head of EMEA Public Policy & Partnerships

The Economics of AI Progress

Our work at OpenAI suggests that as AI advances, progress accelerates and becomes increasingly affordable:

01 **Investing more in AI will continue to make it more capable and beneficial.** The intelligence of an AI model roughly equals the log of the resources used to train and run it. Until recently, scaling progress has primarily come from training compute and data, but we have shown how to make intelligence scale from inference compute, as well. The scaling laws that predict these gains are incredibly precise over many orders of magnitude. It follows that further investment will lead to further gains, and further benefits to society: We believe that the socioeconomic value of linearly increasing intelligence is super-exponential in nature.

02 **The cost to use a given level of AI capability falls by about 10x every 12 months, and lower prices lead to much more use.** We saw this in the change in token cost between GPT-4 in early 2023 and GPT-4o in mid-2024, where the price per token dropped about 150x in that time period. Moore's Law predicted that the number of transistors on a microchip would double roughly every two years; the decrease in the cost of using AI is even more dramatic.

03 **The amount of time it takes to improve an AI model keeps decreasing.** Put another way, AI models are catching up with human intelligence at an increasing rate. The typical time it takes for a computer to beat humans at a given benchmark has fallen from 20 years after the benchmark was introduced, to five years, and now to one to two years—and we see no reason why those advances will stop in the near future.

Developing Responsible AI

At OpenAI, we strive to ensure that our AI technologies align with human values and operate securely. Our commitment to transparency, accountability, and thorough testing helps mitigate risks and maximize societal benefits, setting a high standard for safe AI development.

01 **OpenAI's Preparedness Framework is the safety protocol** guiding how we manage the largest-scale risks when developing and deploying our increasingly advanced AI models. It helps us measure and protect against the risk of severe harm from dual-use frontier capabilities. We are constantly refining our practices and advancing the science underpinning these developments, to unlock the benefits of these technologies while addressing their risks.

02 **OpenAI's Model Spec details our approach to shaping the behavior of AI models to ensure they are useful, safe, and aligned with both user and developer needs.** The Spec provides concrete guidelines for mitigating risks like misaligned goals, execution errors, and harmful instructions, ensuring our models ask clarifying questions when needed and follow a structured hierarchy in decision-making. By publishing this document under the Creative Commons CC0 1.0 deed, we invite public discussion and collaboration, reinforcing our commitment to transparency and continuous improvement in AI safety and alignment.

03 **At OpenAI, we believe in Privacy for the Intelligence Age—a commitment to safeguarding personal data through implementing robust data minimization measures, meaningful transparency, and giving users control over their data.** We pioneered and embedded state-of-the-art privacy safeguards throughout the entire AI lifecycle—from pre-training to deployment—ensuring that privacy is a foundational element of how we build and deploy our technologies.

First Principle

Establish and grow the foundations needed for sustained AI growth - chips, data, energy and talent.

For the EU, the AI era is an unmissable opportunity to drive growth and bolster its strategic capabilities. In the AI era, infrastructure is destiny. Successful nations will turn resources into competitive advantages – and *in this AI age, the critical resources are compute, data, energy, and talent*. Marshalling these in the EU through bold investment will create widespread economic opportunity, and thoughtful harmonized rules will serve as a complement, mitigating risks to citizens and enabling the region to capitalize on its innovative potential.

Critical Resources

We believe the EU should seize the moment and make bold investments so that it has the physical and human infrastructure to support AI at scale. In concrete terms, Europe must be able to produce and access world-class computing power, abundant clean energy, vast and diverse datasets, and a skilled workforce. In the past, Europe's economic strength was built on infrastructure like highways, factories, and universities; in the AI future, it will be about supercomputers, fiber optic networks, power grids, and research centers. If Europe lags in any of these areas, it could find itself dependent on external providers for critical capabilities – something the EU has recognized as a strategic risk.

Compute

Abundant and affordable compute will be the engine of AI development in the EU. The good news is that steps are already being taken – including recent plans to mobilize €200 billion for infrastructure across the EU and €109 billion in investment in France. These come on top of existing EuroHPC and AI Factory initiatives, which are laudable steps to build some of the world's fastest supercomputers and make them available to researchers and startups. Europe's challenge is now to accelerate and expand these initiatives. To complement AI factories initiatives, we recommend an AI Compute Scaling Plan that boosts EU computing capacity by at least 300% by 2030, prioritizing low-latency, geographically spread inference infrastructure.

Access to chips is vital to drive down the cost of compute and help make it abundant in Europe. Our recent submission to the White House Office of Science and Technology for the upcoming US AI Action Plan, recommended that countries that commit to democratic AI principles by deploying AI systems in ways that promote more freedoms for their citizens could be considered Tier I countries. Tier II status should be limited to only those countries that have a history of failing to prevent export-controlled chips and other US-developed IP from being diverted into, or used by Tier III countries.

Data

In AI, data is the raw material that fuels AI model training. The EU has vast amounts of data (from scientific research to industrial data to public sector information), but much of it remains siloed, or underutilized. To innovate in AI, developers need access to rich, diverse, and relevant datasets – while respecting privacy, data protection and security. The EU should assure AI developers and researchers that their models will have the ability to learn from publicly available sources and that more data will be made available in responsible ways. This would also attract infrastructure investment because the more the EU is able to reduce barriers to data access, the more it will be able to convince investors and companies to place their multi-billion euro data infrastructures in its jurisdiction.

Energy

For the EU to support widespread AI, it must ensure sufficient, sustainable energy supply. This aligns well with Europe's green objectives: investments in clean energy (solar, wind, nuclear, and emerging sources like fusion) serve the dual purpose of fighting climate change and powering digital growth. We recommend dramatically increasing European capacity for renewable energy generation and modernizing power grids. Streamlining permitting for new energy projects and cross-border transmission lines (part of the EU's existing efforts under the Trans-European Networks for Energy) will help meet AI's electricity needs in a climate-friendly way.

Talent

Even with the best hardware and data, AI leadership ultimately depends on people – researchers, engineers, domain experts, and informed users - and Europe must focus on retaining and cultivating its considerable human capital. To do this, the EU should significantly expand investments in STEM education, specialized AI scholarships akin to an "AI Erasmus," and networks of AI research excellence centers across the continent.

Equally crucial is large-scale workforce reskilling through apprenticeships and vocational programs, enabling current workers to adapt their skills for AI-driven roles. Broad and equitable access to AI education from primary school onward, similar to Estonia's successful integration of AI in curricula, can prepare Europeans for an increasingly AI-centric world and reinforce Europe's digital values.

Additionally, Europe should attract global AI talent by improving policies like the EU Blue Card and investing in interdisciplinary expertise to ensure responsible and ethical AI deployment.

Our policy proposals

- **AI Compute Scaling Plan:** Increase the EU's computing capacity by at least 300% by 2030, with a clear focus on low-latency, geographically-distributed infrastructure optimized for inference, as a complement to the AI factories focused on training.
- **Green AI Grid:** Fast-track renewable energy projects, prioritizing streamlined approvals for data centers, to deliver carbon-neutral AI infrastructure EU-wide by 2030.
- **EU AI Data Spaces:** Implement sector-specific AI data spaces (health, industrial, environmental, public) under the EU Data Strategy by 2027 to foster trusted data sharing; advancing AI through data sets with privacy-preserving safeguards in place, foster transparency, and accelerate safe innovation.
- **Public Data Availability:** Make the availability of online public data a top priority to fuel AI innovation and model development that is reflective of Europe's diverse languages and cultures.

Second Principle

Ensure that EU rules are streamlined and work in sync to enable AI progress rather than hinder it.

AI is a rapidly evolving technology poised to become integral across society, requiring thoughtful regulation to mitigate risks without stifling innovation. Unfortunately, today Europe's regulatory environment is often too complex – as Mario Draghi noted in his report, the EU now has around 100 tech-related laws and over 270 bodies regulating digital sectors across member states. EU policymakers should focus on reducing barriers within the single market by simplifying the digital rulebook and eliminating redundant or obsolete proposals. The ongoing Omnibus Simplification Package and broader simplification agenda promoted by the European Commission represent critical opportunities to streamline regulations and strengthen EU competitiveness.

AI Act

Any discussion of AI regulation in Europe has to address the AI Act, the world's first comprehensive framework for AI. OpenAI supports the AI Act's core objective – to ensure AI systems are safe – and we have engaged constructively with EU lawmakers throughout its drafting and implementation. It is crucial that this Act strikes the right balance: protecting citizens and giving certainty to businesses, while not unduly burdening innovative developers or pushing research out of Europe.

At OpenAI, we've already taken steps to adopt robust safety measures in our operations. We've been one of the first AI companies to design a detailed framework that guides the development and deployment of our models, including rigorous testing and the publication of safety documentation for each model. We actively engage with global and European initiatives on AI governance – from working with law enforcement on misuse cases to participating in multistakeholder forums on AI ethics. We're also developing tools to help users and third parties identify AI-generated content (through watermarking techniques and other research), contributing to the broader transparency goal.

With an approach that strikes a balance between common-sense guards against risks, and enabling and incentivising innovation, Europe can both protect its citizens and societies while also seizing what AI offers.

Simplification

A house divided against itself cannot stand, and nor can a strategy that nurtures AI growth in some areas while hindering adoption in others.

The Draghi report was right that the sheer breadth and quantity of EU regulations hamper innovation, slow economic growth, and pose an "existential challenge" to the European Union's future. They close the door on the EU's ability to build AI that is of, by, and for Europe.

Now, as policymakers take a step back to evaluate which rules strengthen the EU's AI sector and should be preserved, versus which ones are holding it back and should not—they should make sure they're considering the full picture. This is a one-way door decision that will have huge consequences for Europe's future progress and prosperity. If policymakers get it right, it will open the door to building AI for Europe—and unlocking the progress and prosperity gains the technology offers. If they don't, Europe risks falling further behind in the global competition over AI's future.

Harmonization

The EU must be vigilant against the risks posed by regulatory fragmentation. EU rules should apply clearly and consistently across the single market. Europe's abundant talent and innovative spirit remain constrained by national barriers, limiting the seamless, cross-border growth that is essential for global competitiveness. The EU's strength is its scale: one set of balanced rules for 27 countries will spur competition and innovation, help native AI companies start and scale, and ensure that the benefits of AI reach all Europeans.

For example, a pan-European entity for startups is an area that would benefit from a unified regulatory vision. As underscored by Ursula von der Leyen, while startups in the US benefit from a cohesive market, European companies struggle with fragmented ecosystems. Establishing a unified legal framework—as put forward by EU Inc—will empower startups to secure capital more effectively and realize their potential across and beyond Europe.

Our policy proposals

- **AI Accelerator Fund:** Launch a dedicated €1 billion fund to rapidly finance pilot projects demonstrating clear societal or economic value of AI.
- **A European AI Readiness Index:** To encourage national action, publish a yearly ranking of EU countries based on a methodology assessing AI adoption, skills, infrastructure, and regulation readiness.
- **National AI Readiness Officers:** By 2027, each EU member state should appoint a dedicated "AI Readiness Officer," responsible for accelerating national AI adoption and creating synergies among member states.
- **Pan-European Start-up Entity:** By 2026, launch a unified legal framework that harmonizes corporate law, taxation, and compliance across the EU, enabling startups to operate seamlessly.

Third Principle

Maximize the AI opportunity through widespread adoption across all sectors, regions and society.

As AI models and systems become more capable and versatile, they will increasingly bring prosperity and advantages to people and organizations. Yet translating the potential of AI into concrete economic and social benefits will not happen automatically. These benefits will not materialize without adequate AI adoption at the organizational level and AI literacy at the individual level.

To us, this means that beyond the foundational investments needed in research and in infrastructure, we believe that governments and institutions in the EU, both nationally and locally, have a central role to play in incentivizing adoption in the public and private sector. Moreover, preparing, educating and upskilling individuals would have a high impact on the development of AI throughout the continent.

Europe has enormous assets to build on, including world-class universities and research labs, a highly skilled workforce, and leading industries in sectors like manufacturing, automotive, pharmaceuticals, and finance. By infusing AI into these sectors, Europe can supercharge productivity and create new high-value jobs.

Private Sector

European companies must leverage AI to advance their industries. By encouraging such partnerships and innovation, the EU can ensure that its single market of over 450 million people becomes one of the world's most attractive arenas for AI-driven business.

Every European sector – from agriculture and tourism, to energy and creative industries – stands to gain from AI adoption, whether through cost savings, better services, or entirely new and innovative AI-based offerings.

- **In healthcare,** OpenAI is partnering with Sanofi to use AI for accelerating clinical trials and drug discovery, helping life-saving treatments reach patients faster.

- **In telecommunications,** Orange is fine-tuning our open-source speech model to better understand African regional languages, illustrating how European companies can localize and build on AI to serve global markets. Meanwhile, T-Mobile is working with OpenAI to launch IntentCX, a platform that uses generative AI to transform how businesses handle customer interactions—improving service quality, reducing response times, and empowering support teams with smarter tools.

- **In manufacturing and automotive,** European firms are exploring AI to boost efficiency and create smarter products. Mercedes-Benz is integrating ChatGPT into its in-car assistant, aiming to improve the driver experience. More broadly, industrial players across Germany, Italy, and beyond are using AI for predictive maintenance, supply chain optimization, and design – keeping Europe's industrial base competitive.

- **In creative industries,** Spotify is partnering with OpenAI to integrate advanced language models into its platform, enhancing the music streaming experience through personalized recommendations, dynamic playlist generation, and improved search functionalities.

- **In financial services,** BBVA is leveraging ChatGPT Enterprise to give its 125,000 employees AI tools that enhance productivity and creativity across departments – from credit risk and legal services to customer experience and marketing.

These examples demonstrate flagship European brands operating at the cutting edge of AI adoption and rollout.

SMEs & startups

European small and medium-sized enterprises (SMEs) should be encouraged to adopt AI tools to become more efficient and creative. There is already a thriving European start-up ecosystem with a new generation of AI-fueled companies like Synthflow and Pigment using OpenAI's technology to enhance their platforms and ensure that European companies compete globally.

However, the AI adoption imperative isn't just about the most agile companies – it's also about how AI can boost the broader economy. Despite the EU's ambition to become a global leader in AI, adoption rates remain suboptimal.

According to the latest Eurostat statistics from 2024, only 13,5% of European companies are using AI technology. This figure starkly contrasts with the EU's Digital Decade target of achieving 75% AI adoption among businesses by 2030. At a deeper level, a divide exists between larger organizations that demonstrate higher adoption rate (41%) while SMEs lag behind (11%).

Targeted policy interventions are needed to close the adoption gap.

Public sector

Public sector spending accounts for roughly half of EU GDP, with public institutions managing vast amounts of data and resources. As such, governments are uniquely positioned to lead digital transformation in critical areas such as infrastructure, healthcare, energy, transportation, and a broad range of public services.

Currently, AI adoption in the public sector lags behind the private sector. A European Commission study highlighted that complex procurement processes and concerns about bias in AI decision-making are significant barriers to wider AI implementation in Europe's public institutions. Addressing these issues and encouraging broader AI adoption would deliver more effective and efficient public services, saving taxpayers' money and freeing resources for increased investment in essential areas.

To accelerate this transformation, procurement rules embedding specific technological requirements should be revised if they hinder AI adoption. Instead, these rules should actively encourage the integration of AI solutions. Governments should also prioritize suppliers and vendors who utilize AI tools effectively.

In sectors where the public sector is a primary or significant service provider, governments should remove unnecessary restrictions that prevent sensible and responsible AI use. For example, enabling doctors in public hospitals to leverage AI could improve patient care and extend healthcare access.

Finally, governments should be transparent about AI applications within public services. Clear communication builds trust among citizens and demonstrates the tangible benefits of AI-driven public services.

Education

To date, ChatGPT has become a go-to tool for students globally to personalize their education and advance their personal development. Most ChatGPT users—nearly four in five—are under the age of 35 and the majority of conversations are focused on learning and schoolwork.

By embedding AI literacy and usage in curricula, ensuring equitable access to AI tools, investing in digital infrastructure and providing professional development for educators, we can ensure students will be better equipped as the workforce of the future.

In Estonia, we are proud to work with the Government to provide secure, equitable access to AI across the secondary school system, focused on supporting and empowering teachers.

This is part of our effort to support educational institutions across the EU including ESCP France, IE Madrid, University of Pisa, Gothenburg and WHU Germany that have taken steps to make AI as fundamental to their campus as using the internet.

Educational institutions and governments have a core role in ensuring that all students and teachers have access to AI and receive the necessary training to benefit from its potential. Estonia has fully embraced this challenge and will offer vital learnings to other EU member states.

Society

European governments and industry should work together to scale AI literacy so that users understand these tools and make informed choices. This includes investing in public awareness campaigns and professional training. On a citizen level, libraries and community centers could offer AI familiarization workshops – similar to past efforts on basic digital skills. Such initiatives ensure that all Europeans, not just tech experts, have the chance to benefit from AI and use it on their own terms.

At OpenAI, we recently announced the next phase of the OpenAI Academy with a publicly available, free online resource hub to support AI literacy and help people from all backgrounds access tools, best practices, and peer insights to use AI more effectively and responsibly.

Our policy proposals

- **Introduce fiscal incentives:** Tax credits or subsidies for AI-related investments can help de-risk innovation and make adoption more accessible, particularly for SMEs and local authorities with limited resources.

- **Encourage public-private partnerships:** Governments should bring together industry, research institutions, and civil society to develop tailored AI solutions and support deployment at scale, like the initiative in Estonia.

- **Build a mechanism to share best practices across member states:** especially in sectors like manufacturing, health, or public administration to accelerate learning and diffusion.

- **100 Million AI Citizens:** Train 100 million Europeans in foundational AI skills by 2030 through freely accessible online courses in all official EU languages.

- **AI Literacy Ambassadors:** Mobilize 10,000 European "AI Literacy Ambassadors" by 2030—educators, influencers, and entrepreneurs—to locally promote and teach AI awareness.

- **Build an AI-native education system:** Ensure equitable access to AI in education by embedding AI literacy in secondary and post-secondary curricula, expanding tool access, upgrading digital infrastructure, and training educators.

Fourth Principle

Ensure that AI is built responsibly and reflects European values.

The EU and the AI industry must work hand-in-hand to ensure that, as AI tools increasingly become part of Europeans' daily lives, users can confidently trust that potential harms are being considered and mitigated. Equally important is empowering individuals to meaningfully control and personalize these tools according to their specific needs and preferences. Building this trust will be crucial to ensuring AI is embraced widely and benefits everyone in society.

Protecting young people

Young people deserve special consideration in the AI regulatory agenda. They are avid users of new technologies and will inherit the AI-driven world we create. The EU has long recognized the need to protect children and teens online, and AI should be no exception. At OpenAI, we've made youth safety a priority, establishing robust safeguards in our systems and working with trusted organizations dedicated to protecting young people. We believe AI companies, together with policymakers, should:

- **Prevent exploitation and harmful content.** It is critical to guard against AI being used to generate child sexual abuse material or other exploitative content. We encourage policies and industry standards that prevent the creation or distribution of AI-generated Child Sexual Abuse Material (CSAM). This includes building AI systems with safeguards (e.g., filters, detectors, classifiers) that block attempts to produce such content. OpenAI incorporates protections throughout the AI development life cycle to mitigate misuse. We are collaborating with organizations and law enforcement to identify and block harmful outputs. The EU's proposed legislation to combat child sexual abuse online can intersect with AI governance, ensuring AI tools do not become a new vector for abuse.

- **Promote safe use of AI.** Promote safe use of AI among younger people: AI systems should prioritize the best interests of young people by upholding their rights and supporting their cognitive and socio-emotional development at every stage. This means designing built-in safeguards, including regularly updated content moderation systems that respond to emerging risks and encourage healthy engagement. Transparent practices and appropriate human oversight are essential to ensure AI aligns with young people's well-being.

- **Promote AI literacy.** Governments and industry should work together to integrate AI education into school curricula, equipping young people with the skills to understand how AI works, its limitations, and how to use it safely and critically. These efforts can build on the EU's Digital Education Action Plan and successful models like our partnership bringing AI to Estonian schools and our support for the Beneficial AI for Children Coalition.

- **Coordinate with child protection stakeholders.** Companies like OpenAI should continue to form partnerships with safety organizations globally and in Europe (such as EUROPOL's Cyber Crime Centre or national centers for missing and exploited children), sharing information on emerging threats and working on better reporting and prevention mechanisms. A cooperative approach between AI providers and law enforcement, under clear legal frameworks, will help anticipate and combat new risks to young people that AI might pose – from deepfake abuse imagery to harmful interactions.

With youth safety a core component of AI governance, Europe can ensure the next generation grows up with AI systems that support their well-being and development. European parents and educators should have confidence that AI tools – whether used for education, entertainment, or social connection – come with necessary protections in place. This is both a moral imperative and key to public trust: if families and educators feel AI is safe for children and teens, they will be more open to its adoption in society at large.

Choice and empowerment

AI should be a tool for people. That means users – whether individual consumers, developers, or businesses – should have meaningful control over how they use AI. We think the EU rules should reinforce user agency by allowing people to customize AI tools to their needs.

For example, a user of an AI assistant should have settings to determine the assistant's style or constraints (within safe bounds). Developers building on AI APIs should similarly have freedom to fine-tune models for their use-cases (again, provided safety standards are met).

As AI systems become more adaptive and autonomous, it's essential that regulations do not inadvertently hinder beneficial personalization. On the contrary, EU policy can play a role in encouraging companies to provide users with meaningful control over how their data is used by AI systems—and how outputs are shaped in return. The ability for individuals to tailor their AI experience enhances both the utility and the societal acceptance of these tools.

OpenAI, for example, enables users to customize elements of ChatGPT's within safe and transparent boundaries, and to decide whether their data is used for model training. Putting users in control is not only a product design choice—it's a foundational principle for trustworthy AI.

In essence, a user-empowered approach means AI will be created alongside European citizens, not imposed on them. European consumers should be able to trust that AI services serve their interests, protect their privacy, and exercise choice.

Our policy proposals

- **AI Awareness Day:** Creation of an annual moment on a specific aspect of AI (fairness, AI for science, responsible use…). The theme will change each year and will drive awareness and conversations in schools and organizations.

- **A Responsible AI Innovation Prize:** Inspired by the Sakharov Prize, the European Institutions could award a prize each year to recognise outstanding, innovative work in AI and related disciplines in Europe.

- **Extend Safer Internet Day to AI:** Expand the scope of the EU-supported Safer Internet Day to cover AI, offering resources and guidance tailored to children, parents, and educators. This includes co-developing age-appropriate educational materials, safety guidelines, and awareness campaigns on the responsible and safe use of AI tools. Partner with industry, civil society, and youth organizations to ensure broad reach and relevance.

- **Youth Digital Agency in AI Initiative:** Establish an EU-funded initiative that supports the development of AI applications co-designed with young people, particularly those aged 13–17. This could involve participatory design sprints with schools, youth councils, and child safety organizations to build features that reflect their needs and values.

- **Voluntary EU Code on Child-Centric AI Design and Use:** to establish a collaborative, forward-looking framework for AI developers, platforms, civil society, and policymakers to co-develop, adopt and adapt common standards that protect and empower children in the AI ecosystem—while enabling responsible innovation.

Conclusion

The EU's competitiveness and growth in the AI age depend on embracing innovation and enabling it at scale. By believing in its innovators and giving them the tools and freedom to succeed – with sensible guardrails – the EU can lead in this new frontier. The payoffs will be substantial: higher productivity growth, new jobs and companies, and solutions to challenges from healthcare to climate that improve citizens' lives.

By focusing on four pillars – compute, energy, data, and talent – as well as reducing regulatory burdens, the EU can create an environment where AI innovation flourishes on European soil. The economic impacts of such investment are numerous: it will create direct jobs (e.g., constructing a new chip factory or wind farm), indirect jobs (through ecosystems of suppliers and service providers), and induce innovation-led growth (new startups, productivity boosts for companies that adopt AI, etc.). Boosting local economies through these projects also has a regional cohesion effect if done strategically. For example, AI infrastructure investment can be directed to regions that need economic revitalization – like establishing an AI research center in a city affected by industrial decline, or a data center in a rural area with available space and renewable energy potential.

Yet time is of the essence. The global demand for compute and energy for AI far outstrips current supply; and as mentioned, large pools of capital (hundreds of billions of euros globally) are searching for places to invest in AI infrastructure. Europe should act fast to channel these resources into projects that support a democratic, sustainable AI ecosystem. If Europe doesn't provide the attractive environment for this investment, through clear policies and possibly public-private partnerships, the funds will flow to projects elsewhere, potentially backed by regimes that do not share Europe's values. To put it plainly, infrastructure will define Europe's AI destiny.

We are optimistic about the EU's AI future. In just two years, our technology has become an indispensable and valued part of millions of European citizens' lives. We've opened our first office in Paris; we've based our regional headquarters in Ireland; we've set up an office in Brussels to work with EU officials; and we're now opening a Munich office. We see tremendous promise in the talent and momentum present in Europe and are eager to continue and reinforce our growing partnerships here. We are ready to do our part – whether by collaborating on research, delivering powerful new tools, or working on infrastructure needs – to ensure Europe has access to the best AI capabilities.

To reflect the rapid evolution of AI technology, this Blueprint is a living document. We will continue to update our principles as we learn from working with governments across Europe on how to thoughtfully steward AI technology for the public good. OpenAI strongly believes that with the right vision and execution, Europe can harness AI to solve the grand challenges of our time – from improving healthcare and education to enabling the green transition – and in doing so, deliver wide-spread prosperity and well-being for its citizens.

About OpenAI

OpenAI's mission is to ensure that as AI advances, it benefits everyone. We're building AI to help people solve hard problems because by helping with the hard problems, AI can benefit the most people possible—through more scientific discoveries, better healthcare and education, and improved productivity. We're off to a strong start, creating freely available intelligence being used by half a billion people around the world, including 3 million developers. We believe AI will scale human ingenuity and drive unprecedented economic growth and new freedoms that help people accomplish what we can't even imagine today.

Cover image created with 4o image generation