GOUVERNEMENT

Liberté
Égalité
Fraternité

FRENCH ARTIFICIAL INTELLIGENCE COMMISSION

AI: OUR AMBITION FOR FRANCE

MARCH 2024

SUMMARY

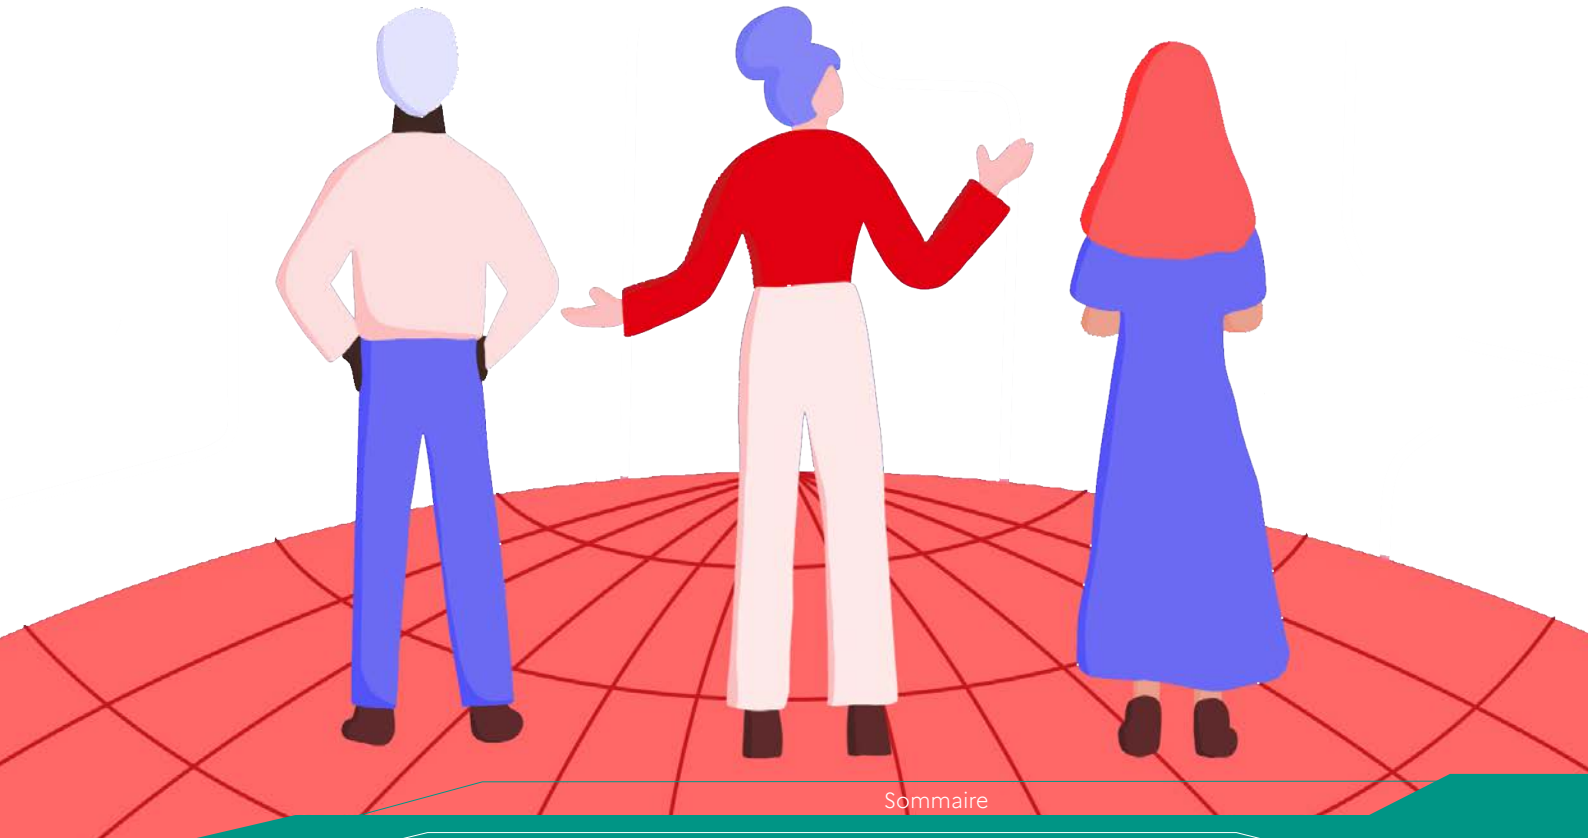

EXECUTIVE SUMMARY, OUR VISION AND KEY RECOMMENDATIONS

Artificial intelligence is an unavoidable technological revolution. The sudden emergence and spread of generative AI marks an important stage in this revolution. We can't help but notice how easy certain tools are to use, how quickly content is generated, how realistic the texts, images and sounds generated are, and more generally the abilities of recent AI models.

This technological revolution affects all areas of activity. It affects the economy, employment, public services, the environment, information, cultural industries... the list goes on. Every aspect of our society is concerned, and will be even more so in the future, given its considerable potential.

We should neither be overly pessimistic, nor overly optimistic. We don't expect mass unemployment or an automatic acceleration of growth. In the coming years, AI will not replace humans, nor will it be the solution to all the challenges of our time. We must neither overestimate the impact in the very short term, nor underestimate it in the long term.

Europe and France are well placed to play a leading role in this revolution, first and foremo st thanks to the excellence of our talent. This wealth and the exceptional dynamism of the French AI ecosystem must not, however, mask a worrying reality. For several decades, the trend has been for our continent to fall behind technologically and economically, jeopardizing its prosperity and independence.

At a time when the United States and China have made mastery of AI one of the pillars of their national strategies, we must rise to the challenge of AI, or risk losing control of our future. We need to overhaul our institutions and public policies, so that AI can play its full part in driving progress.

We propose six main lines of action:

- **Immediately launch a plan to raise awareness and train the nation:** lead ongoing public debates in everyday places on the economic and societal impacts of AI, structure the higher education training offer, bring to scale lifelong training in AI tools, integrate AI as an object and tool of social dialogue;

- **Structurally redirect savings towards innovation and create, in the short term, a €10 billion "France & AI" fund,** to finance the emergence of the AI ecosystem and the transformation of the French economic fabric;
- **Make France a major location for computing power:** secure collective supply on a national and European scale, solicit projects to set up computing centers in France with public guarantees of usage and simplification of procedures, set up an AI tax credit for the training of models;
- **Facilitate access to data: in terms of personal data,** modernize the mandate of the CNIL (the French Data Protection Authority) and its board, abolish certain prior authorization procedures for access to health data and reduce response times; in terms of cultural data, put in place a technical infrastructure to encourage the training of AI models while respecting intellectual property rights;
- **Adopt an "AI exception" principle in public research:** free researchers from administrative constraints, upgrade their remuneration, double the public research resources specialized in AI;
- **Promote global governance of AI:** create a World AI Organization to evaluate and oversee AI systems, an International AI Fund to serve the public interest, and a "1% AI" solidarity mechanism for developing countries.

A collective, massive, immediate and long-term mobilization is imperative. With this in mind, our Commission has set out to draw up an action plan that is as ambitious as it is realistic, serving people, our needs, our values, and our principles. The plan represents an annual public investment of €5 billion over five years. It includes technological investments, but also investments to catalyze at the same time the diffusion of AI in the economy, its deployment in the public interest, its adoption and training of society as a whole.

This investment is significant, but it is necessary if France is to become a leading country in artificial intelligence, and if our society is to reap its full benefits. This ambition is achievable, given France's and Europe's strengths. It is also realistic and affordable for our country: the "AI plan" we're proposing would represent 0.3% of total public spending. The cost of inaction, on the other hand, would be very high. We would forgo major economic and social gains, and risk a historic downgrade. The question set before us is this one: which spending will enable France to take control of its future?

A. ASSERTING A PRINCIPLE OF ACCOUNTABILITY SO THAT INNOVATION CONTRIBUTES TO A COLLECTIVE PROJECT

Societies everywhere are being challenged by the spread of digital technologies. Social networks are undermining political systems. Technological concentration polarizes the distribution of wealth. Algorithms contribute to inequalities in work and employment. The huge increase in uses is accompanied by a growing environmental impact. The power of certain companies limits the ability of sovereign states to act.

Artificial intelligence extends and deepens this trend. Contrary to the horror scenarios put forward by some, today's AI systems will not lead to the end of humanity. They are, however, far from infallible, and come with undesirable side-effects: reproduction of stereotypes, disclosure of confidential information, violation of intellectual property rights, and so on. They open up new possibilities for malicious acts, particularly cyber attacks and disinformation. They are a source of systemic risk, particularly in terms of the potential for technological concentration in the hands of a small number of countries, companies, or individuals.

Faced with these major challenges, we cannot repeat the mistakes of the past. Over the past two decades, France and Europe have done too little, too late, with little commitment to technological innovation and belated regulation. Today, it is up to us to take advantage of AI by putting it in its rightful place: that of a technological means at the service of an ambition for humanity, equality, solidarity, justice, prosperity, and freedom.

Since these issues concern the whole world, several international organization models have been considered in recent years. France was one of the pioneers, co-founding the Global Program for Artificial Intelligence (GPAI) in 2020. International discussion forums on AI abound, with at least fifty to date. To go further and anchor these initiatives in concrete actions, parallels have been drawn with global climate and energy issues. Our Commission notes that these parallels are insufficient: AI cannot conform to a previous model. We also believe that the international community must take advantage of the window of opportunity in 2024 to bring together many initiatives under one umbrella.

To steer AI technologies, we recommend setting up a global governance structure with a coalition of like-minded countries. Our Commission considers three major steps forward. First, a coalition of countries would set up the **World AI Organization**. This international organization would share scientific findings on the workings and effects of AI, and define binding standards for AI systems and how they should be audited. It would be democratically governed, bringing together governments, civil society (researchers, citizens, trade unions) and companies.

Secondly, France could support the creation of an International Fund for Public Interest AI, with an annual budget of €500 million. It would finance public interest initiatives: free and open AI services, independent research projects, innovations (in environment, science, health, etc.). Thirdly, France could promote the "**1% AI" solidarity mechanism**, whereby all international players investing in computing power would commit to offer 1% to developing countries.

The goals of this global governance should also be pursued at a national level. France has the opportunity to position itself as a pioneer in the evaluation of AI systems if it structures its evaluation and market surveillance network. We also need to carry out ongoing and ambitious forward-looking work on AI developments, in order to anticipate its effects on society and prepare for the necessary transformations.

B. AIMING FOR HUMANISM IN AI DEPLOYMENT

The technological revolution of artificial intelligence will have to give more power to citizens and workers. Failure to do so runs the risk of mass rejection of AI. In the past, other technological innovations have affected social cohesion. Above all, innovation only makes sense if it serves the free development of our humanity. In other words, the deployment of AI must aim to be humanistic. To achieve this, our Commission has identified three main pillars: training, social dialogue, and public services.

We recommend immediately launching a plan to raise awareness and train the nation. To achieve this, we must first create the conditions for **collective adoption of AI** and its challenges. This means holding ongoing public debates in our society, encouraging the creation of places for experimentation, public engagement with the technology ("AI cafés"), providing a digital information tool, and launching a competition for positive use cases of AI.

We also need to invest in **training for everyone, at every age**: young people in and after school, specialized and non-specialized students, employees, the self-employed and public-sector workers, and retirees. This means preparing for tomorrow's professions, in particular by structuring a range of hybrid higher education courses, such as "AI + biology" and "law + AI", or by creating AI chairs in design schools. We must also enable the use of AI in today's professions, for example by planning an AI awareness course for all civil servants.

Renewed social dialogue should be the cornerstone of AI deployment. At both national and company level, AI uses need to be built through collaboration with workers. At the same time, we need to invest in analyzing the impact of AI on the quantity and quality of jobs. AI itself can be put at the service of social dialogue, with the creation and deployment of specialized tools.

Finally, AI systems should be used to improve the quality of public services. Artificial intelligence can improve public services, by helping to personalize education, give patients more time, better support and anticipate professional transitions, and reduce bureaucracy. We can achieve these gains only if we transform our institutions. Not only do public administrations need to strengthen their infrastructure and pilot AI projects. We need to accelerate, amplify, deepen and scale AI interest in each and every public service.

C. REALIZING ITS MAJOR ECONOMIC AND SOCIAL POTENTIAL

Provided we deploy and steer AI, it should increase collective prosperity and can contribute to improving the quality of work and reducing inequalities. According to our analysis, France's annual economic growth could double thanks to the automation of certain tasks. After ten years, the increase in GDP would be between €250 and €420 billion. This would be like adding a second industrial sector. However, this increase would be temporary: once AI has been adopted by the entire economic fabric, no further productivity gains could be expected.

More than just improving productivity, AI could have a major impact on prosperity as it seems capable of accelerating innovation. This effect remains uncertain, but it is supported by many recent uses of AI: finding new proteins, identifying new materials, etc. If this effect is confirmed, it would be a remarkable feature of AI: it could induce a permanent increase in the economy's growth rate. In other words, in addition to a temporary effect linked to automation, AI could produce a longer-term effect linked to the emergence of new innovations, new products, new forms of organization, and so on.

In the coming years, AI systems will lead to the transformation of many jobs. Most workers will benefit from the automation of some boring and grueling tasks. The vast majority of jobs will change. Some tasks will be added, some will be automated. Yes, AI automation will lead to the elimination of some jobs and hasten the obsolescence of some skills. It will create a challenge for training on a sectoral or individual level. However, at a national level and despite the uncertainties, our Commission believes that the effects of AI will be generally favorable to employment: AI could generate jobs in new professions, some of which are unknown today, as well as in existing professions.

The mere existence of technology is no guarantee of these economic and social gains. Recent history proves the point. While digital technologies have contributed to American growth, France has benefited little. French activity has grown far less than on the other side of the Atlantic, and no real global digital player has emerged: between 2001 and 2022, wealth per capita grew by 29% in the United States, compared with just 14% in France.

It is therefore essential to put in place a set of appropriate public policies to maximize gains: innovation policy, industrial policy, competition policy... Support for professional retraining and continuing education will also be decisive. The challenge of supporting individual career development is all the more important given that the rapid spread of AI will make transitions difficult.

The main benefits will accrue only to those countries that give themselves the means to master AI. To date, France and Europe are lagging behind.

D. THE SCALE OF THE ECONOMIC CHALLENGE: THE WORRYING WEAKNESS OF FRANCE AND EUROPE

The digital economy is two to three times weaker in Europe than in the US, and AI is following the same path. Of the 100 largest-cap technology companies at the end of 2023, 10 are European. The problem is not only that Europe doesn't produce any digital giants, but also that it doesn't produce any second- or third-tier companies: neither Adobe, Uber, AirBnB, Shopify nor Stripe are European, even though their biggest markets or founders are European. In the digital services sector (software, data processing, etc.), activity is 2.5 times higher in the United States than in the European Union and the United Kingdom combined. We find a comparable order of magnitude in AI. The number of specialized companies financed over the period 2013-2022 is 2.5 times higher in the United States than in Europe.

If this superiority continues or strengthens, France and Europe run the risk of rapid economic decline. This risk is twofold: being largely bereft of companies specializing in AI, and seeing existing companies lose competitiveness. Like the previous wave of digital innovations, we could not only miss out on the AI economy, leading to the increasing capture of our economic value by others, but also see the weakening of other business sectors.

No company is immune. Over the past ten years, the integration of AI into companies has been slower and shallower overall in France than in the USA, the UK or Scandinavia. In the absence of rapid, structural adoption, each of France's companies will face the erosion of its market share, margins and value, but also the risk of being ousted by a new player (disruption). This prospect becomes more acute as AI systems become more accessible and more powerful.

What's more, our lag in the field of artificial intelligence undermines our sovereignty. Weak control of technology effectively implies a one-way dependence on other countries. In the privatized and ever-evolving field of AI, public power appears largely outmatched, limiting our collective ability to make choices aligned with our values and interests.

This weakness in French innovation can be explained by a number of factors. A lack of public engagement with the technological issues underlying AI and its potential impact on society. Our collective aversion to risk, which leads us to avoid unproven technologies and business models. Bureaucracy, which hinders public research in particular.

Lagging behind is not inevitable, and it's not too late to do something about it. Our continent has assets that should not be overlooked. European companies are positioned across the entire AI value chain. French higher education is training world-class engineers and researchers in AI. On the other hand, the technological innovations at the heart of generative AI are recent, and the economic value chain is far from mature. The AI economy is still in its infancy, and most business models have yet to be invented. Since Cédric Villani's report (2018), the French government has also set itself in motion as part of the "*investissements d'avenir*" and France 2030. This action has spurred the creation of specialized training programs in AI, strengthened public AI research and helped private innovation.

E. INNOVATING TO CONTROL OUR FUTURE

If we want to control our future, we cannot just use AI systems developed elsewhere. If France and Europe do not catch up, this will reinforce our dependence on other countries, affect our social cohesion and weaken our economy. Similarly, it's an illusion to believe that we can take a self-sufficient path. So it's up to us to take advantage, right now, of the potential of AI systems, wherever they come from, while creating the conditions for a European AI offer. The emergence of an AI ecosystem depends on three key areas of action.

Firstly, current funding for the AI ecosystem is insufficient to bring about the emergence of world-class players: we recommend redirecting a portion of savings towards innovation. The amounts invested in AI in the United States are now 20 times greater than those invested in France. For comparable wealth, we invest around three or four times less than the Americans, and the gap is likely to widen. In the medium term, a structural increase in the allocation of savings to innovation is essential. Proactive action must be taken quickly in this direction, for example with regard to the taxation of life insurance products, so that in a few years' time we will have a significantly increased capacity for financing.

In the short term, we propose the creation of a "France & AI" investment fund. The fund will mobilize €10 billion in corporate private equity and public support to bring about the emergence of the AI ecosystem and accelerate the transformation of the economic fabric through AI. Alongside financial resources, the fund will be accompanied by the pooling of activity data to drive certain digital projects. Such a scale of resources and the tandem of funding and data are unprecedented in France. Faced with the risk of economic downgrading, boldness will contribute to the emergence of innovative, high-performance solutions and accelerate the modernization of French companies.

Secondly, we won't reap the benefits of generative AI without access to reliable, high-quality data: we therefore recommend rethinking data governance. Firstly, it is essential to **facilitate access to personal data** to enable its use in therapeutic innovations, notably by abolishing certain prior authorization procedures for access to health data and reducing response times from the CNIL. This implies reconstituting the CNIL's mandate to include a focus on innovation, revising the composition of its board and increasing its resources.

We must then **enforce the principle of transparency** of training data for large-scale AI models. Provided for in the European AI Act, this principle must ensure respect for literary and artistic property rights. It must be implemented in the simplest possible way, for AI model developers and rights holders alike, in particular by drawing up standards for the publication of information on AI models and the enforcement of the opt-out clause.

Finally, from a more forward-looking perspective, it's up to us to **devise a new, collective way of managing data**. While data protection is centered today on the individual, with the European General Data Protection Regulation (GDPR), access to data from AI models and the resulting benefits are mainly collective. This dichotomy has so far favored the digital giants, who alone have hundreds of millions of users who generate streams of data every day to drive their models. We need to explore new models of joint data governance, without of course weakening individual protection.

Thirdly, computing power is the other essential ingredient of generative AI: we therefore recommend making France a major hub in this field. This is a prerequisite for strategic autonomy, and public supercomputer capacity, which we must support, will not be enough. On the supply side, we need to secure supplies for the French ecosystem without delay, through a European order for private computing power. At the same time, solicit the installation of private computing power in France and Europe with a public guarantee of usage and support for installation and electrical connection. On the demand side, an AI tax credit would support research and development projects involving the leasing of computing power, subject to the use of a computing center established in the country. Finally, industrial policy could be geared towards the emergence of an electronics industry adapted to AI.

It's not a question of chasing after technological advances, but of creating our own comparative advantages. Targeting and concentrating resources will therefore be key to establishing our superiority in certain segments of the value chain, and thus being able to speak on an equal footing with our competitors and partners. Targeting must go hand in hand with the gradual emergence of AI innovation ecosystems in France and Europe. The path to differentiation could focus in particular on the environmental dimension, by targeting new generations of AI, from hardware architecture to the choice of models, that will consume less energy.

F. RELYING ON THE OPENNESS OF ARTIFICIAL INTELLIGENCE SYSTEMS

The development of AI brings with it a major risk of market domination by a few players. A single (American) company currently holds 80% of the global market share in GPU design. Three companies — all American — capture 80% of the increase in French spending on cloud service infrastructures and applications. These same three companies are combining their cloud services with American generative AI tools.

It cannot be a good thing for the AI technological revolution to be limited to a handful of companies, even more so when none of them is European. We need to ensure the emergence of a diversity of economic players, particularly French and European, on economic and sovereignty grounds. Limiting dominant competitive positions promotes growth and a fair distribution of economic gains.

European competition policy must therefore be fully mobilized to prevent the emergence of dominant positions. In the short term, it is important to make use of the range of actions provided by the European Digital Markets Act. This regulation could also be supplemented to take account of the specific features of the AI value chain. In the medium term, we need to consider a change in competition policy doctrine, moving from a static system (what market shares does this company hold today?) to a dynamic vision (what market shares could this company hold tomorrow, and what companies could enter this market tomorrow?).

Beyond this, our Commission recommends supporting an open ecosystem of AI developers, offering the benefits of transparency, pluralism and competition. We consider that such an ecosystem is a powerful lever for innovation and can contribute to the security of AI systems and the development of benevolent uses, including countermeasures against malicious uses. It also contributes to public confidence and the reduction of certain negative impacts of AI on individuals. We therefore need to provide the ecosystem with legal certainty and quality data, as well as develop model inspection and evaluation capabilities.

Finally, to shape the future of AI, it is essential to free researchers up from administrative constraints: we recommend adopting the principle of an "AI exception" in public research. In the form of an experiment, this principle aims for "zero hindrance for researchers", notably through a commitment on response times to requests and the introduction of an indicator of administrative simplicity. The AI exception should also make it possible to raise the salaries of researchers and teacher-researchers, and facilitate part-time work with companies or other socio-economic players in AI. These actions could go hand in hand with a doubling of public research resources specialized in AI, thus scaling up investment within the context of interdisciplinary institutes and the imminent AI clusters.

From the Enlightenment to the present day, openness has been at the heart of our European continent and our values. Let's embrace this core tradition.

G. WHAT DOES THE FUTURE HOLD IN THE AGE OF ARTIFICIAL INTELLIGENCE?

Generative AI is a milestone in the history of innovation. And it certainly won't be the last. In the months and years to come, we are likely to see further rapid and far-reaching advances. Models will be increasingly capable of being factual, of conducting reasoning, of understanding the physical world around us. AI will accompany people continuously and in all their tasks, perhaps in the form of personalized assistants. Products and gestures will be invented to enable us to take full advantage of these new AI systems. There will also be major advances in robotics.

The societal transformations brought about by these innovations will depend on our ambition and commitment. AI can be harnessed to reduce social inequalities, increase collective prosperity and improve the quality of work. These benefits will not come spontaneously. In the absence of a political vision and collective commitment, AI can, conversely, weaken our democracy, alter our sovereignty and concentrate wealth. Let's be careful neither to underestimate the potential of AI in the medium term, nor to overestimate it in the short term.

Let's give ourselves, collectively and without delay, the means to reap the benefits of AI. Its effects will be all the more beneficial if France and Europe master the technology and its value chain. This mastery is essential. Our Commission therefore recommends closing the French and European gaps and launching a new AI strategy by the end of the first half of 2024. Measures will have to be subject to ongoing evaluation and annual or even biannual review if the pace of innovation remains very high.

It's a race against time. The strategy recommended by our Commission is only the beginning. Going beyond that will require not only consistency in public intervention, but also action to strengthen the adaptability of our organizations, both public and private. A collegial approach to anticipation is also essential, to prepare our country for the effects of the technological revolution.

METHOD

COMPOSITION

Chairpersons	Philippe Aghion	Anne Bouverot	
General Rapporteurs	Arno Amabile	Cyprien Canivenc	
Members	Glles Babinet	Bernard Charlès	Cédric O
	Joëlle Barral	Luc Julia	Isabelle Ryl
	Alexandra Bensamoun	Yann Le Cun	Franca Salis-Madinier
	Nozha Boujemaa	Arthur Mensch	Martin Tisné
			Gaël Varoquaux
Rapporteurs	Marc Auberger	Emilie-Pauline Gallié	Erwan Paitel
	Simon Bunel	Paul Jolie	Timothée Paris
	Philippe Chantepie	Arnaud Mazier	Christophe Ravier
	Eloy Dorado	Vincent Montreuil	Ulrich Tan
			Louis Charles Viossat

AREAS OF EXPERTISE

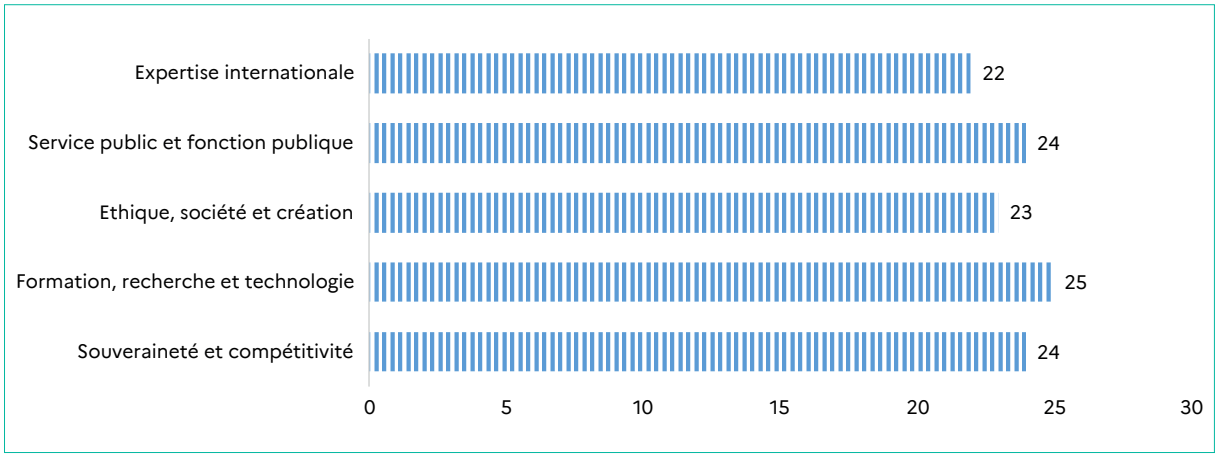

MISSION

In September 2023, the French Government set up the Artificial Intelligence Commission to "help put France at the forefront of the AI revolution". The Commission was thus tasked with putting forward operational, realistic and ambitious proposals supported by a long-term, global and objective vision. It is with this mandate in mind that we wrote this report.

PRINCIPLES AND METHODS

Expertise

The Commission's members and rapporteurs have been appointed, in a personal capacity, for their expertise in artificial intelligence. This expertise is made available to the government to help improve the efficiency of public action.

Collegiality

The Commission's plenary sessions, which have brought together the fifteen members and fifteen rapporteurs on a weekly basis, embody collegiality. The plurality of expert opinions and open, adversarial discussions contributed to the objectivity of the work, as well as to the measure and balance of the recommended action plan. The members of the Commission have given their general approval to the report and recommendations, which represent a majority consensus.

Consultation

The Commission's debates and conclusions were enriched by hearings with experts and stakeholders from a wide range of backgrounds, experiences and geographical horizons. They were enriched by the consultation of citizens, who contributed to the timeliness, realism and pragmatism of the recommendations.

Independence

The Commission has freely defined its program of work and its organizational structure, within the framework of the mission entrusted to it on September 19, 2023. It has conducted its deliberations and drawn up its conclusions independently of the executive and legislative powers.

KEY FIGURES

600 hearings with AI experts and stakeholders

200 spontaneous contacts with the Commission

7 000 contributors to the Agora application

23 plenary Commission working sessions

25 recommendations to the government

INTRODUCTION

Artificial intelligence (AI) is omnipresent in public debate. Its applications are hitting the headlines: text synthesis, music generation, translation and interpretation, photo editing, and more. The same applies to certain AI-based tools whose use is spreading rapidly around the world. Some see this as a source of concern, others of hope. Some observers see or predict a societal revolution, while others see no disruption in their personal or professional lives. But what are we really talking about?

ARTIFICIAL INTELLIGENCE IS NOTHING NEW: IT DATES BACK TO THE 1950S

AI has actually been on the scene for decades. The American MYCIN blood disease diagnosis and prescription system (1970s), the construction of Navlab, the first self-driving vehicle (1986), Deep Blue's victory over world chess champion Gary Kasparov (1997), the Siri virtual assistant built into iPhones (2011), or the defeat of world champion Ke Jie in the game of Go by the AlphaGo machine (2017) have all been described, at one time or another, as artificial intelligence.

The history of AI is indeed over 70 years old. As early as 1950, the British mathematician and cryptologist Alan Turing was interested in the ability of a machine to imitate a conversation. For several decades, this ability was not sufficient to fool a human, who could distinguish a simulated conversation from a real one. We'll come back to this in a moment.

It was a few years later, in 1956, that the term 'artificial intelligence' first appeared. Research in this field gradually gained momentum, with the emergence of several technological approaches. AI first developed in the form of deductive rules of the "if... then" type. This so-called symbolic approach, based on reasoning and instructions, was in the majority until the 1990s.

Although this symbolic approach has not been abandoned, a statistical approach to AI has been gaining ground since the 1990s: machine learning. Unlike the symbolic approach, the human does not determine a set of "if... then" rules. Instead, they ensure that the computer "learns" to identify statistical relationships between data. So there is no explicit human instruction: the machine is trained to recognize links from a set of training data. The machine then applies these links to new data to perform a task.

The success of this second approach relies on two essential ingredients: data and computing power, supported by the emergence of the cloud. The availability of these two ingredients has increased sharply over the last 30 years, under the triple effect of the digitization of our society (producing more data), improvements in semiconductor materials (increasing computing power) and technical progress. These developments have enabled machine learning to make rapid and significant progress. Techniques enabling machines to automatically "learn" rules from data have been diversified and refined.

AI COVERS A RANGE OF DIGITAL TOOLS THAT ARE ALREADY UBIQUITOUS IN OUR SOCIETY.

There is no single, universal definition of AI, not least because the term covers so many technologies: thirty years ago, it seemed far-fetched that a machine could distinguish a cat from a dog. It's also a theoretical notion, because these are the artificial intelligence systems we use every day. So what are these AI systems?

AI systems can make predictions, recommendations or decisions. They respond to a given set of objectives and influence their environment.

AI systems have a huge number of applications in our daily lives, in our economy, and in our public services. Examples include voice recognition on cell phones, industrial robotics, self-driving vehicles, pathology detection in medical imaging, virtual sales assistants, facial recognition on computers, targeted advertising on the Internet, and the identification of financial anomalies to combat tax fraud.

GENERATIVE AI IS A MAJOR TURNING POINT IN THE HISTORY OF INNOVATION.

These numerous professional and personal applications are amplified by so-called generative AI systems. AI is called generative because it can generate new content in the form of text, image, sound, video, or code. This production capacity is a major turning point for AI, for several reasons.

Firstly, the use of generative AI models can be very straightforward. It is indeed possible to use dialog interfaces, whereby a human expresses himself or herself in written or spoken form to control the generation of content. This creates the impression of a conversation or dialogue with the machine.

Secondly, content generation is fast. The most advanced models can produce music or long texts in a matter of seconds, whereas a similar human production would take several days or weeks.

Thirdly, the content generated is realistic. It appears credible to the eyes and ears of a human being, because it displays human characteristics: clarity of speech, logical sequence of words, coherence of images, presence of intonation, etc.

Fourthly, the models have important capabilities. A series of experiments in 2023 showed that generative AI models appear to be able to perform complex human tasks. For example, one experiment concluded that algorithms performed better than 90 % of human candidates on certain bar exams in the USA[1]. Other experiments in the field of medicine have shown that the accuracy of diagnoses made by algorithms is superior to those made by doctors[2]. These performances are promising, even if they are sometimes obtained in situations that are relatively far removed from real-life conditions. We must continue to evaluate these performances rigorously[3] and avoid projecting human intelligence onto them[4].

Realism, simplicity, speed, ability. These characteristics of generative AI enable the automation of a number of tasks that were previously difficult to automate. For example, they facilitate the personalization of commercial offers, simplify the analysis of financial data, speed up scientific research, and so on.

These same characteristics suggest that AI could take over from personal computers, social networks and smartphones as *the* dominant digital platform, the technological layer on which all other new services are built[5]. With each change of platform, the deck is reshuffled and power is redistributed to the companies that control the new platform. IBM, the all-powerful company of the mainframe era[6], has not disappeared, but it no longer has the same centrality since the rise of personal computers. If AI is the next platform, who controls it? The companies that make the models, or those that make AI-enabled products? If chatbots become the new central interface with the digital world, how do we define their conditions and behaviors together?

Generally speaking, there are many uncertainties. Who controls and deploys generative AI systems? What responsibilities do these players have? How will these systems be used? What impact will it have on the economy, work and employment? What impact will the human appearance of generated content have on our relationship with truth and information? How will technologies evolve?

To answer these questions, a more detailed description of the technology and the economic value chain is required.

1. Katz D. M., Bommarito, M. J., Gao, S., and Arredondo, P. D. (2023), "GPT-4 Passes the Bar Exam", *SSRN eJournal*.

2. Caruccio L., Cirillo, S., Polese, G., Solimando, G., Sundaramurthy, S., and Tortora, G. (2024), "Can ChatGPT provide intelligent diagnoses? A comparative study between predictive models and ChatGPT to define a new medical diagnostic bot", *Expert Systems with Applications,* Volume 235

3. In particular, to check that they are not simply the result of including these tests in training data, or that they are only carried out in a highly controlled environment.

4. For a human, passing the bar exam is (reasonably) correlated with more general legal competence. We have no reason to believe that this is also the case for a language model.

5. OpenAI has made no mistake about this, and is striving to become an indispensable platform as quickly as possible. The company quickly opened up tools enabling third parties to create customized versions of ChatGPT, to be purchased from a store.

6. Large, widespread computers from the days when computers were so big and expensive that only a few organizations had them.

HOW DOES GENERATIVE ARTIFICIAL INTELLIGENCE WORK?

As we have seen, generative AI systems are based on machine learning techniques. They are trained on large volumes of data and "learn" the relationships between this training data. This data can take the form of text, images, sound, video or tables of information, and these categories can be cumulative.

Systems can be trained on very large quantities of data to form foundation models (more recently called general-purpose AI models), which can be adapted to many different tasks. These include large language models, which have been trained on vast bodies of texts.

Once trained, the model can be called upon by a user via queries. The generative AI model then responds to the query by producing new data (text, image, sound). The content generated has a certain similarity to the training data, without being identical.

Let's take the particular case of text generation. When the query "Complete the following sentence: France is a major (...)" is submitted, the language model begins by breaking down this query into a series of elementary text units called tokens. A token corresponds to a series of a few letters, not always forming complete words. For simplicity's sake, let's associate a token with each word: "France", "is", "a", and "major".

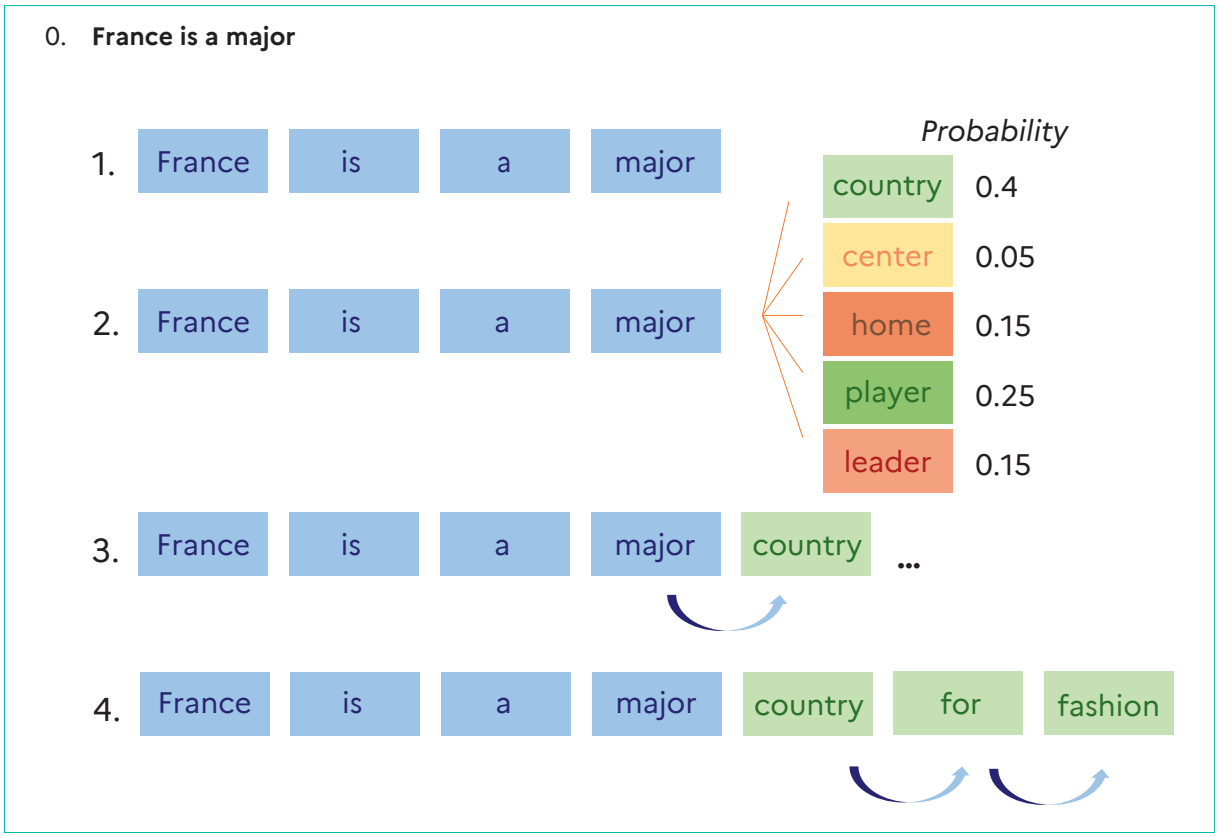

After several technical steps, the model analyzes this succession of tokens in light of the training data. It identifies a set of possibilities for continuing the text: "country", "center", "player", "leader", etc. Each successor token possibility is assigned a probability. The response is generated, token by token, according to the probability of each token. The generated text is then addressed to the user: "country", "for", "fashion", etc.

This operation, described here in very simplified terms, has one essential feature: generative AI models do not understand the meaning of words, images or sound. The meaning of the generated response therefore comes not from the machine, but from humans. Humans project their vision of the world onto the results generated by the machine. The more closely they resemble a text written or spoken by human beings, the more significance they attach to the automatically generated text and voice.

This characteristic therefore calls for a certain amount of vigilance when it comes to using current generative AI models. In particular, they cannot be used — as yet — as reliable sources of true assertions, like encyclopedias. Indeed, AI models sometimes generate erroneous answers. Erroneous responses, often referred to as hallucinations (or confabulations), are one of the areas in which AI systems need improvement.

WHAT IS THE ECONOMIC VALUE CHAIN OF GENERATIVE AI?

We don't need to wait for these technological advances for the applications of AI systems in our societies to be innumerable. These applications offer significant economic value, shared within a long value chain.

In the middle of this value chain are, of course, the creators of the foundation models described above. As these models can neither be trained nor used without two essential ingredients — data and computing infrastructure — the value chain extends upstream to the companies that provide data and computing power platforms. These companies in turn furnish themselves with equipment specific to generative AI models.

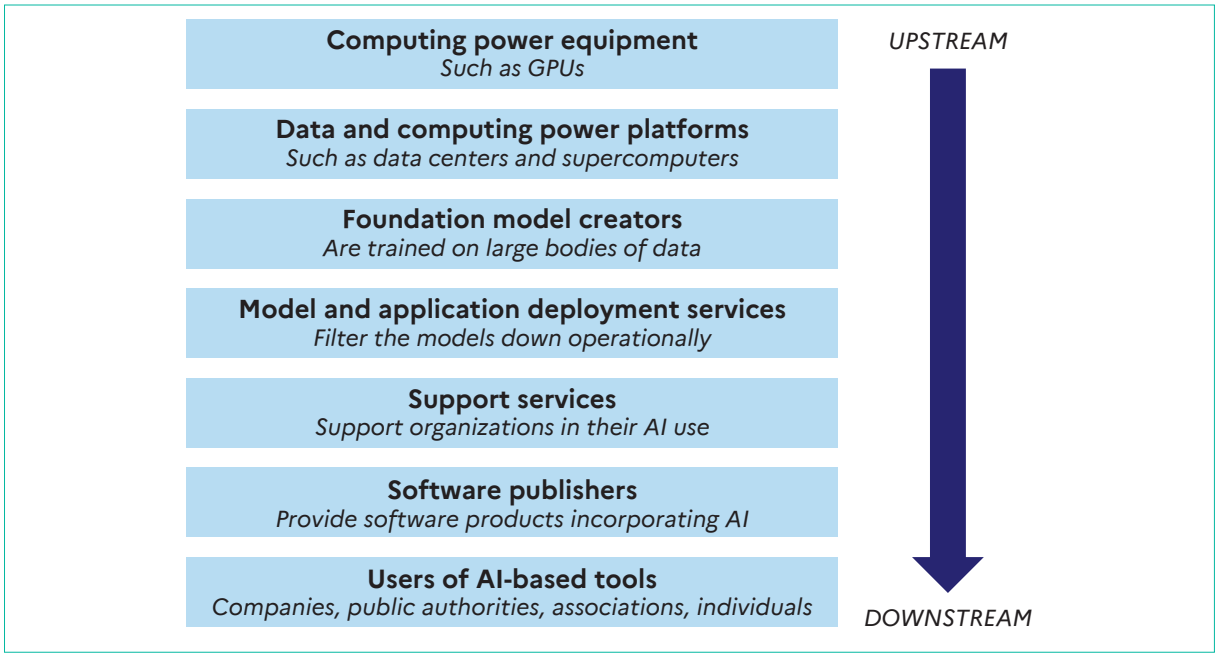

The economic value chain also extends downstream from the foundation models. These can be used directly by end consumers (companies, public authorities, associations, individuals), who then derive an economic benefit, such as productivity or quality gains. Numerous companies are deploying this technology to optimize their production chains and services, in fields as varied as media, finance, law, IT, automotive, and pharmaceuticals.

Most of the time, however, several economic players contribute to the value chain between end-users and foundation model creators. Indeed, foundation models can generally be optimized for specific applications and tasks. Moreover, the deployment of AI-based tools generally calls for a transformation of organizations (adaptation of information systems, evolution of procedures, reallocation of human resources, etc.), which requires support. AI models will also be integrated directly into the software products chosen by end-users, particularly in the office suite (for writing e-mails, meeting minutes, etc.).

All in all, the generative AI value chain is made up of companies whose core business model is AI, as well as companies that use AI and integrate it into a pre-existing or adapted business model. So there's an AI economy and an AI-driven economy. The overall economic value of generative AI is considerable. According to the work of our Commission, the deployment of AI could double our country's economic growth.

Given these economic advantages, but also the potential social and sovereignty benefits of AI, we can expect very strong international competition. The companies — and countries — that come out on top will reap the main benefits.

IN FRANCE AND EUROPE, WE ARE CLEARLY LAGGING BEHIND...

To date, American players largely dominate the upstream end of the generative AI value chain. Let's take a few examples concerning computing power, data platforms and foundation models.

Today, graphics processing units (GPUs) are the most essential computing power equipment for generative AI systems. Yet a single (American) company currently holds 80% of the global market share in GPU design.

Data centers are platforms that enable organizations and individuals to both host data and use AI systems. Worldwide, three (American) companies hold two-thirds of the market share.

The ranking of the most accurate foundation models[7] includes 30 models, created by twelve companies. Of these, the majority are American, and Europe has only three companies ranked, two French and one German.

This American domination can be explained by the fact that the United States invests far more in AI than France and Europe. The amounts invested in AI in the United States are 20 times higher than in France. It's true that the American economy is much larger than ours. However, for comparable wealth, we invest around three or four times less than the Americans.

... BUT WE DO HAVE A FEW TRICKS UP OUR SLEEVE: IT'S NOT OVER YET!

There is no denying that France and Europe are lagging far behind. However, our continent has assets that should not be overlooked.

European companies are positioned across the entire value chain, and some of them are among the very best. For example, the world's leading manufacturer of machinery for the semiconductor materials industry is Dutch.

What's more, Europe, and France in particular, can count on professionals who are precisely trained in artificial intelligence technologies. The quality of French higher education in this field has led many foreign companies to recruit engineers and researchers trained in our country.

What's more, 2023 showed that the lead held by some American players, including OpenAI, was not irrecoverable in terms of AI model production. In 2024, many models are expected to catch up with or surpass GPT-4. The race is not only about technology, but also about business models, products and the ability to deliver them to a large number of users at low cost.

Finally, the technological innovations at the heart of generative AI are recent, and the economic value chain is far from mature. The AI market is still in its infancy, and most business models have yet to be invented. At the same time, Europe should be able to draw on its existing economic fabric to position itself in the many AI-driven markets, i.e. those integrating AI into their business models.

So it's not too late to (re)act. France and Europe have a lot to gain from artificial intelligence. With this in mind, our Commission has drawn up an ambitious action plan.

7. Stanford Center for Research on Foundation Models (2024).

SOME DEBATES AT THE HEART OF OUR PROPOSED ACTION PLAN

To arrive at a proposed action plan, our Commission has raised and debated many questions over the past six months. Here are just a few of them.

- **What weight should be given to control over the creation of AI systems, and what place should be given to foreign technologies?**

  A firm commitment to international AI competition represents a significant investment. However, being one of the pioneers of AI comes with many advantages: greater economic gains, less dependence on foreign countries, mastery of the value repository underlying the AI system, the ability to adapt the technology to anticipate the effects of its deployment, etc. We know that the United States derives major benefits from its mastery of previous waves of digital innovation (computers, the Internet, smartphones).

- **To be able to compete internationally, do we need public policies targeted at AI, or more cross-functional action to encourage innovation in general?**

  It's very difficult for the State to target its support on the next successful innovations, because it's not necessarily the most far-sighted player, but also because we're generally surprised by disruptive innovation, which arises from the unexpected intersection of research fields and business models. This argues in favor of public policies that provide horizontal support for innovation. Even so, AI has a number of specific features and prerequisites — data, computing power, particular skills — which may justify targeted actions. In any case, public support should not aim to play endless catch-up (i.e., constantly chasing after pioneers), but to achieve superiority in certain links in the value chain (i.e., choosing a few areas of competition to win).

- **How much trust should be placed in the private players behind AI systems, and what should national, continental and international public authorities supervise and regulate?**

  In the years to come, AI systems will have an impact on our societies, from the economy and the organization of work to our relationship with information and truth. These effects will be far-reaching, whether or not France is at the origin of these AI systems. These effects justify public intervention, to ensure that AI is deployed in the service of the common good, to promote a fair distribution of economic and social gains, to limit the harmful uses of AI, or even to limit the dominant positions of a few companies. However, excessive or ill-sized public intervention could disproportionately harm innovation and deprive us of its benefits.

- **What is the right balance between data protection and accessibility?**

  The machine-learning technologies at the heart of AI systems rely on vast quantities of data. Consequently, restricting access to data means restricting innovation and its benefits, or leaving it in the hands of a few players capable of collecting more data and bearing the cost of regulation. Nevertheless, restricting access to data is justified by numerous protection objectives: privacy, intellectual property, and so on. For example, limiting access to researchers' health data ensures a certain confidentiality of sensitive data, but also slows down the discovery of new therapeutic treatments.

- **What are the benefits and risks of open access to AI systems?**

  The increasing power of AI models raises fears of a proliferation of malicious uses. Open access to these increasingly powerful models — to which people with malicious intent can therefore gain easy access — can be a cause for concern. At the same time, open source clearly facilitates the development of benevolent uses, including countermeasures to malicious uses. Open access also broadens the base of contributors to the development of AI systems, which can help to make them safer. Similarly, the protection of cultural creations helps to promote French culture by ensuring a degree of economic independence for the industry, but also limits the presence of the French language in AI models.

- **What levers do we need to pull to master AI?**

  For France and Europe to take full advantage of the AI technological revolution, several levers are obvious: access to quality data, availability of computing power, investment capacity, and expert personnel in the field. But these few levers are far from sufficient. What actions will enable us to steer innovation in line with our political objectives? Which ones will increase confidence in AI tools? Should public service evolve? Do we need to change the legal framework? As we can see, the social, economic and legal dimensions cannot be considered in isolation, and many levers need to be activated simultaneously to develop and guide AI projects... but according to which priorities?

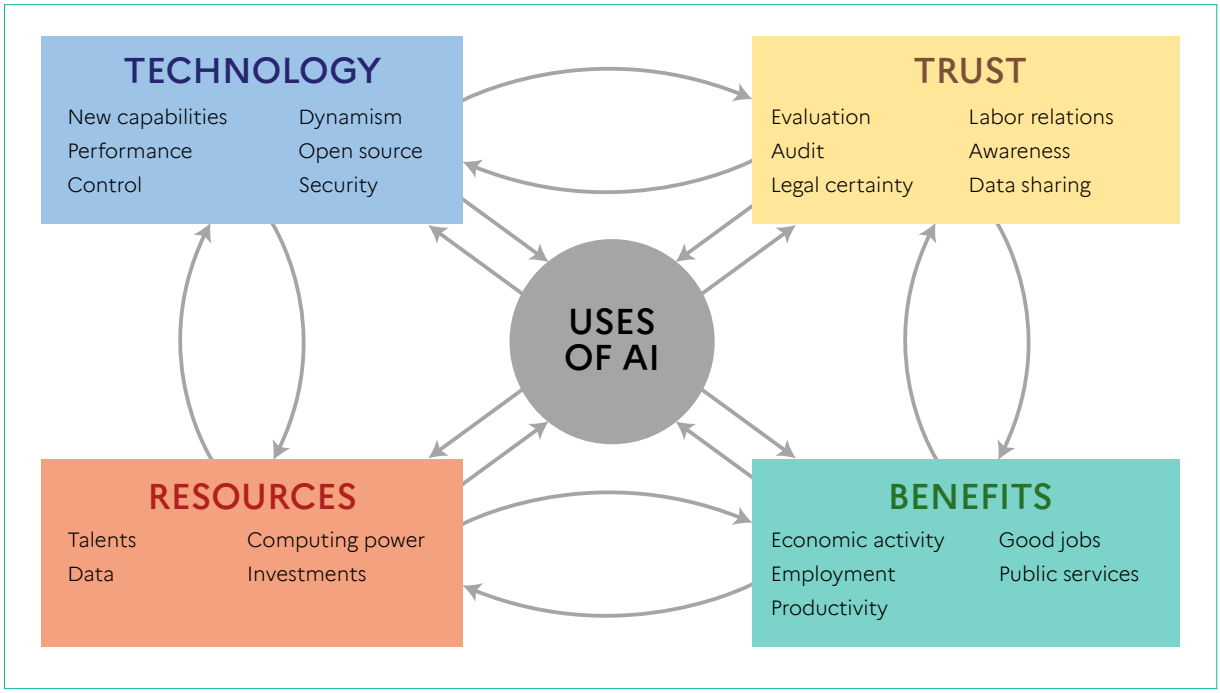

The virtuous circle of AI between resources, technological development, trust, and economic and social benefits.

This report sets out the terms of the debate and provides a set of answers. These reflect the convictions of our Commission. These convictions, forged between September 2023 and February 2024, are based on the expertise of our members and rapporteurs. They are underpinned by consultations with 7,000 people and hearings with 600 experts and stakeholders.

THE INEVITABLE DEBATE: WHAT WILL TOMORROW'S WORLD LOOK LIKE?

Artificial intelligence is less than a century old. Recent technological developments, which have taken many by surprise, have been dazzling. It's certain that new innovations will follow the wave of generative AI we are experiencing today. What will these next innovations be? Will they be an extension of generative AI, or a technological breakthrough? What do we need to do to be in the running for the next stages of innovation?

AI systems are already omnipresent in our societies, and we can expect their effects to become increasingly pervasive. But how far will these effects go? At what pace? Which sectors of society will be the first or most affected? How can we maximize the benefits of innovations, while minimizing their negative effects? What can we do to prepare our society for the upheavals to come? And beyond that, what kind of society do we want?

In 1881, the first incandescent lighting was introduced at the Paris Universal Exposition. At the 1900 Universal Exposition, the Electricity Fairy triumphed. Yet who could have anticipated the scale of the changes the world would undergo as a result of this new form of energy?

At the same time, the car was appearing in the homes of a few wealthy individuals. Who could have anticipated that the democratization of the car would affect the environment, redesign cities, transform our relationship with distance and thus alter our social interactions?

Unfortunately, our Commission has no crystal ball. So we don't pretend to see right where so many others have seen wrong in the past. We do hope, however, that we have asked the right questions and raised the right issues. Some of our proposals will undoubtedly prove to be wrong, so uncertain are the next technological and societal evolutions. However, we believe that the recommended action plan will enable France and Europe to resolutely enter the international AI competition, and place innovation at the service of our principles, values and interests.

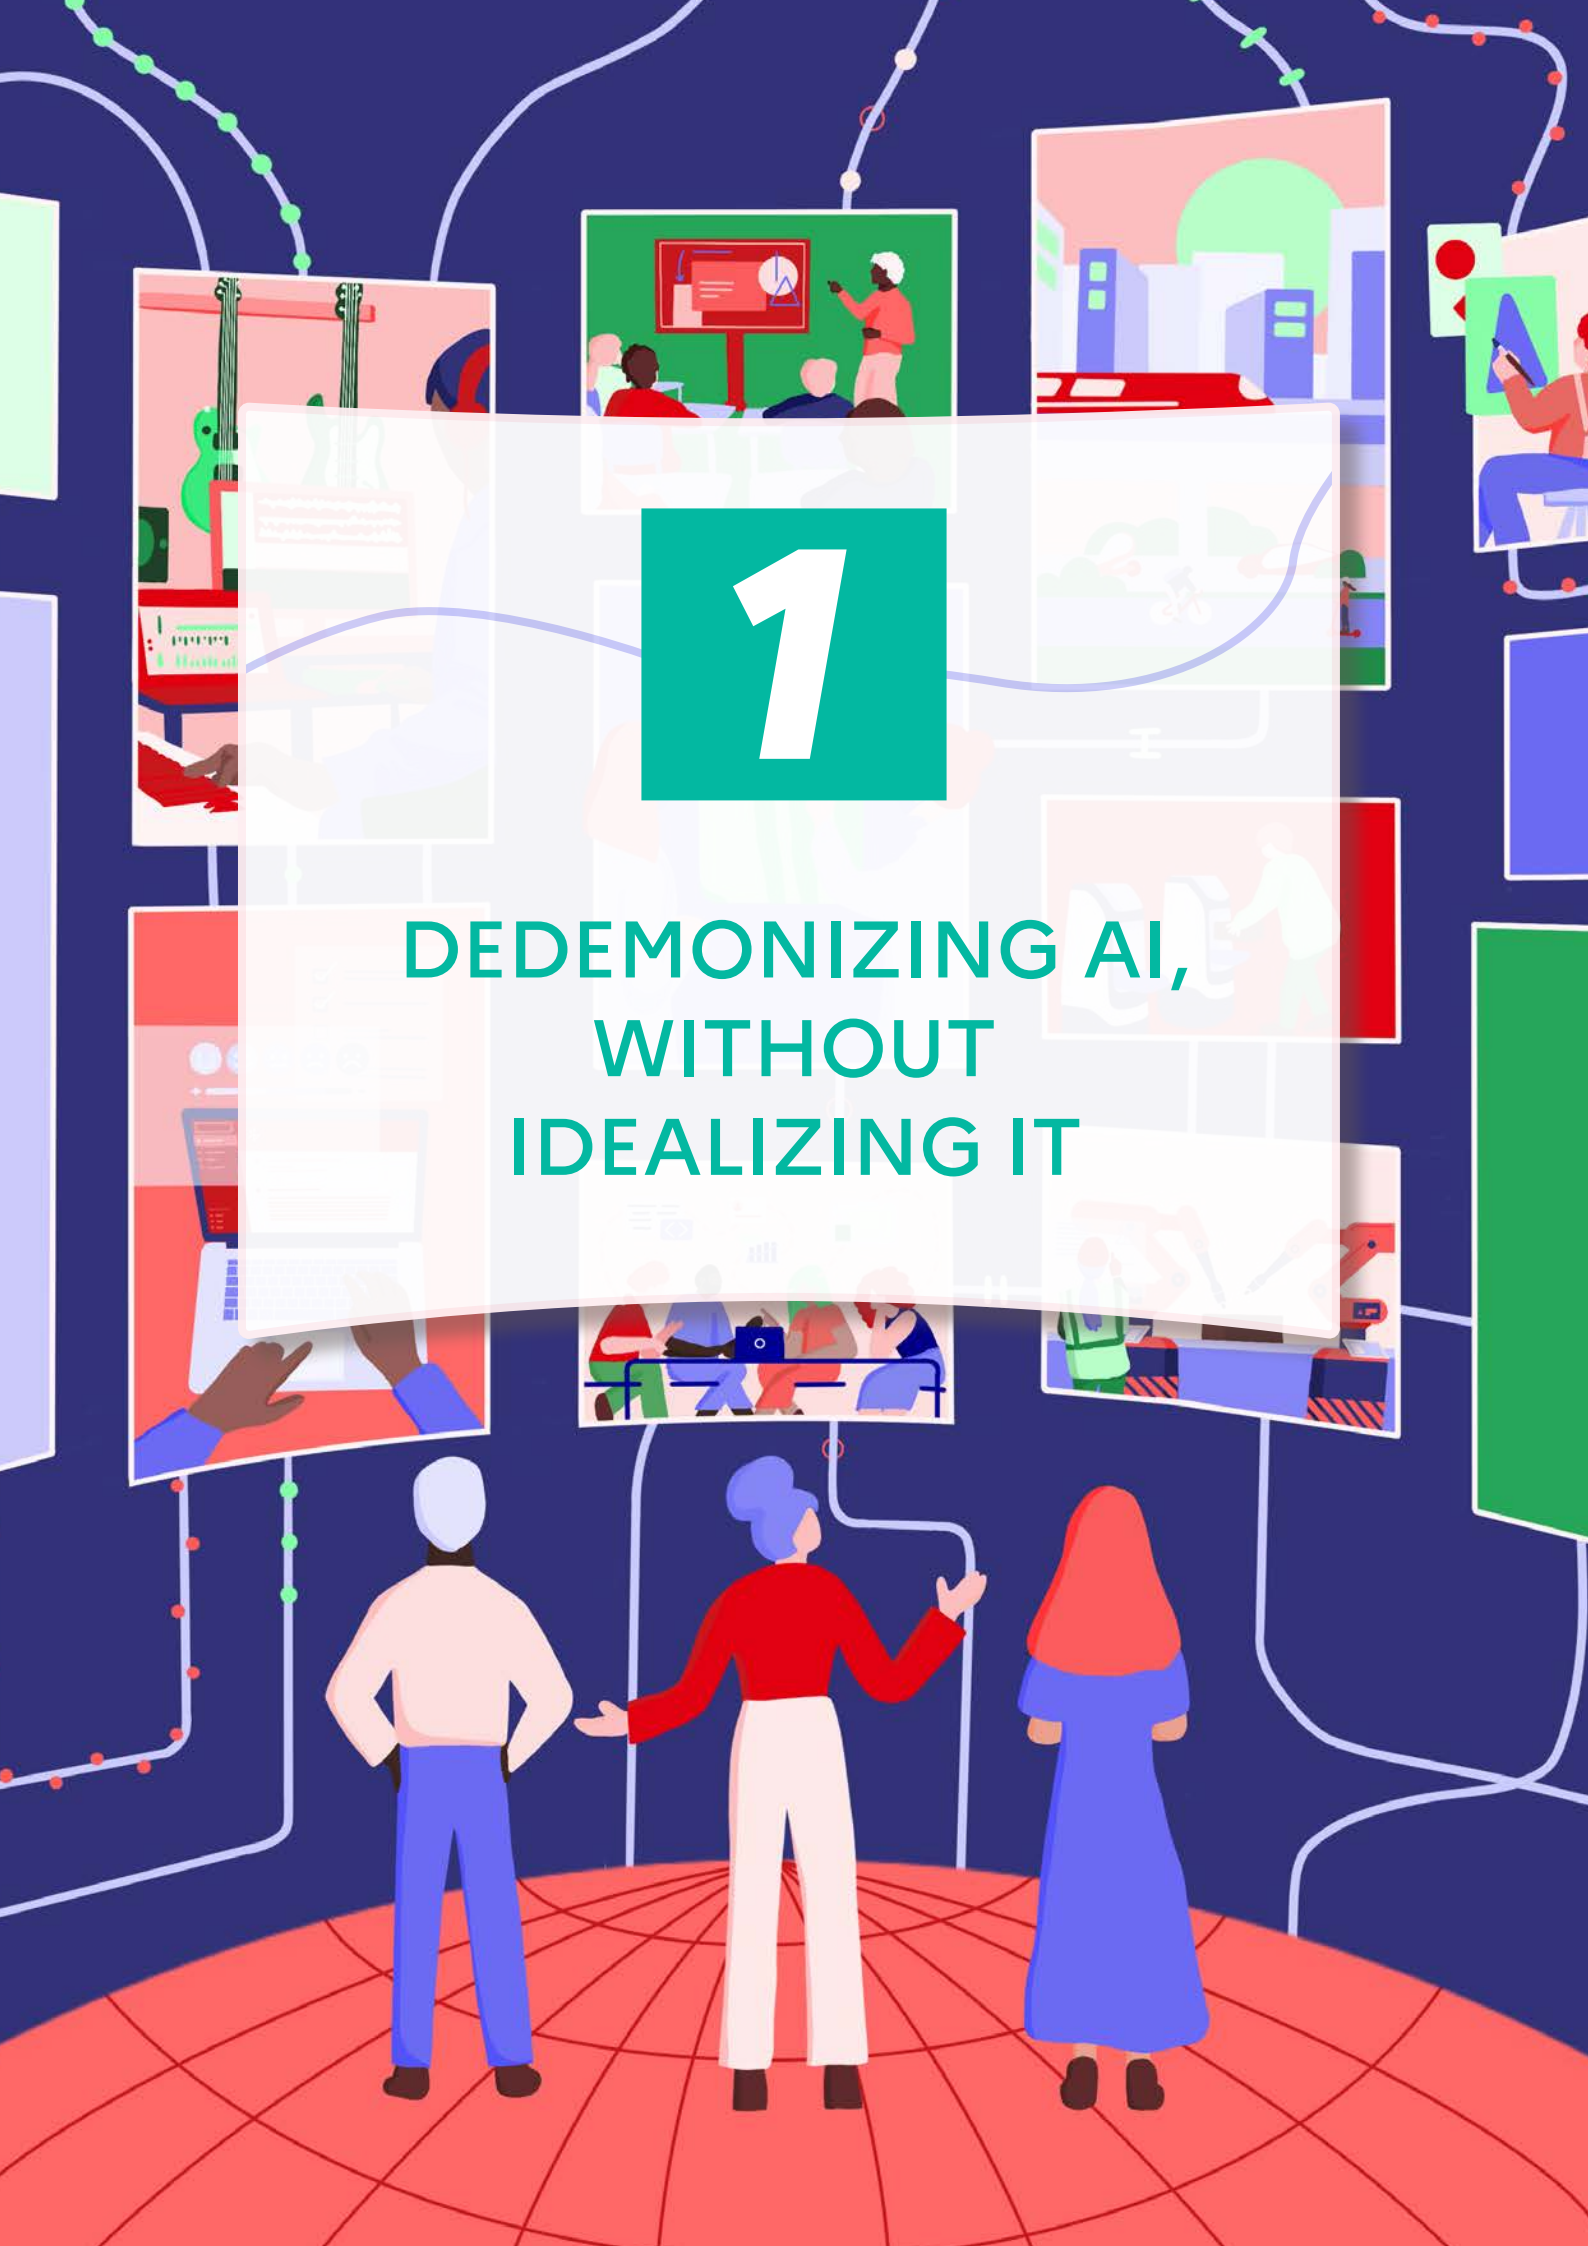

1

DEDEMONIZING AI, WITHOUT IDEALIZING IT

1.1 DOES AI CONCERN ME?

> **Yes,** AI concerns us all. As users, because we already use AI-enabled services in our daily lives. And as citizens, because we will have to decide how we want to use these technologies. Between 1920 and 1930, most Western cities and countries decided to separate the flow of pedestrians from cars, in order to avoid accidents and get around as quickly as possible by car. This decision changed our cities and our daily lives for at least a century. The many choices surrounding AI will affect us just as much.

The launch of ChatGPT at the end of 2022 put AI at our fingertips. One year later, 55% of French people say they are familiar with ChatGPT, and 28% say they have heard of it but don't really know what it is[8]. But ChatGPT is only the tip of the iceberg. Many observers point out that AI didn't appear with conversational robots like ChatGPT, and they're absolutely right: research into artificial intelligence began as early as the 1950s; it is, in fact, almost as old as computer science.

Over the years, AI's progress has resulted in some resounding feats in activities previously thought to be the preserve of humans, such as chess and Go. At the same time, AI has become an integral part of many everyday activities. Our phones use it to unlock by recognizing our face or fingerprints, it helps with foreign language translation, image recognition, automatic video subtitling, fraud detection, product recommendations, and more. When an ad appears on our screen, it's often chosen by artificial intelligence. With ChatGPT, we can interact (or even play!) with the AI and imagine new use cases.

8. Online survey conducted by Ipsos for Sopra Steria in October and November 2023 among 1,000 people representative of the French population aged 18 and over.

The importance of AI is not only linked to what it can do, although beating the best chess or Go players remains impressive. It is the rapid spread of AI across a wide range of products and services that gives it such transformative power. In this sense, it is often compared to other innovations which, in their time, have profoundly transformed our lives, our economies and the very functioning of our societies, such as electricity or the telephone. AI, integrated into our communication and information systems, will also transform our democracies.

It took us a few decades to get to grips with the previous technological revolutions, so we need to design a society with AI without delay.

Recommendation No. 1

Create the conditions for collective appropriation of AI and its challenges, in order to collectively define the conditions under which it will become part of our daily lives.

1.2 SHOULD WE BE AFRAID OF AI?

> **No, but as with any tool, we must be vigilant.** Today's AI will not lead to the end of humanity. On the other hand, AI systems already come with a set of risks that need to be managed.

In the spring of 2023, 60 AI experts and world-famous personalities signed a declaration that caused quite a stir[9] : they warned that preventing the risk of human extinction caused by out-of-control AI should be a global priority, on a par with preventing pandemics or nuclear conflict. Indeed, the rhetoric surrounding AI focuses as much on its risks and dangers as on its potential to improve humanity's everyday life, and often it's the most extreme risks that are highlighted. The extinction of mankind is no mean feat!

This ambivalence in the treatment of AI influences the way we perceive it. Fascination goes hand in hand with fear. One year after the launch of ChatGPT, 77% of French people say that AI is a real revolution[10], but 68% think that we should pause in the development of AI[11]. And we're not alone, since 79% of Chinese and 74% of Americans are in favor.

In this respect, the AI situation is not very original: all the technologies that have revolutionized our daily lives have, in their time, aroused fears, some imaginary, others very real. The fear that the speed of trains would blind their passengers has proved to be entirely unfounded, but the development of railroads has also been the source of incidents, sometimes serious, which have required a response from public authorities: rail tunnels have long been considered an unhealthy, even dangerous environment. By 1900, even the Electricity Fairy was raising fears of electrocution in the street.

Anti-electricity drawing from 1900.

9. *"Pause Giant AI Experiments: An Open Letter"* published by the Future of Life Institute on March 22, 2023
10. Online survey conducted by Ipsos for Sopra Steria in October and November 2023 among 1,000 people representative of the French population aged 18 and over.
11. Online survey conducted by Ipsos for AXA in May and June 2023 among 3,226 experts in 50 countries and 19,000 people in 15 countries, representative of the national population aged 18 and over.

In the same way as trains and electricity, which provide us with immense services and are part of our everyday lives, AI presents risks. These risks must not be ignored, and they call for a response. This response must be proportionate and must not deprive us of the benefits of AI. All the more so as there will be not "one AI", but many tools integrating AI functionalities. For example, an AI tool in healthcare does not involve the same risks as an AI tool in online advertising, so different precautions need to be taken.

The risks associated with the spread of AI can be grouped into three broad categories.

Risks of imperfection. Many AI systems work with probabilities. This is what gives them their flexibility, but also their capacity to be "wrong". When I ask a generative AI model to explain to me "how a car works", it may give me an answer that is very well formulated, but wrong. When I ask it to draw me a car, it may forget to add the car doors. This is called "hallucination". For example, an American lawyer didn't realize that ChatGPT had provided him with examples of completely fictitious cases. Some of these imperfections will disappear with technological progress, while others will have no significant impact. Some will persist, notably when training data is biased or false. In sensitive areas, such as healthcare or law enforcement, the use of AI must therefore be carefully evaluated and supervised.

Risks of malicious use. Cybercriminals have not waited long to use generative AI to produce particularly convincing forgeries, just as their predecessors seized the car 100 years ago to escape after a robbery. However, there is nothing to suggest that AI will permanently change the balance of power between cybercriminals and those charged with protecting us, provided the latter are able to seize these technologies. Nor does generative AI appear to make the production of physical, chemical or biological weapons any easier than an online search[12].

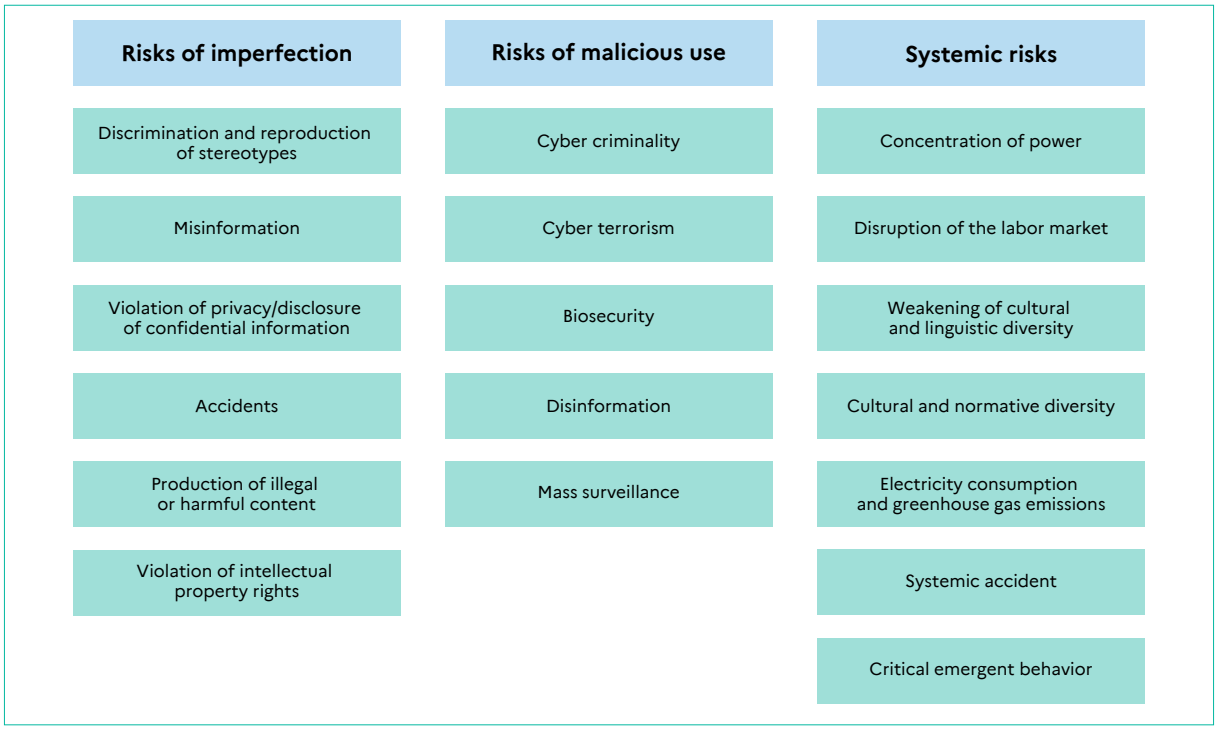

Risks associated with generative AI systems
Source: French AI Commission

12. Mouton, C. A., Lucas, C., and Guest, E. (2024) "The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study, *RAND Corporation.*

Systemic risks. The development of AI can be a source of risks for society as a whole, and even for humanity as a whole. While the risk that AI will result in massive job destruction seems limited, we must be prepared for certain professions to be significantly transformed, or even disappear (see 1.4 *AI: creator or destroyer of jobs?*). Other risks arise from the concentration of the development of the most advanced AI systems in the hands of a small number of countries, companies and individuals. They call for a determined response in terms of industrial policy and competition. On the other hand, no AI is yet capable of surpassing human intelligence in all tasks, let alone posing an existential threat to humanity as a whole. These prospects, which are still hypothetical and may never materialize, cannot constitute the alpha and omega of our approach to AI. They do, however, call for vigilance. It is therefore essential for our country to equip itself with the capacity to evaluate the most advanced AI systems, so as to anticipate the emergence of new risks, and with the governance to respond to them, both nationally and internationally.

1.3 WILL AI HELP US TO PROSPER?

> **Without doubt, because AI will make us more productive.** The scale of these gains and how they will be distributed across society are uncertain and not defined *a priori*. In 10 years, they could increase GDP from €250 to €420 billion, equivalent to the added value of industry as a whole..

Since the 'Thirty Glorious Years' of the post-war boom, the French economy, and more generally the economies of most developed countries, have experienced a decline in their growth rate. Some economists have concluded that this weak growth was inevitable, even speaking of "secular stagnation".

AI could help us return to high growth rates in our developed economies through two effects: (i) by increasing our productivity, i.e. the speed at which we produce goods and services; (ii) by increasing our ability to generate new ideas, and therefore new innovations, new products or new forms of organization.

AI SIGNIFICANTLY BOOSTS PRODUCTIVITY

AI can increase our growth potential by automating tasks in the production of goods and services. It thus contributes to increased productivity, as happened with mechanization in the agricultural sector, the invention of the assembly line in industry, or more recently the digitization of a significant part of the economy. Productivity gains on these tasks will contribute to higher growth rates.

A recent U.S. study[13] looks at the effects of generative AI adoption on the productivity of a company's customer service employees. The company progressively deployed an AI tool helping employees responsible for responding to customers via online chat by offering automatically generated responses. In this study, the effect of deploying this tool was substantial: the productivity of employees who had access to the AI assistant increased by 25%, including 14% in the first month of use. The effect is immediate and persistent over the five months of the study.

13. Brynjolfsson, E., Li, D., and Raymond, L. (2023), "Generative AI at Work", *NBER Working Paper.*

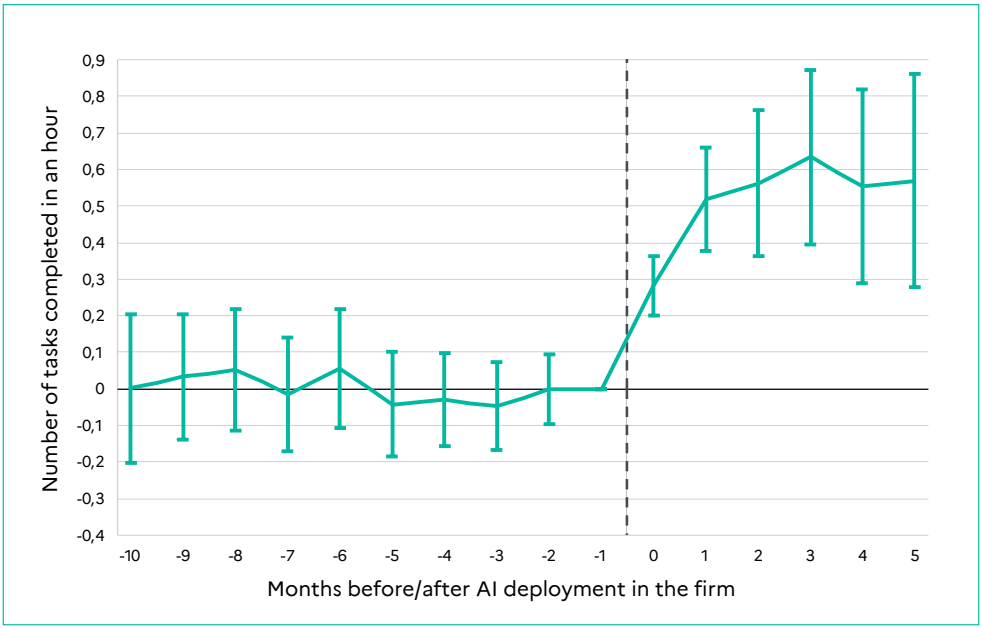

Figure 1: Effect of generative AI adoption on the productivity of employees in a customer service department.
Source: Brynjolfsson, Li and Raymond (2023)
Interpretation: Employees with access to AI see their productivity increase more than those without, while their productivities were evolving similarly in the 10 months prior to the introduction of AI.

Of course, these results only apply to a specific type of profession within a given company. However, two studies[14] focusing on highly qualified individuals in the USA (consultants, managers, etc.) show that using ChatGPT increases the productivity of tasks typical of these professions by between 25% and 35%.

This seems to show that productivity gains are observed across a wide range of occupations, with different levels of qualification. These effects are expressed by the workers themselves. In France, a wide-ranging survey by Pôle Emploi, "*Les employeurs face à l'intelligence artificielle*" (June 2023[15]), also highlights a positive effect on productivity: 72% of employers using AI mention a positive impact on their employees' performance, in particular by making it possible to reduce tedious tasks (63%) or the risk of error (51%).

If we look beyond a single company to estimate the overall impact of AI on the economy, two questions arise: when will we see the economic gains of AI, and how big will they be? We can draw a parallel with the effects on productivity of past technological revolutions. In the United States, as in Europe, in the case of electricity, productivity gains materialized some twenty years after the technology was invented.

To understand this time lag, let's look at the case of electricity. At the beginning of the 20th century, the adoption of electricity in business was still limited. Factories maintained an internal organization similar to the one they had adopted when powered by water mills: with a central

14. S. Noy, and Zhang, W. (2023), "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", *Science*. // Dell'Aqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani (2023), "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality", *Working Paper.*
15. Survey conducted by Pôle Emploi among 3,000 establishments with 10 or more employees in June 2023. The study does not specify whether the AI used is generative or non-generative, whereas American studies focus on generative AI.

drive shaft turned by hydraulic power. Neither the advent of steam power with the first industrial revolution, nor that of the dynamo at the start of the second industrial revolution, led to any change in the internal organization of factories. However, the presence of this drive shaft meant that similar machines had to be placed side by side. We had to wait until the 1910s to see the productivity gains linked to electricity, thanks to the invention of the electric wire and the miniaturization of electric motors. Each machine became autonomously powered by electricity. This innovation eliminated the drive shaft and enabled machines to be arranged more efficiently: the invention of the assembly line marked the rise in factory productivity.

The time lag will also exist for AI, as its adoption will require both a change in work organization within companies and additional investment. However, the time lag should be less marked with generative AI, which is much easier to deploy in the economy, and applies well to service jobs in our economy. The video game industry, for example, deploys generative AI rapidly to generate the outline of a game in 2 months instead of 6 months, and generate many different ideas for a character. What's more, the proliferation of AI innovations dates back ten years or so. So we could soon start seeing productivity gains. .

This leads us to a second question: how big are the economic gains to be expected? If we consider that the productivity gains enabled by the AI wave over the next decade will be comparable to those of the electricity wave in the 1920s in Europe, then productivity growth would increase by 1.3 percentage points per year from 2024[16].

If we prefer to take the digital technology wave of the late 1990s and early 2000s in the USA as a point of comparison, the increase in productivity growth would be around 0.8 percentage points per year. By comparison, France's potential productivity[17] is now estimated at 0.5% per year over the medium term. A closer look at the growth differential between France and the United States over the period 1997-2006 reveals that it is the sectors producing or heavily using digital technologies that account for almost all of the observed gap: the development and deployment of a new technology as decisive as digital technology is the main factor explaining the prosperity gap between the United States and France over this period.

All other things being equal[18], the gains generated by AI would significantly boost France's growth rate, estimated at 1.35% per year in the medium term. Such productivity gains over ten years would lead to an increase in GDP of between €250 billion and €420 billion in 2034, equivalent to the added value of industry as a whole!

On the other hand, this increase in growth would only be transitory: once AI has been adopted by the entire economic fabric, the productivity gains linked to this adoption and the transformations engendered cease, as shown in the graph below.

16. Estimate based on data from Bergeaud, A., Cette, G., and Lecat, R. (2016), "Productivity Trends in Advanced Countries between 1890 and 2012", Review of Income and Wealth, 62(3), pp 420-444. *Review of Income and Wealth*, 62(3), pp 420–444.
17. Stability program for the period 2022-2027.
18. By transposing these productivity gains to potential GDP; without a corresponding change in employment, a point discussed in detail in 1.4. *AI: creator or destroyer of jobs?*

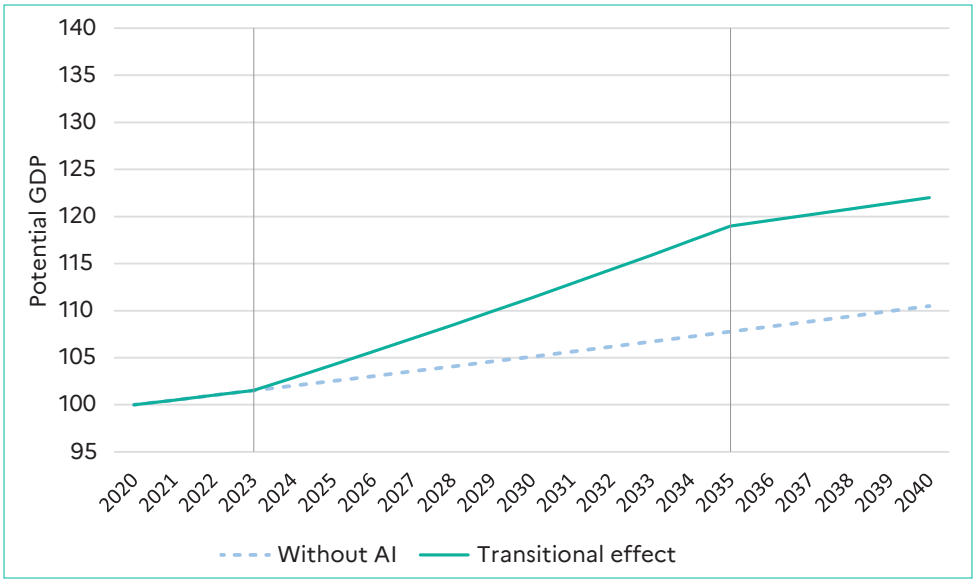

Figure 2: Expected transitional effect of AI adoption on growth
Source: French AI Commission

However, this prediction will be considered too pessimistic for some, too optimistic for others. The former will argue that AI can also automate the production of ideas, thereby generating additional growth, and this on an ongoing basis. The latter will point to the existence of obstacles to growth, notably the lack of competition in various segments of the AI value chain.

AI SHOULD ALSO MAKE IT EASIER TO GENERATE NEW IDEAS

AI could automate the generation of new ideas, or at least make it easier. It will thus help us generate new inventions and solve complex problems, as in the example of AlphaFold, which helps find new proteins[19], or GNoME, which suggests new materials that could be used in our vehicles or everyday objects[20].

The impact of AI on science and innovation is difficult to quantify. All the more so as AI's ability to generate new ideas could come up against practical difficulties: it's not enough to identify 2.2 million potential new materials to produce them, they still need to be validated experimentally. At the very least, AI will make researchers' work easier. If AI tools gradually support humans in identifying new hypotheses, creating protocols and carrying out experiments, then the production of relevant ideas will increase. This uncertain prospect is, however, on the horizon. Nearly one in ten research articles already mentions the use of AI[21].

So what effect would AI's ability to generate new ideas have on prosperity? Once again, let's draw a historical parallel to illustrate how an innovation can have a long-term impact on the rate of productivity growth. In the 17th century, the invention of infinitesimal calculus enabled gigantic advances in physics, notably in understanding the movements of projectiles or planets. Similarly, advances in glass polishing techniques led mankind to see smaller and smaller, and to

19. Jumper, J., Evans, R., Pritzel, A., et al. (2021) "Highly accurate protein structure prediction with AlphaFold". *Nature 596*
20. Merchant, A., Batzner, S., Schoenholz, S.S. et al. (2023) "Scaling deep learning for materials discovery". *Nature 624*
21. Analysis of the *Nature* journal on the Scopus database.

discover the previously unknown world of microbes. In this way, the microscope has enabled crucial advances in medicine. In the same way, AI opens up a field of possibilities difficult to imagine. These effects lead to a permanent increase in the rate of productivity growth. The magnitude of this effect, however, is impossible to quantify.

Let's now project ourselves into a future that combines the transient effects on growth of automating the production of goods and services, with the permanent effects on growth of automating the production of new ideas. This is the blue scenario in the graph below, with potential gains even greater than the €500 billion by 2034 presented above.

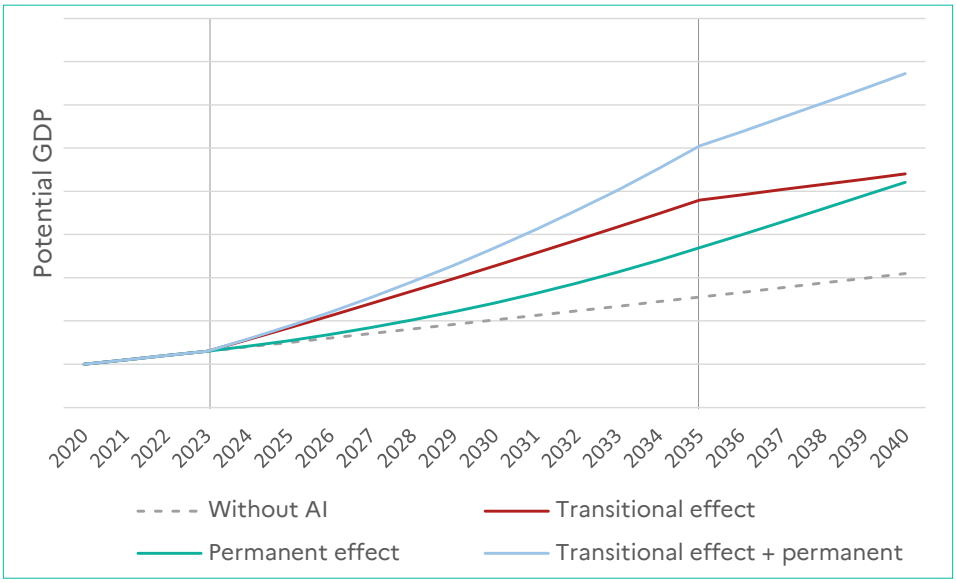

Figure 3: Total expected impact of AI adoption on growth
Source: French AI Commission

THE IMPORTANCE OF INSTITUTIONS AND COMPETITION

Aren't we being over-optimistic in our growth predictions? After all, the digital revolution, too, was supposed to lead to accelerating growth. Yet since the early 2000s, developed countries, starting with the United States, have experienced a sharp decline in their growth rates. How can we explain our low growth rates despite major innovations that have considerably changed our daily lives: the computer, the smartphone, social networks, etc.?

For some economists, this is simply a measurement problem, while for others it's a sign that these digital innovations have above all enhanced our entertainment[22]. A more convincing explanation is that the information and communication technology (ICT) revolution has fostered the emergence of "superstar" companies, notably the GAFAMs (Google, Amazon, Facebook, Apple, and Microsoft). Born of the ICT revolution, the GAFAMs initially contributed to the observed increase in productivity growth during the decade 1995-2005. However, an overly lax competition policy allowed the GAFAMs to grow to the point of controlling entire sectors of the US economy, ultimately discouraging the entry of new, innovative companies, with negative effects on the growth of the economy as a whole[23].

The difference between the ICT revolution and the AI revolution is that this time, the GAFAMs are dominant from the outset, and can therefore immediately discourage the entry of new, innovative companies. The lack of competition is particularly pronounced in the upstream segments of the AI production chain, namely access to data and access to computing power: these segments are dominated by a small number of giants, including the GAFAMs. Hence the importance of adapting our institutions, and in particular our competition policies, so that the AI revolution can fully act as a growth driver (see 2.3.3. *Avoiding dominant competitive positions*).

CONVERSELY, A RISK OF ECONOMIC AND GEOPOLITICAL DOWNGRADING

The deployment of artificial intelligence systems in our economy and society is not without risks. So what would happen if France and Europe chose to remain on the sidelines of this technological revolution? First of all, it must be stressed that it is virtually impossible to completely prevent the spread of AI systems in society. In particular, it would mean isolating our country from the global flow of information and imposing very strict controls on the population. With the exception of North Korea, few countries are taking this route.

So if total isolation is unenviable, what would happen if the spread of innovation were held back? To answer this question, let's turn to the history of industrial revolutions. China is a case in point. For centuries, the Middle Kingdom was by far the world's leading power. From the second half of the 17th century, the Ming and Qing dynasties chose to oppose international trade for fear of new technologies and their effects on the economy and society[24]. State control over society and the economy was tightened, even going so far as to order inhabitants living along the southern coast to move 18 miles inland.

22. As investor Peter Thiel said of social networks in 2022: "We wanted flying cars, instead we got 140 characters".
23. Aghion P., Bergeaud, A., Boppart, T., Klenow, P., and Li, H. (2023), "A Theory of Falling Growth and Rising Rents", *Review of Economics Studies,* 90(6).
24. Acemoglu D. and Robinson, J. (2012), "Why Nations Fail: The Origins of Power, Prosperity, and Poverty", *Crown Business.*

This isolation from innovation changed the course of China's history. By the middle of the 17th century, per capita GDP in China was only 10% to 20% lower than in the UK or France[25]. Then China did not benefit from the effects of the industrial revolutions unfolding in Europe. Two centuries later, per capita GDP in China was five times lower, and ten times lower in 1900 (Figure 4). The economic stagnation led to a collapse in the Empire's sovereignty, with foreign powers imposing unequal treaties, the "rush on China".

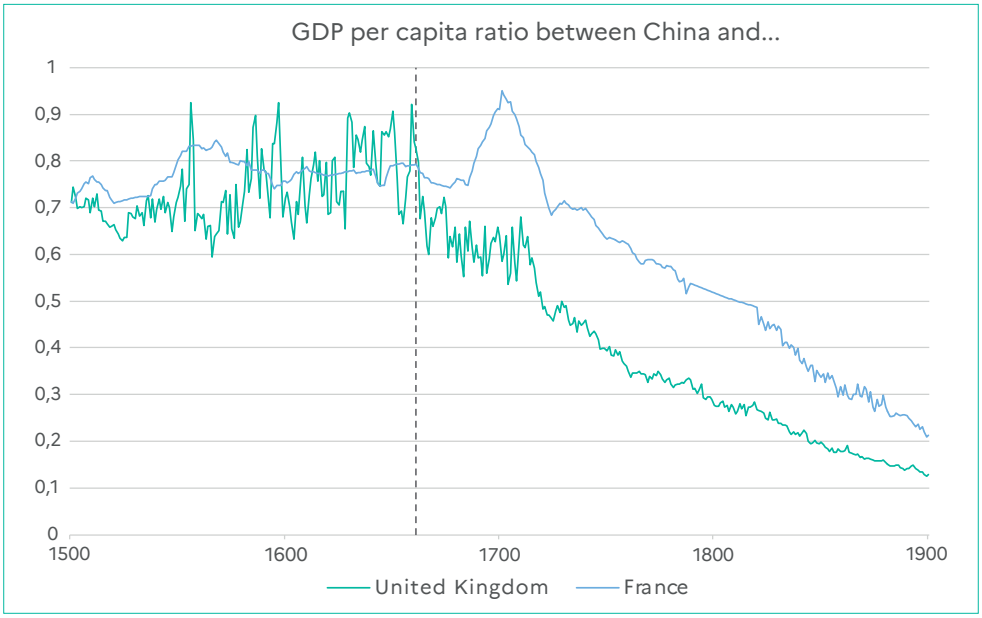

Figure 4: Comparison of China's GDP per capita with that of the United Kingdom and France
Source: French AI Commission based on data from the 2020 Maddison Project Database.

Numerous other examples show that failure to participate in a technological revolution leads to economic and societal downgrading, which hastens foreign domination and fuels international greed. Preventing this deleterious trajectory does not require us to seize AI with no compass, but rather to organize ourselves to master it and steer it in line with our political objectives.

25. Based on 2020 data from the Maddison Project Database, the reference database for historical GDP estimates.

1.4 AI: CREATOR OR DESTROYER OF JOBS?

> **Our own empirical analysis suggests a positive effect of AI on employment in companies that adopt AI, as it replaces tasks, not jobs.** In 19 out of 20 jobs, there are tasks that AI cannot perform. Jobs that can be directly replaced by AI would therefore represent just 5% of jobs in a country like France. On the other hand, the spread of AI will create jobs, not only in new professions, but also in old ones. All in all, some sectors or fields could experience net job losses, which need to be supported by the public authorities, but this does not imply that AI will have a negative effect on national employment in France.

Hollywood screenwriters and actors fearing for their jobs, a French media monitoring company triggering a job-saving plan under the pretext of AI: news and pronouncements are raising fears of the end of work and mass technological unemployment.

Artificial intelligence, in particular, enables the automation of tasks, which is a key driver of economic growth (see *1.3. Will AI help us to prosper?*), and implies two opposing effects on employment. On the one hand, automation displaces certain tasks from human labor to machines, which tends to destroy jobs: this is the **crowding-out effect**. On the other hand, automation increases the productivity of individuals, leading to an increase in the value for money of products offered to consumers, and therefore to higher demand and, ultimately, more hiring and the creation of new tasks: this is the **productivity effect**.

To understand the effect of a new technology on employment and the labor market, we need to understand which of these two effects prevails. To this end, two main approaches have been adopted. Firstly, by directly studying the effects of AI adoption within companies or sectors. Secondly, by studying the effects of AI on the different tasks making up the economy.

FIRST APPROACH: THE EMPLOYMENT EFFECTS OF AI ADOPTION BY COMPANIES

The approach is not new, and has already been used to measure the employment effects of other technological revolutions. For example, a study[26] based on French data and focusing on the adoption of industrial machinery (machine tools, robots, etc.) shows that companies that adopt more of this type of new technology lower their prices, increase their sales and employ more people than their competitors who have not adopted automation technologies.

What about AI? A survey carried out annually by Insee studies the effects of AI adoption by companies in France[27]. We find that total employment in companies that have adopted AI is increasing more than in companies that have not, whereas these two groups were following a similar previous trend (Figure 5). The effect is mainly due to the creation of new jobs, rather than a greater retention of existing ones. We also note that there are no differentiated effects on jobs held by men compared to those held by women: the adoption of AI has comparable effects for male and female employment.

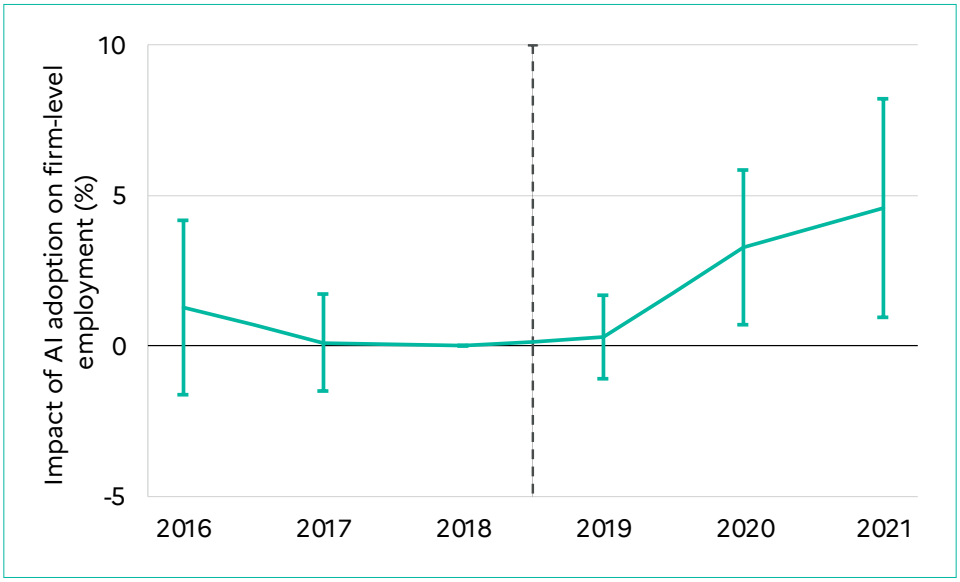

Figure 5: Effect of AI adoption on total employment within companies in France
Source: French AI Commission
Interpretation: Companies adopting AI are employing more people than those not adopting AI, whereas they were evolving similarly in the previous 3 years.

However, this effect of AI on total employment is not uniform from one profession to another. In particular, certain levels within the company or certain professions are likely to experience net job reductions. Indeed, Figure 6 shows that companies that adopt AI for administrative management or marketing see their employment in "intermediate administrative and commercial professions" decline.

26. Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2023), "Capital industriel moderne, demande de travail et dynamique des marchés de produits : le cas de la France", *Insee Working Paper.*
27. Among companies that hadn't already adopted AI in 2018, we can compare employment trends between the 321 companies that did so between 2018 and 2020 and the 897 companies that didn't.

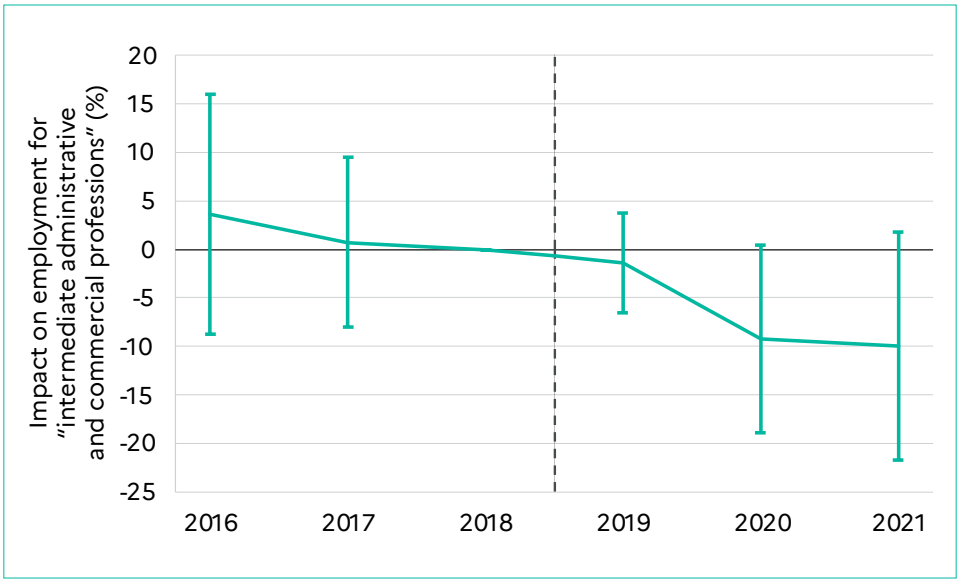

Figure 6: Effect of adopting AI for marketing or administrative management on the employment of intermediate administrative and commercial professions within companies in France
Source: French AI Commission
Interpretation: Companies adopting AI for marketing or administrative management are reducing their employment more than those not adopting it, whereas they were evolving similarly in the 3 previous years.

These results are new and important, but they are not sufficient to capture all the potential effects of AI on the labor market. In particular, it is likely that an innovative company that adopts AI will become more productive than companies in the same sector that do not, and that as a result the innovative company will gain market share at the expense of competing companies that have not adopted AI. In terms of employment, this will result in a net creation of jobs within the innovative company, but to the detriment of companies that have not adopted AI: these competing companies are likely to suffer job losses. All in all, what is the aggregate effect of AI on employment when we consider the labor market as a whole?

A first element of the answer to this question is provided by recent studies, focused on the USA, which analyze the effects on employment of the adoption of non-generative AI by companies. These studies suggest a positive effect of AI on employment. In particular, a US study[28] published in 2023 suggests that AI adoption is associated with increased employment and sales at the sectoral level, not just at the level of the company converting to AI.

So far, we've focused on the effects of non-generative AI. But what's new with generative AI is that certain knowledge, strategy and creativity professions (doctors, teachers, lawyers, journalists, artists, etc.), once seen as crucibles of human intelligence, could be affected by a reduction in the total number of jobs.

Recent work, carried out at a very microeconomic level, focuses precisely on the effects of adopting generative AI. In particular, the study by Brynjolfsson et al.[29] (see *2.3. Will AI help us to prosper?*) focuses on the case of the adoption of a generative AI-based tool to assist customer service agents. While this study does not conclude that there is an effect on total employment, it does highlight the fact that the probability of an employee quitting their job in the current

28. Babina, T., Fedyk, A., He, A., and Hodson, J. (2024), "Artificial Intelligence, Firm Growth, and Product Innovation", *Journal of Financial Economics*.
29. Brynjolfsson, E., Li, D., and Raymond, L. (2023), "Generative AI at Work", *NBER Working Paper*.

month decreases by 8.6 percentage points after AI adoption, while at the same time increasing employee productivity. The effect is even stronger for employees who have only recently joined the company (within the last six months). All in all, this study suggests a positive effect of generative AI on employment.

However, another study[30] shows that the arrival of ChatGPT has had a negative effect on the employment and remuneration of American freelance workers The study focuses more specifically on a platform offering assignments targeting small or medium-sized tasks within several professions (data entry, graphic design, software development, marketing, etc.) and highlights a decline in the number of jobs and income, including for workers with more experience, following the launch of ChatGPT in November 2022.

To sum up, these initial studies suggest that the productivity effect dominates on average for employees in companies, while the crowding-out effect seems to be greater for self-employed individuals who have to perform mostly tasks that can be more easily replaced by AI.

Beyond the effect on employment, what about the effect on inequality? The first series of studies focusing on non-generative AI[31] highlights the fact that companies adopting AI subsequently hire more highly educated and technical profiles, particularly in favor of so-called "STEM" jobs (science, technology, engineering and mathematics). This would lead to a rise in inequality, as these complementary profiles of AI adoption have higher-than-average salaries. However, the second wave of studies on generative AI[32] overturns this view by highlighting the fact that the employees who are initially the least qualified or least productive are those for whom the productivity gains offered by the use of AI are the greatest. This could then give these employees leverage to renegotiate their pay upwards, and thus reduce inequalities within companies.

Finally, the direct approach of comparing, *a posteriori,* companies or sectors that have adopted AI with those that have done so little or not at all, does not allow us to look back very far in time, particularly in the case of generative AI: the release of the first mass-market generative AI tool dates from November 2022, and these studies do not yet allow us to conclude on an effect over a horizon of a few years. In addition to the direct approach, it is also worth taking a more forward-looking, task-based approach.

30. Hui, X., Reshef, O., and Zhou, L. (2023), "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market", *Working Paper*.
31. Babina, T., Fedyk, A., He, A., and Hodson, J. (2023), "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition", *NBER Working Paper*.
32. Brynjolfsson, E., Li, D., and Raymond, L. (2023), *op. cit.* // Noy, S., and Zhang, W. (2023), "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", *Science*.

SECOND APPROACH: THE EFFECTS OF AI ADOPTION ON TASKS

Like the previous one, this approach predicts a positive effect of AI on employment. A recent study by the International Labour Organization (ILO)[33] shows this. It gives each task performed in the economy a probability of being replaced by AI. The study distinguishes between two cases. If a significant proportion of the tasks that make up an occupation can be performed by AI, then that occupation has the potential to be replaced by AI. On the other hand, if an occupation is made up of a few tasks that can be automated, but a majority of tasks that are difficult to automate, it has potential for improvement by AI: automating some tasks frees up time for others. For example, office workers would face the risk of being replaced, while managers would have potential for improvement.

The study concludes that, worldwide and including in developed countries, the number of jobs with high augmentation potential (13.4%) is much higher than that with high automation potential (5.1%). The study also notes gender inequalities: 3.5% of jobs mainly held by women have high automation potential, compared with 1.6% of jobs mainly held by men[34]. A study by the International Monetary Fund (IMF)[35] adopting a relatively similar methodology comes up with higher figures: 60% of jobs would be highly exposed to AI, with around half of these jobs able to benefit from high augmentation potential, while the other half would face high automation potential.

Regarding gender inequalities, the study notes that in developed countries, women are more exposed than men, both for jobs with the potential to be replaced by AI and for those with the potential to be enhanced by AI. It concludes that women face both greater risks and greater opportunities in the face of AI.

A recent note[36] adapts the ILO methodology to the case of France in order to provide an overview of the expected effects of AI on the 222 occupations listed by the Ministry of Labor's nomenclature. The vertical axis (see Figure 7 below) represents an occupation's exposure to AI in general: the higher an occupation is on this graph, the more exposed it is to AI overall. The horizontal axis represents the proportion of tasks deemed difficult to automate: the further to the right an occupation is on this graph, the greater the proportion of tasks that cannot be easily replaced by AI. The size of the circles is proportional to the number of people working in that occupation.

The occupations at top left are the most vulnerable to AI: both highly exposed and with few tasks that are very difficult to replace by AI. Those at top right combine high exposure to AI with a high proportion of tasks unlikely to be replaced. We can therefore expect major transformations in these jobs: some better-trained workers may be able to take advantage of the new possibilities offered by generative AI, while other workers may face greater competition, thus increasing pay inequalities within these professions. Occupations in the lower part of the graph, particularly in the lower right-hand quadrant, appear to be largely unaffected by the impact of artificial intelligence.

33. Gmyrek, P., Berg, J., and Bescond, D. (2023), "Generative AI and jobs: A global analysis of potential effects on job quantity and quality", *ILO Working Paper* 96.
34. For example, customer service, administrative and communication professions are highly feminized.
35. Pizzinelli, C., Panton, A., Mendes Tavares, M., Cazzaniga, M., and Li, L. (2023), "Labor Market Exposure to AI: Cross-country Differences and Distributional Implications", *IMF Working Paper*.
36. Bergeaud, A. (2024), "Exposure to generative artificial intelligence and employment: an application to the French socio-professional classification", *Working Paper.*

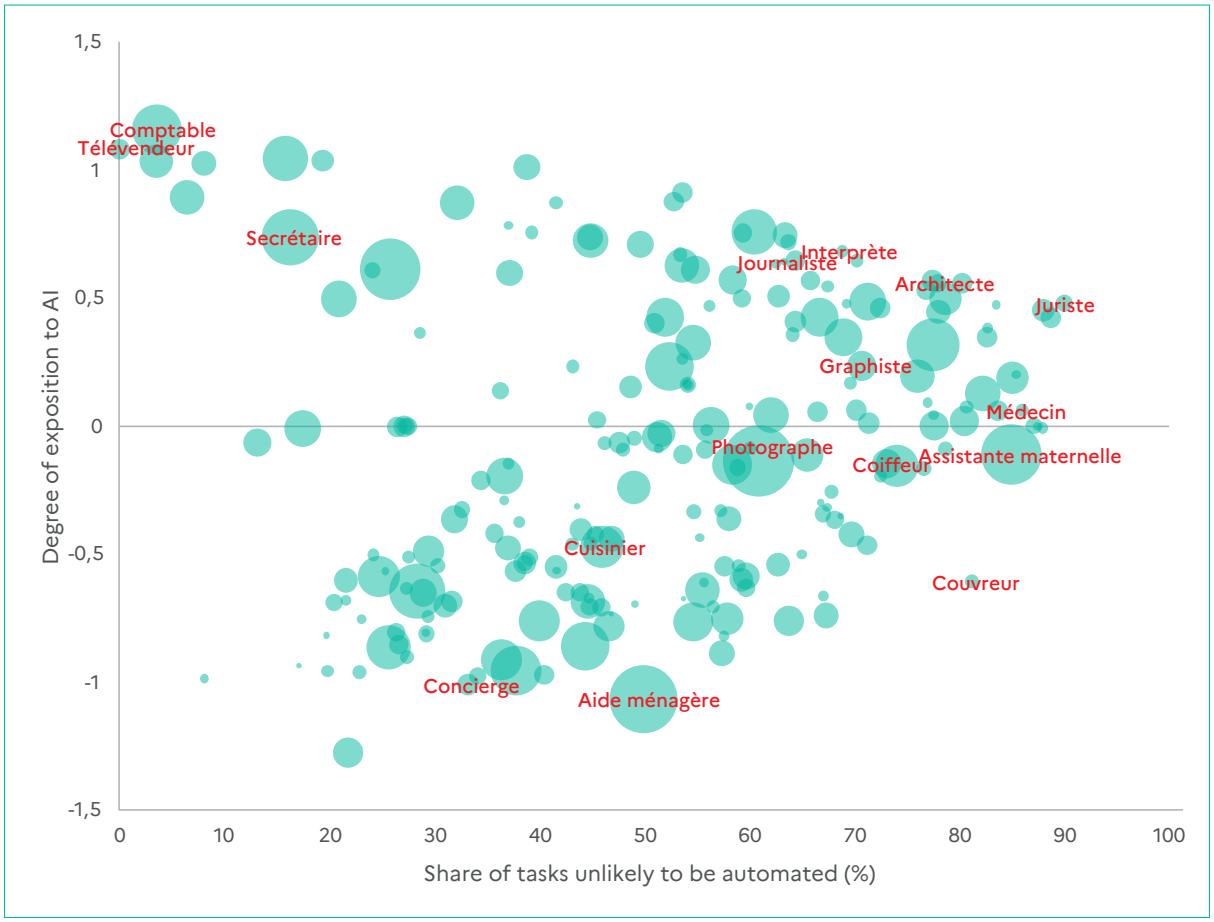

Figure 7: Expected impact of AI on professions in France
Source: A. Bergeaud (2024)
Interpretation: The higher up the graph the professions are, the more exposed they are to AI. The further to the right of the graph, the lower the proportion of their tasks likely to be replaced.

This approach, based on task exposure to AI, has the advantage of making it possible to estimate aggregate effects at the level of the economy as a whole, and to enable comparisons between countries. However, it has several limitations. Here are the two main ones. On the one hand, it is a static approach: the studies are based on existing tasks, and therefore take no account of tasks that might be created as a result of the development of AI. To draw a parallel, it was hard to imagine that the job of data scientist could take on such prominence at the start of the digital revolution in the 2000s, or that the job of household appliance engineer could exist before the mass adoption of electricity in the mid-20th century. On the other hand, this approach is based on an estimate of the probability of AI replacing various tasks (see box).

Why was a 2013 study on the effects of automation wrong?

A study[37] published in 2013 caused quite a stir when it concluded that 47% of American jobs would be threatened by automation in the next ten to twenty years. Ten years on, we have to admit that this prediction was wrong.

The study was based on the subjective idea that 70 occupations would be fully automated, i.e. that all the tasks performed by these occupations could be carried out by machines. Then, based on the characteristics of these occupations (types of tasks performed, skills required), the study calculated a probability of automation of between 0 and 1 for the remaining 632 occupations. Choosing a threshold of 0.7 to conclude that an occupation was threatened by automation, the authors deduced a figure of 47% of jobs threatened.

Several errors crept into this line of reasoning. The first was to place too much emphasis on subjectivity, and to jump to the conclusion that certain professions, such as truck and cab drivers, will be automated, betting on the rapid development of autonomous cars as the perfect alternative to drivers. The second was to confuse exposure to automation with the risk of replacement by automation. The aforementioned ILO and IMF studies avoid both these pitfalls.

We can see that the two approaches converge on similar conclusions: all in all, the deployment of AI in the economy should have a positive overall effect on the number of jobs. Catastrophic forecasts about the end of work are no more credible than similar predictions made in the past. Especially as even the task-based approach represents an upper bound for the impact of AI, since it makes the assumption that it is profitable to automate all automatable tasks. But this assumption is far from being true today. The falling cost of AI systems and the possibility of distributing the same AI system to a very large number of users will be key factors in determining the impact of AI on tasks and jobs.

Moreover, this overall effect will cover a variety of situations. Some professions may experience net job losses. In culture and media, for example, the proportion of jobs exposed to AI is higher than in other sectors. For the majority of workers, this will involve changes in skills and tasks. For others, high exposure to AI means high complementarity with AI, leading to new specializations, new or enhanced forms of expression or technical skills. However, we must also expect certain professions to disappear or to see their numbers significantly reduced. The complexity of situations and professions makes it impossible to provide a uniform, global response to the challenges of AI. More precise studies are needed to address the variety of sectors, value chains and take statuses into consideration. Beyond this, it is essential to organize access to and development of initial and lifelong training.

37. Frey, C. B., and Osborne, M. A. (2013), "The future of employment: How susceptible are jobs to computerisation?", *Oxford Martin School Working Paper.*

1.5 WILL AI DEGRADE OR IMPROVE QUALITY OF WORKING LIFE?

> **AI can increase quality of working life, even for middle-class workers.** Some AI users say they are more fulfilled and more successful[38], because they can get rid of routine tasks and improve the quality of their work. However, there are risks (surveillance, discrimination, increased stress, etc.). The consequences of AI on quality of working life will depend on our collective choices and the quality of social dialogue with regard to it.

The impact of AI on quality of working life can be seen from three perspectives: job content, the risk of losing one's job, and working conditions.

As far as job content is concerned, there is no doubt that the majority of professions will evolve. Most tasks will be transformed, others will be eliminated, and new ones will appear[39]. These transformations may concern relatively ancillary, low value-added tasks (such as a manager writing a job description for a recruitment process), as well as the high value-added tasks that make up the core business.

Within the same profession, productivity gains so far seem to benefit the least productive workers. For example, the introduction of AI to help cab drivers find customers by suggesting routes on which demand will be high increases the revenue of the least skilled drivers, without improving that of the most experienced. This result is also found in the case of customer service employees[40] and consultants in a large consulting firm[41].

In a way, where industrial machines and "classic" computer programs are a compendium of explicit knowledge, an AI system can be a compendium of implicit knowledge. Machines and programs follow well-defined rules, and they have led to the automation and disappearance of professions that consisted in following these rules precisely (pouring such and such a quantity of material when the temperature reaches such and such a threshold, cutting a steel board in a

38. OECD (2023) "Employment Outlook": 63% said enjoyment in their job improved and 80% reported improved results
39. See in particular the ILO and IMF study cited in section 1.4.
40. Brynjolfsson E., Li, D., and Raymond, L. (2023), "Generative AI at Work", *NBER Working Paper*.
41. Dell'Aqua F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani (2023), "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality", *Working Paper*.

regular way, etc.). AI can detect regularities where we don't know how to explain them. It can, in certain cases, give every worker access to the intuition of the best expert, even when the latter would be unable to explain where his intuition comes from, and therefore to pass it on. AI can therefore not only relieve routine and irregular tasks, but also enrich work and improve the quality of work that anyone can do independently.

Conversely, misuse of AI systems can also introduce mental overload, which can lead to cognitive exhaustion if the time freed up by the machine results in stress for the worker and an excessive increase in complex tasks. These risks call for reflection on the redistribution of the gains enabled by AI, which could in particular enable a reduction in working hours and a better work-life balance.

The risks also justify the establishment of governance rules for the interaction between workers and machines. The draft regulation on artificial intelligence (AI Act) specifically identifies AI for recruitment purposes, including the filtering and selection of applications, as a high-risk application, and lays down strict obligations for such applications. High-risk AI systems will also have to be designed in such a way that individuals can control their operation: the human operator will have to remain in charge of the final decision.

What's happening in the rest of the world? The example of the United States[42]

In the United States, the use of AI by employers is regulated by various federal-level legislations, which govern the use of AI in recruitment processes and aim to combat discrimination. For example, Law 144 in New York, which came into force in July 2023, requires audits to be carried out to verify the absence of bias in AI tools for recruitment, and that applicants be informed about the use of such tools.

On the other hand, we are only beginning to better understand the different interactions of workers and work with AI and the conditions increasing workers' skills and competencies and improving work organizations. The deployment of AI systems in organizations is not the end point of innovation processes, but rather a new starting point[43].

When it comes to keeping one's job, the spread of AI has its advantages. It can actually have a positive effect on worker envy: employees in customer service and qualified service occupations[44] (consultants, managers, etc.) who use AI are more likely to stay in their jobs, as AI increases their positive feelings about their work. However, surveyed consultants and managers using AI also have a higher fear of losing their jobs. They see for themselves the technology and the tasks it will be able to perform for them. Added to this is the fear of seeing their company disappear if their sector is disrupted.

As regards working conditions (work organization, managerial practices, labor relations, compensation, health and safety), there is little doubt that the deployment of AI will lead to the emergence of new forms of organization and coordination, like previous technological

42. Source: French Treasury Department, Economic Service in the USA.
43. LaborIA (2023) "Study of the impacts of AI on work. General synthesis of the LaborIA Explorer survey report".
44. Noy S., and Zhang, W. (2023), "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", *Science*.

revolutions. The steam engine concentrated industry in a few factories, bringing large numbers of workers together in one place. The electric motor enabled the installation of numerous lines and facilitated the organization of assembly-line work. Communication technologies facilitated telecommuting and the organization of international production chains, notably through the relocation of activities.

In the field of occupational health and safety, advances in artificial intelligence open up interesting prospects in epidemiology and accident research, as well as new possibilities for supervising a work environment, a construction site, or an industrial site, for example, notably through the use of predictive maintenance[45].

Conversely, misuse of AI tools can exacerbate occupational and psychosocial risks. The development of "algorithmic management" could lead to a loss of autonomy in the workplace, dehumanizing subordination to the machine, excessive surveillance of workers, isolation of workers, and a loss of a sense of community.

One major unknown remains the impact that AI will have on salaries. Because technologies not only transform professions, they also change bargaining power and the value of certain expertise, which becomes less necessary or less rare. The industrial revolution facilitated the production of certain goods, thereby reducing the wages of craftsmen. On the other hand, it created new needs, particularly for those capable of operating machines and organizations. These new jobs were at the heart of the "middle class" in industrialized countries. In turn, information and communication technologies have revolutionized working conditions. On the one hand, they have facilitated the automation of certain routine manual or cognitive tasks, thereby threatening certain middle-class jobs. On the other, they have facilitated the circulation and processing of information, increasing the centrality of the most qualified people and concentrating decision-making power in their hands. The last thirty years have thus seen a polarization of wages and, more broadly, of working conditions, including for people with higher education qualifications.

Will AI extend the same dynamic? Initial results suggest the opposite, since they show that generative AI increases the productivity of the least qualified. Beyond the empirical results, it's a hope that we can formulate[46]: by condensing intuition and implicit knowledge, AI could help workers with a certain amount of experience and basic training to carry out higher value-added work. This would replenish the ranks of jobs offering good working conditions. It would not, however, destroy the value of expertise. After all, just because anyone can buy plumbing and heating tools doesn't mean we don't lack those skills. A tool only has value when combined with certain skills.

The positive consequences of AI for quality of working life are a scenario, not a forecast. They will depend on how it is introduced into organizations.

To steer the deployment of AI in a positive direction, a series of initiatives has recently been launched: commitment to the development of employment and skills on a regional scale in Hauts-de-France (Cité de l'IA), and on a national scale ("Perspectives IA"), development of support guides, self-diagnostic tools, etc. But many authorities, industries and companies have yet to take the time to analyze the consequences of AI on production processes and their work organization. The studies carried out by LaborIA[47] and other research teams are valuable, but still too few in number.

45. INRS (2022) "L'intelligence artificielle au service de la santé et de la sécurité au travail. Challenges and prospects for 2035".
46. Autor, D. (2024), "Applying AI to Rebuild Middle Class Jobs", *NBER Working Paper.*
47. A joint initiative of the Ministry of Labor, Full Employment and Integration and INRIA created in 2021.

Recommendation No. 2

Invest in observation, studies and research into the impact of AI systems on the quantity and quality of employment.

To avoid certain negative consequences, the legal framework defines an inescapable foundation of rights (labor law, personal data protection law, etc.), which for the moment appears sufficient to ensure a worker-friendly deployment of AI. With the current framework, the French Data Protection Authority (CNIL) was thus able to fine Amazon €32 million for "employee monitoring". On January 23, 2024, it considered that measuring every few seconds' interruption of employees' scanning or measuring the speed of scanner use when putting things away was excessive, even in view of the delivery time stakes. The same framework could help prevent any abusive "algorithmic management". Above all, we need to ensure that it is effective.

The rapid development of AI systems, particularly generative AI, in work organizations, as well as the unsupervised use of "algorithmic management", will have to be addressed by the social partners and through specific social dialogue in companies and administrations. In addition, labor inspectorates need to be modernized and strengthened. Training must be organized for its staff, as well as for those responsible for preventing health and safety risks in the workplace.

Recommendation No. 3

Make social and professional dialogue a tool for co-constructing the uses and regulating the risks of AI systems.

1.6 IS AI A THREAT TO ARTISTIC CREATION?

> **Like other technologies before it, AI is being integrated into creative processes to serve human creation.** However, it is also undermining the sector, as it poses a huge challenge to designers, their skills, careers and remuneration.

From Hollywood to Paris, from Europe to India, creators and their organizations are worried. Generative AI models suggest that artistic creation is in danger, so rapid has their adoption been for text, image and video generation. The creative process is faced with the challenge of AI which, from query to query, asks software (Dall-E, Midjourney, Adobe Firefly, etc.) to refine a created text or image, to create or summarize text, or even to write or depict "in the style of...".

This spraying of "creative AI" seems to pose an immense challenge for creators, just as the democratization of personal publications on the Web alongside so-called traditional media did two decades earlier. This earlier mutation made the boundaries between amateur and professional less watertight. AI, particularly generative AI, is also lowering the barriers to artistic creation. But they do more than that.

In terms of jobs, careers and income, the emergence of AI is not just a new form of competition. For many creators, authors, artists, translators or actors, it poses an existential challenge, which may result in a reduction in activity, or even job substitution. A drop in income is looming, particularly for supplementary income, which raises the question of the start of a creative career. Even if the movement is not general, it is precisely for this social issue that the effects of AI are being debated. As of now, without appropriate training, starting out and progressing in one's career seems a very delicate matter.

Furthermore, the use of protected human creations to train generative AI, which is then likely to produce competing content, raises questions of authorization and fair remuneration. Numerous lawsuits have been filed in the USA for the unauthorized (and therefore unpaid) use of copyrighted content when training generative AI.

However, AI does not jeopardize the originality of creation itself, nor its selection processes. The production and production capacity of AI systems is no more than a new material, accompanied by the software evolutions that have become customary for music and images. The panoply of creative techniques is expanding. Artists are already using AI, and have been for a long time in some cases, to stimulate thought and imagination, and enable expression that would not otherwise be possible. Some (currently few in number) have already chosen these paths of discovery: Robbie Barrat, Justine Emard, Gregory Chatonsky, Pierre Giner, Benoît Carré, to name but a few. From this point of view, training will be essential, so as not to leave out a major part of the creative sector.

In many fields — music, architecture, multimedia, etc. — generative AI does not represent a break with the past. It adds to the long history of adopting technological layers for composition, arrangement, design, and production. It is part of the creative process, which moves from all possible universes, through multiple selections and iterations, to a work of art. The use of AI can accelerate creation times, open up new spaces, and encourage, nurture and enhance creativity. What's more, the development of specific, singular and personal AI for artists could enable them to identify and hone their own style.

By reducing repetitive or low-value tasks, and lowering barriers to entry, AI enables us to concentrate on the most essential elements of the imaginary or its output, where the human element, its creation, becomes even rarer and more distinctive in a regime of "megabundance" of AI productions. This opens up two new avenues. On the one hand, there is the need to recognize human creation, to identify it, to recognize its irreducible character, which stems from its originality and deserves distinction and protection. On the other hand, fair remuneration of human creation, essential to the harmonious reception of AI in the cultural sector.

1.7 CAN AI HARM THE QUALITY OF INFORMATION?

> **Information, at the heart of democracy, is being disrupted by the rise of AI.** Action is needed to preserve trust in and quality of information.

Information is not just another good or service. It is an essential element of democracies, presupposes freedom of communication, and shapes opinion, judgment and citizenship. To ensure that the information space is trustworthy, democracies have designed, established and protected major responsibilities: for press publications (defamation, insults, etc.), for sources (journalistic presence and ethics, reliability and fact-checking, etc.), and for the readership.

However, AI has the potential to call into question the activity of media companies, and consequently the fundamental role they play in the production of reliable, pluralist information. New AI-only media, with no respect for these responsibilities, are already seeking to position themselves as competitors to "traditional" media. Sites are providing unreliable, AI-generated information, often bearing names designed to make it appear as if the content has been produced by journalists. In addition, plagiarist robots improperly use content published by traditional media to produce articles, without crediting their sources and without remuneration.

The democratic risk stems from technical issues and the availability of quality data: as the production of quality information weakens in a world where AI promotes the multiplication of media, there is a risk that AI models will become increasingly trained on posts from social networks or doctored news sites. Today, the risk is perceived as that of a vicious circle based on the progressive eviction of professionally produced and verified information in favor of low-cost information based on false information.

The polarization of information quality is also at work. On the one hand, we can see that the biggest news players (news agencies, world-famous publications, media groups) can enter into agreements with AI providers, or sue them for previous unauthorized use (in violation of copyright) for AI training purposes (e.g. Getty Images, New York Times, etc.) because they measure the economic value of their content and their news mission. On the other hand, many media cannot install specialized AI on their own content. The risk, therefore, is to accentuate a two-speed information system: on the one hand, high-quality, paid-for information from major media groups; and on the other, mediocre, sometimes inaccurate information based on generalist AI models.

Finally, AI reinforces the propagation of information that isn't information, and enables the creation of personal assistants, promoting the fragmentation of information and the creation of information "bubbles".

In the face of such challenges, the world's information players have set to work: creating charters (*Reporters sans frontières*), identifying images and texts produced by AI (Google, Meta and others), etc. Similarly, the fight against parasitism and copyright infringement, the identification of unreliable information and the fight against deepfakes all require ongoing investment. We also need to ensure the traceability of information, in line with the fundamental responsibilities of information in a democracy. These requirements are fraught with difficulties, starting with the interaction between the requirement for veracity and traceability, since content labelled as "AI-generated" is automatically seen as less credible[48]. Because they are difficult, they require collective investment.

48. Not all AI-generated content is fake, and not all fake content is AI-generated. Wittenberg, C., Epstein, Z., Berinsky, A. J., and Rand, D. G. (2023) "Labeling AI-Generated Content" *MIT Topical Policy Brief*

1.8 SHOULD WE SPREAD OTHER PEOPLE'S AI OR CREATE OUR OWN?

> **France and the European Union must mobilize to disseminate AI systems, but cannot simply use those developed elsewhere.** If we are to reap the benefits of AI and control its risks, we need to play an active role in the ongoing technological revolution.

t was the Babylonians who first used multiplication tables. Should we use their multiplication tables or create our own? This question seems far-fetched to us, and no one in France is thinking of announcing the creation of sovereign multiplication tables. How is AI different? AI systems are far from being a neutral technology, like multiplication tables, which are the same for everyone.

Firstly, AI models are effectively imbued with the data they are trained on, and the cultural referents present in that data. A French national will tell you that Clément Ader made the first airplane flight, while an American will reply that it was the Wright brothers. ChatGPT, on the other hand, will answer the Wright brothers, whether questioned in France or the USA: 93% of GPT-3's training data comes from English-language texts[49]. AI-based tools will have a growing influence on our society, so it's important to master their cultural references.

Secondly, the deployment of AI systems in sensitive sectors (defense, energy, research, etc.) raises a question of sovereignty. Knowing how models are trained, how they work, their weaknesses and strengths, is a prerequisite for deploying AI in these sectors with confidence.

Finally, if all government agencies, businesses and individuals in France use foreign tools, a growing proportion of our wealth will benefit these suppliers, worsening our trade balance in particular. This is what happened with the previous wave of technological innovation (computers, software, internet): our digital trade balance shows a deficit of at least €22 billion (in 2019, pre-COVID year[50]), compared with an overall French trade deficit, all sectors combined, of €23 billion!

49. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry G. and Askell, A., (2020) "Language models are few-shot learners" *Advances in Neural Information Processing Systems.*
50. Data from the "OECD Digital Transformation Toolkit".

On the other hand, it's from the very rapid deployment of AI tools that we'll derive a lot of value, for our public services as well as for our businesses. France cannot wait until it has European tools at its disposal to reap the benefits of AI, in particular the 1.5-point annual gain in productivity presented earlier (see *1.3 Will AI help us to prosper?*). In fact, it is unlikely that our continent will be able to fully master the technology in the medium term.

All in all, without opting for an autarkic approach to AI that would be detrimental, France and Europe must be players in the ongoing technological revolution and its value chain. In other words, let's not close ourselves off to AI created by others, but at the same time commit to creating our own differentiated and internationally competitive technological links and business models. We need to have a dynamic vision: as the years go by and our expertise grows, our sourcing can become increasingly European.

1.9 SHOULD AI REMAIN IN THE HANDS OF A FEW PLAYERS?

> **It cannot be a good thing for just a few companies to master AI.** Supporting an open ecosystem of AI developers has benefits in terms of transparency, pluralism and competition, without causing specific risks.

Progress in AI is the product of both massive efforts by private companies and a very open research community where researchers and developers together contribute to the building blocks that enable these systems to be constructed. With regard to the foundation models that are the basis of generative AI, a seminal paper was published by a Google team in 2017[51], while OpenAI was created with the explicit aim of advancing open AI research, but neither of them publishes their models openly anymore. For their part, Meta and Mistral have opted for a partial opening of their models. In China, Alibaba and 01.AI continue to publish their models, while Baidu has stopped publishing its Ernie model[52].

This panorama reveals a tension within the ecosystem: on the one hand, openness widens the circle of contributors, democratizing AI and its dissemination; on the other, it reduces control over the technology and increases the risk of inappropriate or malicious use. Some even advocate a "registry" of AI systems to limit their "proliferation".

This is nothing new in IT. It lies at the heart of the open source debate. Open source software is software whose code is public, allowing anyone to use, inspect, modify or share it. Open source protocols make it possible to read and receive e-mail, and most of the world's servers are equipped not with Windows, but with the open source operating system Linux.

Open source is therefore present in many of the digital tools we use. Outlook or Gmail are based on open source bricks, but they are not themselves open source, both for economic reasons and to ensure a certain degree of control over the user experience. Open source and proprietary solutions are therefore complementary in many ways.

While the debate is not new, AI has two particular characteristics. Firstly, AI could cause significant risks, which would justify being more cautious than with messaging. Secondly, the openness of an AI system is not binary, but occurs along a "gradient" of openness.

51. Vaswani A. et al, Attention Is All You Need (2017), 31st Conference on Neural Information Processing Systems.
52. Baidu is the Chinese search engine, Alibaba an e-commerce leader and 01.AI a start-up founded by Kai-Fu Lee.

Should we encourage the development of open foundation models in AI? To answer this question, we need to look closely at the risks mentioned, the comparison with closed models, and the benefits of openness. We also need to bear in mind that some of the safeguards built into closed models today can be circumvented[53].

As far as biological and cyber risks are concerned, there is no evidence to suggest that open models pose any greater risk than closed ones. A model provides access to information on virology and bacteriology, but does not appear to facilitate the production of a biological weapon any more than a simple search engine[54]. On the other hand, open models reduce the cost of producing disinformation by 70%[55] and facilitate the creation of non-consensual images. However, these capabilities are already within the reach of current models, which are widely distributed on the Internet. Restricting the development of open models will therefore not reduce the risk.

Overall, our Commission considers that open foundation models do not pose any significant additional risk compared with closed models. However, like closed models, they do require investment in countermeasures, such as cybersecurity and disinformation.

Furthermore, we must not overlook the benefits of open models for society, including in terms of risk management. Transparency facilitates evaluation by third parties and mobilizes the community to improve and secure models. Openness also contributes to the spread of AI[56], since it facilitates the customization of models, and thus their adaptation to different contexts. It reduces the environmental footprint, since it avoids the need for everyone to retrain their own foundation model. Last but not least, openness lowers barriers to entry, enabling new companies to enter the market. Countries that contribute more to open source create more start-ups, and these are of higher quality[57]. Several major technology companies (MongoDB, Huggingface, Confluent, etc.) have made open source the core of their strategy, while monetizing their technology in other ways.

How can we realize the benefits of opening up AI models? We need to ensure that the ecosystem can actually use, inspect, modify and share systems. This will depend first and foremost on how models are accessed, how transparent they are, and how well third parties can evaluate them.

For a system to be reusable and shareable, you need to be able to download not only the code that makes up the model, but also the weights that are the parameters resulting from training, and a permissive license[58]. To be able to inspect the system or modify it appropriately, the model code is not enough. You need the tools to modify it, as well as knowledge of the techniques and data used to train it[59]. Access and transparency thus form a gradient of openness for AI systems.

Being in favor of openness does not mean opposing closed models, but rather supporting the open ecosystem to produce quality AI bricks and systems, reused because of their quality. Tomorrow, open AI systems, partially open systems, closed systems, open AI systems integrated into proprietary interfaces, etc. are likely to coexist.

53. Henderson P. et al, "Safety Risks from Customizing Foundation Models via Fine-Tuning" (2024), *Stanford HAL Policy Brief*
54. Mouton C. A., Caleb L. and Guest, E. (2024) "The Operational Risks of AI in Large-Scale Biological Attacks", *RAND Research Report*.
55. Musser A. (2023) "A Cost Analysis of Generative Language Models and Influence Operations", *Computers and Society.*
56. Ding, J. (2023) "The Diffusion Deficit in Scientific and Technological Power: Re-assessing China's Rise", *Working Paper.*
57. Wright N. L. et al. (2023) "Open source software and global entrepreneurship", *Research Policy.*
58. Google's PaLM 2 is thus closed, and GPT-4 is only accessible on a limited basis via an API. Meta's Llama 2 is downloadable with a license that authorizes commercial reuse except in services with more than 700 million users.
59. The *Stanford Transparency Index* lists 100 transparency indicators for AI models.

What the open AI ecosystem needs is not a protected garden, but legal certainty (to avoid individual contributors being liable for a company's use of their model), rare commodities like rich data, and certain bricks that no company will want to publish openly. Public support should focus on bricks that are not very capital-intensive or of no interest to private players, such as refining in rare languages.

Recommendation No. 4

Develop a strategy to support the open AI ecosystem internationally by supporting the use and development of open AI systems and third-party inspection and evaluation capabilities.

1.10 DOES AI POSE A THREAT TO THE PLANET?

> **It all depends on how AI is used.** The use of models implies an increase in energy consumption, even if this is partly offset by leaps in their energy efficiency. However, AI models can accelerate green innovation and thus combat global warming.

The energy consumption required to train large language models has attracted a great deal of attention. In 2019, a study showed that the greenhouse gas emissions associated with training one of the first large language models in the USA were of the same order of magnitude as those of a flight between New York and San Francisco[60]. Increasing the size of the models rapidly increases these emissions as soon as the energy supply is polluting. On the other hand, increasing efficiency gains in microelectronics and training of models reduce these emissions for equal performance[61].

The carbon footprint of AI models comes not just from their training, but from their entire lifecycle: equipment manufacture and transport (20-30% of the footprint); model development and use (50-60%: training and inference); and waste (10-20%). The growing use of AI also increases its energy impact. In the case of a model like ChatGPT used by over 10 million daily users, while the energy consumption of each query remains relatively low, the total energy consumption of inference exceeds that of training after a few weeks of use[62].

In total, AI could consume 85 to 134 TWh of electricity in 2027, equivalent to that of Argentina or Sweden[63]. These figures, based on an indirect approach, should be treated with caution. Above all, energy consumption represents such a cost in training and inference that all players are seeking to improve their energy efficiency. In 2023, a Google query augmented with AI would cost 10 times more in energy than a conventional query. It's hard to imagine this technology being deployed to Google's 5 billion users. It's even certain that the cost of each query will be a key factor in the competition for AI systems. In this respect, specialized models today consume less energy than general-purpose models, and many avenues are being explored, whether to improve models, find new architectures or new models. Increasing transparency on

60. Strubell, E., Ganesh, A. & McCallum, A. (2019) "Energy and Policy Considerations for Deep Learning in NLP".
61. Luccioni, S. (2023) "Towards Measuring and Mitigating the Environmental Impacts of Large Language Models".
62. Luccioni, S. et al. (2023) "Power Hungry Processing: Watts Driving the Cost of AI Deployment?"
63. Vries, A. (2023), "The growing energy footprint of artificial intelligence", *Joule*

the energy impact of AI models would enable users to make informed choices, and weigh up beyond economic incentives.

As far as the non-energy environmental impact is concerned, this comes mainly from the production of the processors used in computing power, and in particular from the extraction and use of water, silica and rare earths in production. However, the high price of specialized AI processors should not obscure the fact that they make up a very small proportion of processors produced worldwide. In 2022, they represented less than 1% of sub-7nm chips, and less than 0.00026% of all chips produced[64].

The environmental impact of AI must be weighed up against its potential benefits. Thanks to its ability to optimize complex processes, AI could significantly reduce greenhouse gas emissions in many sectors: energy, transport, agriculture, housing, etc. Furthermore, AI could accelerate the ecological transition by reducing dependence on the innovation path[65]. A recent study[66], based on patent data, highlights just such an effect. Concrete examples of innovation acceleration are emerging[67], but still need to be supported.

Recommendation No. 5

Make France an AI pioneer for the planet by strengthening environmental transparency, research into low-impact models, and the use of AI to serve energy and environmental transitions.

64. Heim, L., and Pilz, K. (2024). "What Share of All Chips Are High-End Data Center AI Chips?" *blog.heim.xyz*
65. A company that has innovated in polluting technologies in the past is more likely to continue to innovate in these technologies, because it has acquired advantages (skills of the research team, mastery of industrialization, etc.).
66. Andres, P., Dugoua, E. and Dumas, M. (2022), "Directed Technological Change and General Purpose Technologies: Can AI Accelerate Clean Energy Innovation?", *LSE Working Paper.*
67. AI could, for example, enable aircraft to reduce the contrails they leave in their wake by 54%, which are themselves responsible for 35% of emissions by volume. Google (2023), "How AI is helping airlines mitigate the climate impact of contrails".

1.11 IS THERE A BUBBLE IN GENERATIVE AI?

> **Probably, but the bubble is helping to attract investment in risky projects.** When the bubble bursts, it won't be proof that generative AI is good for nothing, but that it isn't good for everything. It will then be a matter of maintaining the course of investment in AI to ensure the emergence of a solid European ecosystem.

In terms of financing, 2023 was an exceptional year for generative AI start-ups. They raised $22 billion in 2023, compared with $4.3 billion in 2022, $4.5 billion in 2021, and less than $2 billion a year previously[68]. In other words, two-thirds of all investments received by generative AI start-ups took place in 2023! At the same time, venture capital investment in start-ups continued to fall, reaching $224 billion in 2023, compared with $655 billion in 2021.

This influx of funding reflects the effervescence of the sector, which is accompanied by a host of start-ups with uncertain products or business models. The tendency is to see this as a "bubble", i.e. an unreasonable rush by investors and entrepreneurs driven by mimicry.

From a financial point of view, a bubble exists when the price (of a stock, a company, a sector) greatly exceeds its "fundamental value". To find out whether there is a bubble in generative AI, we would therefore need to estimate the future value of the companies and products being financed today. However, we're still in the early stages of generative AI deployment. It's likely that many of the generative AI products funded today will be neither useful nor profitable.

However, this euphoria is normal and necessary at the start of a technological revolution. It corresponds to the "settling-in phase" of a new technology[69], when everything seems possible, the benefits seem limitless, industries fear disruption, and there is much to learn. This phase enables the first innovations to be financed despite the uncertainty.

The excitement surrounding generative AI in 2023 has been compared several times to a "gold rush". It's much more akin to the boom that accompanied the development of the railroads in the 19th century, or the development of online uses during the dotcom bubble.

This installation phase is often followed by a "turning point", at which point the relevant uses of the technology become clearer. Then the bubble bursts, company valuations fall, and investment is redirected to useful, profitable companies and projects. Over the next few years, we're likely

68. Dealroom data
69. Carlota Perez (2002), Technological Revolutions and Financial Capital

to see some spectacular bankruptcies in the field of generative AI. Some very well-funded companies won't find their market and their gambles won't pay off. At the same time, other companies will succeed and find their market.

The bursting of the bubble will not be proof that generative AI doesn't work, but that certain uses don't work, or don't work yet. After the dotcom bubble burst, Europeans turned their backs on digital. Graduates left for the financial sector and governments didn't make digital a priority, while Silicon Valley continued to invest in digital. There's probably a bubble in generative AI, but let's not make the same mistake when it bursts: while letting unviable companies disappear, let's stay the course in AI investment.

1.12 SHOULD WE PREPARE FOR AN AI THAT'S SMARTER THAN US?

> **Yes, we can prepare for a future in which machines will surpass humans in many, many areas.** This superiority will arrive gradually and over an uncertain timeframe. It's up to us to act collectively to limit the risks associated with this evolution and reap all the benefits for humanity.

Generative artificial intelligence is a milestone in the history of innovation. It is far from being the last. In the months, years and decades to come, we are likely to see further rapid and far-reaching advances. Models will progressively be able to be factual, to adapt easily to increasingly sophisticated applications, to generate voice and video in any language with precision, to conduct reasoning, to do mathematics, to understand the physical world around us.

These developments will be made possible by the continuing massification of available data, due in particular to the multiplication of sensors on connected and embedded objects (cars, robots, etc.), but also in space (satellite constellations) and in the oceans. They would also be accelerated by the falling cost of computing power, contributing to both the widespread use of AI systems and the growing accuracy of models. They would probably be facilitated by a combination of technological approaches, mixing machine learning and the symbolic approach of AI.

By the end of the decade, AI systems are likely to be supporting humans continuously and in every task, personal or professional. In particular, this support could take the form of powerful personalized assistants, which will perform tedious tasks, support thinking and decision-making, and speed up group work.

We can also expect robotics to make giant leaps forward. However, the complexity of the real three-dimensional world, the difficulties of interacting with the environment and societal inertia mean that we cannot expect robots to be used on a massive scale in the near future. The self-driving car, whose democratization is constantly being postponed, illustrates the twofold difficulty of removing technological barriers and effectively deploying automated systems in real-life situations. Just because machines can perform tasks that seem complex to us today, doesn't mean they'll be able to do what's simple to us tomorrow. This is the paradox underlined since 1988 by Hans Moravec: "the most difficult thing in robotics is often what is easiest for man".

As technology advances, machines will overtake humans in an ever-increasing number of areas. This overtaking by machines is sometimes referred to as general artificial intelligence. The notion is much debated and far from universally accepted. It seeks to describe high-performance AI systems with capabilities that are both broad (hence the term general AI) and sharp[70]. In all likelihood, there will be a gradual improvement in AI systems, and superiority will not be sudden. Our personalized assistants should become progressively more competent. Their use will become increasingly easy and unobtrusive, for example by replacing today's familiar interfaces (screen and keyboard) with more natural ones.

The societal transformations brought about by these innovations will depend on our ambition and commitment. AI can be used to reduce social inequalities, promote collective prosperity, improve the quality of work by eliminating the most thankless tasks, and promote scientific progress for the benefit of humans (their intelligence, health, nutrition, democratic life, etc.) and their environment (optimization of industrial processes, new forms of energy, new decarbonization technologies, etc.). In this version of our future, humans would gradually gain access — using powerful, personalized digital tools — to a range of knowledge, goods and services that previously seemed beyond their reach.

These benefits will not come spontaneously. They can only be achieved through a political vision and collective commitment. Because, in contrast to the gains outlined above, a dystopian future may also be in the offing. A future in which information bubbles and cognitive influences weaken our democracy. A future in which many workers find themselves out of place in the face of ever more competent machines. A future in which the concentration of the most advanced technologies in the hands of a few players alters our sovereignty and absorbs most of the value produced by our economy.

The path that France and Europe will take has not yet been mapped out. It's up to us to define a political project for society, and to forge AI applications in line with it. It's up to us to take advantage of AI, by investing strategically in mastering the technology and its value chain. The action plan recommended by our Commission focuses on marking out the first few miles of the road ahead. Going beyond this will require constant action, the plasticity of our institutions (public and private), and continuous anticipation and preparation.

70. Ringel Morris M., Sohl-Dickstein J., Fiedel N. et al. (2024) "*Levels of AGI: Operationalizing Progress on the Path to AGI.*"

2

HUMANISM, SOVEREIGNTY, RESPONSIBILITY: INNOVATING, DEPLOYING AND MASTERING AI

2.1 HUMANISM: PUTTING AI TO WORK FOR US

2.1.1. MAKING SOCIAL DIALOGUE AND CO-CONSTRUCTION THE CORNERSTONE OF AI USE

The spread of AI is caught up in a face-off between "techno-sceptics" and "techno-enthusiasts" that is easy to caricature. The former fear that AI will reinforce inequalities, degrade the quality of working life and ultimately benefit only a minority. They therefore demand to be able to define how AI systems will be designed and deployed. The latter say they need to move fast, to experiment to find the right uses, the right products — in short, to have a free hand[71].

At this stage, two things are certain. On the one hand, we need to experiment, to feel our way with AI in order to find how to reap its full benefits. All the more so if we want to find the "right" ways of deploying AI so that it improves workers' everyday lives. On the other hand, this experimentation is not just a question of tools, but also of training and work organization. Mechanical cotton spinning only increased productivity in the textile sector when the industry was reorganized around large textile mills that could take advantage of these new machines[72]. Electricity only increased productivity in factories when they reorganized (see 1.3 *Will AI help us to prosper?*).

To harness the full potential of AI, we need to find ways of reconciling the rapid pace of experimentation with the necessarily slower pace of our skills and organizations. Social dialogue is essential to encourage the use of AI, to discuss the aims and meaning of technological transformations, to develop the learning capacity of organizations and to design appropriate training plans. The participation of all stakeholders is an essential condition for the deployment of new technologies with a view to emancipation, empowerment and improved working conditions, notably through the reduction of thankless tasks[73].

The spread of AI will be a key asset in boosting companies' competitiveness and employment (see 1.4 *AI, creator or destroyer of jobs*). It will not happen without social dialogue, based on mutual trust, experimentation and co-construction. Yet, while the effects on the world of work

71. Uber founder T. Kalanick went so far as to make this his mantra: "*move fast and break things*".
72. Juhász, R., Squicciarini, M. P. and Voigtländer, N. (2020), "Technology Adoption and Productivity Growth: Evidence from Industrialization in France", *National Bureau of Economic Research Working Papers.*
73. International Commission Olivier Blanchard and Jean Tirole (2021), "Major economic challenges".

of previous digital waves are profound, workers and their representatives are today little involved in technological and organizational choices in workplaces and at national level.

The right to information[74] and the informed opinion of employee representatives on transformations in the workplace are little used in practice by companies and administrations. This is due to the fact that artificial intelligence and digital technology in general are presented as primarily technical issues, which are difficult for workers and employers to grasp. Social players are not sufficiently informed or trained in these issues and tools. The Information Systems Department (ISD), on the one hand, and the Human Resources and Social Relations Department, on the other, often work in silos. The weakness of co-construction can become a source of anxiety and even rejection for workers, and increase their sense of insecurity and fear of demotion.

What's happening in the rest of the world? The example of Canada[75]

In Canada, the use of AI systems by employers is becoming increasingly sensitive, and has led to demands in the context of large-scale social movements during 2023: a strike by federal civil servants, a strike by port employees, etc. In particular, dockworkers at British Columbia ports and employees of the Metro distribution chain have demanded stronger supervision of the use of automation. In addition to their wage demands, employees of the three major automakers (Ford, General Motors and Stellantis) are also making demands concerning automation and subcontracting.

In September 2023, the federal government presented a voluntary code of conduct aimed at the development and responsible management of advanced generative AI systems. The code of conduct is built around six principles: accountability, safety, justice, transparency, human oversight and reliability. Application of the code is voluntary, and several companies around the world have already signed up to it.

For social dialogue to be able to integrate AI issues, and facilitate its experimentation and dissemination, two characteristics need to come together. Firstly, the social partners must be trained[76] and active participants in the bodies where AI deployment will be discussed. Secondly, this technological social dialogue must be part of an iterative process that characterizes AI projects.

To take this a step further, AI itself could be put to use in social dialogue. Tools based on generative AI can be developed with the social partners to help employees better understand technical debates, be they IT, financial or legal. These AI tools, for example in the form of a simple dialogue forum, could integrate a part common to all companies (labor code, etc.) and a part specific to the organization in which each worker is placed (collective agreement,

74. European agreements on digital transitions signed by the social partners in 2020 for companies and in 2022 for administrations.
75. Source: French Treasury Department, Economic Service in Canada.
76. See, for example, the "Dial IA" project, led by the Institut de recherches économiques et sociales (Ires), which aims to deploy a methodology that makes technological social dialogue at work an operational lever for digital transformation.

organization's internal regulations, etc.) and to its union representatives. The tool could help to increase knowledge of rights and understanding of ongoing transformations, or improve preparation for meetings (boards of directors, works council, elections of representatives, etc.). During labor negotiations, AI can also help to analyze and exploit vast quantities of data, and thus support negotiations.

2.1.2 TRAINING: IMMEDIATELY, ON A LARGE SCALE AND CONTINUOUSLY

2.1.2.1 Initial training

Current and future needs in AI require a vast training plan for everyone and at every age. More specifically, the training challenges cover three different needs: training people capable of designing and developing AI solutions, training people capable of deploying these AI solutions within their companies, and more generally raising the general population's awareness of the culture and understanding of the main operating principles of AI.

As we mentioned earlier (see *1.4. AI: creator or destroyer of jobs?)*, companies investing in AI skills are hiring more highly qualified, more technical profiles, particularly in favor of so-called "STEM" jobs (science, technology, engineering and mathematics). This applies equally to companies that design AI solutions and companies that adopt them. More specifically, a study[77] based on Danish data examines the need for adequate training for AI production on the one hand, and AI adoption on the other. It shows that AI-producing companies, which sell a product or service "containing AI", recruit more students with backgrounds in computer science, mathematics or physics, while companies that deploy AI without developing it in-house target more applied "STEM" profiles with backgrounds in chemistry, biology or biotechnology in particular.

In terms of demand for AI-specific skills, a recent study by the OECD[78] provides an order of magnitude. This study notes that online job offers requiring AI skills represent 0.35% of the offers posted in France. Of these "AI job offers", 69% relate to the IT and specialized activities sectors, and are therefore mostly linked to the development of AI solutions, while the remaining 31% relate to the deployment of AI within other sectors. This second category covers people with training in adapting AI to the specific uses of a discipline such as health, law or physics, sometimes referred to as "X + AI" profiles. In addition, the total number of "AI job offers" increased by around 45% between 2019 and 2022.

What does this mean for France? If we project a similar trend for the next ten years, and assume that the split between development and deployment will remain similar, job vacancies in AI development and AI deployment should represent 1% and 0.5% respectively of all vacancies in 2034. If we project a similar trend in manpower requirements over the next ten years to that seen over the last ten, we would end up needing around 56,000 positions per year in AI development and 25,000 positions per year in AI deployment ("X + AI").

77. Humlum, A. and Meyer B. (2020) "Artificial Intelligence and College Majors", *Working Paper*.
78. Borgonovi, F. et al. (2023), "Emerging trends in AI skill demand across 14 OECD countries", *OECD Artificial Intelligence Papers*.

It is therefore necessary to calibrate initial training provision to the AI skills needs of today and tomorrow. In 2021, a report by the Cour des Comptes audit office estimated that there would be 16,687 places in specialized AI training courses at graduate level. To meet AI development needs, this figure would have to at least triple over the next decade.

To deploy AI in the economy, the need for 25,000 people per year by 2034 corresponds to training around 1.5% of all higher education students each year in "X + AI" skills, either by creating specific courses, or by creating an advanced AI module within the various courses, as already recommended in the Villani report.

A second type of profile, indirectly linked to AI, seems necessary for the deployment of AI within companies: that of people in charge of the information system infrastructure, with the knowledge essential to AI, particularly on data collection and processing, to best deploy AI solutions within companies. The term "MLOps" is sometimes used to describe this type of profile. In 2023, while Pôle Emploi indicates a labor requirement of 12,180 positions in skilled IT services, 16,959 students are enrolled in engineering courses in the "IT and computer sciences" field. In order to meet this training need, it would therefore be necessary at the very least to train all students on specialized IT courses in the AI issues relevant to their activity. Given that not all the needs of IT departments are related to AI, and that not all of them are going to disappear, we should probably aim to increase the number of students enrolled in these courses by 25%, to reach 20,000 students.

Added to all the specific training needs for AI development and deployment is the need to raise awareness among everyone, adults and younger generations alike. What's more, the low proportion of women among STEM graduates (31% in 2019), but even more so in the "computer science and information technology" field (19% of enrolees in 2023), brings issues relating to employment and gender pay inequalities to the table alongside the spread of AI. This makes it all the more important to raise awareness upstream, to attract female students to these fields.

The French government has encouraged higher education establishments to develop training programs in this area through two calls for expressions of interest (CEIs). The "*Compétences et métiers d'avenir*" and "*IA Cluster*" CEIs (the latter is currently open) aim to structure the AI training sector in order to consolidate around ten clusters of excellence and triple the number of students trained in AI. These CEIs also include an element of continuing education and the deployment of introductory AI modules to reach students from a variety of disciplines. These investments will enable a significant rise in the number of profiles trained in AI over the next few years, but this will only be effective if the training courses on offer are attractive and find their audience.

However, the CEI-funded training programs have one limitation: they do not cover all the students in a given generation, notably because they are not distributed throughout the country, and there is a shortage of trainers specialized in AI. A three-year assessment of the results of these CEI programs will be necessary to verify the reality of the ramp-up. However, it is already essential to pursue efforts to reach as wide an audience as possible. This could be achieved by sharing courses online, or through an ambitious training plan for teacher-researchers in all disciplines. The Commission also recommends greater reliance on the sector's professionals, who are probably ready to participate in the training effort, provided that it includes a significant practical component, which is crucial in this sector.

In addition, to ensure that the courses on offer attract enough students, it is necessary to act upstream of higher education to acculturate students to the challenges of AI as they progress through their schooling. Some school curricula already include explicit learning about AI: the draft curriculum for junior high school technology, the science syllabus for the core curriculum of the general baccalaureate, and the curriculum for the engineering science specialization of the general baccalaureate. However, these learning elements do not form a progressive learning curriculum for all students. We need to identify the contribution of each discipline taught, so that all students (whether vocational, technological or general) can benefit from learning related to AI.

Recommendation No. 6

Generalize the deployment of AI in all higher education courses and acculturate students in secondary education to make specialized courses accessible and attractive.

2.1.2.2 Ongoing training

In addition to initial training and higher education, as well as recruitment and work reorganization policies, continuing vocational training will be an essential tool for coping with the profound transition of professions that artificial intelligence will entail. This trend has been reinforced by generative AI, but is not new. By way of illustration, La Poste Groupe started digital training in 2015 and has made it part of the company's social pact. A specific training budget (€500 million over 5 years) has been established and the training catalog has gradually been enriched on AI. The new 2020-2030 strategic plan extends this action. Joint work must be undertaken to anticipate changes in professions and the content of professions, as well as to identify the necessary training.

The main lesson from our citizen consultation: the need for training

From mid-December 2023 to mid-January 2024, our Commission conducted an online citizen consultation, to better understand expectations and fears when it comes to AI. We collected 6,917 responses[79]. Among the key results of the survey, the need for information and training on AI in the professional environment stands out in particular.

The majority of participants do not fear seeing their jobs disappear or devalued by the emergence of AI (note that a higher proportion of women, blue-collar workers and intermediate professions share this concern). However, the majority of respondents indicated difficulties in expressing themselves on the tangible effects of AI in their jobs. Many of them expressed the need for a better understanding of the benefits they could derive from AI, and the concrete applications for using these tools at work. After the training, respondents stressed the importance of social dialogue.

AI also looks set to play an important, if not dominant, role in the continuing professional education sector. Firstly, in terms of instructional design, to structure content and organize ideas for a training plan, or to help design pedagogical tools. During training, where artificial intelligence can act as a virtual assistant to learners, helping them to individualize their learning paths. And finally, for post-training support and tutoring. For example, some software engineering courses already incorporate language models to learn how to code.

For the time being, according to the OECD, job vacancies requiring artificial intelligence skills represent only a limited number of all job vacancies (OECD, 14 countries), with less than 1% in the USA in 2022, the country with the highest proportion. Nevertheless, they are growing rapidly in almost all countries. The OECD also points out that demand for jobs linked to artificial intelligence is highly concentrated in terms of sectors and professions, with differences between Europe and Anglo-Saxon countries.

While being proactive, public authorities need to show humility in this area. Companies have an interest in training their employees in new skills. If they aren't doing so, or not yet, it's also

79. The consultation was conducted on the Agora application with the support of the interministerial center for citizen participation. Participant profiles were diverse, but cannot be considered fully representative of the French population (for example, more men than women responded).

because there's still a lack of clarity about the skills and continuing professional development issues involved in the professions transformed by the introduction of AI. This vagueness will only be lifted gradually. Where public support will be useful, it will be to support the ongoing mapping of training needs and modalities, to bring clarity to a plethoric offer, and to ensure training for jobseekers.

Recommendation No. 7

Invest in continuing vocational training for the workforce and in training schemes around AI.

In the culture and media sector, the impact of AI on employment (4% of total employment) is already analyzed as significant, especially as 40% of artistic and cultural jobs fall within the rest of the economy (luxury goods, advertising, industrial design, automotive, etc.). Overall, the level of exposure to AI is higher than for all professions, particularly for these artistic and cultural jobs outside the cultural sectors.

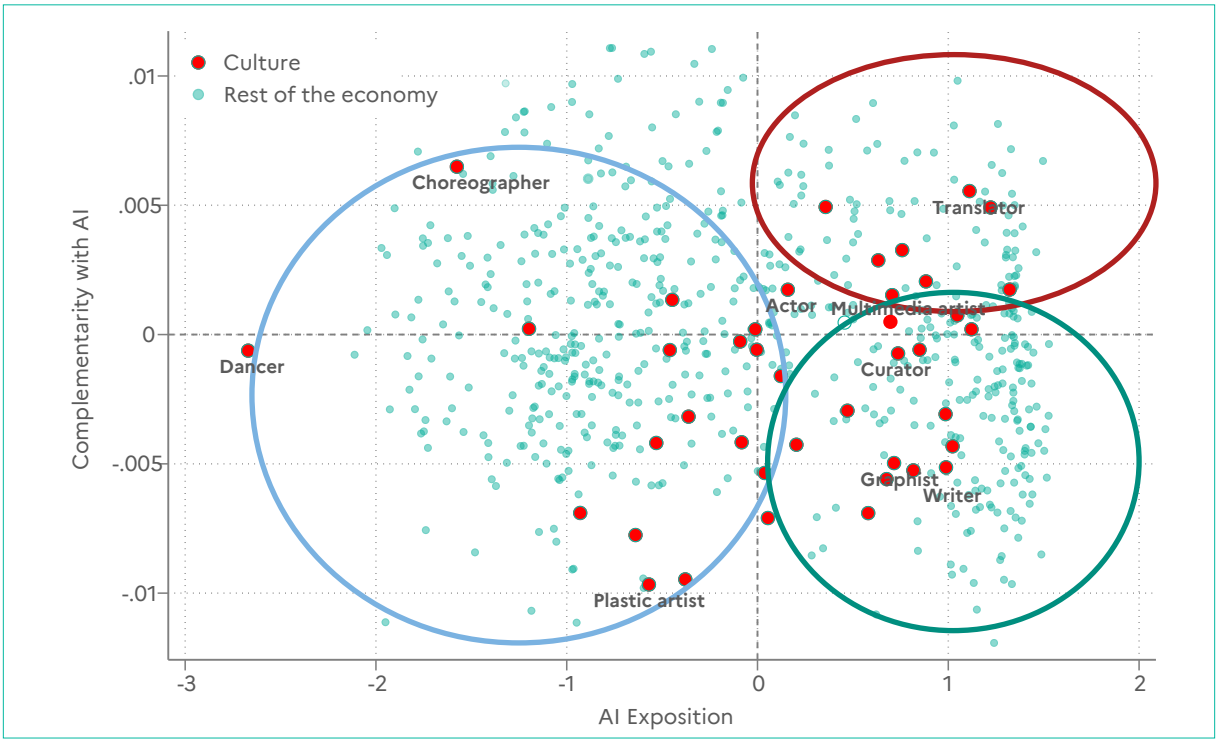

Figure 8: Exposure of artistic professions to AI and complementarity with AI.
Source: calculations by A. Bergeaud for the French AI Commission (2024)
Interpretation: the further to the right of the graph, the more tasks are likely to be automated by AI. The higher up the chart, the more complementary they are to AI, as they mix automatable and non-automatable tasks.

Substitution effects already exist (translators, voice artists), with forms of downgrading (loss of author status, transfer of activities due to lower skills and risk of lower income). They are likely to extend — at some unknown date — to many jobs where automation is possible: artistic professions in non-cultural sectors. In the artistic and cultural sectors, complementarity may play a more important role, with less risk of substitution: activities "augmented" thanks to AI and time freed up for more creative and distinctive activities. Complementarity with AI is the key to the polarization of professions in the already competitive world of talent.

To a large extent, career distinctions and job stability should depend on status (salaried/self-employed; private/public) and employer size. The conditions of access to vocational training to support careers will also be determined. The stakes are immediate. In the longer term, it's a question of career paths based on specific durations and varied statuses, where the complementarity of jobs is essential. Some professions are already facing this challenge (freelancers and graphic designers)[80]. They are in fields where AI is deploying fast, such as the press[81] or communications, where tasks can be substituted by AI systems that will perform them more cost-effectively.

Favoring the adaptation of artistic and cultural employment means very quickly creating the conditions for training, particularly for the younger generations, to make France and its creators a place of excellence. In a world of AI, original creation and cultural specialization have every chance of becoming a rare resource as "small" models develop, or as "large" models become more widely available. We need to strengthen initial professional training in specialized higher education and research, by building bridges between creation, research, technology, economic, social and cultural projects.

The reconfiguration of the economy as a whole by the introduction of a technological regime dominated by AI places the creative sector in a crucial position. Already highly competitive and much sought-after, the creative sector will have to position itself at the crossroads of rapid technological change and evolving business models.

Recommendation No. 8

Train creative professions in AI, from the early years of higher education and on an ongoing basis.

80. Reshef, O., Hui, X., and Zhou, L. (2023) "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market", *CESifo Working Papers*.

81. Beckett, C., and Yaseen, M. "Generating Change A global survey of what news organisations are doing with AI", *LSE JournalismAI*.

2.1.3 EQUIPPING PUBLIC SERVANTS: AN OPPORTUNITY TO TRANSFORM THE ADMINISTRATION

Digital technology coming to the rescue of public service is nothing new. While the dematerialization of online procedures has made good progress, digital transformation remains incomplete. The deployment of AI is an opportunity to relaunch this transformation, provided we remove the brakes that continue to prevent a real digital transformation of public service. For these are the same obstacles that will prevent us from taking advantage of AI and banking on a profusion of initiatives.

The civil service should be one of the first to benefit from digital technology. In fact, from the citizen's point of view, the dematerialization of public services and procedures has made great strides in recent years. More than three-quarters of the 250 most frequently used procedures have been dematerialized[82], and 79% of French people believe that the dematerialization of public services makes their lives easier[83]. However, for some, this has been accompanied by a feeling of dehumanization and alienation from the public service[84]. For public servants, this dematerialization has ambiguous effects, particularly as 51% still find their digital environment average (32%), poor (14%) or mediocre (5%)[85].

All too often, digital transformation has stopped at the dematerialization of procedures, without any in-depth transformation of information flows or the processing of requests. Promises to personalize (and therefore humanize) the civil service, speed up processing and simplify the work of agents have not been kept. The formalities to be completed when a child is born are numerous and complex, and it remains frustrating to have to bring to each public service the documents proving that a child has been born, and that it is indeed yours. While there are some promising initiatives underway ("Tell us once", "proactive administration", etc.), the frustration is all the greater given that, in comparison, our digital lives are becoming increasingly integrated.

AI is an opportunity for public services to take their transformation one step further. It promises to personalize public services and make them more efficient, and generative AI promises to streamline communication with users. A generative AI could soon re-explain several times in accessible language the steps to be taken to enroll your child in school, or to file your tax returns. An agent could even do it for you. At the Rectorat de l'Académie de Lyon, since September 2023, Cassandre AI has been providing the 40 human resources managers in the teaching personnel department with answers to questions asked by the 45,183 teachers in the academy about the assignment of trainees and transfers within the academy. Surveys show that the tool is widely appreciated, and its use will be extended in 2024.

82. Observatory of online procedures.
83. Online survey conducted by Ipsos for Sopra Steria in September 2019 among 6,000 people in 6 countries, representative of the national population aged 18 or over
84. Défenseur des Droits (2019), "Dématérialisation et inégalités d'accès aux services publics".
85. Direction interministérielle du numérique (2021) "Baromètre numérique de l'agent".

What's happening in the rest of the world? The example of Sweden

The Platsbanken.se platform, Sweden's largest online job board, managed by the Swedish equivalent of the France travail governmental agency, has evaluated the AI-enabled implementation of targeted job recommendations, based on applicants' job search history[86]. Jobseekers with access to recommendations clicked and applied more to recommended offers, leading to a positive effect on the re-employment rate of around 0.6%. Although it is difficult to extrapolate and transpose, a similar increase in France would lead to the hiring of just over 12,000 jobseekers. But the main gain probably lies in the time freed up for agents, who can focus their support on other areas.

AI, and in particular generative AI, can free agents from repetitive or time-consuming tasks, while improving service quality. Thanks to its ease of use, generative AI offers the opportunity to unleash agents' creativity, enabling them to experiment with technology at their own level, without always needing a specific system. A teacher can already use generative AI to help her design a training program or vary an exercise.

The civil service can set itself two levels of objectives with regard to AI. The first involves deploying AI systems that fulfill a specific objective: responding to users, simplifying a message, summarizing a video conference, making a financial analysis, etc. The second aims to rethink public service based on its missions, users' needs and the capabilities offered by AI. For example, by imagining a companion to help us with all the "paperwork" for our child.

The first level corresponds to the dematerialization of the digital transformation of the last 20 years. The second is the profound transformation of our public services.

To successfully deploy AI in public services, the first level, public services will need the same ingredients as for digital transformation:

- **a clear vision of the service's objectives,** enabling AI's contribution to be assessed, systems to be continuously adjusted in response to feedback, risks to be managed, and roles to be allocated between public developments and the role of private players;
- **the confidence of agents and users to experiment,** which will require a sense of responsibility without passing the buck to the tool;
- **people** capable of designing, piloting, producing or purchasing these AI systems;
- **data,** to train or re-train a model, but also more simply to deploy AI-based tools in the work process;
- **robust infrastructures,** facilitating data circulation, common repositories, application security and regular updates.

Efforts have been made in recent years, notably to strengthen the State's digital skills and invest in a cloud infrastructure. These efforts must be continued, as they will only have an impact in the long term. The fragmentation of public digital hosting infrastructures is costly, reduces flexibility in the event of peak workloads, prevents access to certain development, testing and integration tools, reduces the attractiveness of the public service for technical talent, and

86. Le Barbanchon, T., Hensvik, L., and Rathelot, R. (2023), "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", *Working Paper*.

makes it more difficult to improve applications. And, crucially for AI, it hinders the collection and circulation of data. The transformation of public services by AI will not advance without a leap in public digital investment efforts, both in quality and quantity[87].

These ingredients do not, however, form the basis of a strategy, which must avoid two pitfalls. On the one hand, the "big AI project", designed to do everything, to replace everything, developed far from agents, users and the reality of public service. On the other hand, the "all ChatGPT", in which a commercial and foreign universal conversational robot would become the only use of AI in public service.

Between these two pitfalls, public services need to chart a course that combines technological mastery, cost control, abundant experimentation and profound transformation. The Commission proposes four pillars: (i) clarify objectives and the division of roles between public services and private providers; (ii) pool investments in AI; (iii) strengthen steering and execution capacity; (iv) empower agents and citizens to become involved in this transformation.

From 2024, public services will have to decide whether to use off-the-shelf AI solutions, enter into partnerships with companies, or redevelop their own tools. Following the "public platform" logic already applied in Health or Education, work on the "why" (their mission and objectives) is essential in order to deduce the "what" (the bricks, services and datasets to be developed) and the "how" (the distribution of roles between different public and private players, the involvement of agents and users, consideration of ethical issues)[88].

Off-the-shelf solutions have the advantage of performance and simplicity, as they are immediately available, and can even be integrated directly into agents' tools (office suite, search engine). But they do present risks, chief among them data leakage[89]. This is not an insignificant issue. Some ministries have already banned the use of IT development tools such as Github's Copilot or GPT-4. The IT developers recruited (with difficulty!) to reinternalize and master digital technology within the State are thus deprived of tools that will soon be indispensable to any coder. Public services should therefore rapidly adopt charters for the use and contractualization of off-the-shelf AI solutions, to encourage public servants to appropriate these tools[90], all the more so when they are free for occasional use.

However, off-the-shelf solutions will have their limits, either because they don't fit in with existing tools, or because they aren't adapted to certain sensitive uses. It would be tricky to ask ChatGPT to summarize a memo to a minister, for example. But it would be absurd for every ministry and local authority to redevelop or buy an AI capable of summarizing a memo without leaking the data. The generative AI production chain, made up of re-training and refining models, lends itself particularly well to mutualization.

The Albert project, launched in autumn 2023, is a good illustration of this dynamic. This generative AI can summarize text in administrative language, and will soon be integrated into existing collaborative tools. Starting in 2024, Albert will be tested at local public service counters — Maisons France services — to help users with their paperwork.

Albert is a building block that can then be refined with proprietary data, inserted into interfaces specific to each public service. Together, public services could identify other functionalities to

87. Digital investments accounted for just 3% of total State investments in 2019: Cour des Comptes (2020), "La conduite des grands projets numériques de l'Etat".

88. Digital Roadmap for Health, 2019-2022 and 2023-2026, Digital Strategy for Education 2023-2027

89. A 2021 circular already prohibits the use within the State of Microsoft 365 on a Microsoft cloud, which will prevent the use of the "Copilot" module based on OpenAI technology. Local authorities and hospitals are not affected.

90. The Quebec government's Digital Transformation Academy, supported by Université de Laval, is an interesting example of how to develop skills and spread a digital culture within public services.

be pooled, particularly where they have a need that is shared by several services, and relatively specific compared to the private sector. Simplifying administrative language and providing assistance with procedures could be the first steps in this direction.

In order to provide widespread access to this generative AI service, set up an AI-adapted infrastructure and avoid duplication of investment, our Commission also recommends that the capacity for technical steering and execution in public services be strengthened. At inter-ministerial level, a real technological department should be able to provide not only doctrine, but also quality infrastructures (hosting, computing power, software factory, digital identity[91], etc.), expertise and a transformation budget. At a time when the State is seeking to reinternalize skills and circulate them between ministries, it needs to strengthen its ability to produce quality digital products.

In all public services, digital and AI must be taken to the right level by administrations and politicians, beyond simple support functions or ad hoc projects. In the same way that every public service has in mind the constraints and possibilities of the legal framework, the vector of digital and AI must become inescapable in the development and conduct of public policies. Too often, high-level decisions focus on the design of an interface rather than on the speed of the application, its resistance to peak loads or the measurement of its quality. Too often, development or implementation decisions are made without thinking about the right digital vector, when better solutions exist and are equivalent from the point of view of the public policy objective[92]. The COVID-19 testing and vaccination campaigns would not have been possible without providing for data feedback at national level, and without cooperating upstream with software publishers in the liberal professions to enable them to communicate data.

Beyond steering, the arrival of AI could be an opportunity to give public servants the ability to transform their own work, rather than having it transformed from above. Today, there are tools that make it possible to design a digital tool without coding. "*Démarches-simplifiées*", for example, has enabled public services to dematerialize 32,468 procedures without coding, both in the State and in local authorities. The public service would benefit from equipping agents with configurable AI that they can deploy themselves, whether these solutions are off-the-shelf or specific to the public service. If agents take ownership of AI, the uses will abound and will enable us to identify more quickly where AI has value. For example, the *LlaMendement* project of the Direction Générale des Finances Publiques, which greatly facilitates the processing of amendments to the finance bill, was suggested by an agent involved in digital transformation issues.

This empowerment of the agents themselves would also enable us to better define where to place the "human element in the loop". It is essential that a public service user can always turn to a person or department responsible for a decision. But there are often gains to be made both in terms of quality and processing efficiency by automating certain tasks. Bureaucracy, made up of rules, processes and equal treatment, sometimes comes into tension with the objectives of humanizing and personalizing public service. By making it possible to automate bureaucratic tasks and free up agents' time for their public service mission, AI offers a path to reconciliation.

Citizens themselves could be brought in to contribute to AI-enabled public services, whether to define how they operate or to participate in their construction. Their involvement is crucial to

91. As in the case of the *France Identité Numérique* program, it is not necessary for the interministerial level to operate all these shared infrastructures, as long as the operator knows how to take into account the needs of other administrations.

92. This problem is not specific to France. See Pahlka, J. *Recoding America*

prevent the transformation of public services by AI from reinforcing bureaucratic, inexplicable and distant centralization. In Taiwan, for example, "Alignment Assemblies" have been convened to define rules for AI deployment and behavior in the public service. Initiatives of this kind could be deployed in France, extending the concertation experiments already carried out in the digital health and education sectors, for example.

The dematerialization of public services has sometimes been synonymous with dehumanization, both for users and for agents forced to fit their service into the precise boxes of a rigid process. AI, and in particular generative AI, can be an opportunity to rehumanize public services, by bringing public services closer to users and enabling agents to be involved in improving their work.

Recommendation No. 9

Strengthen the technical capacity and infrastructure of public digital in order to define and scale a real transformation of public services through digital and AI, for agents and at the service of users.

2.1.4 BETTER CARE THROUGH AI: INDIVIDUALIZED PATIENT SUPPORT

The transformation of the healthcare system by AI is not coming, it has already begun. AI is already present in many medical devices to improve analysis and clinical processes. Advances in computer vision have led to progress in surgical robotics and X-ray analysis, to the point where some thought the profession of radiologist would disappear.

Over the past ten years, the performance of AI systems has already improved the accuracy of diagnoses, as well as the targeting and speed of treatment processes. In the field of medical research, AI is enabling the discovery of new treatments. For example, AI recently led to the discovery of a new antibiotic against staphylococcus aureus, after 60 years of unsuccessful research[93]. In public health, AI can improve predictions of the evolution of epidemics[94].

These advances have triggered ethical reflections[95] and an evaluation of these systems in real-life situations, in order to distinguish between potential and actual effect. For example, AI-based diagnostic systems are configured to be highly sensitive and not miss any signals. On the other hand, they are more likely to spot signals when there are none. They may therefore require double checks, which they were designed to reduce[96,97].

With this already long history, what do the latest technological advances change? With their ability to identify elements in text and voice, to generate text and voice, and to adopt a wide range of tones, these new systems are radically changing the way we interact with caregivers and patients.

On the caregivers' side, digital transformation requires them to constantly structure their notes, thoughts and medical decisions in software that can only disappoint: either its interfaces are too complex, or the options it offers are too limited. The multiplication of addressees (specialists, GPs, social security, mutual insurance companies, etc.) reinforces the feeling of overload. This is where generative AI holds great promise.

On the patient side, the ability of AI to communicate credibly with a patient raises many questions. Some users will turn to generalist AI for diagnoses, despite warnings[98]. In the context of a text conversation, specialized AI may be judged more empathetic than a human doctor[99]. Outside of diagnosis, AI may soon be better able to hold a long conversation with an Alzheimer's patient than a human. Whereas 5 years ago, it seemed clear that the "relational" part of care would remain beyond the reach of AI for a long time to come, this boundary is now blurred.

While the study cited above[100] highlights a better diagnosis by the AI than by the doctors in the majority of cases, it points out that this result was obtained under specific conditions: the doctor-patient relationship was reduced to an online chat, a far cry from usual medical practice, which could have enabled the doctor to detect elements impossible for the AI to perceive and

93. Wong, F., Zheng, E.J., Valeri, J. A., et al. (2023) "Discovery of a structural class of antibiotics with explainable deep learning" *Nature*.
94. Olawade D. B., Wada, O. J., David-Olawade, A. C., Kunonga, E., Abaire, O., and Ling, J. (2023) "Using artificial intelligence to improve public health: a narrative review" *Front Public Health*.
95. Comité consultatif national d'éthique pour les sciences de la vie et de la santé (2023). "Diagnostic Médical et Intelligence Artificielle : Enjeux Ethiques".
96. Antun, V., Renna, F., Poon, C., Adcock, B., and Hansen, A. C. (2020) "On instabilities of deep learning in image reconstruction and the potential costs of AI" *Proceedings of the National Academy of Science*.
97. Shen, Y., Shamout, F. E., Oliver, J. R., et al. "Artificial intelligence system reduces false-positive findings in the interpretation of breast ultrasound exams" *Nat Commun*
98. A ChatGPT diagnosis on a child today is wrong 83% of the time according to Barile, J., et al (2024) "Diagnostic Accuracy of a Large Language Model in Pediatric Case Studies" *JAMA Pediatrics*.
99. Tu, T. et al (2024) "Towards Conversational Diagnostic AI".
100. Tu, T. et al (2024) "Towards Conversational Diagnostic AI", *arXiv preprint*

thus improve its diagnosis. This raises the question of the optimal framework for the use of AI by caregivers. Another study[101], focusing this time on the case of radiology, comes to the following conclusions: if the doctor is very confident in his or her diagnosis, then AI assistance is detrimental to the quality of that diagnosis, as it sometimes leads to doubt on the part of the doctor, and thus to him or her wrongly changing his or her mind. On the other hand, if the doctor is uncertain, then AI assistance improves the quality of the diagnosis on average. These studies illustrate the complexity involved in deploying AI systems for healthcare: at this stage of the technology, they need to assist caregivers with recommendations, without giving the illusion of providing the most appropriate diagnosis in every situation.

In France, digital and AI-enabled transformation in healthcare has seen significant progress in recent years, catalyzed by the COVID pandemic. The roadmap for digital healthcare has clarified objectives, strengthened confidence in traditional digital tools and invested in certain infrastructures, notably common repositories.

To develop the interaction capabilities of AI-based tools in the healthcare field, rich data will need to be made available on dialogue between caregivers and patients, failing which this AI will only be trained on non-French-speaking data and in different healthcare systems. The same data will be used to develop AI capable of handling some of the administrative tasks involved.

To develop AI capabilities in prevention, diagnosis and treatment, healthcare data will be just as crucial. Certain AI tools will facilitate the collection and structuring of data, but it will also be necessary to improve the digital infrastructures of hospitals and the availability of this data to researchers. Healthcare databases exist and are of a high quality, thanks to the centralized healthcare system. However, access to these databases remains limited, both because of the way they are financed (each hospital tries to value this data individually, including to pay for its digital infrastructures) and because of the way they are authorized (the scientific and ethical committees of hospitals are not large enough to authorize every request for access to databases).

So far, AI has mainly transformed care for a small proportion of the population (those with certain diseases, or with complex-to-read X-rays). New advances promise to transform the care experience for all, freeing up caregiver time. To achieve this, we need to collectively accept better data circulation and demand in return that it be protected, and that a discussion take place on the evolution of the care system we want.

Recommendation No. 10

Facilitate the circulation of data, the sharing of practices and evaluation to reap the benefits of AI in care, and improve the offering and the daily lives of caregivers.

101. Agarwal, N., Moehring, A., Rajpurkar, P., and Salz, T. (2023), "Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology", *Working Paper*.

2.1.5 BETTER EDUCATION THROUGH AI: INDIVIDUALIZED SUPPORT FOR STUDENTS

Education is one of the areas where the impact of generative AI could be greatest. However, education does not lend itself well to digital transformation. It is made up of human relationships, where the content is not very standardized, despite the programs, and where the teacher must have the autonomy to adapt to the context, to the students. So it doesn't lend itself well to the standardization that digital technology requires to collect and exploit data.

This is why most innovations in digital education have taken place less in the classroom than around it, to make textbooks more accessible, particularly online, or to facilitate administrative procedures. Outside schools, totally digital training courses have been created, but mainly without teachers, for language learning in particular. Even in the most technophile countries, the disruption of education by digital technology has not really taken place. Everywhere, the influence of screens on pupil concentration calls for caution, although a clear distinction needs to be made depending on how screens are used.

In this world, which is difficult for machines to understand, generative AI could be a game-changer. Because generative AI can produce text and images, modify or classify them, and because it is possible to interact simply in natural language with this AI, it reopens the field of possibilities. Its uses are manifold: supporting the production of educational content and sequences, notably by building bridges between subjects; personalizing this content for students; tutoring students; assessing learning, by automating part of the correction process; guiding students; and even training teachers themselves.

To take the case of tutoring alone, it is now well established that personal or small-group tutoring is a highly effective learning method[102], as it helps students to get to grips with the subject matter. Despite its very positive results, it remains costly. Generative AI could provide each pupil with a tutor adapted to his or her level and course, capable of helping the pupil to reason, available 24 hours a day, giving the teacher information on misunderstood elements of his or her course.

Before deploying a generative AI tutor to all students, there are still a few steps to be taken. These systems are costly and slow, not all students have a terminal capable of running generative AI, generative AI models are not very good at spotting errors in reasoning, and above all, we don't yet know how students would grasp such a tool. Perhaps tutoring is only effective when an adult is leaning over the student's shoulder, prompting them to think. Deployed at scale, this solution would raise many questions: should AI be used for homework, the amount of which we're trying to limit in France? Would it require the presence of an adult behind each student?

The technology industry is already working to make AI cheaper, more efficient, capable of reasoning, and able to identify flaws in reasoning. Beyond technology, national education and higher education are the places where teachers and students will be able to experiment to discover precisely when and how AI can be useful for education and training.

Our Commission therefore recommends moving forward on three fronts simultaneously. In the short term, it is crucial to encourage and secure the use of AI by teachers in the preparation and organization of their lessons, but also to help students make intelligent use of AI in their learning.

102. Nickow, A., Oreopoulos, P. and Quan, V. (2020) "The impressive effects of tutoring on PreK-12 learning: a systematic review and meta-analysis of the experimental evidence", *NBER Working Papers*.

In the medium term, it is essential to experiment with a deep integration of AI in education, by rigorously evaluating its contribution, the necessary equipment and the evolution of pedagogies. The Digital Strategy for Education 2023-2027 is a good support for this effort, not least because it organizes collaboration with private players and even provides a "resource account" for educational teams to use commercial tools.

As with public services, five ingredients will be needed to achieve this transformation: clear objectives, trust, skilled people, data, and infrastructure.

Building the confidence of students, teachers and parents in education with AI will require transparency, and evaluation of the tools, starting with those currently being tested ("MIA seconde" for math teaching or "Jules" for homework help for secondary school students)[103]. If students are to make AI their own, they need to be trained in its technical, sociological, historical and philosophical dimensions from secondary school onwards. Local authorities could play a key role by developing projects linking school and extracurricular activities. This appropriation, whether it takes place thanks to or in spite of the national education system, will necessarily lead to a discussion on the use of digital technology and AI at home. After all, students are already using commercial tools to revise with videos or learn math with image recognition. An official stance on the subject would help students and teachers to get to grips with these tools.

As far as people are concerned, the success of AI in education will depend on the involvement of and complementarity with teachers. The teaching profession will be shaken up by the development of AI, which could facilitate the evolution from a "knowing" teacher (disciplinary expert) to a teacher "accompanying" the student, outside the historical paradigm of monodisciplinary transmission. This will require today's teachers to be trained and to take possession of these tools. It may also require a review of the current[104] organization and recruitment methods.

It is vital that the French education system encourages, secures and rewards teachers who experiment with AI. It is reassuring to see teacher training networks, notably the Canopé network, quickly seizing on the subject to train teachers in the creation of educational resources based on AI systems, and to set up pilot projects in real-life situations. Other systems have published guides for parents, pupils and teachers on the use of generative AI, which would benefit from being emulated in France[105].

When it comes to data and infrastructure, the requirements for AI in education are twofold. On the one hand, content data, to ensure that AI in education is relevant and appropriate to the context, the student and the teacher. On the other hand, data relating to the activities of students and teachers (timetables, school life, etc.). This data will only be acquired and made accessible with a major effort on infrastructure and student equipment. The stumbling block remains the fragmentation of content formats (paper or digital textbooks) and platforms, for both students and teachers. Collaboration with the private sector could be a useful way forward, for example, by making annotated copies available to them in order to develop AI that can support teachers in marking.

103. As such, it is imperative that the relevant ministry teams have access to the usage data for these tools.
104. In this respect, the Danish model seems interesting for AI: a group of teachers is responsible for following a class through its entire first cycle (9 years, the equivalent of primary and junior high school in France), teaching several subjects and creating a group spirit.
105. Back in February 2023, the state of North Rhine-Westphalia published a guide for schools and students to incorporate generative AI into their practice.

In terms of objectives, the transformation of education will require us to recall and question the missions assigned to education. The evaluation of AI will require precise, shared protocols to determine the conditions in which digital technology and AI are useful not only for disciplinary learning, but also for the development of critical thinking. We also need to collectively overcome the fear of teachers being replaced. There is no doubt that they will remain central to the transmission of knowledge and the development of critical thinking. Complementarity between teachers and AI should enable us to multiply the impact of each teacher on student learning, by estimating the benefits, costs and evolution of each option[106].

Beyond apprenticeships, AI could be mobilized for guidance support. Parcoursup sees 917,000 applicants pass through its doors every year, guiding them to the 23,000 courses offering state-recognized diplomas. This is a rare instance where a single platform can be used to support the introduction of AI, bringing together in one place processes, information on student profiles and wishes, and choice of courses.

Recommendation No. 11

Encourage the individual use, large-scale experimentation and evaluation of AI tools to strengthen the public education service and improve the day-to-day lives of teaching teams.

106. The operating costs of AI are such that widespread use of a "teacher + RN" system today would not be financially sustainable either. However, costs decrease very rapidly at constant performance. Evaluation and transparency on benefits and costs at scale will be crucial to ensure the confidence of parents and teachers in these tools.

2.2 SOVEREIGNTY: INVESTING IN OUR STRATEGIC AUTONOMY

The digital economy is dominated by a handful of companies, such as Alphabet/Google, Meta/Facebook, Microsoft and Amazon, which also have a strong presence in AI. European companies specialize in traditional industries, while the USA is spearheading the digital technology revolution[107].

This gap has a significant impact on research and development (R&D) spending. Each year, the European Commission draws up a list of the world's top 2,500 R&D investors[108]. In the field of information technology and telecommunications, the latest ranking is striking: the United States has more than four times as many companies as the European Union, and these American companies invest more than six times as much in R&D as European companies. Europe also lags far behind China in this respect[109].

The economic supremacy of foreign players raises the question of our sovereignty, a question that is reinforced by the development and potential of generative AI: if the European economy is largely dependent on foreign companies, what will happen if international relations become strained with its supplier states? Aren't we in danger of seeing a significant part of our economic added value, and even our knowledge, captured by foreign companies supplying AI systems? Since AI contributes to the way we perceive the world, how can we preserve French and European culture if our tools are built on foreign frames of reference?

107. Coste, O., and Coatanlem, Y. (2023) "Tech : quand l'Europe s'éveillera", *Commentaires*
108. European Commission (2020). *The 2020 EU Industrial R&D Investment Scoreboard*.
109. China has more than three times as many companies as the European Union, and these Chinese companies invest more than twice as much in R&D as European companies.

New U.S. measures risk slowing the development of AI capabilities[110] in the rest of the world, including Europe.

On October 30, 2023, the Biden administration issued an executive order on the safe, secure and trustworthy development and use of AI. The text includes reporting obligations for American cloud service providers. They — and their resellers — will have to notify the US government of transactions with any foreign actor relating to the hosting of large-scale AI models whose capabilities could be used for malicious purposes in the cyber domain. At this stage, the scope of these reporting obligations would include the training of any foundation model above a certain threshold of complexity.

These provisions do much more than raise questions about data confidentiality and the protection of business secrecy for our European companies. They contribute to reinforcing American domination, by hindering the development of foreign AI capabilities and facilitating economic intelligence.

To answer these questions, let's define the concept of sovereignty. A sovereign state is one that is not subject to any other state. A radical approach to sovereignty could therefore lead to autarky, in order to control the entire AI value chain. Several countries around the world are moving in this direction, notably Russia and China.

Our Commission considers that an autarkic approach to sovereignty is not appropriate, either in terms of our values or our interests. On the one hand, France promotes an open, democratic framework for society. On the other hand, it would be technically and financially impossible to control all the goods and services making up the AI value chain.

We believe that France should choose the path of interdependence, the right balance between total dependence (no sovereignty) and autarky (no dependence). In concrete terms, this means giving France comparative advantages by positioning itself on a few technological bricks and links in the value chain. To achieve this and strengthen European sovereignty, we recommend developing a dynamic and attractive AI ecosystem. Four pillars are needed: funding, computing power, access to data and talent.

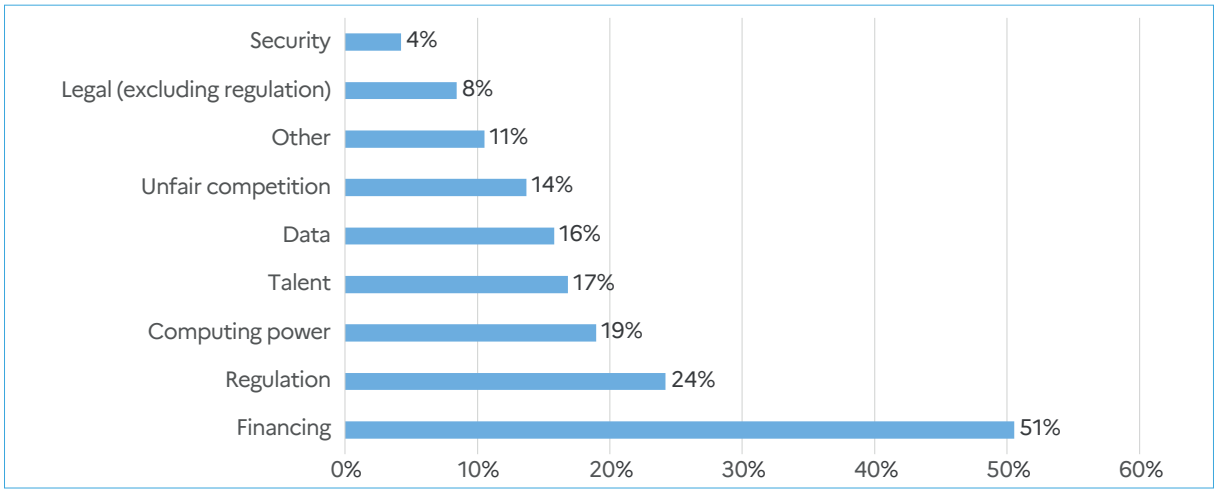

Figure 9 : PMain challenges identified by European generative AI start-ups.
Source: Survey conducted as part of the Generative AI in the European Start-up Landscape 2024
Note: start-ups can identify several challenges

110. Source: French Treasury Department, U.S. Economics Department.

2.2.1 SUSTAINABLE FINANCING FOR INNOVATION: THE NEED TO SCALE UP

2.2.1.1 Investment in AI at least three times too low

Technological innovations are driven either by existing companies or by start-ups. AI is no exception. However, we have just shown that French and European companies are poorly positioned in the digital economy, and therefore less inclined to invest in the emergence of disruptive AI solutions.

In France, the development of a fabric of new companies specializing in AI is therefore essential. To achieve this, specific financing must be mobilized. In particular, this is the purpose of "venture capital", which finances the creation or development of risky but high-potential companies.

However, venture capital investments in AI are far from sufficient. In 2022, investment in France will amount to $2.8 billion[111], compared with $56.8 billion in the United States. To have an investment comparable to that of the United States, France would need to invest between $8.4 and $10 billion per year, i.e. at least between $5.6 and $7.2 billion extra each year. These differences are often explained by cultural differences and the low risk appetite of French players.

Our Commission therefore believes that France should at least triple its investment in AI. This amount is a minimum, as it does not take into account the evolution of financing needs. This trend is achievable. Between 2018 and 2022, venture capital investment in AI in France increased fourfold, from $0.6 billion to $2.8 billion. So we need to stay the course.

What's more, the majority of investments from European funds are small (rarely more than €30 million), whereas the needs announced for certain companies run into the hundreds of millions or even billions of euros. It is essential to remedy this situation, as the absence of European investment means that financing provided by non-European players contributes to the relocation of activities.

The increase in French and European investment in AI does not mean that we should aim for 100% national or European funding. Indeed, the diversity of funding sources contributes to the successful development of start-ups. In addition, foreign investment is testimony to the attractiveness of the national innovation ecosystem. In the United States, 63% of AI start-ups are financed by domestic investors. This figure rises to 45% in France, 25% in the UK, and 24% in Germany.

111. OECD AI Observatory

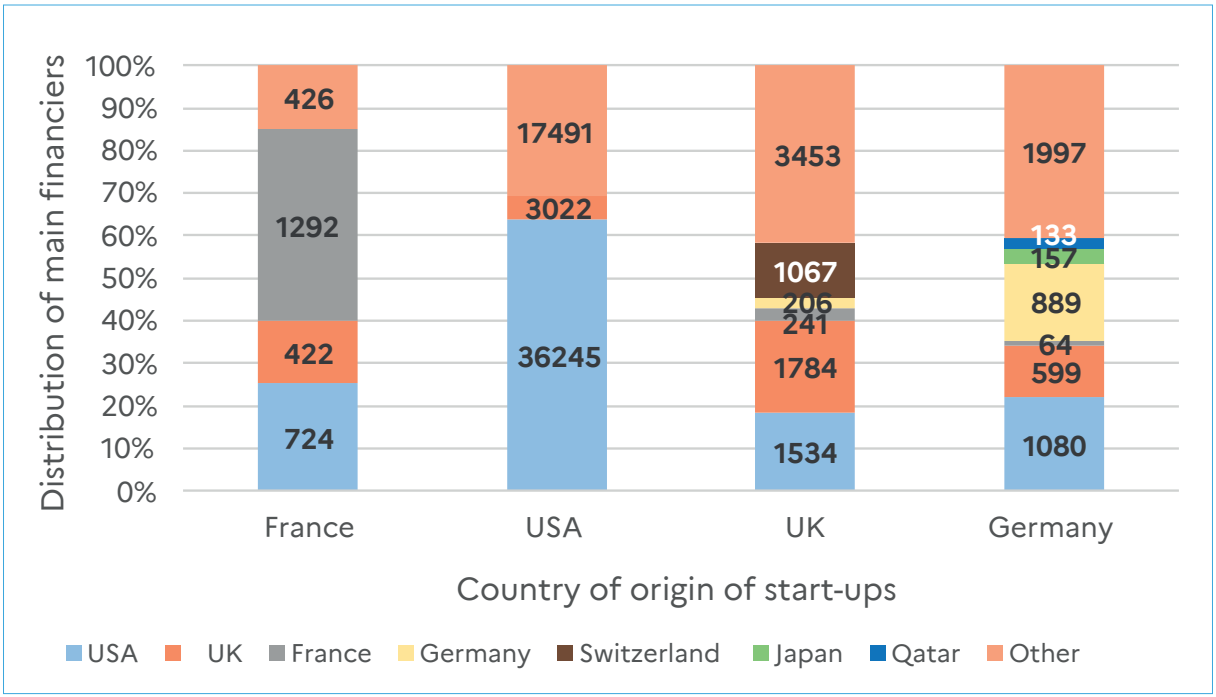

Figure 10: Origin of start-up financing by country in 2022 (all sectors)
Source: OECD.
Interpretation: More than 60% of American start-ups are financed by American venture capital.

2.2.1.2 An investment ambition that requires societal choices to be made

Our Commission estimates that between $5 and $8 billion more needs to be invested in AI every year, tripling the sums currently committed. While this dynamic is a minimum, we recommend that, in this phase of emergence of new technologies with great potential (see part 1), where the first investments are essential to position and take up positions in the value chain, France should set AI funding targets commensurate with its ambitions to be among the leading AI economies. An annual investment target of $15 billion could be targeted.

Scaling up involves redirecting part of private savings, mobilizing between 0.1% (catch-up approach) and 0.3% (pioneering approach) of French household savings[112].

Societal reforms can be envisaged to enable savings to be channeled differently in the long term. Countries with the most numerous and dynamic technology companies offer several possible paths. France could change the tax incentives for life insurance policies, so as to direct funds more towards innovation. It could also change the way supplementary pensions are managed, so as to take on riskier, longer-term investments.

The deepening of European integration is another way of financing innovation. The fragmentation of capital markets in Europe reduces the size of investments. To remedy this, the European Union should establish a true capital markets union, i.e. allow and encourage the free circulation of savings and investments. The first European action plan dates back to 1999... but so much remains to be done! This could be a priority for the next European Union legislature, starting in June 2024.

112. By 2023, French people's savings will total just over €6,000 billion.

A third objective is worth highlighting: enhancing the attractiveness of foreign investment funds, so that they establish themselves in Paris and not just in London. This is not an alternative to the two objectives mentioned above (changing the way French private savings are channeled, strengthening the integration of the European capital market), but rather a complementary solution that will enable the French innovation ecosystem to develop more rapidly.

Our Commission believes that these three projects should be carried out in parallel. They will bear fruit in the long term, and beyond the field of artificial intelligence alone.

To accelerate the development of the AI ecosystem in the short term, we recommend creating a "France & AI" investment fund by the end of 2024. The fund would aim both to support the emergence of start-ups specializing in applied AI and to facilitate the transformation of the economic fabric of SMEs and ISEs. It would mobilize €7 billion in corporate private equity and €3 billion in public support, according to several intervention modalities (a first envelope of funds of funds, a second of co-investments, a third of debt financing of digital transformation projects). In addition to financial resources, the fund will be accompanied by the pooling of business data for certain digital projects. The scale of the resources involved is unprecedented in France, and would represent a high average investment for each company. If each of the 250 major French companies[113] not under foreign control were to participate, this would represent an investment of around €25 million each. We could therefore envisage raising the total €10 billion envelope in two stages. Faced with the risk of economic downgrading, boldness is the key to the emergence of innovative, high-performance solutions and the acceleration of the modernization of French companies.

Recommendation No. 12

Invest massively in digital companies and business transformation to support the French AI ecosystem and make it one of the world's leaders.

The goal of tripling (at least) investment in AI should not be understood as a purely quantitative target. The quality of investment is just as important. Targeting and concentrating resources are key to establishing France's and Europe's superiority in certain segments of the value chain, and thus being able to speak on an equal footing with our competitors and partners. This path to differentiation could focus in particular on the development of open source components, but also on the environmental dimension, by aiming for new generations of AI, from the hardware architecture to the choice of models, which will consume less energy. This differentiation goes hand in hand with the gradual and constant emergence of AI innovation ecosystems in France and Europe.

113. Companies with at least 5,000 employees, or at least €1.5 billion in sales and over €2 billion in balance sheet total.

2.2.2 SOVEREIGN COMPUTING CAPABILITIES: A PREREQUISITE FOR STRATEGIC AUTONOMY

Computing power is an essential ingredient in generative AI. It forms the foundation of the generative AI value chain (see diagram in introduction to report). In concrete terms, it consists of a mix of distributed electronic hardware and associated services.

In fact, the calculations required to train or use ("infer") an AI model are carried out by networks of assembled semiconductors, forming servers which in turn are grouped together in supercomputers or data centers. When a company, public authority or private individual uses an AI model, it therefore calls on the computing services of several supercomputers or data centers.

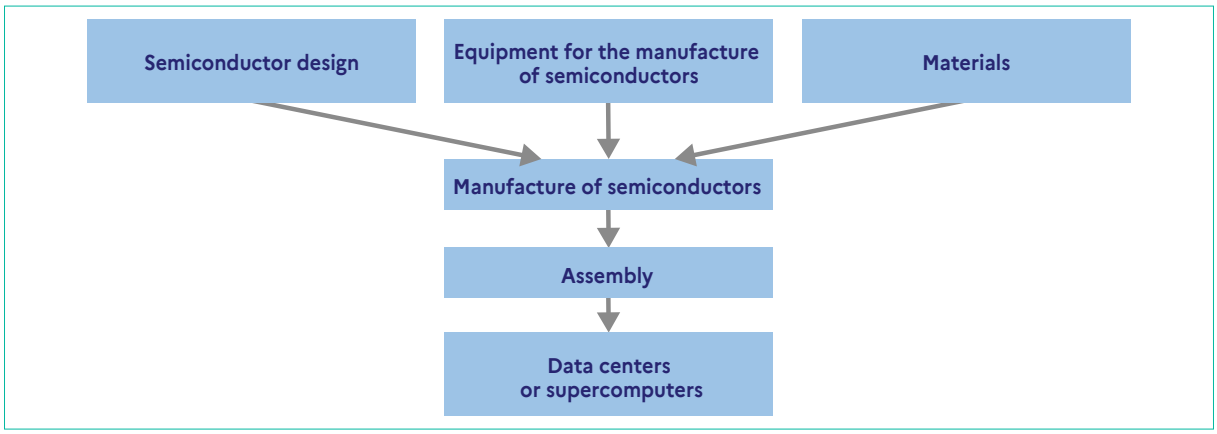

Today's computing power requirements are growing very rapidly as a result of three dynamics. Firstly, the training and use ("inference") of increasingly large models is demanding in terms of computing power. Conversely, major efforts are being made today to reduce the computing power required for training and inference, as this is one of the main cost items in the construction and use of an AI model, and one of the obstacles to its dissemination. The competition is currently as much about model performance as about the ability to develop AI at an acceptable cost. Since 2020, the power required to infer AI models has been increasing faster than the cost of inference (see Figure 11). Finally, even if the training or use of models becomes more economical for the same performance, this could accelerate their spread, and therefore the total computing power required. Despite the uncertainty surrounding these three dynamics, our Commission believes that needs will continue to grow rapidly in the short term, as AI spreads across all economic and social spheres.

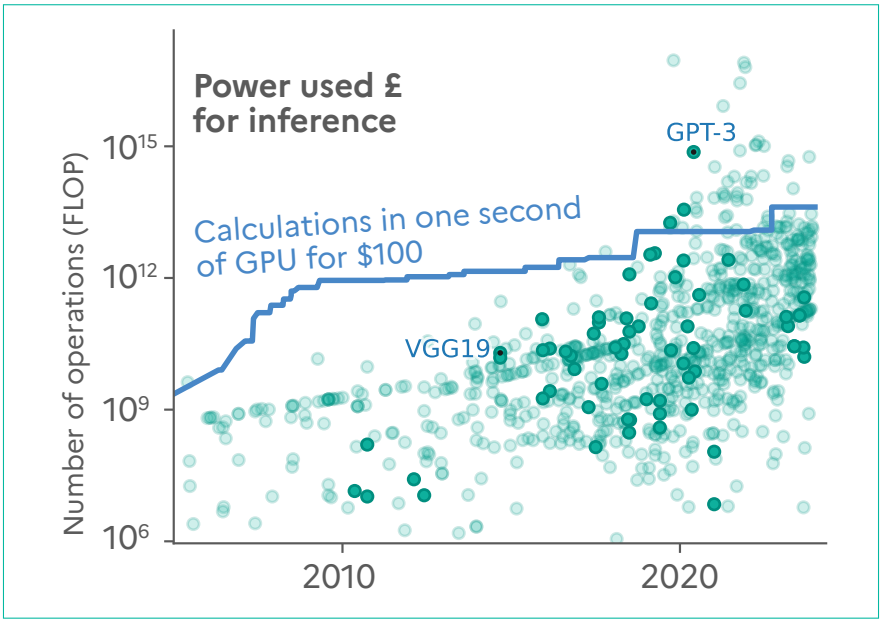

Figure 11: Evolution of the power and cost of AI model inference
Source: French AI Commission calculations (2024).
Interpretation: Each point represents an AI model and the power required for inference. The curve shows the evolution of the number of calculations achievable per second of GPU use for $100.

Today, however, France and Europe are a long way from being able to meet this growing demand. The market is dominated by a handful of companies, mainly American, holding individually or collectively up to 80% of the world market share in certain business segments, whether in the design of semiconductor components specialized in AI, or in the supply of computing power to train or use AI.

This dependence will prevent us from benefiting from a significant share of the added value of AI, which will be captured by foreign players. More broadly, it raises issues of sovereignty: difficulties in protecting data, difficulties in drawing up standards appropriate to poorly mastered technological issues, exposure to geopolitical tensions, etc. It is therefore imperative to remedy this situation.

What's happening in the rest of the world? The example of the United States[114]

In the United States, the Executive Order of October 30, 2023 requires cloud computing capacity providers to declare to the US government the location and capacity of their computing infrastructures above a certain threshold[115].

114. Source: French Treasury Department, Economic Service in the USA.
115. Note that these thresholds, measured in computational operations, are a poor measure of a model's "power" or "capabilities", especially as it is possible to cut the model to train it "piecemeal" or, on the contrary, to exceed the threshold simply because a model has been refined to make it more specific.

2.2.2.1 Investing in AI-optimized semiconductor components

To date, the semiconductor component supply chain for generative AI systems is highly concentrated. When it comes to the design of certain key semiconductors, the American company NVIDIA has a virtual monopoly. In addition to component design, this company also offers a software layer (known as CUDA) that dominates the market.

The manufacture of semiconductor components specifically for generative AI is equally monopolistic, with Taiwanese company TSMC holding 85% of the sub-7 nm semiconductor market. It should be noted, however, that this Asian manufacturer relies on equipment produced worldwide and assembled in machines designed by a Dutch company called ASML.

This monopolistic position is not inevitable. On the one hand, the semiconductor market is set to grow rapidly, driven by demand for computing power. On the other hand, current electronics are not fully optimized for training and, above all, for using generative AI systems (see box below). This dual context of market growth and segmentation is conducive to the entry of new players (Google and Amazon are already designing their own specialized AI components).

Which electronics for generative AI models?

The various stages in the creation of a generative AI application require different networks of semiconductor components to perform the necessary calculations. The learning and specialization stages require massive computation to be carried out intensively, and are performed by GPUs (graphics processing units), which were originally created for video games, not AI. The inference stage, on the other hand, requires much less computing power. It can call on GPUs or more widespread semiconductor components: CPUs (central processing units).

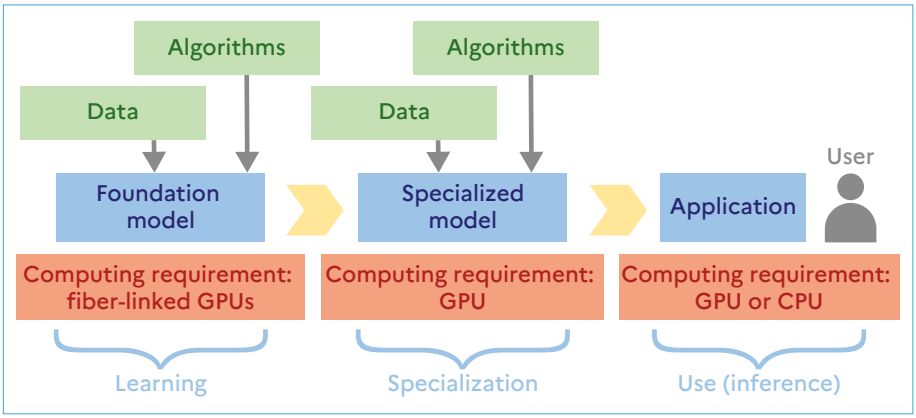

For the most part, the semiconductor components currently used for the various design stages of a solution mobilizing generative AI have multiple applications. A minority of these components are optimized for AI, whether in terms of speed or energy consumption. These criteria are particularly important in view of the massive deployment of AI-enabled solutions in data centers and everyday objects.

In addition to the type of semiconductors to be used, we also need to consider the number of players involved in each phase, in order to measure overall computing power. Indeed, a foundation model enables the creation of multiple specialized models in several fields, each of these specialized models then being used by multiple users. All in all, training and inference create a very significant need for computing power throughout the value chain.

AI therefore represents an opportunity for Europe to position itself in the next generations of semiconductors, which will form a strategic technological backbone. The aim is not to be self-sufficient, but to benefit from the growth of new economic activities and restore the strategic balance of power.

In view of the scale of the investments required, the associated risks and the timeframes inherent in the development of a semiconductor industry optimized for AI, public intervention would appear to be necessary. Such intervention could target specialized chips for inference on the one hand, and chips adapted to embedded AI systems (edge systems) on the other. The framework for such support could be found in the European Semiconductor Regulation, adopted in 2023, which aims precisely to reduce the European Union's dependence in this field.

On the other hand, public support should not be aimed at creating competition between players on the technologies currently in use. In fact, the scale of investment is such that public support would not be sufficient without market traction or an alliance with a digital giant.

To this end, France must continue to support public research and provide financial backing in the form of subsidies and equity capital, in order to encourage the emergence of such an industrial sector. In addition to the design and manufacture of components, European players will also need to support the promotion of open source solutions for the software layer, in order to move away from dependence on the proprietary software layer, which is the predominant one used today.

Recommendation No. 13

Accelerate the emergence of a European semiconductor component industry adapted to AI systems.

2.2.2.2 Increasing the presence of data centers in our territory

Alongside the semiconductor industry, the services associated with supercomputers and data centers form the essential second brick of computing power for AI. In 2023, the combined power of the world's ten most powerful supercomputers increased by 72% compared to 2022[116].

Europe is lagging behind when it comes to public computing power, and lags far behind when it comes to computing power installed by private players.

Europe's public computing power is significantly lower than on the other side of the Atlantic. It currently hosts 143 of the world's 500 largest supercomputers — most of them public — (compared with 161 in the USA), representing 24% of the corresponding computing power (compared with 53% in the USA). However, the European Union and its Member States are investing in a program called EuroHPC (European High Performance Computing). Future installations of massively parallel, very large-scale "exascale" supercomputers are planned for Germany in 2024 ("Jupiter") and France in 2026 ("Jules Verne"). At the same time, however, the USA and China have already begun work on so-called "post-exascale" supercomputers.

The computing power of private players is largely based in the United States. This predominance is twofold. Firstly, data centers are generally much more present in the USA than in Europe: the USA has 2,109 data centers (representing 37.8% of the world's data centers), compared with 1,244 data centers in the European Union as a whole (representing 22.3%)[117]. What's more, the biggest providers of services from these data centers (cloud) are American. Three major players account for 80% of the increase in spending on infrastructure and applications for public cloud services in France.

This imbalance tends to worsen when it comes to computing power dedicated to AI. In 2023 alone, the combined purchases of Meta, Microsoft, Alphabet and Amazon will amount to 400,000 GPUs, compared to a few thousand units for the European players. In 2024, Meta alone plans to buy 300,000 latest-generation graphics cards specialized in AI training (NVIDIA's H100 model or equivalent), while the company already has 300,000. Admittedly, some of these will be installed on infrastructures in Europe. However, to exploit the full potential of these machines today, they need to be located in large-scale infrastructures, which are less present in Europe than in the United States.

It would be tempting to consider that Europe's lag is not a problem, that it's simply an extension of a deficit in hosting services. In fact, in a globalized economy, it is not necessary for all AI used on European soil or by European consumers to be trained and deployed on strictly European infrastructures.

However, Europe's lag in computing power poses three difficulties. Firstly, there are sensitive AI applications for which it is not feasible to train models or deploy them on non-European computing power. Today, these uses are certainly limited to a few players, notably the military, who have their own infrastructures. This will no longer be the case in the future, whether we're talking about tax, healthcare or justice. Secondly, in a world where AI will be integrated into many aspects of our social life and economy, resilience and strategic autonomy require a minimum amount of computing power on our soil. To draw a parallel with electricity: it's not necessary for all the electrons powering our economy to be produced in France, but it would be dangerous to be forced to import 80% of our electricity. Computing power is a control point in the AI supply chain, because it can be detected, prohibited, quantified and concentrated. As a

116. Based on Top500 data for November 2023. The source is identical for the following paragraph.
117. Based on DataCenterMap data from January 2024.

result, some already see it as a neuralgic point for governing AI[118]. Finally, the rise in demand for computing power will lead to a deterioration in our trade balance, if France does not simultaneously increase its supply.

What's happening in the rest of the world? Examples from Sweden, Taiwan and Japan[119]

In Sweden, the highly decarbonized energy mix (fossil fuels account for 1% of electricity production) and the lower price of electricity are an advantage for the installation of supercomputers. German translation services start-up DeepL, for example, chose this country for its new Mercury supercomputer, which is now considered the most powerful computer in Sweden.

In Taiwan, data centers and computing power for AI must be located on Taiwanese soil, for resilience reasons. On the other hand, there are no restrictions on the nationality of computing power providers. In fact, the three American hyperscalers (Amazon Web Services, Google Cloud and Microsoft Azure) provide almost all the country's hosting and computing power.

In Japan, the government subsidizes the use of computing resources to support the development of basic models for generative AI. It also supports the development of data centers by subsidizing a portion of the infrastructure costs for companies providing these services. Public funding is higher in rural areas.

Our Commission considers it imperative that France and Europe become a major hub of installed computing power with three objectives: (i) provide public computing power for sensitive use cases; (ii) provide accessible and affordable computing power to stimulate research and development for AI start-ups; (iii) be able to train and use the most advanced AI models on European soil.

For sensitive use cases, research and start-up support, investment in public supercomputers should continue. Access to supercomputers financed by Europe or France is free for public or private players who contribute to open research. This approach enables them to benefit from publicly funded infrastructures in return for a contribution that advances collective knowledge. This dynamic must be pursued to enable European research players to remain at the technological frontier.

For private or non-sensitive uses, it is imperative to develop data centers in France. Setting up a data center in France raises issues of land management, coordination of players, administrative procedures and access to competitively priced electricity. In this respect, electricity consumption represents a significant cost for market players: it is therefore a key factor in the decision to locate. Faced with the financial support offered by certain US states, which are keen to attract data centers, France could strengthen its tax competitiveness. To ensure that our players developing and using AI have access to affordable computing power, we also need to take note of the very tense context in which the market for specialized AI processors finds itself. Access to computing power will be a major differentiator in the next few years, and cost a crucial parameter for the competitiveness of players developing AI.

118. Sastry, G., Heim, L., Belfield, H., et al. (2024) "Computing Power and the Governance of Artificial Intelligence" *Working Paper*
119. Source: French Treasury Department, Economic Service in Sweden and Japan.

What targets should be set for computing power in France?

In a fast-growing market, with uncertain technological parameters, our Commission has attempted to estimate computing power requirements using several methods. Computing power requirements vary widely, depending on whether we consider a company training foundation models, refining more specialized models, or using AI models in its operations.

By 2026, France could set itself the target of being able to offer computing power on its soil to support 5 companies training next-generation foundation models in a reasonable time, 50 companies refining specialized foundation models and 1,000 companies using AI models in their operations.

Based on the data published for Llama 2, training three versions of a state-of-the-art model in 2023 required the equivalent of 100 latest-generation GPUs in 2024 ("H100 equivalent") for just over four months. Next-generation models (from 2024) will probably require around ten times as much computing power for training, or 1,000 H100 equivalents, and the same for inference. Companies refining models will need 10 times less power for training (100 GPUs) and inference (100 GPUs). Companies using models only will need 100 times less computing power (10 GPUs), concentrated on inference.

In 2024, France could therefore set itself the target of securing at least 30,000 H100-equivalent GPUs from private players. This would represent 3% of worldwide production forecast by NVIDIA for 2024, which corresponds to France's share of world GDP. If this target were met, France would have 21 MW of private computing power, or almost 2 exaflops[120].

With a utilization rate of 60%, these 30,000 GPUs would consume 110 GWh of electricity, equivalent to the annual consumption of 50,000 people, or 0.22% of our electricity exports in 2023. With this kind of low-carbon electricity production capacity, France could double or triple this target in order to host computing power for its European neighbors. At European level, the same reasoning would lead us to set a target of 150,000 H100-equivalent GPUs (or 105 MW).

Finally, when the "Jules Verne" supercomputer is delivered, France will have just over one exaflop of public computing power. Given the timeframe involved, we need to start work right away to increase public computing power after "Jules Verne".

This target does not include computing power requirements for a digital giant. Their needs are such that simply hosting one or two AI computing centers for one of them would easily double the number shown here. France could even enter into discussions with one of them to host one of their computing centers. Generally speaking, these figures should be taken as short-term projections, given the rapid growth in demand.

120. Unit of measurement used to calculate the power of a processor, corresponding to a quintillion floating-point operations per second.

With this in mind, we propose to act simultaneously on the supply and demand sides of computing. On the supply side, we recommend speeding up the expansion of French and European "exascale" supercomputers, launching a group purchasing operation for the ecosystem in the short term, and setting a target for the establishment of computing centers in Europe, with a public guarantee for the use of computing power, as well as support for implementation and electrical connection. This call for tenders aims both to increase computing power — to meet growing needs — and to diversify hosting solutions — to encourage the emergence of European solutions. It is therefore not closed to solutions from non-European suppliers. On the demand side, an AI tax credit would support research and development projects in the rental of computing power, on condition that they use a computing center established in France. The aim here is not so much to subsidize computing in general as to encourage the installation of computing power in our territory.

Recommendation No. 14

Make France a major center of computing power, in both the short and medium term.

2.2.3 ACCESS TO QUALITY DATA

AI needs data in huge quantities, as does generative AI, even if it needs cultural data more than personal data. The evolution of these increasingly precise models, the desire to save computing power and the depletion of available data (for text, the largest models have already used most of the existing corpora) will probably lead to future emphasis on data quality rather than quantity. In Europe, in France, the challenge is not just technical, but cultural, in terms of the presence and discoverability of language, images and videos.

2.2.3.1 Personal data

Artificial intelligence makes it possible to grasp the mass of available data, which human intelligence can no longer embrace. For example, over 5 million scientific articles are published every year, half of them in the field of medical research alone. Of course, it's impossible for any single researcher or team of researchers to read them, let alone evaluate and analyze them.

Conversely, data is an indispensable ingredient in recent developments in artificial intelligence. For example, the discovery of a new antibiotic to combat staphylococcus aureus[121] (see *2.1.3. Better care through AI*) was only made possible after training an AI system with almost 40,000 known results of molecular structures of existing antibiotics. This figure is very low compared to the learning and processing capacities of such systems, and is due in particular to the quality of the data used for this training. To put it another way, the quality of the data used for training is at the heart of an AI system's reliability.

This data is not necessarily personal, but it's clear that a lot of the data of interest for training AI is personal. In healthcare, of course, but not only. Even generative AI, a priori more interested in cultural data, may need it to develop a specific interaction capability. In education, training a model capable of interacting in a credible and relevant way with a student will probably require training on student-teacher dialogue data, which is personal data.

Exploiting the potential of artificial intelligence and deploying it in the service of human beings therefore requires that researchers, developers and innovators have access to massive, reliable and easily manipulated data, whose representativeness and quality can be assessed. Against a backdrop of rapid technological change and heightened competition, this access must also be open to them rapidly, and the data used without undue constraints, at the risk of further favoring the players in place or seeing others appropriate our research and innovations, getting ahead of us in their experimentation and dissemination.

As things stand, however, the difficulties of accessing data and the constraints on its use, regularly regarded as excessive, are observations widely shared by AI players of all kinds (companies, researchers, laboratories, public and private institutions, associations, etc.). These difficulties are of two kinds.

Firstly, **certain French rules and practices are more restrictive than the European framework** when it comes to processing personal data. The current framework is defined by the General Data Protection Regulation (GDPR), which came into force in 2018.

The GDPR has completely reversed the logic of the law that has prevailed in France since the 1978 Data Protection Law (*Loi Informatique et Libertés*). Whereas the possibility of processing personal data was based on prior authorization or declaration procedures with the administration, the GDPR has established the principles of freedom and responsibility: players

121. Wong, F., et al. (2023) "Discovery of a structural class of antibiotics with explainable deep learning", *Nature*

are free to create and implement personal data processing, subject to ensuring themselves that such processing complies with the principles and rules set out in the European regulation. In particular, they must analyze the specific risks that may be created by the most sensitive processing operations, and take appropriate measures to remedy the situation. In return for this freedom, instituted with the precise aim of fostering innovation, personal data protection requirements have been strengthened, as have the *a posteriori* control and sanction powers of the authorities in charge of data protection. In France, this authority is the CNIL (Commission Nationale de l'Informatique et des Libertés).

In France, this development has not been implemented fully. There are still some prior authorization procedures not provided for under European law[122]. This is particularly the case for access to health data for research purposes. A simplified procedure for declaring compliance with reference methodologies does exist, but it is far from widespread. In practice, the simplified procedure remains the exception to the[123] prior authorization procedure, since the slightest deviation from these methodologies requires prior authorization, which can involve up to three levels of prior authorization. Similar cumbersome procedures can be found in the fields of public order, security and justice[124].

Secondly, there is a **growing discrepancy between the logic of protecting the individual and the evolution of collective use of data**. Whereas the law on the protection of access to and use of data has, since its origins, focused on the individual (personal data), the development and use of AI systems is focused on massive, aggregated and dynamic data.

From a technical point of view, several key notions of the GDPR are thus awkward to handle given the way AI works. This is particularly true of "data controller", where the division of responsibilities between the developer who has trained a generative AI and makes it available to third parties, and the end-user of the system for his or her own needs, does not necessarily appear to be self-evident. The notion of the purpose of processing, which determines the type of data that can legally be used, and on which the consent of data subjects is based, is also more complex to grasp, given the many possible uses of a generative AI once it has been trained.

In terms of principles, the very notion of personal data, which is the key to applying the GDPR, raises questions in a context of increasing use of collective data. As it stands, legally, only a process of anonymization of personal data allows one to "opt out" of the GDPR's personal data protection regime. Yet technology is increasingly opening up possibilities for re-identifying anonymized data. Furthermore, more collective data management (currently in its infancy in the GDPR) could improve the protection of interests and the exercise of rights, particularly in the face of global data players. Indeed, whether through associations, trade unions or any other organized collective, agreements relating to data processing and the use of AI systems could make it possible to increase the effectiveness of the guarantee of everyone's rights.

Recommendation No. 15

Transform our approach to personal data to protect while facilitating innovation to serve our needs.

122. The national legislator has mobilized leeway opened up by the GDPR to Member States in the direction of maintaining or even strengthening prior authorization procedures for health data.
123. *Fédérer les acteurs de l'écosystème pour libérer l'utilisation secondaire des données de santé*, report by the mission led by Jérôme Marchand-Arvier (December 2023).
124. A specific procedure for the authorization by decree of the Conseil d'Etat of all processing carried out on behalf of the State in the fields of public order, security and justice has been maintained, making any change, however minor, in the way personal data is processed in this field particularly burdensome.

Our Commission recommends that we continue to modernize our approach to data, by better combining protection and innovation. A number of changes are underway. In 2023, for example, the CNIL set up an AI and publication department, and a few months ago published "AI factsheets" designed to guide developers of AI systems. In this way, the authority is consolidating its mission to support innovation players.

To go further, we recommend eliminating prior authorization procedures for access to health data, and reducing the CNIL's response times. This move should be accompanied by a reform of the CNIL's mandate, to include a focus on innovation. This will entail adjusting the composition of its college, so that a broader range of skills is represented (innovation, research, etc.), and strengthening its operating resources. More broadly, it will be necessary to strengthen the coordination of all regulators working in the field of data and digital technology.

Beyond that, it's important to find the path to collective governance of data that could, right now, make use of the under-exploited legal leeway provided by the GDPR and, in the long term, pave the way for an evolution of the legal framework that would take better account of the changing ways in which data is used. In calling for improved governance of public sector open data and the immediate implementation of the new European data regulations, our Commission is calling for the full exploitation of a formidable resource of general interest and economic growth.

2.2.3.2 Heritage data

Since the early 2000s, France has had an outstanding collection of digitized works. This asset has been cultivated since the time of François I, as a result of an ongoing policy of universal legal deposit (books, magazines, press, images, sounds, maps, cinema, video games), extended to television and radio, and through acquisition and conservation policies. Linguistically, artistically and culturally, this corpus represents a very broad intellectual and emotional landscape of the French-speaking world.

What's happening in the rest of the world?
The examples of India and South Korea[125]

In India, AI is identified as a lever for strengthening social and economic inclusion. The country intends to be at the forefront of applied AI solutions, and is focusing on gathering quality data. In particular, the "Bhashini" program aims to provide companies with audio and text datasets from India's 22 national languages, so that they can develop innovative solutions.

In South Korea, the Ministry of Culture and the National Institute of the Korean Language plan to build a high-quality Korean data corpus of around 120 million phrases to support the development of "K-ChatGPT that speaks Korean well", then expand it to 1 billion phrases by 2027.

The appropriate mobilization of this unique heritage in the training of artificial intelligence systems therefore represents an issue of cultural diversity and sovereignty. At present, AI models perform less well in French than in English, having been trained mainly on English corpora.

125. Source: French Treasury Department, Economic Service in India and South Korea.

Meeting this challenge calls for rapid change. While a large proportion of content is already available, digitized and accessible via APIs, there is another part that warrants accelerated funding to make it available, and to free up any rights that may exist, such as photographers' rights. Most institutions, such as archives, are destined to be open and available digitally. However, there are infrastructure issues that need to be addressed to speed up access.

Legal obstacles remain. For example, there is the legal deposit, which is fed on a daily basis by rights-holding publications from publishers of printed matter, phonograms, films and audiovisual programs. This deposit cannot be used as an open fund for making AI available. Moreover, our international commitments prohibit this. Metadata, on the other hand, could be opened up. More broadly speaking, economic models still need to be devised for research and commercial uses.

To remove some of these barriers, a vast plan for making content available is needed, with funding on a scale commensurate with this objective. This plan will have to take into account the particularities of international standards for each category of content.

Three avenues need to be explored. Firstly, it is important to create a framework of trust for access to rights-free content, notably by making metadata available. Secondly, we need to prepare for the creation of a register of available content (including all relevant information, such as access conditions), for which the Bibliothèque nationale de France, the Institut national de l'audiovisuel and the Ministry of Culture could act as trusted third parties. A platform could then bring together the major holders of public cultural data, or even private cultural data holders who so wish, and AI developers. Finally, the aim is to build public tools available to private players, for example to combat misinformation by enabling verification of past audiovisual information.

Recommendation No. 16

Set up a technical infrastructure to bring together AI developers and holders of cultural heritage data.

Beyond that, **public data remains insufficiently open**. Worldwide, the movement to open up public data began in the late 2000s, peaked in the mid-2010s and then faded. We need to revitalize this movement, in the new context of generative AI.

The public sector is indeed a source of both abundant and high-quality data. The reuse of this data, in particular through AI, can enable the creation of many new services, whether for commercial or general interest activities. The still too limited opening of public data[126] deprives France of many benefits.

Examples include health data for research, road traffic data to optimize the use of means of transport, agricultural data to improve yields and reduce the use of inputs, and so on. The circulation of data within the public sector simplifies administrative procedures, avoiding the need for multiple administrative files and the production of supporting documents, and improves the effectiveness of public policies. By way of illustration, it enables us to identify people who meet the conditions for receiving social benefits, and thus avoid the phenomenon of "non-application" due to a lack of knowledge or sufficient technical skills on the part of beneficiaries.

When it comes to public data, the obstacles encountered are practical rather than legal. The principle in force in Europe since 2003[127] is indeed that of opening up and freely re-using data produced or held by all public bodies and belonging to the public sector. The very recent amendments to the European regulation on data governance[128] aim to extend the scope of this principle and facilitate its application, notably by providing a single European platform for the re-use of public data.

2.2.3.3 Data protected by literary or artistic property rights

A global movement protesting the use of protected data for AI training is emerging on both sides of the Atlantic.

In the United States, the use of protected data is subject to the authorization of the rights holder, unless the use is considered "fair use". The application of this exception is contested by authors and artists, who have brought numerous lawsuits against model publishers. The outcome of these legal actions will be decisive, but also uncertain, as the assessment of "fair use" is highly contextual. In these conditions, some AI players — the same ones who initially contested the application of copyright — are already seeking to contractualize the use of protected content by entering into licensing agreements with rights holders.

In Europe, the use of protected data to train a generative AI system is governed by Article 4 of a 2019 European Directive ("copyright and related rights in the Digital Single Market"; note that this predates the proliferation of generative AI models). This Directive imposes an exception for Text and Data Mining (TDM), which will be transposed into French law in 2021[129]. The monopoly is therefore suspended for these operations, for all purposes (including commercial ones) and to the benefit of all players.

126. See, for example, E. Bothorel's mission report, *Pour une politique publique de la donnée*, December 2020.

127. Directive 2003/98/EC of the European Parliament and of the Council of November 17, 2003 on the re-use of public sector information. The principle is now enshrined in Directive (EU) 2019/1024 of the European Parliament and of the Council of June 20, 2019 on open data and the re-use of public sector information.

128. Regulation (EU) 2022/868 of the European Parliament and of the Council of May 30, 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Regulation)

129. Order No. 2021-580 of May 12, 2021 transposing Article 2(6) and Articles 17 to 23 of Directive 2019/790 of the European Parliament and of the Council of April 17, 2019 on copyright and related rights in the digital single market and amending Directives 96/9 and 2001/29/EC.

However, certain conditions must be met in order to benefit from the exemption. Firstly, the data must have been obtained lawfully. It has now been proven that generative AI models have been trained on databases containing pirated content. Secondly, in return for the infringement of their exclusive rights, data owners can opt out of the use of their data. This return to exclusivity may enable the holder to regain negotiating power with a view to obtaining remuneration.

Faced with the growing use of their data, French rights holders are multiplying these opt-out strategies, with worrying consequences: weakening of the reliability of the results produced by AI, absence of French content and more generally of authentic creations, which can lead to the production of stereotyped and mediocre results by generative AI models accessible to the public.

The current situation is therefore contentious. Rights holders are demanding that their rights be respected, but their actions have been hampered by the lack of transparency regarding the training data used. The European regulation on artificial intelligence (AI Act[130]) provides for such transparency both upstream, with regard to training sources, and downstream, to combat misinformation.

Upstream, rights holders must have the opportunity to be informed about the data used, in order to ensure that it has been obtained lawfully. Contrary to the current trend, the training of generative AI models should not be a new, sustainable outlet for companies to pirate protected data by "scrapping" or "crawling". Transparency is also essential to enable rights-holders to exercise their opt-out rights with discernment, to authorize against remuneration where appropriate, and to monitor compliance with the parties' commitments. Without this principle, the right conferred on them in return for the TDM exception is deprived of its meaning, as it cannot be made operational.

In practical terms, this means that AI publishers must make public the fact that they have used data protected by copyright, and indicate from which entity the rights holder can obtain information. Such data may be made public or communicated on request to persons with an interest in the matter, if it is necessary to preserve confidentiality. The European AI Regulation does not stipulate that this obligation applies to publishers of specialized models (but only to general models). It will therefore be necessary to regularly assess the relevance of this scope of application of the transparency obligation, in the light of technological developments and model training methods.

The contact platform identified above could usefully integrate any opposition from rights holders and, where applicable, the conditions of use for training (license, remuneration, etc.). Ultimately, the idea is to be able to easily build up quality data corpora for training AI models while respecting rights. Not all such corpora have the same value, and the marginal value of a particular work or author can be quite low. On the other hand, their collective value can be significant. This is the case of the legal deposit of the Institut National de l'Audiovisuel, which represents 47.85 Petabytes of quality multimodal data, or 100 times the size of Common Crawl, a partial "backup" of the Internet.

Downstream, AI-generated cultural products and information must be clearly labeled. AI-generated content must be recognizable. This requirement, which is gaining ground internationally, is essential for tackling parasitism, identifying unreliable information and combating deepfakes. It is also essential for informing consumers who, as with food, have a right to information on how, where and what ingredients are used in the manufacture of products. Finally, downstream transparency is essential in relations between customers and

130. Political agreement of the proposed regulation voted in Coreper in February 2024.

suppliers, to manage any potential liabilities. In this respect, we support a broader form of transparency than that set out in the AI regulation.

Recommendation No. 17

Implement and evaluate the transparency obligations set out in the European AI Regulation by encouraging the development of standards and a suitable infrastructure.

The quality, diversity and depth of data from European culture are clearly an asset for artificial intelligence publishers. The implementation of transparency rules should enable the emergence of a market for "white" data whose use by software publishers is legal and secure from a legal point of view. This could be an advantage over the United States, whose "fair use" rules will be challenged in the courts for many years to come. A strong collective mobilization in Europe would be essential to encourage the creation of public and private data warehouses, accompanied by rights mapping and remuneration models. Only this mobilization will make it possible to avoid the three risks before us: concentration (if only a few very well-funded AI publishers can access cultural data at a high price), cultural submersion (if AI is trained on purely American data), or circumvention of property rights (if a "black" or "gray"[131] cultural data market develops).

131. Without any agreement between rights holders and AI publishers, some are already talking about the possibility of producing synthetic databases reproducing the characteristics of illegally obtained cultural databases, without directly containing protected works.

2.2.4 ATTRACTING TALENT TO BUILD THE TECHNOLOGIES AND USES OF TOMORROW

2.2.4.1 Bringing the world's most qualified people to our territory

Talent is one of the major pillars for accelerating the development of AI, creating a dynamic ecosystem and contributing to France's sovereignty. The notion of talent is taken in its broadest sense as any person with rare skills who can contribute to the development of AI, whether in its technical, scientific or commercial dimensions. Talent is the main resource for both public research and companies in the field, accounting for around a third of their costs.

However, France is in strong competition with the United States, and to a lesser extent with the United Kingdom, and French ecosystems are not (yet) as renowned. Many young French graduates, particularly from engineering schools, are moving abroad to work there. The AI departments of American companies, for example, include a large number of French people, underlining the quality of French training, but also our country's lack of appeal.

What's happening in the rest of the world? The example of the United States[132]

In the United States, the Executive Order of October 30, 2023 implements a set of measures aimed at attracting the world's most qualified AI professionals. Firstly, it simplifies the visa application process and prioritizes the processing of these applications. Secondly, it revises the criteria for granting fast-track visas (J-1). Finally, it implements a comprehensive program to identify and attract foreign AI talent to the US.

The situation has improved in recent years, with American companies setting up research laboratories in Paris. These companies have recruited young French people in Paris, who have gone on to set up start-ups in France. This dynamic has contributed to the emergence of a French AI ecosystem and its attractiveness. A foreign company specializing in AI code generation, for example, is in the process of setting up in Paris.

Nevertheless, the momentum is insufficient. Our Commission considers that of the three to five thousand highly qualified international profiles likely to have a significant impact on the growth of the AI ecosystem, France needs to attract between 10 and 15%. The skills required are diverse. We need to attract the talent capable of developing major language models and future generations of AI. We also need to attract the best researchers and engineers in the field to create new products and services, as well as leaders of innovative growth companies. The same applies to qualified product engineers in the marketing field.

A diversity of profiles will not only foster innovation, but also limit the risk of bias, which can occur when teams have profiles that are too similar. This diversity issue is particularly important

132. Source: French Treasury Department, Economic Service in the USA.

in France, where the proportion of foreign engineers is low[133] compared to the situation in Silicon Valley or London. The attractiveness and diversity of the innovation ecosystem will therefore also depend on attracting foreign students from reputable degree programs abroad, so that they can continue their studies and then start their careers in France.

Recommendation No. 18

Attract and retain international talent with scientific, entrepreneurial and managerial skills in the field of AI.

Attracting talent goes hand in hand with attracting companies. The latter can play a catalytic role in the French ecosystem, in particular by attracting qualified foreign profiles to our country. In addition to a range of measures designed to support the growth of French start-ups, measures to attract foreign groups to France can therefore contribute to the sovereignty of France and Europe.

Our Commission recommends the creation of an "AI talent attractiveness" mission, inspired by the one set up to attract financial profiles to France in anticipation of the UK's exit from the European Union (Brexit). This mission consists of a targeted, personalized approach to identifying companies and talent abroad. It needs to be implemented quickly to benefit from the current momentum of the French AI ecosystem and its growing reputation abroad.

Apart from public research, recruiting individual talent is the responsibility of companies. However, the State must create the conditions to facilitate the installation of qualified profiles, notably through assistance with administrative formalities: visas, children's schooling, information on the tax regime for impatriates, etc. These conditions will also be offered to employees of foreign companies who settle here, and to students who have been canvassed.

133. The proportion of foreign engineers in the Paris Region is estimated at less than 10%. In Silicon Valley, the figure is 68%.

2.2.4.2 Enabling researchers to research

It is essential to have high-level public research in AI in France. Public research enables us to explore longer-term avenues that do not necessarily meet immediate economic objectives. It also guarantees the quality of training for future engineers and researchers who will join the private sector. It is a factor of attractiveness.

Following C. Villani's report[134] (2018), which highlighted this issue, several measures have been taken. Several laws[135] have facilitated the use of fixed-term or open-ended contracts, made bonus schemes more flexible, simplified the declaration of multiple activities[136], strengthened relations with companies, etc. In particular, the possibility of offering fixed-term or indefinite-term contracts with remuneration conditions higher than those offered by the State has facilitated the recruitment of young researchers or renowned foreign researchers.

However, many practical obstacles remain. When it comes to remuneration, for example, the researchers interviewed by our Commission feel that they almost always end up finding solutions for additional remuneration, but at the cost of complex, tedious and time-consuming arrangements. At a time when public-sector research cannot compete with business when it comes to salary conditions (salaries are commonly ten times higher), the simplicity of using additional remuneration is essential: red tape discourages researchers.

Enhancing the attractiveness of French public research is not only imperative, but also achievable. France's public research sector still offers a number of conditions that are appreciated by researchers, such as access to high-quality students and academic freedom for researchers (a point that distinguishes France from the USA). In the field of artificial intelligence (taken broadly), a salary in excess of €80,000 per annum, more or less depending on age and reputation, should be achievable without any administrative burdens. In addition to existing researchers, we also need to provide additional remuneration for doctoral and post-doctoral students, who make a major contribution to public research and whose quality is an appeal for more mature researchers.

Countless other disproportionately time-consuming administrative complexities need to be tackled. Validating hirings, signing contracts with companies, putting together administrative files in response to calls for projects, or even setting up a training program all take up at least a third of a public-sector researcher's time. These activities are time-consuming, laborious, discouraging and offer little added value in terms of researchers' skills. They therefore slow down the production of new knowledge. In a world of AI innovation that is progressing daily, these complexities are unthinkable.

Finally, working conditions are not up to international standards. By way of illustration, foreigners recruited into the French public research sector do not benefit from comprehensive support to facilitate their move (relocation, assistance with administrative formalities, hiring assistance for spouses, school enrolment for children, etc.). The *Choose Paris Region* scheme has made some progress, but more needs to be done.

Although France has initiated a process to simplify administrative measures in research, the timetable is not in line with the global development of AI. For this reason, we are proposing to move faster for artificial intelligence by setting up an "AI Exception". In the form of an experiment, this exception aims for "zero hindrance for researchers", notably through a commitment on response times to researchers' requests. The AI exception should also make it possible to raise

134. Cédric Villani, March 2018, For a meaningful artificial intelligence: towards a French and European strategy.
135. In particular, the research programming law and the law for the growth and transformation of businesses.
136. Despite simplification, many establishments still require authorizations that have to go through numerous layers of signatures.

remuneration and facilitate part-time work with companies or other socio-economic players in AI. Beyond this, it is essential to at least double funding for public research in AI to join the world's top 5 in terms of production and publication in AI.

Recommendation No. 19

Assume the principle of an "AI Exception" in the form of an experiment in public research to boost its attractiveness.

2.2.5 WIDESPREAD DEPLOYMENT OF ARTIFICIAL INTELLIGENCE IN OUR ECONOMY

2.2.5.1 Using AI to boost productivity and competitiveness

The French and European ecosystem specialized in artificial intelligence must develop in parallel with the widespread adoption of AI systems by the companies that currently make up our productive fabric. We need to focus not only on the AI economy, but also on the economy with AI. The potential for productivity gains is indeed immense[137]. The particularities of generative AI (realism, simplicity, speed, aptitude) open the way to rapid, almost immediate productivity gains. However, the main benefits will require transformational action (in terms of data infrastructures, organization, human resources, etc.).

The impact of previous major innovations on French productivity is often considered moderate. One of the main reasons for this is the delay in adopting these technologies compared with other countries, notably the USA. To reap the benefits of AI, French companies therefore need to adopt AI systems quickly. If they are slower in this adoption, they will lose market share to companies that use AI faster. The opposite is equally true: if France mobilizes faster, it will be able to gain market share both domestically and in export markets.

What's happening in the rest of the world? The spread of AI in Europe and abroad

AI is a general-purpose technology, with no single measure of its use (as with electricity). It is therefore difficult to compare the speed at which AI is spreading through the European economy, a fortiori generative AI. And 2023's publications on the subject don't always distinguish between experimentation and production projects. Be that as it may, the hearings conducted by the Commission are convergent.

Firstly, the dynamism of the AI ecosystem is greater in the USA than in Europe (activity on contribution platforms is 50% higher there). But this difference is mainly due to the size of the US technology sector.

Secondly, the adoption of generative AI is roughly similar in the USA and Europe when compared sector by sector. It is very broadly linked to cloud adoption in businesses. The major cloud providers all offer AI and generative AI services, facilitating their experimentation. However, France is one of the European countries where companies use the cloud the least (30% of companies, compared with over 70% in the Nordic countries, 60% in Ireland, and 45% in Germany).

137. See previously "Will AI help us to prosper?"

It is therefore up to each French and European company to identify the most relevant AI application cases. The simplest cases are likely to be those where AI errors have few consequences, and can be easily detected and corrected. Given that the costs of inference are still substantial, these will also be cases where the value of the service rendered by AI is fairly high. It is not the role of the State to identify these use cases on behalf of economic players, but it can encourage the creation of a collective dynamic that will enable these players to share experience and information. This could take the form of a "*Convention des Entreprises à l'*Ère *de IA*", similar to the "*Convention des Entreprises pour le Climat*", enabling companies in the same region or industry to get informed, experiment and share best practices, and establish roadmaps for transformation with AI.

An impetus to create a collective dynamic is necessary, but will not be sufficient to encourage widespread adoption of AI by all economic players. This will also largely depend on the profitability of AI systems. Beyond the general lowering of the cost of these technologies, facilitating data collection or large-scale use of the same AI system will be crucial for it to become profitable. To facilitate the adoption of AI while avoiding it being driven by only a few players on a very large scale, the collective dynamic needs to be complemented by a more sector-specific (aeronautics, luxury goods, automotive, agriculture, etc.) or more functional (human resources, finance) approach. The public authorities, led by the State and regional councils, could provide greater support for joint projects involving consortia of private players, bringing together AI start-ups and companies from the traditional economic fabric.

What's happening in the rest of the world? The example of the United Kingdom[138]

In the UK, the BridgeAI program was launched in April 2023 to stimulate the adoption of AI in sectors of the economy with low productivity gains and thus boost their productivity. The £100 million program is particularly targeted at the agriculture, construction, transport and creative industries sectors.

Lastly, the French state should play a greater role as an early adopter of technology, i.e., contracting with companies, particularly small ones, that are developing innovative solutions capable of having a major impact, but which are not yet stabilized due to a lack of customers. Such an approach also contributes to the development of the ecosystem and enables the State to set an example for companies in the adoption of AI.

Recommendation No. 20

Encourage, facilitate and amplify the use of AI tools in the French economy by promoting the use of European solutions.

138. Source: French Treasury Department, Economics Service in the UK.

2.2.5.2 Accelerating the deployment of artificial intelligence in culture

Value chains in the cultural and national media sectors will have to make a rapid, and sometimes major, transition. The transition has already begun in the press (selection, documentation, verification, labeling of content); it is underway in the major news agencies (Associated Press, Reuters, Agence France Presse) and photo agencies (Getty Images); appropriation is proceeding apace in the major communications groups (Publicis). It is still being deployed in technology-intensive cultural sectors, such as animation and, of course, video games (Ubisoft or Mac Guff). The use of AI has focused mainly on repetitive tasks, generating gains that can be invested in greater information quality, creativity and differentiation.

The uses of AI in other sectors are varied. Some are less sensitive — think of the performing arts, apart from the technical aspects (sound, lighting). Most can be used in creative processes, as in architecture and already in video games, with new interactions with players. AI is already involved in production and post-production for visual and sound effects, and its use is set to increase in the audiovisual, film and music industries (mixing, arranging, etc.). However, entire activities and jobs are obviously affected by the rise of AI: translation, dubbing, extras, photography, etc.

The search for productivity gains is spreading. In book and music publishing, it is tending to develop (selection of manuscripts or music, talent spotting). As in the rest of the economy, it can increase the efficiency of physical distribution (books, records). Coupled with online distribution, it is the relational and individualized marketing segments that seem most promising for AI tracking tastes and preferences.

This use of AI is not without risk in terms of cultural and informational diversity and discoverability. For it is in the commercialization and marketing segment that the most important changes are taking place. AI can lead to the rarefaction and homogenization of creation. It is accompanied by a strong polarization of players: on the one hand, a small number of powerful players offering blockbusters to vast audiences; on the other, a more fragile fabric of small and medium-sized players, often at the heart of the renewal of creativity.

The challenge, then, is to make the shift to AI a success by accelerating and structuring an ecosystem. The current ecosystem for designing specialized AI exists, but is still in its infancy: 3D crowd generation and simulation (Golaem), Face Engine (Mac Guff) for processing facial images, colorization of film archives (Composite Films), ChatBox for cultural mediation (Ask Mona), image transformation (Photoroom), DeepVoice voice applications from Ircam. These new tools are growing in the cultural sector and in all sectors that rely on creativity (advertising, luxury goods, the textile industry, design). We need to support this ecosystem of research-based players, but also take into account the industrial fabric of cultural SMEs and ISEs, to avoid polarizing the culture and media economy.

Recommendation No. 21

Facilitate the appropriation and acceleration of AI uses in culture and the media to limit polarization between large groups and small players and combat misinformation.

2.3. RESPONSIBILITY: CONTROLLING, AUDITING, PROTECTING

2.3.1 BUILDING THE INTERNATIONAL GOVERNANCE THAT IS CURRENTLY LACKING

Widespread awareness of the lightning progress being made in the field of AI has given rise to a profusion of international initiatives on the subject of AI. In all, there are no fewer than fifty initiatives, most of which involve France!

2023 was particularly rich in debate. At the UN, a high-level advisory committee on AI to the Secretary-General was announced. Within the framework of Unesco, a new world summit on the ethics of artificial intelligence was prepared and held at the beginning of February. Within the G7, the so-called Hiroshima process led to several joint declarations on the development and supervision of AI systems. In addition, the UK organized a summit on AI risks in November 2023.

These events extend an initial body of work on the development of this technology. Within the European Union, the High-Level Expert Group on AI set up in 2017 laid the foundations for the Artificial Intelligence Regulation. Within the OECD, a network of experts and an AI policy observatory were set up in 2018, leading to the establishment of guiding principles on AI (2019), a framework for classifying AI systems (2022) and a platform for monitoring AI-related incidents (2023).

States have also supported joint projects. In 2020, the Global Partnership for Artificial Intelligence (GPAI) was launched under the impetus of France and Canada, to develop robust scientific expertise on AI and make concrete recommendations for the development of responsible AI systems that respect human rights. The GPAI brings together a solid community of stakeholders, with 29 member states from all continents and all levels of development, as well as experts from industry, academia and civil society. For example, the GPAI has worked on the integration of AI systems in organizations and their effects on employment, and on the use of AI by small and medium-sized enterprises.

Should we go further than these initiatives? Yes, because they do not constitute true international governance of AI. On the one hand, the work carried out to date concerns analysis, declaration or recommendation, and is not binding. On the other hand, the initiatives are scattered, and none of them has the resources or legitimacy to ensure effective AI governance at international level.

Our Commission believes that global governance must now emerge from this concert of initiatives, and recommends that a World AI Organization be set up at the AI summit hosted by France. Let's take a look at four key questions related to this organization: its composition, its missions, its functioning, and its temporality.

Recommendation No. 22

Structure a coherent and concrete diplomatic initiative aimed at founding a global governance of AI.

Recommendation No. 23

Structure an open national AI governance ecosystem now.

Whose governance? The participation of States is of course essential, as they have the capacity to conclude treaties and make them applicable in their territory. However, progress in AI research and the worldwide dissemination of technology are very much the work of very large private companies with a global reach. Whether one finds it regrettable or not, the legitimacy of the interstate dimension is therefore not sufficient, on its own, to conceive of fully effective AI governance on an international scale. Moreover, in the digital age, civil society is indispensable. It is not only a major user, but also a contributor to the development and dissemination of AI, with a set of principles and cultural references that underpin open source, for example.

We have therefore ruled out several existing governance models. A governance model based exclusively on inter-state relations, such as that inherited from the international organizations of the 20th century, would be incomplete. Similarly, AI governance based solely on the components of international civil society, even if widely understood (associations, non-governmental organizations, companies, research structures, etc.), would be largely ineffective. In particular, the weight of a few major AI companies would make it impossible to draw inspiration from the model used in the air transport sector, which would see, on the one hand, an organization with standardization powers bringing together only these private commercial players — the International Air Transport Association (IATA) — and, on the other hand, an interstate organization, under the aegis of the UN, with regulatory powers — the International Civil Aviation Organization (ICAO).

AI governance must therefore be of a new kind. To ensure the full legitimacy and effectiveness of its action, governance should be made up, in equal parts, of representatives of States and inter-State organizations, on the one hand, and individual representatives divided into four colleges corresponding to the key AI stakeholder groups (research, private general-interest structures, companies, territories), on the other.

For what missions? The treaty establishing the World AI Organization would entrust it with a range of responsibilities. These would be of three kinds. Firstly, the organization would be tasked with establishing binding standards for AI systems, including standardization of AI auditing processes. Secondly, like the IPCC, it would be tasked with assessing the state of knowledge on the evolution of AI systems and their impacts. Thirdly, it would decide on strategic orientations for projects of global general interest, for example in the field of AI tools for environmental transition. The financing of projects as such would be entrusted to a specific international fund (see below).

How would the organization function? The Organization's General Assembly would bring together all state and individual representatives. A small collegiate executive body (maximum 10 people) would be responsible not only for running and managing the organization, but also for carrying out its projects.

Funding for this organization could be provided on a permanent basis in equal shares by States and all non-state components, and would be sufficient to guarantee effective participation by non-state players and to ensure strong, recognized expertise.

To ensure the confidence of all players, standardization and normalization missions, particularly in the field of auditing, could be entrusted to a specialized body, endowed with a high degree of autonomy and strong guarantees of impartiality. The governance of this specialized structure could be made up of trusted individuals elected or appointed by the various components and colleges of the main organization, in the manner of an Anglo-Saxon "trust", and national or regional technical professionals. In addition to its own expertise, this body could call on representatives of national or regional technical expert professions (CEN-CENELEC, etc.).

When? The AI summit to be hosted by France in late 2024 or early 2025 provides the opportunity to lay the foundations for this global governance of AI. With this in mind, we need to start building this project now, with France's diplomatic partners as well as the non-state actors called upon to participate in this future organization. Before and after the summit, a series of actions could be organized to mobilize civil society in France and abroad.

The meaning of a French bid to host the World AI Organization

While not leading the international AI competition, France is contributing to the development of the technology and its business models. A training ground for leading contemporary AI researchers, France has been stepping up its research and training efforts since the 2018 report by Fields Medal winner C. Villani. At the same time, its AI business fabric is growing, with, for example, two companies appearing in the ranking of the most accurate foundation models (out of twelve worldwide[139]).

In recent years, France has also committed itself to laying the foundations for global governance. Co-founder of the GPAI with Canada, it also hosts the headquarters of the OECD in Paris, which is carrying out important work in the field of AI.

Furthermore, within the European Union, during the negotiation of the European AI Regulation, France expressed a position aimed at reconciling the two imperatives of protection and innovation. At the London summit in November 2023, which saw discourses based on existential risks prevail, France also promoted a vision that recognized the risks of AI, but also emphasized the benefits of this technology. This line could help build a space of international convergence.

Finally, in the current context of heightened economic competition between the United States and China, France could appear as a point of equilibrium, favoring the rapid emergence of a form of governance that includes elements of regulation, standardization and auditing.

Alongside this World AI Organization, France could promote the creation of an **International Fund for Public Interest AI** (IFPAI). This fund would contribute to the emergence of a range of projects beneficial to humanity: free and open source AI services, independent research projects, innovations (in the environment, science, health, etc.). The fund's budget could be around €500 million a year.

We suggest an International Fund, distinct from the World Organization presented above, for four main reasons. Firstly, it seems necessary to separate — in order to prevent potential conflicts of interest — the mission of setting norms and standards, which involves political issues and would therefore fall within the remit of the World Organization, and the mission of financing initiatives of global general interest (independent research, ecosystem construction, etc.), which would fall within the remit of the International Fund. Secondly, an independent Fund would bring together expertise in financial evaluation, support and monitoring of innovative projects. In addition, the Fund will rely on contributions from donors, while the World Organization will rely on statutory contributions from member states and third parties. Last but not least, the creation of the Fund will enable the involvement of non-state actors to be stepped up, with a multi-stakeholder Board of Directors.

Au-delà de la cristallisation de la gouvernance mondiale, le sommet sur l'IA qu'accueillera la France pourrait être l'occasion de faire progresser la coopération internationale dans quatre directions.

139. Stanford Foundation Models Research Center (January 2024).

Putting AI to work for the common good. France could propose the creation of a mechanism for access to computing power for developing countries, for example within the framework of a "**1% AI**": all players investing significantly in computing power could commit to allocating 1% to these countries.

Fostering dynamic innovation and an open ecosystem. In view of the risk of concentration in the AI ecosystem, France could promote international funding of AI commons, particularly in terms of data. Best practices in competition policy (see *2.3.3. Avoiding dominant competitive positions*) could also be identified and implemented on a global scale.

Preparing the future of work with AI. Following on from the work of the International Labor Organization and our Commission, France could lead an international project to assess the impact of AI on employment and work, in order to facilitate social dialogue and the orientation of technology towards job quality. It could also initiate a reflection on certain professions now at the heart of the AI value chain, such as those in human-enhanced learning.

Promoting safe and secure AI systems. The summit could lead to the principle of international convergence of standards and evaluation methods for AI systems, to avoid the fragmentation of safety and security rules around the world. These rules are intended to protect users, but also to create a framework of trust favorable to the development of AI business models.

2.3.2 PROVIDING FRANCE WITH THE CAPACITY TO EVALUATE AI SYSTEMS

The performance of AI systems often amazes us, but their limitations are almost as obvious as their qualities: a few exchanges with a conversational robot driven by generative AI is enough to understand that an AI system can respond in an incomplete and misleading way. Specialists use the term "hallucinations" (or confabulations) to refer to erroneous results or imaginary facts stated by a generative AI system. These are far from being the only shortcomings of these systems: tests have shown that, if poorly controlled, they can produce discriminatory or sexist results, reveal confidential information present in their training data, create child pornography content, or threaten their users.

To avoid the risks that can arise from the use of generative AI systems, the companies developing these systems subject them to extensive testing, and use a variety of techniques to improve the quality of the results. This may involve, for example, providing the machine with human feedback during the training phase, or intervening directly on the algorithm to control the results produced. These companies need to be able to verify and demonstrate that the precautions taken really do reduce the risks to an acceptable level, but also, more simply, that their systems are effectively fulfilling their intended purpose. In this respect, it's not easy to improve the upstream quality of a general model, as it then has to be fine-tuned for specific uses.

Symmetrically, users of AI systems and public authorities want to be sure that these systems offer a sufficient level of reliability and security. This is particularly important when the system is to be used in a sensitive area, such as human health, law enforcement or recruitment. Public authorities also need to be able to monitor the performance and biases of AI systems over time, in order to anticipate new risks. This need for evaluation according to shared standards goes beyond simple performance, since it also concerns social biases (such as discrimination) or environmental impact. All this means that we need to be able to evaluate AI systems holistically, reliably and in a way that is recognized by all.

The demand for assessments will come from the ability of AI system vendors to differentiate themselves through these assessments, and from regulatory requirements. Supply, for its part, will be facilitated by the clarification of evaluation standards and norms.

In terms of regulatory obligations, a framework for assessing AI systems is currently being put in place. The European regulation on artificial intelligence, currently in the process of being adopted, requires a conformity assessment to be carried out on so-called "high-risk" AI systems before they are placed on the market. This assessment is in addition to those provided for by numerous sector-specific regulations already applicable: an X-ray machine incorporating AI is first and foremost a piece of medical equipment; it must be assessed under the same conditions and offer the same guarantees of reliability and safety as a "conventional" X-ray machine. Similar initiatives are underway in other parts of the world.

Who will implement the AI regulation?

The aim of the European AI Regulation is to provide a favorable framework for the development of AI in the European Union, taking advantage in particular of the size of the single market, while providing solid guarantees for the protection of fundamental rights. The regulation lays down a set of common rules for the marketing and use of AI systems, but differentiated according to the degree of risk associated with these systems. It provides for complex governance.

At European level, four bodies have been created. The European AI Office is the main body for implementing the regulation. A scientific panel of independent experts supports the office in carrying out its tasks. A European AI Committee brings together states to coordinate and share best practices. Finally, an Advisory Forum brings together economic and research players to provide technical expertise to the Union and Member States.

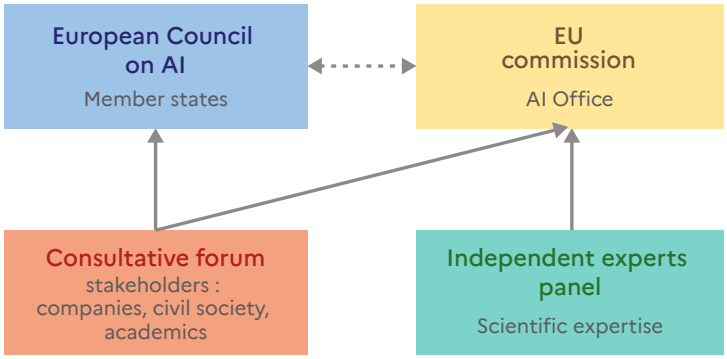

At national level, implementation of the AI regulation relies on two sets of players. On the one hand, the so-called notifying authorities, who are responsible for designating ("notifying") and monitoring the bodies that assess the conformity of AI systems placed on the European market. On the other hand, market surveillance authorities are responsible for monitoring issues specific to AI (e.g., supervision of testing) or sector-specific issues that intersect with AI (e.g., financial markets, data protection, fraud control).

Against this backdrop, our Commission stresses the importance of appointing market surveillance authorities based on competence and knowledge of the market in question, without seeking to designate a single authority. It will therefore be necessary to support the development of the skills of these bodies, in particular by making pooled resources available.

If standards exist, they are still evolving in this still nascent ecosystem. These standards must translate into technical solutions the requirements imposed on AI systems, for example in terms of robustness or reliability. The aim is to provide concrete answers to questions such as: what is a reliable AI in the medical field? What is an unbiased AI in the recruitment field?

The rapid definition of standards enabling the evaluation of AI systems is in the interest of everyone, including the industry, as it is needed both to meet regulatory obligations and to create trust and differentiation. In the absence of standards for effectively evaluating the performance of AI systems, many markets could remain closed to AI solutions. However, the definition of these standards is not only an economic and legal issue, it is also a political one, as

it involves giving concrete content to the values that AI systems must respect. This issue cannot be left entirely in the hands of companies.

In Europe, the harmonized standards on which the AI regulation is based are set to play a central role. The European Commission has mandated the European Committee for Standardization (CEN/CENELEC) to draw them up, with the ambitious aim of being ready by the time the regulation comes into force in 2026. On a global level, a number of bodies are working to define the standards and methods needed to evaluate AI systems, whether they be standards bodies (ISO, worldwide, NIST in the USA, BSI in the UK, etc.) or "private" forums.

Despite their technical nature, the Commission considers that several issues need to be raised to a high level.

Firstly, standards should be harmonized worldwide, despite the proliferation of initiatives. This issue should therefore be taken into account in the international governance to be put in place (see 2.3.1). At European level, the application of the AI Act should be harmonized to avoid any race to the bottom between member states.

Secondly, the standards to be defined will have to be evolving: as AI moves into new fields, the needs for AI system evaluation will evolve in parallel. Above all, it is unlikely that a definitive answer to the questions raised by the evaluation of AI systems can be found at the outset. It is more likely to be an iterative process, as we have seen in the field of cybersecurity. Work is still needed on many subjects. For example, the use of independent teams to carry out adversarial testing (red teaming) is seen as a very promising evaluation method, already used by some companies, but there is as yet no consensus on what constitutes an effective and exhaustive approach in this field. Other subjects are still in the research stage, such as the development of machine unlearning techniques, enabling a generative AI system to "forget" data or information ingested during training.

Thirdly, standards, and in particular the definition of documentation and evaluation practice obligations, will have to be defined taking into account the reality of AI system development practices. Today's developers use numerous tools to automatically document and evaluate the performance and biases of their systems. Measurement grids need to be defined and harmonized, but they would benefit from being easily integrated into these tools. To put it plainly, compliance with the AI Act can become a burden by forcing cumbersome manual documentation for small companies, or it can be an example of technological regulation.

To ensure that the implementation of the AI Act in France is as close as possible to usage, scalable and technological, the Commission proposes to anticipate the implementation of the AI Act in France, and in particular to designate supervisory authorities as close as possible to the markets, supported by strengthened mutualized assessment capabilities. Concentrating the entire implementation of the AI regulation in France in a single authority (like a French "AI Safety Office") would reduce the resilience of the French ecosystem and limit the rise in competence of sector-specific authorities. If AI is indeed a general-purpose technology, all sector authorities need to address it.

On the supply side, we need to develop a supplier-independent AI system audit capability. To this end, in support of the European regulation on AI, we need to help the bodies involved in conformity assessment of AI systems (known as "notified bodies"), whose designation is provided for in the regulation, to become more competent. It also seems appropriate to encourage the

certification of AI systems, which can be a competitive advantage for companies that use them, and for which a number of programs are already available[140].

Finally, evaluation and auditing capabilities need to be complemented by the development of solutions for securing AI systems. These include, for example, the ability to detect cyber attacks (from AI models or with AI models) or to carry out protection or remediation actions.

Recommendation No. 24

Equip France and Europe with a public and private evaluation ecosystem for AI systems that is as close as possible to usage and the latest technological developments.

140. Afnor Certification and LNE, in particular, have set up AI-related certifications.

2.3.3 AVOIDING DOMINANT COMPETITIVE POSITIONS

The early 2000s saw the emergence of digital giants such as Alphabet, Apple, Meta, Amazon and Microsoft, which today dominate most segments of the digital technology value chain. This state of affairs has been made possible by the lack of responsiveness of competition policy in both the USA and Europe. However, the development of AI confirms and reinforces the risk of concentration in the digital value chain.

Limiting dominant competitive positions promotes growth and a fair distribution of economic gains. It is therefore crucial to study, but even more so to anticipate, the competitive evolution of the AI value chain, to enable competition to be exercised properly. European institutions will therefore have to make AI-related competitive issues one of the priorities of their actions, to ensure that markets remain contestable[141], i.e., that barriers to the entry of new companies and to the exit of existing ones are limited. Competitive issues arise at different levels of the AI value chain.

Upstream, the largely preponderant weight of a single player on the market for the design of graphics processing units (GPUs), which are an essential building block for training and optimizing large language models, generates risks of anti-competitive behavior. Furthermore, as the French competition authority indicated in its June 2023 opinion[142], the cloud sector, also upstream in the AI value chain, is dominated by three major players, the "hyperscalers" (Amazon Web Services, Google Cloud Platform, and Microsoft Azure). By 2021, these three players will account for 80% of the growth in spending on public cloud infrastructures and applications in France. These companies, with their financial clout and rich ecosystem of digital services, are in a position to hinder effective competition across the value chain.

Downstream, bundling is a worrying prospect. This consists of players present in one segment of the value chain proposing a commercial offer made up, in addition to the main product, of additional products (software, for example) in other segments of the value chain. This issue is not unique to the AI value chain, as it has existed since the beginning of the digital revolution (think of the marketing of the Microsoft Office suite, for example). While this commercial practice is not inherently anti-competitive, it can lead to reprehensible behavior, and it is important to identify it.

At the heart of the value chain, we need to be vigilant when it comes to concentration. Indeed, the most promising American companies in the development of the largest, so-called foundation models, are already linked to historical players in the digital revolution (albeit in different ways): OpenAI to Microsoft and Anthropic to Amazon and Google. Even if these companies are not formally controlled by these historical players, funding from the digital giants could constitute a form of control and would therefore raise competitive issues.

Public authorities are not helpless in the face of these challenges. The Digital Markets Act (DMA) is due to come into force in the first half of 2023, with the aim of ensuring that markets in the digital sector remain contestable and fair. More specifically, this regulation aims to prevent a major platform enjoying a "gatekeeper" position visàvis a large number of users from abusing this position by preventing potential entrants from gaining access to these users, and therefore to this market.

141. A market is contestable if it is possible for new producers to enter and existing producers to exit under bearable conditions, i.e., if barriers to entry and exit are acceptable.

142. Opinion 23-A-08 of June 29, 2023 on the competitive operation of cloud computing.

The DMA defines a clear framework that will address most of the competitive issues facing the "historic" digital sector. It should make it possible to tackle anti-competitive behavior downstream in the value chain, such as tying, or the re-use of personal data collected through the use of a first product (for example, the collection of personal data via a messaging system for targeted advertising on a social network belonging to the same platform). The DMA should also make it possible to address problems of concentration on certain digital markets, a situation in which a small number of players or even a single player dominates the market, thanks to the obligation placed on "gatekeepers" to declare all their acquisitions, including those of small companies whose acquisition has hitherto passed under the radar of the European Commission.

However, the DMA responds to these challenges through a well-defined prism, that of platforms. Yet the AI value chain, still emerging and subject to numerous evolutions, will not necessarily converge towards an exclusive "online platform" mode of operation. This is why the DMA as it currently exists may not be effective in addressing all the competitive challenges of the AI value chain, but rather certain points (search engine referencing, data portability, interoperability of personal agents, for example).

Thus, the DMA is able to provide solutions to the competitive issues present on the AI value chain as long as these issues are linked to a platform problem. It will probably not be sufficient to cover all competitive issues. The regulation therefore needs to be supplemented to take account of the specificities of AI. Finally, in the medium term, we need to envisage a change in competition policy doctrine, moving from a static system (what market shares does this company hold today?) to a dynamic vision (what market shares could this company hold tomorrow, and what companies could enter this market tomorrow?), enabling us to anticipate mergers rather than waiting to see them happen.

Recommendation No. 25

Monitor developments in market concentrations and rapidly put in place the regulations needed to prevent abuse of dominant positions.

OUR RECOMMENDATIONS

The plan proposed by our Commission includes 25 recommendations, representing an annual commitment of around €5 billion over the next five years. The resources can be divided into five main categories: collective appropriation, training and research; deployment of AI at the service of citizens; technological and industrial investment; dissemination of AI in the economy; French, European and global governance.

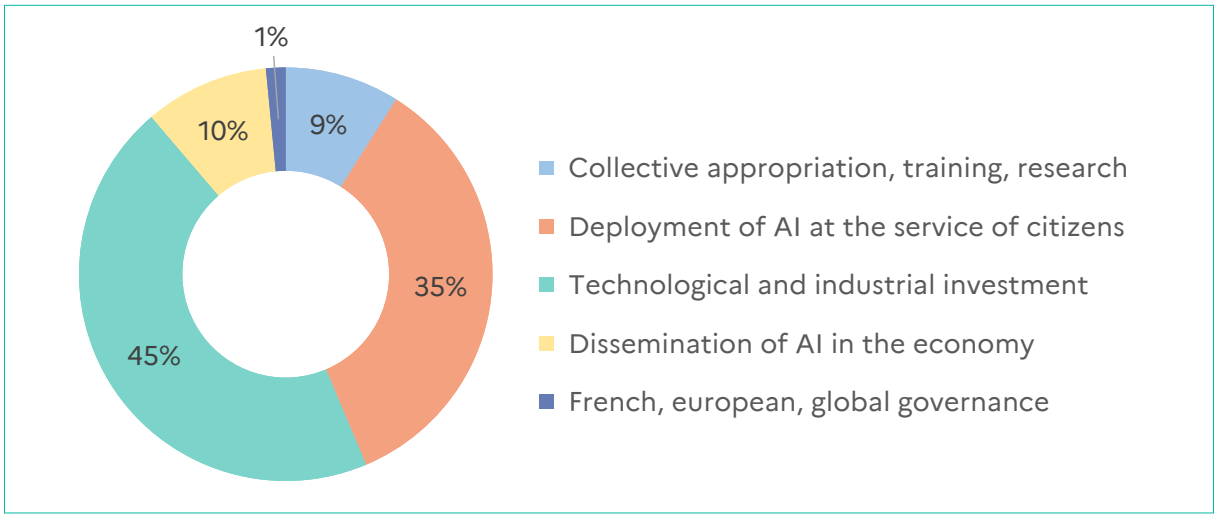

	Recommendations	Financial impact over 5 years (M€)
1	Create the conditions for collective appropriation of AI and its challenges, in order to collectively define the conditions under which it will become part of our society and our daily lives. *Leaders: Ministry of Territorial Cohesion*	10
2	Invest in observation, studies and research into the impact of AI systems on the quantity and quality of employment. *Leaders: Ministry of Labor and Ministry of the Civil Service*	5
3	Make social and professional dialogue a tool for co-constructing the uses and regulating the risks of AI systems. *Leaders: Ministry of Labor*	0
4	Develop a strategy to support the open AI ecosystem internationally by supporting the use and development of open AI systems and third-party inspection and evaluation capabilities. *Leaders: Ministry of Higher Education and Research; Ministry of the Economy*	0
5	Make France an AI pioneer for the planet by strengthening environmental transparency, research into low-impact models, and the use of AI to serve energy and environmental transitions. *Leaders: Ministry of Europe; Ministry of Research*	100

	Recommendations	Financial impact over 5 years (M€)
6	Generalize the deployment of AI in all higher education courses and acculturate students in secondary education to make specialized courses accessible and attractive. *Leaders: Ministry of Higher Education and Research, Ministry of Education*	1,200
7	Invest in continuing vocational training for the workforce and in training schemes around AI. *Leaders: Ministry of Labor and Ministry of the Civil Service*	200
8	Train creative professions in AI, from the early years of higher education and on an ongoing basis. *Leaders: Ministry of Higher Education and Research; Ministry of Culture*	20
9	Strengthen the technical capacity and infrastructure of public digital to define and scale a real transformation of public services through digital and AI, for agents and at the service of users. *Leader: Ministry of Transformation and the Civil Service*	5,500
10	Facilitate the circulation of data and the sharing of practices to reap the benefits of AI in care, and improve the offering and the daily lives of caregivers. *Leader: Ministry of Health*	3,000
11	Encourage the individual use, large-scale experimentation and evaluation of AI tools to strengthen the public education service and improve the day-to-day lives of teaching teams. *Leader: Ministry of Education*	1,000
12	Invest massively in digital companies and business transformation to support the French AI ecosystem and make it one of the world's leaders. *Leaders: Prime Minister's Office; Ministry of the Economy and Finance*	3,600
13	Accelerate the emergence of a European semiconductor component industry adapted to AI systems. *Leaders: Prime Minister's Office; Ministry of the Economy; Ministry of Higher Education and Research*	7,700
14	Make France and Europe a major center of installed computing power. *Leaders: Prime Minister's Office; Ministry of the Economy and Finance*	1,000
15	Transform our approach to personal data for better innovation. *Leader: Prime Minister's Office; Ministry of Justice*	16
16	Set up a technical infrastructure to bring together AI developers and holders of cultural heritage data. *Leader: Ministry of Culture and its operators*	35
17	Implement and evaluate the transparency obligations set out in the European AI Regulation by encouraging the development of standards and a suitable infrastructure. *Leader: Ministry of Culture*	0
18	Attract and retain world-class talent with scientific or entrepreneurial and managerial skills in the field of AI. *Leader: Ministry of the Economy; Ministry of Foreign Affairs*	10
19	Attract and retain world-class talent with scientific or entrepreneurial and managerial skills in the field of AI. *Leader: Ministry of the Economy; Ministry of Foreign Affairs*	1,025

	Recommendations	Financial impact over 5 years (M€)
20	Encourage, facilitate and amplify the use of AI tools in the French economy by promoting the use of European solutions. *Leader: Prime Minister's Office; Ministry of the Economy*	2,600
21	Facilitate the appropriation and acceleration of AI uses in culture and the media to limit polarization between large groups and small players and combat misinformation. *Leader: Ministry of Culture*	60
22	Structure a coherent and concrete diplomatic initiative aimed at founding a global governance of AI. *Leader: Ministry of Foreign Affairs; Ministry of Culture; Ministry of the Economy*	300
23	Structure a powerful national AI governance ecosystem now. *Leaders: Prime Minister's Office, Ministry of the Economy, Ministry of Research*	5
24	Equip France and Europe with a public and private evaluation ecosystem for AI systems that is as close as possible to usage and the latest technological developments. *Leaders: Ministry of the Economy and Finance*	15
25	Anticipate market concentrations across the entire artificial intelligence value chain. *Leader: Ministry of the Economy*	0
	Overall plan	27 bn€

APPOINTMENT LETTER

Le Premier Ministre

Paris, le 28 FEV. 2024

Madame, Monsieur,

L'arrivée progressive de l'intelligence artificielle (IA) au cœur de notre quotidien révèle chaque jour un peu plus son potentiel et soulève de très nombreuses questions, notamment dans les domaines de l'éthique, de l'économie, de la productivité, du travail, de l'organisation des entreprises ou encore de la souverainetés industrielle et numérique des États. Cette innovation constitue un point de rupture dans nos sociétés modernes, venant bouleverser nos modes de pensée, nos modes de production, nos modes de consommation, en bref, nos modes de vie.

Les bouleversements technologiques et sociaux qui en découlent prouvent que la bataille pour la souveraineté passe par la maîtrise de ces technologies. Dès 2018, la première phase de la stratégie nationale pour l'IA, lancée par le Président de la République, a permis de développer une filière de rang mondial. Avec le lancement du plan France 2030 en octobre 2021, notre pays a rehaussé son ambition en la matière.

Aujourd'hui, face à ces nouveaux défis, nous devons aller plus loin, notamment en renforçant nos formations pour développer davantage de talents en France, en nous assurant que notre tissu économique ait les moyens de tirer le maximum de cette technologie, en investissant pour favoriser l'innovation française sur la scène internationale et en définissant une régulation adaptée des différents secteurs pour protéger des dérives.

C'est la raison pour laquelle ma prédécesseure, Mme Elisabeth Borne, a souhaité installer la Commission de l'intelligence artificielle, le 19 septembre 2023 à Matignon.

Sous votre présidence, cette Commission réunit des acteurs de différents secteurs (culturel, économique, technologique, de recherche), pour contribuer à éclairer les décisions du Gouvernement et faire de la France un pays à la pointe de la révolution de l'intelligence artificielle. Elle doit présenter au Gouvernement des propositions concrètes permettant d'adapter notre stratégie nationale.

Vous animez cette Commission composée de treize experts de l'intelligence artificielle : Gilles Babinet, Joëlle Barral, Alexandra Bensamoun, Nozha Boujemaa, Bernard Charlès, Luc Julia, Yann Le Cun, Arthur Mensch, Cédric O, Isabel Ryl, Franca Salis-Madinier, Martin Tisné, Gaël Varoquaux.

Monsieur Philippe AGHION
Professeur au Collège de France
Madame Anne BOUVEROT
Présidente du conseil d'administration de l'École normale supérieure

Chacun des experts a été désigné *intuitu personae* par le Gouvernement et ne peut donc en aucun cas se faire représenter. Les présidents, les membres et les experts de la Commission sont par ailleurs tenus à la confidentialité des débats et travaux auxquels ils participent ou assistent. Vous veillez à ce que cette obligation de confidentialité et de discrétion soit respectée.

Vous êtes responsables de l'avancée des travaux, dont je souhaite qu'ils s'articulent autour de cinq thématiques clés : impacts économiques, souveraineté industrielle et numérique, éthique et impacts sociétaux, enjeux culturels et service public. Pour chacun de ces thèmes, vous veillez à la diversité des personnes auditionnées qui vous permettent d'étayer les constats et les propositions établis par la Commission.

Pour mener ces travaux thématiques, la Commission de de l'intelligence artificielle bénéficie de l'appui de treize rapporteurs issus de l'administration : Marc Auberger, inspecteur général des finances, Simon Bunel, administrateur de l'Insee, Philippe Chantepie, inspecteur général des affaires culturelles, Eloy Dorado, administrateur général, Emilie-Pauline Gallié, inspectrice générale de l'éducation, du sport et de la recherche, Paul Jolie, ingénieur général des mines, Arnaud Mazier, ingénieur général des mines, Vincent Montreuil, inspecteur général de l'éducation, du sport et de la recherche, Timothée Paris, maître des requêtes au Conseil d'État, Christophe Ravier, ingénieur général des mines, Erwan Paitel, inspecteur général de l'éducation, du sport et de la recherche, Ulrich Tan, ingénieur des mines, et Louis--Charles Viossat, inspecteur général des affaires sociales.

Afin d'assurer la coordination d'ensemble des travaux de la Commission, vous êtes également accompagnés de deux rapporteurs généraux mis à votre disposition, Cyprien Canivenc, conseiller référendaire à la Cour des comptes, et Arno Amabile, ingénieur des mines. Vous pouvez également vous appuyer sur l'expertise de Guillaume Avrin, coordonnateur national pour l'intelligence artificielle.

Je souhaite que vous me rendiez régulièrement compte de l'avancée de vos travaux, jusqu'à la remise de vos conclusions définitives début mars 2024. Les propositions devront être opérationnelles, réalistes et ambitieuses soutenues par une vision long terme, globale et objectivée.

Je vous remercie et vous prie de croire, Madame, Monsieur, à l'assurance de notre considération distinguée.

Gabriel ATTAL

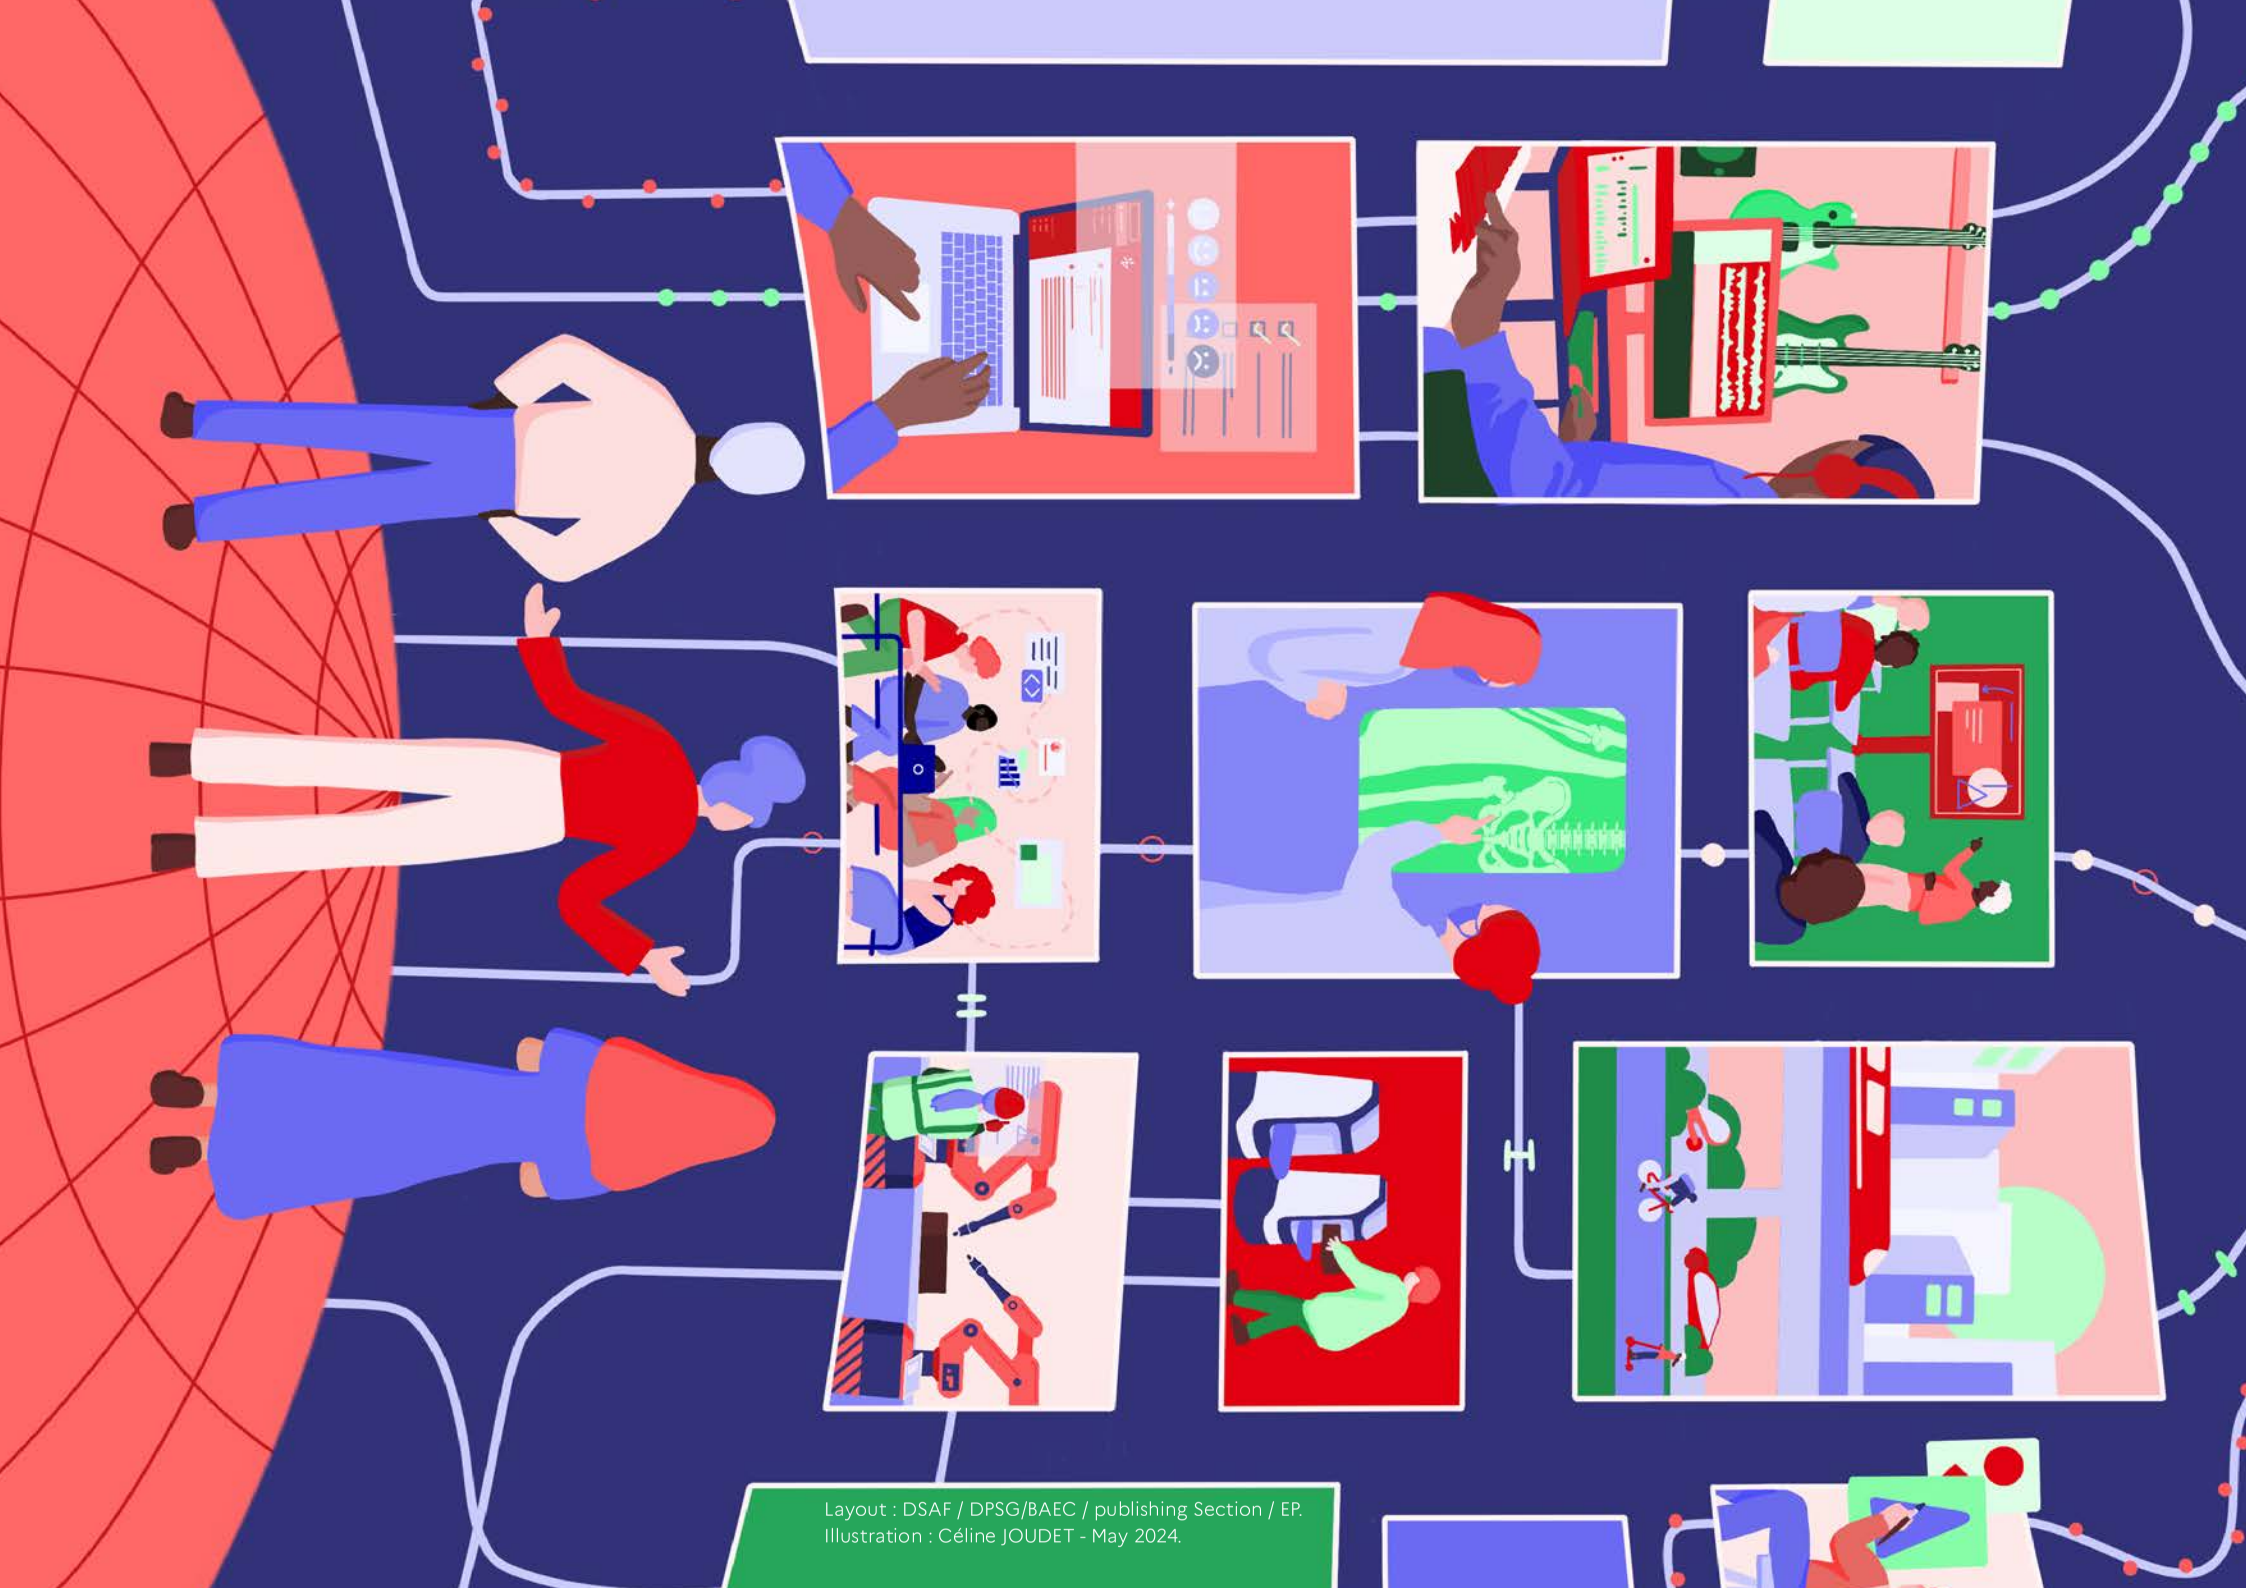

Layout : DSAF / DPSG/BAEC / publishing Section / EP.
Illustration : Céline JOUDET - May 2024.