Discussion Paper for Breakout Session

Theme 4: The Future of Work

Skills for the Modern Economy

Authors:
Giorgio Metta[1], Massimo Maresca[2], Giuseppe Attardi[3],
Salima Benhamou[4], Nicolas Vayatis[5]

[1]Istituto Italiano di Tecnologia, Via Morego, 30 – 16163, Genova, Italy,
giorgio.metta@iit.it
[2]University of Genoa, DIBRIS, Via Opera Pia, 13 – 16145, Genova, Italy,
massimo.maresca@unige.it
[3]Univ. of Pisa, Dept. of Computer Science, Largo B. Pontecorvo, 3 – 56127 Pisa,
Italy, attardi@di.unipi.it
[4] France Stratégie, Department Jobs and Skills, 20 av. de Ségur 75007 Paris,
France, salima.benhamou@strategie.gouv.fr
[5]Ecole Normale Supérieure Paris-Saclay, 61, av. du Président Wilson 94 235
Cachan cedex, France, vayatis@cmla.ens-cachan.fr

G7 Multistakeholder Conference on Artificial Intelligence
December 6, 2018, Montreal

Executive summary

In this paper, we highlight some of the issues of the impact of Artificial Intelligence (AI) on the future of work. First, we acknowledge the advances of AI-based technologies as well as the anisotropy of progress along the various sectors of human activities. Pushed by technological advances and massive investment, AI is introducing asymmetries which are transforming the job market in content as well as in location.

We consider the technological aspects in depth as often the common perception of AI is misaligned with the effective technical readiness level. This is typically the case with any new technology. We devote a section to the analysis of the major limitations of present-day AI including the need of large amounts of annotated data, of massive computational power and the effort required to apply AI to each new problem/domain. We describe the potential of AI technology and potential (sometimes actual) application areas ranging from data analytics, robotics, engineering, genetics, climate change analysis, etc.

With respect to impact to the job market, it is clear that AI already brings tangible benefits for certain businesses and improved service levels to consumers. On the other hand, there is not a unique interpretation of how things will evolve. Predicting technology diffusion and, especially its impact, hinges on a variety of exogenous factors including further technological advances to overcome existing limitations but also political and economic factors (regulations, availability of funding, cost vs. benefit aspects, etc.). Our analysis about the risks and benefits of AI on work and employment relies on the assumption that in most contexts, we will have to consider the challenges of achieving true complementarity between human and machine through organizational choices and continuous learning.

Under any possible scenario, the deployment of AI systems will lead to rethinking work organization ensuring that individuals develop continuously their learning capacities.

We conclude with a series of recommendations for the effective management of AI, which can be summarized as follows:

- Ensure that people receive appropriate training and education in the potential issues related to AI: we underline the fact that this is not only a question of training the developers of AI systems but all strata of society;

- Support research and innovation in AI: the rapid evolution and diffusion of AI requires appropriate actions to support research and innovation both at the government level and clearly at the international level;

- Democratize AI: favoring openness and accessibility to AI will speed up the capacity to train individuals at a large scale, as well as facilitate the emergence of certification processes and explainability of AI algorithms.

Introduction

About ten years ago, Artificial Intelligence (AI) technologies started to achieve incredible progress in a surprising variety of applications because of three fundamental technological advances:

1. Increases in processing power that allowed building larger models trained with machine learning algorithms;

2. Availability of large amounts of annotated data for training these large models;

3. Progress in machine learning theory and, consequently, improvements in learning algorithms.

The impact of AI is magnified by the fact that the "digital economy" has reached a global scale, so that "countries that become leaders in the field of AI will not only capture much of the value of the systems that they transform, but also control these same systems, calling into question the independence of other countries"[1]. AI is a cross-sectorial technology that is poised to disrupt all aspects of life, economy and society. This scientific and technical revolution offers a horde of promises, but it also raises important issues.

Setting the stage: AI and the future of work

Several studies (such as Frey and Osborne[2]) predict that a large percentage of jobs will change considerably through automation, AI, and computerization. Predictions range from the "risk of vanishing" to more moderate "modification of existing jobs". A distinction is to be made between automation of a task versus the complete automation of entire jobs. Some economists argue that the lost jobs will be replaced by new jobs as it happened with previous technological revolutions. Conversely, Erik Brynjolfsson of MIT Sloan School warns that we "should not just blindly extrapolate from the past, this is a different technology that is going to affect different groups of people"[3].

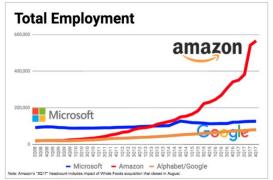While it is hard to make predictions on occupations that are yet to be invented, we can at least safely say that:

1. The phenomenon is not affecting a single sector as in the past (e.g. transportation), but it will impact all sectors (transportation, retail, healthcare, science, to name a few);
2. The transformation is happening at a much higher rate than in the past[4]. It may be difficult for workers to retrain for a new job;
3. New jobs can be created although the distribution across countries, sectors, companies is difficult to predict.

Cédric Villani[2] says that "it looks increasingly certain that most sectors and companies will be widely reshaped. We are entering a new era of major technological transition and history shows us that previous periods of transition did not always run smoothly. Indeed, they sometimes involved drastic political readjustment, which often hit the most fragile portions of the population the hardest".

On the other hand, the following graph shows that the phenomenon is not necessarily uniform across companies or sectors. For example, for companies like Microsoft and Google the increase has been modest. Conversely, Amazon has been steadily growing worldwide, due to the heavy logistics of its business model. The predominantly digital services can grow without increasing manpower, because intense automation of production and delivery allows handling a huge customer base at a global scale.

In this scenario, no job is exempt from the risk of de-localization, except for those that involve local delivery, must be carried out next to the consumer premises, or where consumers need to interact in person: these include public services like healthcare, justice and schools or services like construction or maintenance of facilities, public transport, tourism and food. Manufacturing may concentrate in large-scale production facilities as a few global conglomerates manage distribution. Therefore, work may not need to be outsourced where labor is cheaper, but rather to companies that can perform effectively and efficiently exploiting perhaps a *highly specialized technological infrastructure*.

AI advancements will make robots operate more flexibly in tasks that so far required human skills. In the longer term, fully autonomous robots will become more ubiquitous and will drive supply chain innovation and tighter integration of the whole manufacturing process. As AI continues to advance, problem solving, and learning abilities will enable robots to be adaptive and responsive with minimal human feedback.

Is AI more disruptive than the other past revolutions?

It is debatable whether AI, in combination with the plethora of automation technologies, will be more disruptive than past technological revolutions. It is clearly very difficult to make predictions as future impact depend on many factors including political, economic, market evolution, technical readiness, etc.

On the one hand, we already noted certain key elements of the ensuing AI-driven revolution. First and foremost, the phenomenon is not affecting a single sector as in the past, but it seems to impact all service and production sectors. Simultaneously, the transformation is happening at a much higher rate than in the past, which leaves little time for retraining to new jobs (e.g. machine translation has gone from research to production in just two years: faster than ever before). The distribution of new jobs – there will surely be new jobs, whose extent is difficult to foresee – may be highly uneven. We already see a concentration of new jobs related to AI in a few global companies, with a possible net loss in countries that cannot keep up with the pace of investment and innovation.

On the other hand, we may reason that the history of the past two centuries has been punctuated by disruptive technological changes. The steam engine in the nineteenth

century changed the landscape of agriculture, in particular in large territories where hundred thousand manual workers were replaced by harvesting machines. Railways made horse-based transportation disappear, the telephone, and later the Internet, killed the postal service, aviation destroyed passenger ships, TV impacted on cinema and movie distribution and so forth. Such an impact was by and large simultaneously negative and positive. Lots of existing jobs were lost and lots of new jobs were created. So far, the economic balance was positive. Quality of life improved considerably as well as life expectancy. Work conditions and work-life balance increased leading to, for example, the concept of free time (and hence tourism). These were not small changes and we need to factor them in into the equation.

Potential and limitations of AI

AI progress to date

Artificial Intelligence (AI) is a field with a varied tradition, covering different aspects of intelligence, such as: representation of knowledge and reasoning, machine learning, natural language processing, vision and speech, robotics with sensing, control of movement and manipulation, planning and coordination. Each of these areas developed its own specific methods – e.g. symbolic, probabilistic, evolutionary, neural networks, etc. – in order to address the complexities of each specific task.

In the evolution of the field, the following two waves[5] can be recognized:

- *Handcrafted knowledge*. Engineers create sets of rules to represent knowledge in well-defined domains. It enables reasoning over narrowly defined problems. It has no learning capability and poor handling of uncertainty.

- *Statistical learning*. Engineers create statistical models for specific problem domains and train them on data. The resulting systems exhibit nuanced classification and prediction capabilities but unfortunately poor contextual capability, minimal reasoning ability and, in general, though not always, poor explainability.

A *third wave* is still expected to emerge, which will be capable to build on both the perceptual abilities of statistical learning and the abstraction and reasoning capabilities required by more intelligent systems.

Deep learning seems promising in relation to the creation of systems that combine perception with abstract reasoning capabilities, since it allows the development of end-to-end applications, able to perform complete tasks that were once divided into simpler tasks. Deep learning is also applicable across fields: e.g. both speech recognition and machine translation can be dealt with the same type of models.

The combination of deep learning with reinforcement learning is another attempt to increase generality and the ability to deliver end-to-end applications, from perception to decision (and action in the case of robots). The design of bespoke algorithms for robotics is another frontier direction which exploits the ability of robots to both explore

the environment through action, sense it in various ways (vision, touch, etc.) and therefore automatically collect data for training. Clearly, this bears promise of more flexible and efficient production, new autonomous means of transportation, logistics, robot helpers, etc.

Limitations of AI

In essence, AI technologies (still) suffer from the following limitations:

- High computational costs in training (though not in prediction);
- Large annotated collections required for training (input data, categories, etc.);
- Accuracy measures not adequate to assess real applicability (evaluation criteria);
- Designing the architecture for applications is still an art that lacks a theoretical framework;
- Artificial General Intelligence still far away (though maybe not essential).

Exploiting deep learning involves huge computational requirements during training. For example, according to Jangquing Jia[6], director of engineering for Facebook's AI platform, "to train one typical ImageNet model takes about one exaflop of computing". Achieving human brain capability would require several orders of magnitude increases in computational power. According to Amodei and Hernandez[7], the amount of compute used in the largest AI training runs has been increasing exponentially with a 3.5 month-doubling time (by comparison, Moore's Law had an 18-month doubling period). On the other hand, AI applications based on deep learning models can be deployed on the edge[8] on cheap devices or embedded processors as long as they do not need to be trained on the fly to function.

Training large neural networks requires suitably annotated data from humans, which in turn needs significant effort to produce. For example, data from healthcare providers can be an invaluable source for AI systems capable of performing diagnoses or analyzing past patient history to discover better treatment strategies. The effort goes into making data suitable for training purposes by cleaning them up, annotating them and converting to a processable format. Also, guaranteeing that data is free of biases requires tremendous human effort both in data collection and successively in testing "in the field". Reducing the reliance on large training collections would require developing new theories of learning. For example, self-supervised learning methods or online acquisition of training data as in the case of robots is a concrete possibility which has been explored in the literature[9].

Each AI application must be tailored to the specifics of the problem to be solved and according to the data that represent it. The process of developing an AI application is in fact laborious, because it involves performing many experiments, repeating trials and verifications, until a sufficient level of accuracy is reached. This process requires having a suitable accuracy metric, necessary both for learning and for verifying the final performance. Therefore, difficulties and failures must be taken into account since the beginning.
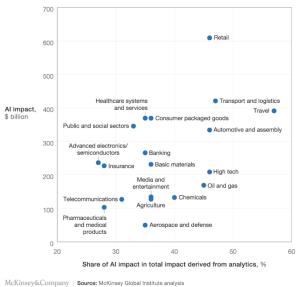
According to Michael Jordan[10] from UC Berkeley, the community may be focusing entirely on the wrong problem, that is "imitative AI", the copying of human-level intelligence, while there is ample room for using the existing statistical learning methods to set up a planetary scale learning and inference engine to address problems in the above-mentioned application areas. One example that illustrate the need for a large infrastructure is the domain of transportation, where the introduction of autonomous driving would require something closed to air traffic control rather than scattered human-type intelligence on each vehicle. Similarly, in healthcare, a large-scale medical infrastructure exploiting the ability to correlate data acquired from a plethora of sources such as images, genetic data, overall therapeutic outcomes, nutrition data would be key to improve cures and discovering general patterns when applying medical policies at the nation-wide scale.

Potential of AI

As we mentioned earlier, automation and AI are *per se* not new. The recent progress is nonetheless moving the boundary of what we can do. This offers a tremendous opportunity for addressing some of the most difficult societal challenges (sustainability, health care, demography). An example of application that has recently become technologically "feasible" is autonomous (driverless) transportation. To put it in numbers, according to the McKinsey Global Institute the applications of deep learning may account already for as much as 3.5 to 5.8 trillion dollars only in data analytics techniques (see figure below for additional details)[11]. In fact, as the pro capita productivity is declining (0.5% negative) from a positive 2.4% during the boom of the Internet, AI has the potential to reverse the trend in spite of a declining workforce due to the demographic megatrend (fertility index reached 2.4 average worldwide with advanced economies well below the 2.1 threshold needed to maintain the population stable).

The impact of AI is not limited to transportation and manufacturing. Text and data mining have enormous potential for scientific discovery and the development of new expertise. The very design process of complex system can be revolutionized by the application of advanced AI-based optimization methods (e.g. evolutionary design). The field of engineering can simulate not only the existing solutions but conduct "what if" searches corroborated by access to large databases of scientific literature, providing the designer with new tools to avoid "reinventing the wheel".

In the field of genetics, AI can help the design of better models. As the cost of sequencing DNA dropped even faster than Moore's law, the ability to process large quantity of data is becoming even more fundamental. The dogma of the so-called

precision medicine lays in the ability to link DNA to RNA to proteins possibly for each single patient. Applications of machine learning are blooming in this sector[12].

They range from the ability to interpret genetic variations, gene editing, to determining patient's risks for developing certain genetic diseases (prevention). In the research sector, understanding the role of the non-coding portions of DNA (only 2% of the human genome encode proteins) is yet another case where data analytics may have a large impact. Adjacent to the direct study of the human genome, there is the large field of genetics for agriculture. Improving soil quality, crop yield and understanding the link between diet and disease at the genetic level are all problems amenable to AI-based solutions. Research on climate change is also benefiting by the availability of data (new and more accurate environment measures) and, consequently, the data analytics needed to cope with this data deluge. AI here is replacing handcrafted models leading to more accurate predictions on how human activities for example affect the environment.

In summary, AI and automation can boost productivity and consequently global growth in GDP. They can help scientific discovery, generating a positive loop between new development and further growth. Seizing the opportunity while managing possible workforce transition, building ethical policies for technology application, and in general avoiding polarization of resources is a big challenge for governments to address.

Risks and benefits of AI on work and employment

Although it is difficult to quantify the number of jobs that can be created or lost by AI, we know that the deployment of AI will profoundly transform the content, organization and conditions of work. These changes will in fact affect the evolution of skills and trades in many sectors of human activities. In order to anticipate these transformations but also to "evaluate" the risks of machine substitution, one approach is to analyze them in terms of the transformation of tasks.

Any business consists of several tasks, some of which are "peripheral" or low added-value in nature, while others constitute the "heart" of the business and/or have high added-value. Depending on the type of job and task, AI does not have the same impact. Thus, the first question we need to ask is whether AI leads to substitutions or it is rather complementary to human work[13].

Substitution or complementarity

First of all, AI enables the execution of tasks that were previously impossible to carry out as they were either too time-consuming or not economically viable (control of 24/7 services such as health monitoring, etc.). In this case, there is no substitution with current work: these are new tasks supervised by automation. For example, in the banking sector, detection of transaction anomalies is facilitated by AI-based systems. In urban transport, autonomous shuttles may be[14] extending areas and times in which services are provided. These AI-based systems seem complementary to human tasks.

A second category is automation of tasks previously performed by human beings. This is by no means a new field: robotization in the automotive industry and digitization of

banking operations have been with us for a long time now and have led to repositioning workers to supervisory tasks. AI is very much part of such an evolution. Automatic highway driving, in the case of road freight transport, falls into this category. All the less complex based on predefined rules such as organizing, planning, control and information management activities are prone to be automatized as well.

Support functions are particularly concerned. These functions are transversal to many sectors such as retail, banking, insurance, marketing, legal services, and healthcare to name a few. Both skilled and less qualified workers in support functions stand high chances of being affected by automatization. The challenge is to anticipate what occupations will be mostly affected – i.e. with automation in the majority of their tasks.

The third type of task transformation concerns assistance in decision-making. Here the human task is not conceptually modified but the worker can draw on systems that can help boosting performance: diagnosis and therapeutic recommendations, customer service in the banking sector, etc. Human intervention is solely needed because of either limitations of technology or "acceptability reasons".

In the first case, the technology is not mature enough as for example, autonomous driving in heavy traffic conditions (in bad weather for example) or the detection of multiple pathologies on a patient whereby the analysis, collection and processing of data is complicated. In the second case, human intervention may be needed for reasons of "social acceptability" such as in interacting with patients or when decision making has individual consequences whereby existing regulations prohibit automatization. Human contact in domains requiring social interaction is often indispensable and therefore not replaceable by AI.

In addition, tasks that are related to the design and deployment of AI itself are at this point still a human endeavor including research and development, production, deployment, maintenance, control, data annotation, testing, validation, etc. These activities involve several types of qualifications from the scientist, to the researcher, down to the data annotator. While AI research requires highly qualified individuals (e.g. design autonomous driving vehicles), other activities required to create AI systems may simply require to be proficient at "doing the task" (e.g. driving a car). In the same category we find updating software learning system and algorithmic processes, detecting and correcting classification errors, collecting and digitizing data.

Uncertain effects on qualifications

AI technologies are often heralded as being well suited to tediously repetitive tasks, obeying rules that a machine can "learn". This is only half of the truth. AI can also handle complicated tasks – as for example autonomous driving. When AI manages basic tasks, this may result in an increase in workers' qualifications. In customer care services, if simple requests are handled by AI, complex cases are then handed to human advisors. This transfer may lead to an increase in skill levels and a growing need for social skills, required for management of complicated interactions with customers, patients or users. This in turn may result in work intensification, as simple routine cases, previously in the majority, are replaced by cases requiring greater attention.

For example, in the retail sector, AI may lead to further worker's specialization, in order to provide the customer with up-to-the-minute expertise and, simultaneously, to a more general set of skills to refer customers to the right specialist. Conversely, in the health sector, AI may bring about relative deskilling through the automation of complex tasks traditionally performed by humans. In contrast, nurses and radiographers may see their performance increased by the assistance provided by AI-based diagnostic systems.

This two-fold effect is not specific to AI, however, arising more generally from automation and digitalization, leading to workers taking on the role of "supervisors", which of course assumes their proficiency in the required digital skills.

More complex tasks at work can lead to a cognitive exhaustion

If the time freed up by AI is allocated in the same proportion to complex activities, there is a potential risk that some specialists qualify as "professional cognitive exhaustion" which can even lead to burn out. The question of balancing time spent on demanding versus that on less demanding tasks is thus an important question concerning the psychosocial well-being of workers.

For example, in the healthcare sector, time can be devoted to the dialogue and relationship with the patient, which is important since it influences positively the adherence to therapeutic protocols and healing process. The dialogue between doctor and patients could also be placed on a completely different level. As the need for transparency and explanation will become more important thanks to better access to AI-enabled knowledge, the doctor will have to give more and more reasons for having or not following the recommendations provided by AI. In short, he will have to justify his protocol and explain it.

Alternatively, time can be allocated to learning, promoting greater satisfaction and sense of purpose in the workplace. A physician, for example, may spend less time interpreting routine data while retaining the most complex cases. This reinforces the capacity for continuous learning and enrichment by placing higher focus on the value-added tasks related to the core of the profession.

In summary, even if the evolution of qualifications is uncertain, i.e. specialists vs. generalist, we can safely say that the introduction of AI "at work" will promote soft skills such as autonomy, communication, ability to listen, negotiate, and influence our peers.

Transformation of the organization of work

Deployment of AI systems leads to rethinking the organization of work. Once digitalization is underway, AI systems can improve operational performances by contributing to information management, activity planning and actor coordination. Predictive maintenance in transport infrastructures and activity planning in hospitals are both emblematic examples. However, if it is to be fully efficient, such organization assumes that users can provide feedback on system operation. For example, assistance systems in the banking sector have had counterproductive results when responses were not adapted. This phenomenon is simply a continuation of the administrative overload

that many organizations complain of, when activity coding takes precedence over the activities themselves.

As it is likely to foster better coordination of organization, AI may also lead to greater isolation of workers. Such is the case, for example, when automatic systems convey instructions to workers who become simple "performers". We should therefore not underestimate the risks attached to deployment of AI tools as far as working conditions are concerned (loss of autonomy, work intensification, etc.).

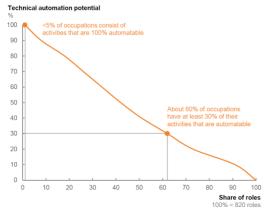**Interaction with AI systems: who is concerned**

Artificial intelligence affects three major categories of workers. First of all, production of AI services generates highly qualified work: researchers, data scientists, engineers and other specialized technicians. Having such qualified staff is crucial to successful integration of AI although only a limited number of jobs are involved. The main problem in this segment is that of training and retaining talents. Secondly, large numbers of workers will be using AI-based systems without necessarily knowing that AI is involved: the issue here is that of training them to use such tools correctly. Finally, AI is set to affect the relative importance of certain skills. Interpersonal skills, empathy, artistic and creative sensibilities and certain manual tasks will be positively reassessed due to the introduction of AI. As a result, activities that do not use AI will become proportionally more attractive.

In customer care, with technical skills being the purview of AI systems, interpersonal skills and dedication to customers, along with sale skills, are likely to see their importance increased. Just as digitization of society and the economy now affects all sectors, AI is likely to affect all professions, directly or indirectly, as regards their nature and/or the conditions under which they are carried out. Various forms of evolution are all the more probable in that successful deployment also depends on various non-technological factors.

Some numbers

Overall, the impact of AI may be clearly be positive as well as it is certainly unavoidable. It allows for discovering and building systems that did not exist before and would not even be possible just through manual design. A note of caution is obligatory as clearly any

Technical automation potential

automation activity bears the risk of job loss/displacement. A study by McKinsey Global Institute analyzes the automation potential across a variety of jobs (820 different roles) and finds what may be the technical impact of automation (see graph below). Clearly, while only less than 5% are totally automatable, there is though 60% of jobs whose 30% of tasks can be automated. On the other hand, adoption of AI and automation technologies is affected by technical feasibility, cost of development but also factors such as the labor market dynamics, overall benefit (cost-performance tradeoffs) and the regulatory and social acceptance.

Conclusion and recommendations

AI promises to revolutionize, and it is already revolutionizing, processes and paradigms in all application domains. Health care, environment, intelligent transportation systems, security, industry, public administration, marketing are just a few examples. As in the past, in each of these application domains some jobs will become irrelevant and rapidly disappear and new jobs will appear. New skills and new attitudes will be required to participate in the construction of the new world deriving from technological disruption.

AI is central in the current debate on social transformation. First and foremost, expected changes in the job market generate two contrasting attitudes. Some people proclaim their optimism in the face of a technology that ensures productivity gains and is therefore a source of wealth, promising to eliminate the most tedious and dangerous tasks. Others make pessimistic prophecies on the inevitable disappearance of whole realms of activity and corresponding jobs. The public debate is polarized an unproductive as it fails to highlight factors of transformation or levers for action. If we bet that AI will not mechanically lead to the disappearance of work, it is because such a scenario is based more on fiction than on science and aims to feed more anxiety and fantasies than to engage in real reflection and constructive dialogue on the subject. In order to shed light on the issue, this paper aims to be realistic and avoid fiction: identifying the potentiality and the limitations of AI is, in our opinion, a necessary step to identify appropriate public policies. The effects of AI, like all technologies, will depend on the use we make of it and the collective objectives we set for ourselves. It is the "encounter" between the uses and the technology itself, the major determinant of the future, whether disruptive or progressive. However, this does not mean that the effects of AI will not be significant and that we will work tomorrow as we do today.

In summary, we would like to put forward an optimistic view of the evolution of the work market deriving from the diffusion of AI. We recommend governments and international organizations to work at three levels to promote *training and education (also continuous)*, support *research and innovation* while making sure *AI remains "democratic"* in nature (accessible and available to everyone).

Who to address: a multi-layered approach

We propose a multi-layered training and education scheme specific to AI assuming that basic/mass education is a "first principle" of any of the G7 countries.

The first layer is that of scientists and scholars, i.e. the people who lead the development of AI's science and technology. These people typically live and operate in universities and public and private research centers. The center of gravity of this layer is undoubtedly located in North America, where a unique ecosystem based on top universities, companies and capitals has been consolidating for almost a century. All countries, and the international community, need to activate ambitious investment plans to play an active role in the development of the AI discipline. Higher education in AI and related areas must receive appropriate funding. A cultural mindset must be instilled in the young generations. They have to think that anyone can have a primary role in technological development and not merely the role of user. While funding of universities

and more in general of public research is the main instrument to achieve such an objective, other more specific initiative should be identified. One of these initiatives is the creation of joint international research centers devoted to AI, aimed at bringing together researchers coming both from the public sector and from the private sectors to stimulate the implementation of pilot projects and more in general innovation.

The second layer is that of domain experts, i.e., the people who master the application domains of AI. While in AI research the issues are the same as in any other scientific area, namely excellence in research and higher education, on the contrary domain expert needs to pay special attention. Here the problem is how to leverage the experience and the knowledge of domain experts who know little or nothing of AI. It is a very delicate and ambitious objective as it requires the development of courses and material aimed at mediating between technology and domain processes and organizations. While such mediation will take place independently from specific political actions, an appropriate governance will make it more fluid and effective. For each relevant application domain, a task force should evaluate the potential of AI and define the skills needed by the domain experts to manage, govern or supervise its deployment.

The third layer is that of school teachers and educators, at all levels. New course programs must be developed to accompany the evolution of young students during their paths from elementary school up to high school and college. The complexity of this effort is huge, as it impacts on school organization which typically tend to resist to change. Governments and international organizations must take an active role in promoting such an evolution.

The fourth layer, which is the largest in number, is concerned with people in general. Not all people are supposed to become AI expert, but they must be helped to develop a sensitivity towards data. They must be helped to analyze phenomena quantitatively, through tables, charts and more in general numbers. It is only through numbers that AI can deliver its potential on education. AI can be a great opportunity to move from "training", which is typical of manual tasks, to "education", which is typical of human cognitive skills. People must not be "trained" to push buttons to control unknown actions, no matter whether the buttons are mechanical or appear on a touch screen. People must be educated to master and understand what they are required to do. In this case, governments and international organizations must take an active role to promote, fund and monitor such type of education, through universities and professional schools.

How to address individuals: continuous learning

While AI does not necessarily bring concern about employment, which itself changes over time, there is a profound question about learning, especially what to learn and how to learn it. AI emphasizes the importance of "learning" to learn. Taddei's report[15], like many others on the subject, already stresses the need to develop "learning" societies. The study by Benhamou "the future of work in 2030"[16] emphasizes that the job market will require individuals to be highly adaptable and highly autonomous. The key to success, according to the author, is the ability to integrate and compose new knowledge rather than the initial mastery of specific knowledge. Working in a team, coordinating information from a variety of sources, transmitting organizational knowledge are skills that will be increasingly required the job market. Other skills will also be crucial for

companies operating in an increasingly innovation-oriented environment. The ability to solve complex problems, be critical and creative are the skills that will be increasingly required by companies. The recent World Economic Forum's report on the future of jobs[17] ranked complex problem-solving, critical thinking, creativity and managerial skills among the top four skills that will be most in demand by 2020. Cognitive and behavioral skills – so called *transversal skills* – are also those that will best ensure the complementarity between AI and human intelligence.

One of the major challenges that will therefore emerge is the ability of individuals to learn continuously and to develop their transversal skills "on the market". The evolution of the initial training system is an important issue insofar as it can promote the development of these skills "upstream" via the education system. The continuous training system is also essential to support workers in continuously developing their skills and enhancing them (either through certification, validation of professional experience). Other levers will also have to be mobilized to ensure complementarity in work organizations and to support people's learning capacities, as highlighted in Section 3. AI can also foster new innovation towards the so-called "learning" organizational models. In this case, the challenge of integrating technology becomes an advantageous learning at the organization level, which becomes fluid as the structure of the organization itself can be learnt as technology is deployed. Such organizations enable response not only to technological issues but also to societal expectations as regards greater autonomy and wellbeing in the workplace[18].

As AI typically work on "historical data", it may stifle innovation and creativity by encouraging conservativism in decision making. Progress requires creativity and risk taking. In addition, algorithms often follow high-level rules predetermined by the designer that do not necessarily correspond to the complex reality of the user. It is therefore imperative that the human worker stays in control, that we encourage critical thinking. Individuals will have to be able to question the recommendations of AI systems in order to enrich them with their preferences and innovative interpretations of results. One of the major challenges of the initial education system and continuing training will be to equip individuals with skills that enable them to move beyond pre-established "frameworks and standards" and to develop a "systematic thinking" to analyze complex phenomena. Initial and continuing training systems will have to find a good balance between the acquisition of formal and academic knowledge but also knowledge based on "experiential" learning based on trial and error, creativity and risk-taking.

Secure impacted sectors

A continuous transition scenario would appear to be the easiest to manage. It corresponds to an evolution where skills, organizations and workers make use of AI tools in order to improve efficiency, free themselves from arduous tasks and perform new tasks that were not possible before AI. This form of evolution is not the only scenario, however. Disruption may occur if advances in AI come sooner than expected. For example, drivers may disappear altogether. If the transition is progressive, retirements, and professional training in new transport professions may be enough to accompany it. If it happens that a safe automatic truck becomes available, authorized by public authorities, accepted by road users and economically viable, major professional conversion problems may arise relatively quickly.

The "skills blocks" approach that has been developing for several years now within the training system and among certifying bodies may well provide an answer to this major issue: a qualified or certified individual may only have to adapt his/her skills via a complementary training "module" without having to complete a whole qualifying or certifying program[19]. This approach, would have a dual advantage for organization of the continuing training system, with regard to responsiveness in the face of technological changes as well as to financial cost[20]. It would involve acceleration of the process of breaking down certifications into skills blocks. The "block" approach would also enable more rapid adaptation of skills reference frameworks (in as much as it would only be partial): even though it still remains to be tested, this hypothesis would contribute to ensuring the relevance of frameworks in the face of changes resulting from AI.

Support research and innovation in AI

The rapid evolution and diffusion of AI requires appropriate actions to support research and innovation both at the government level and clearly at an international level. Actions should address basic and applied research, innovation (technology transfer), and stimulate the creation of higher education programs, at the Master and Doctoral level, in AI and, more in general, in computer science and engineering. In short, government and international organizations should support:

- *Basic research* that aims at the development of theory, models and methods for knowledge extraction and processing as well as the software and hardware architectures;

- *Applied research* in coordination with experts from the application domains of AI utilizing real-world metrics to assess the quality of results including the definition of "grand challenges" or "competitions" in well-targeted domains [21];

- *Innovation actions* especially geared to small-medium enterprises, startups and in general to support advanced industrial initiatives in the areas of AI;

- The creation of *open and rich data sets* through large scale pilots and taking into account the issues of privacy and copyright;

- The creation of *forums* (in the form of panels, academies, associations, etc.) that bring together all stakeholders from both the private and public sector;

- Innovation through *private capital* (e.g. bank, private equity, venture funds) *to invest in AI* projects, with appropriate mechanisms to minimize risk.

Finally, governments and international organizations should pay attention to the legal issues associated to the deployment of AI-based solutions as well as to the general issues of privacy. The requirements are to promote widespread use of AI "with care", without stifling the development of technology or imposing unnecessary burdens to the market. This is a delicate balance that needs to be studied in depth.

Promote vectors of AI democratization

A means to favor training the population at all levels is to make widely available the components of AI technologies. The access to data is crucial for researchers and scientists involved in the development and the analysis of AI-based methodologies, while the access to algorithms is essential to field experts or anyone interested in the experimentation of such techniques. Furthermore, the adoption of AI should be helped by lowering the bar of technological access (cost, time, resources). This is a matter of policies as well as technology.

It may be useful to implement a "one-stop shop" for AI per application domain, where anyone could find (or gain access) to software, hardware, data, knowledge, example applications, rather than just simply purchase AI solutions. The more the "shop" is open (in the sense of Open Data, Open Science, Open Source), the better for the citizens. Users' communities may group and thrive around such a platform as it has been often demonstrated with Open Source communities (refer for instance to the development of open source operating system such as Linux). As AI still requires enormous resources, democratizing it may have to go through the development of methods that are as energy-efficient as the human brain – a few dozens of watts versus the megawatts of present-day computer farms – learning from smaller data and requiring less operations per data sample.

We recommend authorities and public institutions to promote and favor the access to such open platforms, while providing a framework and governance to face possible issues of responsibility and competitiveness (misuse of data, off-shoring).

---

[1] Cedric Villani. AI for Humanity: French Strategy for AI. 2018. https://www.aiforhumanity.fr/en/

[2] Future of Employment, C. Frey and M. Osborne, Oxford.

[3] Erik Brynjolfsson. What AI can do? Workforce implications. April 2018.

[4] Steve Lohr. A.I. Will Transform the Economy. But How Much, and How Soon? NYT 2017

[5] https://www.darpa.mil/attachments/AIFull.pdf

[6] https://www.techrepublic.com/article/four-ways-machine-learning-is-evolving-according-to-facebooks-ai-engineering-chief

[7] https://blog.openai.com/ai-and-compute/

[8] https://www.st.com/content/st_com/en/about/innovation---technology/artificial-intelligence.html#edge

[9] https://project.inria.fr/paiss/files/2018/07/zisserman-self-supervised.pdf

[10] https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7

[11] https://www.mckinsey.com/featured-insights/future-of-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

[12] https://www.techemergence.com/machine-learning-in-genomics-applications/

[13] S.Benhamou and L. Janin (2018): Artificial Intelligence and Work, France Stratégie Report, mars 2018. This mission was carried out in parallel with that entrusted to Member of Parliament Cédric Villani, which, given its wider scope, covers questions of research on AI, industrial development of AI and its applications in the public sphere, along with ethical issues and social acceptance.

[14] Beware of alternative predictions: http://rodneybrooks.com/bothersome-bystanders-and-self-driving-cars/

[15] Taddei.F, C. Becchetti-Bizot. C and G. Houzel (2017), Rapport sur la recherche et développement de l'éducation tout au long de la vie for the French Minister of National Education, Higher Education and Research

[16] S.Benhamou (2018): The world of Work: in 2030: Four scenarios in Work in the Digital Age: Challenges of the Fourth Industrial Revolution, Ed. Max Neufeind, Jacqueline O'Reilly, Florian Ranft, Rowman and Littlefield Intl., London-New York.

[17] World Economic Forum Report (2016): The Future of Jobs, Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution, January 2016.

[18] S.Benhamou (2018): *op-cit*

[19] Employment Skills Network (REC) (2017), Transferable and crosscutting skills: what detection, recognition and valorisation tools for individuals and companies?), France Stratégie, April

[20] S.Benhamou and L. Janin (2018) :*op-cit*

[21] https://en.wikipedia.org/wiki/DARPA_Grand_Challenge/

Draft – For discussion purposes only; does not represent the views of the G7 or its members.

16