Santander X Innovation Xperts

Transformative potential of generative AI:

innovation, impact, and implications for the future workforce

JUNE 2023

By Santander

Authors

Esteban Moro

is data scientist at the MIT Connection Science at IDSS and a tenured associate professor at Universidad Carlos III. Previously, he was a visiting professor at MIT Media Lab and researcher at University of Oxford. He holds a PhD in in Physics and is an affiliate faculty at Joint Institute UC3M-Santander on Big Data at UC3M and the Joint Institute of Mathematical Sciences (Spain). He has published several peer-reviewed, widely-cited studies on the interaction between technological advances and the explosion of data in our cities and work environments.

Morgan R. Frank

is an assistant Professor in the Department of Informatics and Networked Systems and the Department of Information Culture and Data Stewardship in the School of Computing and Information University of Pittsburgh. He is also a MIT Connection Science Fellow and a Research Affiliate at MIT's Media Lab, as well as a Digital Fellow at the Stanford Institute for Human-Centered Artificial Intelligence and Stanford Digital Economy Lab, among other affiliations. His research focuses on embracing the complexity of artificial intelligence (AI), the future of work, and the socio-economic consequences of technological change.

Table of contents

01	Intoduction	04
02	Generative AI tools	06
03	Occupational exposure to generative AI	09
04	Demographics of exposure to generative AI	18
05	Substitution vs. complement	20
06	Roles, tasks, and sector implications	24
07	Opportunities and future directions	30
A1	Annex 1 | Glossary	34

01

Introduction

Generative Artificial Intelligence (GenAI for short) has experienced significant advancements in recent years, becoming a potent instrument for transforming numerous sectors and industries. These advancements have opened up new avenues for innovation, from the creation of novel content such as music, art, and text to the discovery of new molecules in the pharmaceutical industry.

Beneath the umbrella of generative AI, there is a set of new algorithms capable of generating seemingly new realistic context —such as text, images, or audio—from training data. The most powerful generative AI algorithms are built on top of foundation models that are trained on a vast quantity of unlabeled data in a self-supervised way to identify underlying patterns for a wide range of tasks. While generative AI has existed for a long time, Large Language Models (LLMs) such as GPT (Generative Pretrained Transformer) or text-to-image models like Stable Diffusion have significantly advanced the capabilities of generative AI, opening the door to solving unprecedented complex problems, creating art, or assisting in scientific research. The key component of those models is their creativity, i.e., they are able to generate new, realistic content because they capture the underlying structure or pattern of the input data.

Generative AI has taken creative industries by storm, revolutionizing the way we interact with these models and how we work, communicate and create. ChatGPT, the LLM tool by OpenAI, is the fastest-growing consumer application in history, with more than 100 million active users in 2 months. Midjourney, another GenAI tool for generating images from text prompts, has between 1.5 and 2.5 million active users at any given time. Big companies such as Microsoft and Google are rapidly adapting their search and workplace tools to new generative AI technologies. Even in its early stages, generative AI is already shaping the future across various domains, and its influence on our lives, companies, economy, and society is set to grow exponentially in the next few years.

Like other technologies in the past, generative AI will likely affect the economy in many ways, potentially boosting economic growth and changing the way people work. At the same time, however, many are concerned about implications for employment and wages. Reports indicate that generative AI has the capacity to raise global GDP by a staggering 7%[1]. The broad impact of generative AI tools, combined with the rapid pace of improvement in recent months, has also sparked public interest in understanding whether and how generative AI will reshape our occupations and workplaces. To uncover the difference between positive and negative employment outcomes for workers, we need to identify the exposure of occupations to generative AI applications while modeling how workplace activities may adapt. This paper shows the current understanding of how different skills, jobs, workplaces, and sectors will be more affected by generative AI. Section 2 describes the current technology of generative AI and potential avenues for its evolution. Although it is very early to see the impact of generative AI on our labor markets, in Section 3 we present key ideas about which jobs, economic sectors, demographic groups and countries are more exposed to the impact of generative AI, building on previous research on the impact of automation on labor markets. The paper concludes with recommendations for policymakers and stakeholders on how to prepare for the impact of generative AI on labor markets and industries.

02

Generative AI tools

Generative AI systems fall under the general category of machine learning tools, but the revolution in the recent years has been in applying these tools to general creativity in language and image generation. OpenAI's ChatGPT builds on rapid advances in so-called large language models (LLMs). LLMs are trained through unsupervised learning using vast amounts of text data. They utilize neural networks, called transformers, with self-attention mechanisms to understand language patterns and contextual relationships. The training involves pre-training the model on a large corpus of text and fine-tuning it for specific tasks or domains. LLMs are not limited to text, but can also be trained on other types of information, such as images, videos, audio, software, protein structures, or financial data. In these contexts, the model is able to generate answers from specific inputs using data patterns and contextual relationships between different parts of the data.

The process of text generation in language models such as GPT is based on probability: given a starting sequence of words or "prompt", the model estimates the probabilities of the next word or sentences and chooses the one with the highest probability. That selection is not always deterministic. It is controlled by the "temperature", a parameter that influences the randomness of the model's output to make the output text more diverse or creative. Higher temperature values result in more random and diverse outputs, while lower values make the outputs more focused and deterministic. In other words, the higher the temperature, the more exploratory the text and the more creative the elements. However, it can also lead to the generation of nonsensical or hallucinatory responses that are sometimes factually incorrect.

Generally trained models like those in ChatGPT lack the kind of focus required for specific purposes like health and financial applications. But personalization or fine-tuning of LLMs, for example, can be achieved by providing them with more specific examples to learn from. The idea is to upgrade pre-trained models by further training them with questions and answers pairs relevant to our context or task. For example, an LLM model can be fine-tuned for specific tasks such as recognizing disease symptoms in medical texts, providing more objective answers to general questions, being used as a chatbot in customer assistance, etc. Another approach is to train the model from scratch incorporating domain-specific datasets. An example of this is the recently announced BloombergGPT, targeted to financial language. Finally, generative AI tools like LLMs can be fine-tuned with our own data to provide a hyper-personalized experience.

Moreover, initial training of these GPT models was done with data until September 2021, and thus they have a "knowledge cutoff" that prevents them from learning new information that occurred after that date.

However, LLM tools can now connect to the internet to retrieve information through plugins. This allows for more factual and up-to-date responses in LLMs like ChatGPT or better search responses in tools like Bing (Microsoft) or Bard (Google). In the near future, any generative AI tool might be able to incorporate near current data from the internet or local documents/data.

Most of the generative AI tools are today in the hands of a small set of private companies, raising concerns about monopolization of the technology. The cost of training and maintaining tools like Midjourney or ChatGPT are out of reach for most individuals and companies. ChatGPT costs $700.000 per day to operate, while GPT-3 required $3.2 million in computing resources alone to be trained. This could lead to a further digital divide between Big Tech companies and wealthy institutions in the global north and those that cannot (nonprofits, small startups, etc.). But both the cost of accessing these tools and training them are falling. For example, the cost of AI training dropped 100-fold between 2017 and 2019. According to some reports, recent advancements in computing and hardware have plummeted the training cost of an LLM with a performance level similar to GPT-3 to $450.000[2]. At the same time, more pre-trained LLM models have become available, including some open-source options like Meta's Llama. Running one of these models on your laptop is now possible. Similar to Moore's Law, some experts predict that the cost of training and running these models will be minimal by 2030. Google has recently released a solution that allows LLMs to be deployed natively on Android devices, suggesting that we could see extensive use of generative AI tools on a myriad of devices and platforms in the coming months, including highly personalized generative AI with our own personal information.

Finally, most of these models are trained on internet text data, from social media to books, blogs, etc. Similarly, image generation models such as Stable Diffusion or Midjourney use the same image corpora to be trained. Therefore, they may reflect and even amplify the biases present in these data. These include, but are not limited to, racial, gender, and political biases. This also raises some questions about privacy, property rights, and copyright as they are trained on the public text and are re-trained with new questions or images posted online by millions of people every day. Regulators are already taking care of these concerns. Both the U.S. and Europe are working on what could be the first set of rules to globally govern general purpose AI systems (which include systems like ChatGPT) set to be approved later this year.

Although it is impossible to predict exactly what tools and uses the new generative AI will bring, developments in computing and hardware, as well as the general character of these tools, will make them reach everyone, everywhere, and in everything we do, write, work, or communicate today and in the future. Most of the generative AI tools we use today and in the next few years will be commoditized and easily accessible either through APIs (application programming interfaces) or natively in our devices. Thus, we can expect a pervasive, easy-to-use and hyper-personalized deployment of this technology in all aspects of our lives.

03

Occupational exposure to generative AI

The rapid advances in technology—particularly in the areas of AI, automation, and robotics—are causing a paradigm shift in parts of the workforce and raising questions about the future of work. The extent to which occupations are exposed to technology is becoming a key metric for understanding the susceptibility of jobs to being replaced or transformed by these technological innovations. Identifying the extent to which various tasks, skills, and responsibilities within an occupation can be automated allows academics, policymakers, and organizations to anticipate disruptions and develop strategies for worker retraining, economic diversification, and social support that will foster workforce resilience and adaptability in an increasingly technology-driven era.

Previous technologies automated repetitive tasks without compromising creative work. In contrast, generative AI can execute various creative tasks with minimal human guidance. This capability requires new methods to measure how much a job is exposed to generative AI. An occupation is exposed to generative AI if current AI applications can perform any of its current workplace activities more efficiently than workers. In the presence of new AI tools workers must shift their work activities in order to complement the work that can be done by AI. Workers who fail to adapt their activities may lose their jobs. Thus, exposure does not directly predict employment loss (i.e., job separations or unemployment), but instead indicates where the most change in the workforce is likely to occur and which workers are most likely to need to adapt.

Although generative AI tools are new, several efforts to estimate workplace exposure have emerged[3]. Some of these preliminary analyses are concerning: for example, two-thirds of U.S. occupations are exposed to some degree of automation through AI[4]. Here we use exposure estimates developed by Felten[5] at the level of both workplace abilities and whole occupations. These estimates are designed for workplace abilities in the O*NET database, which is used by the U.S. Bureau of Labor Statistics to represent the skill requirements of over 700 occupations in the U.S. economy. Occupations are identified according to their Standard Occupation Classification (SOC) code. Using the estimations of what skills will be impacted by generative AI, we can calculate the exposure of these jobs to these tools. Table 1 shows the top ten jobs in the US most exposed to language modeling and image generating technologies. As we can see, white-collar jobs in finance, education, and law are the most exposed to generative AI. In some sectors, most occupations are heavily exposed to generative AI. For example, *Financial Analysts,* *Financial and Investment Advisors,* and *Financial Managers* are all in the top 10% of most exposed occupations. This does not mean that these occupations will disappear, nor that employment or wages will decrease. Rather, their exposure highlights where things could change in the economy. On the other hand, *Dancers and Choreographers, Crop Farm Laborers,* and *Building Construction Laborers* are among the least exposed occupations.

Table 1.

Top 10 occupations more exposed to change provoked by language modeling or image generation in the U.S. Only occupations with more than 50,000 workers are shown.

	Ocupation	Exposure to Language modeling
01	Management Analysts	0.70
02	Human Resources Specialists	0.69
03	Financial Managers	0.69
04	Market Research analysts and Marketing Specialists	0.69
05	Lawyers	0.68
06	Accountants and Auditors	0.68
07	Medical Secretaries and Administrative Assistants	0.68
08	Middle School Teachers, Except Special and Career / Technical Education	0.68
09	Secondary School Teachers, Except Special and Career / Technical Education	0.67
10	Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	0.67

	Ocupation	Exposure to Generating images
01	Architects, Except Landscape and Naval	0.60
02	Interior Designers	0.60
03	Civil Engineers	0.60
04	Operations Research Analysts	0.59
05	Art Directors	0.59
06	Aerospace Engineers	0.59
07	Financial Examiners	0.59
08	Accountants and Auditors	0.59
09	Chief Executives	0.59
10	Financial Managers	0.59

Further insight might be gained by calculating the exposure by economic sector. For this purpose, we aggregate the share of workers in each occupation in that sector, weighted by the exposure of each occupation to generative AI. The results are presented in Tables 2 and 3. Again, we find that *Finance and Insurance, Real State, Professional Services*, and *Accommodation and Food Services* sectors are the most exposed to language models or image generation, as they are not only the most exposed, but also have the relatively largest number of workers.

Table 2.

US industries (i.e., two-digit NAICS codes) ranked by exposure to AI language modeling according to 2020 employment data from the US Bureau of Labor Statistics.

Industry	Exposure to Language Modeling (%)
Finance and insurance	58.020
Real estate and rental and leasing	53.709
Accomodation and food services	52.398
Professional, scientific, and technical services	52.220
Utilities	51.642
Wholesale trade	51.144
Management of companies and enterprises	50.455
Other services (except public administration)	49.299
Arts, entertainment, and recreation	49.183
Construction	48.524
Mining, quarrying and oil and gas extraction	48.104
Government	47.901
Educational services; state, local and private	46.500
Administrative and support and waste management...	46.412
Information	46.353
Federal government	44.555
Healthcare and social assistance	42.927
Agriculture, forestry, fishing and hunting	7.2827

Table 3: US industries (i.e., two-digit NAICS codes) ranked by exposure to AI image generation according to 2020 employment data from the US Bureau of Labor Statistics.

Industry	Exposure to Generating Images (%)
Real estate and rental and leasing	50.555
Accomodation and food services	50.553
Finance and insurance	50.345
Utilities	50.089
Construction	49.868
Wholesale trade	48.958
Mining, quarrying and oil and gas extraction	48.818
Arts, entertainment, and recreation	46.986
Professional, scientific, and technical services	46.640
Other services (except public administration)	46.505
Government	45.308
Administrative and support and waste management...	45.235
Management of companies and enterprises	44.752
Information	42.273
Federal government	41.843
Educational services; state, local and private	41.058
Healthcare and social assistance	39.559
Agriculture, forestry, fishing and hunting	7.297

We calculate exposure scores for other countries' economies using an official crosswalk between SOC codes and International Standard Classification of Occupations (ISCO) codes. This enables us to assess the exposure (percentage of workers affected) to generative AI based on their occupational employment distributions according to EUROSTAT (for European countries, see Figure 1) and ILOSTAT for a more global assessment (see Figure 2).

Figure 1:
Mapping exposure to generative AI in European countries.

(A) Employment exposure to AI language modeling. (B) Employment exposure to AI image generation. (C) Relative to other European countries, employment in blue countries is more exposed to AI language modeling than to AI generated images.

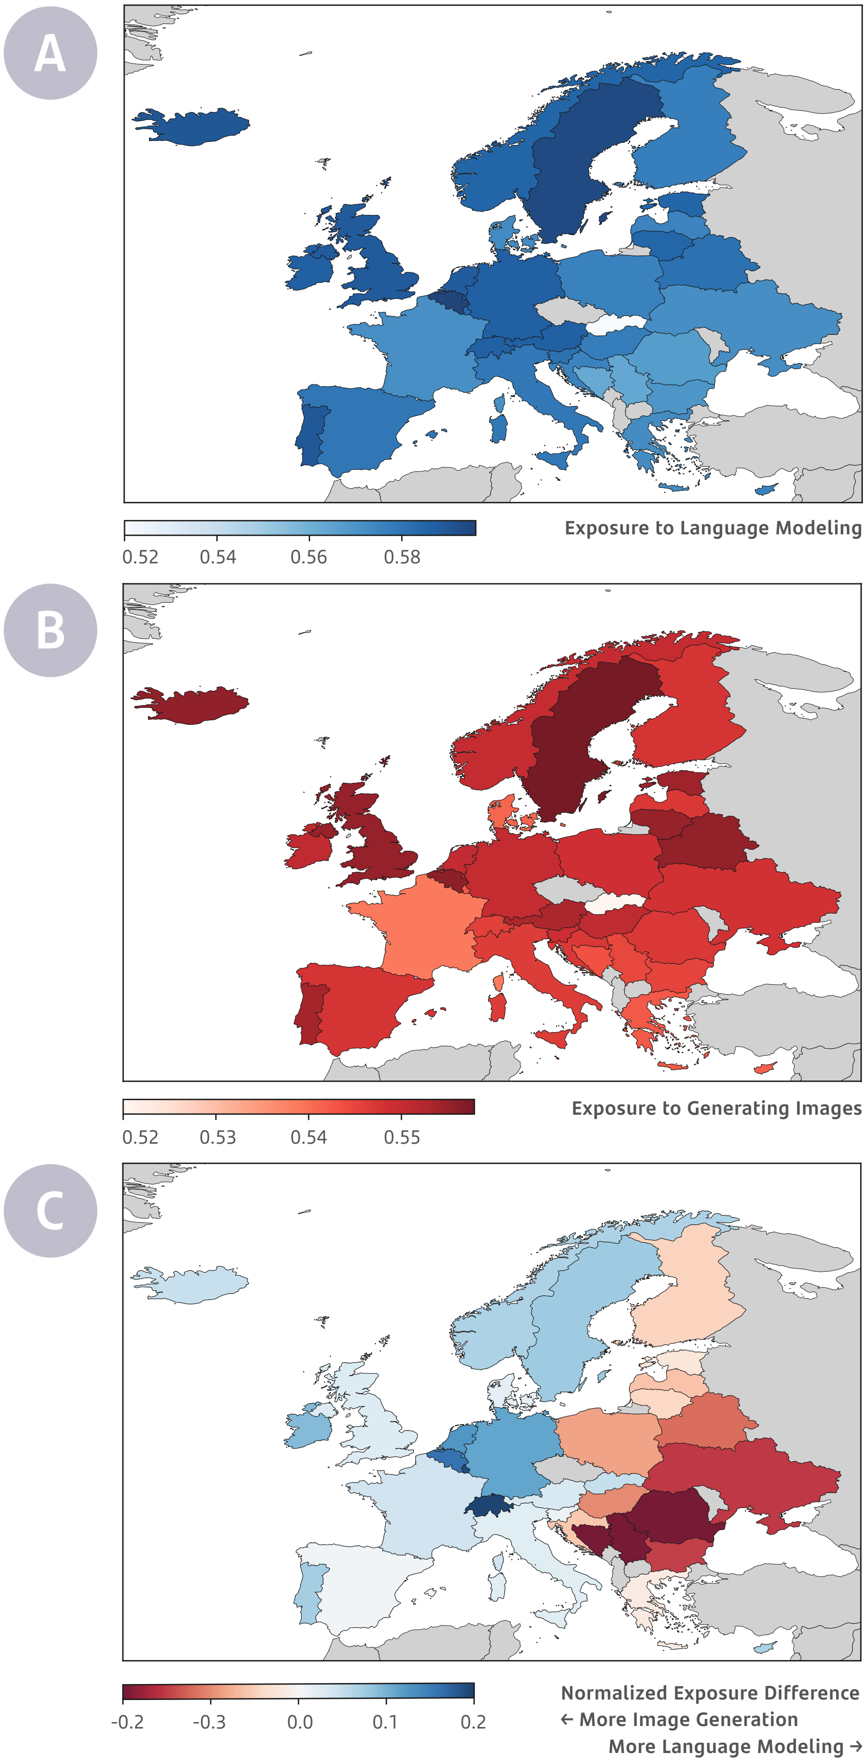

Figure 2: Mapping exposure to generative AI across the world.

(A) Employment exposure to AI language modeling. (B) Employment exposure to AI image generation. (C) Relative to other countries, employment in blue countries is more exposed to AI language modeling than to AI-generated images.

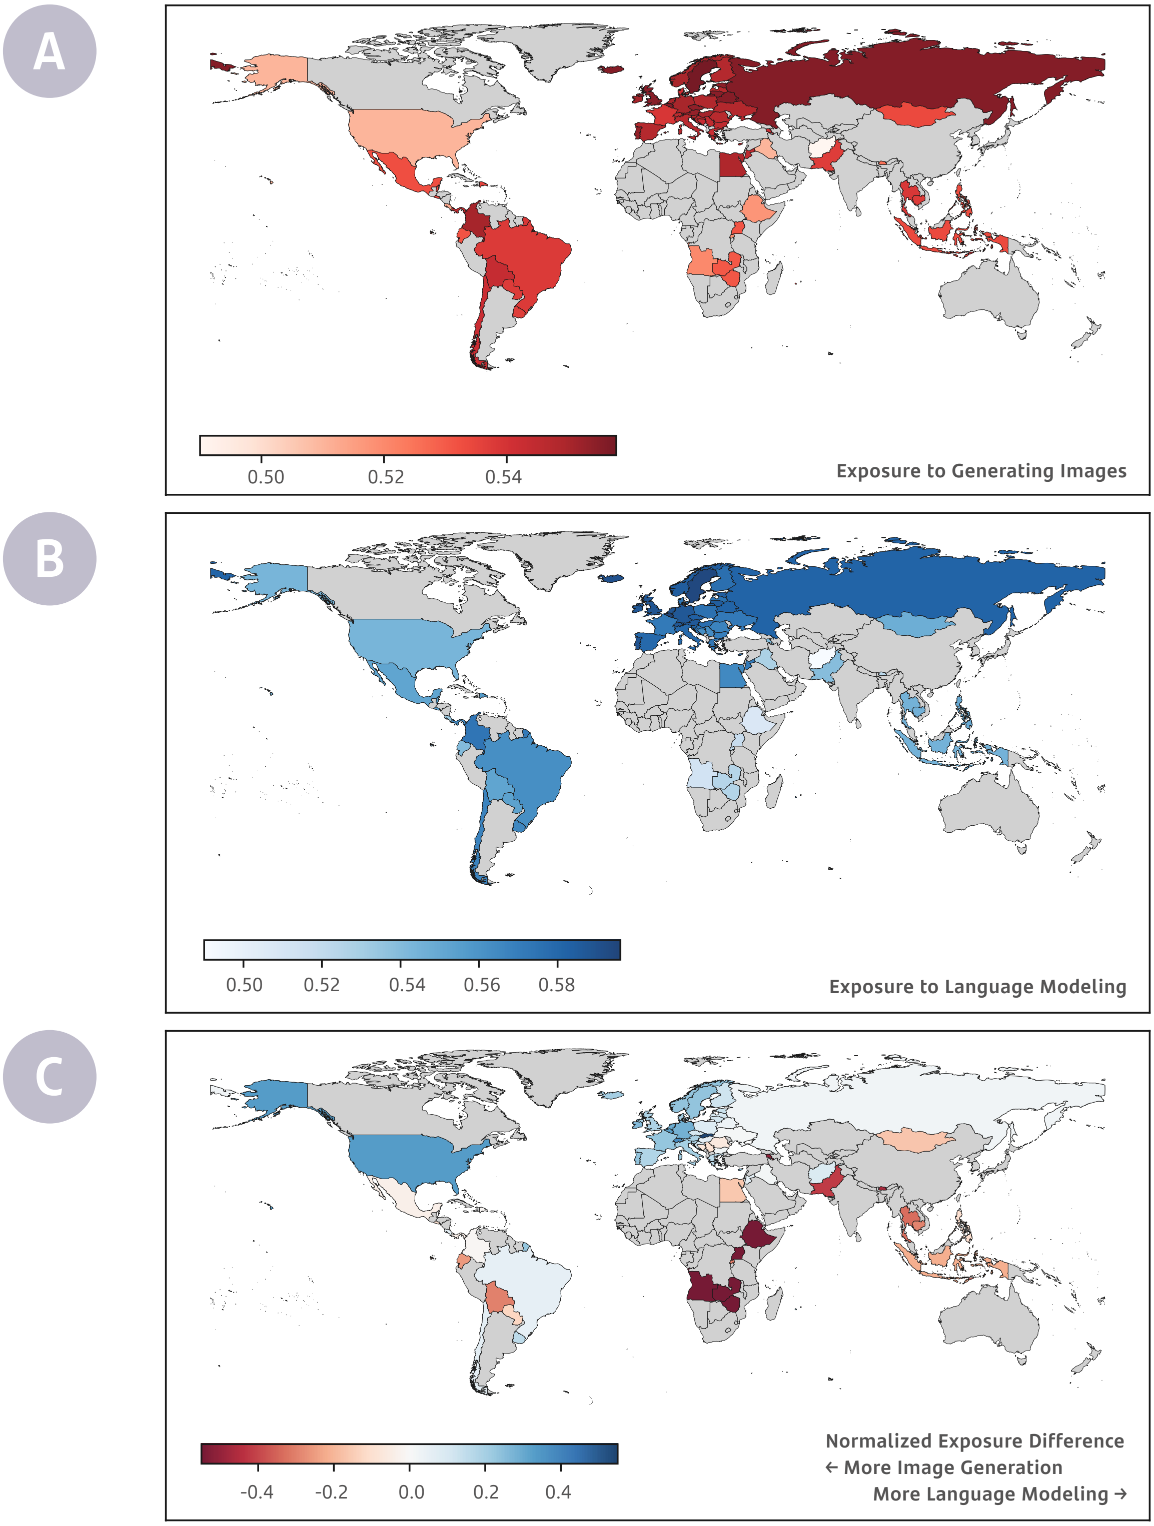

In Figures 1 and 2, we color countries by their share of employment that is exposed to AI language modeling (e.g., ChatGPT or Bing) or AI-generated images (e.g., DALL-E or Midjourney). All countries exhibit significant exposure to both language modeling and image generation. National scores around 0.50 suggest that demand for half of all workplace abilities may shift in response to generative AI. However, in general, wealthier countries show relatively higher exposure to language modeling than to image generation.

One way to understand this heterogeneity is to compare differences across countries in the distribution of employment across occupations. For example, in panel C in Figures 1 and 2, we color countries according to their relative exposure to AI language modeling or AI-generated images to quantify which generative AI application produces the greatest exposure in each country compared to the average exposure across all countries[i]. In summary, we found that countries with larger exposure to language models tend to have a greater proportion of in white-collar occupations associated with *Services, Professionals* or *Managers.*

i. Specifically, for each AI application we normalize the distribution of countries' exposure scores relative to the mean (i.e., we calculate z-scores) and subtract a countries relative exposure to image generation from their relative exposure to language modeling.

In focus

Examining GenAI impact: case studies of selected economies

Although it is individual occupations and workplace activities that are exposed to generative AI, the economic context is important in explaining national exposure. To see this, we compare specific pairs of countries with different exposure to generative AI (see Figure 3). For example, France is less exposed to AI language modeling than Belgium, in part because *Managers* and *Professionals* are more exposed and account for a larger share of the Belgian workforce than in France. At the same time, *Agricultural/forestry* and *Service & sales workers* are less exposed to AI language modeling and are more abundant in France than in Belgium. Interestingly, there are some occupations that contradict this general trend: *Elementary occupations* are less exposed to AI and are more common in Belgium while *Craft & related trades* occupations are more exposed and more abundant in France. We provide a similar analysis in panel B to explain why Sweden is more exposed to AI image generation than Denmark based on their employment distributions. These analyses highlight the parts of national economies that need the most attention as workers adapt to new generative AI tools and can help to explain why some countries will adapt more easily than others. For example, Sweden should investigate AI image generation and *Clerical support workers* if it wants comparable outcomes to Denmark.

Figure 3: Explaining differences in AI exposure.

(A) Comparing exposure to AI language modeling in France and Belgium. (B) Comparing exposure to AI image generation in Denmark and Sweden. In both plots, occupations on the right side have more AI exposure than the less exposed country in the example. For instance, *Managers* have more exposure to AI language modeling than France when taken on whole.

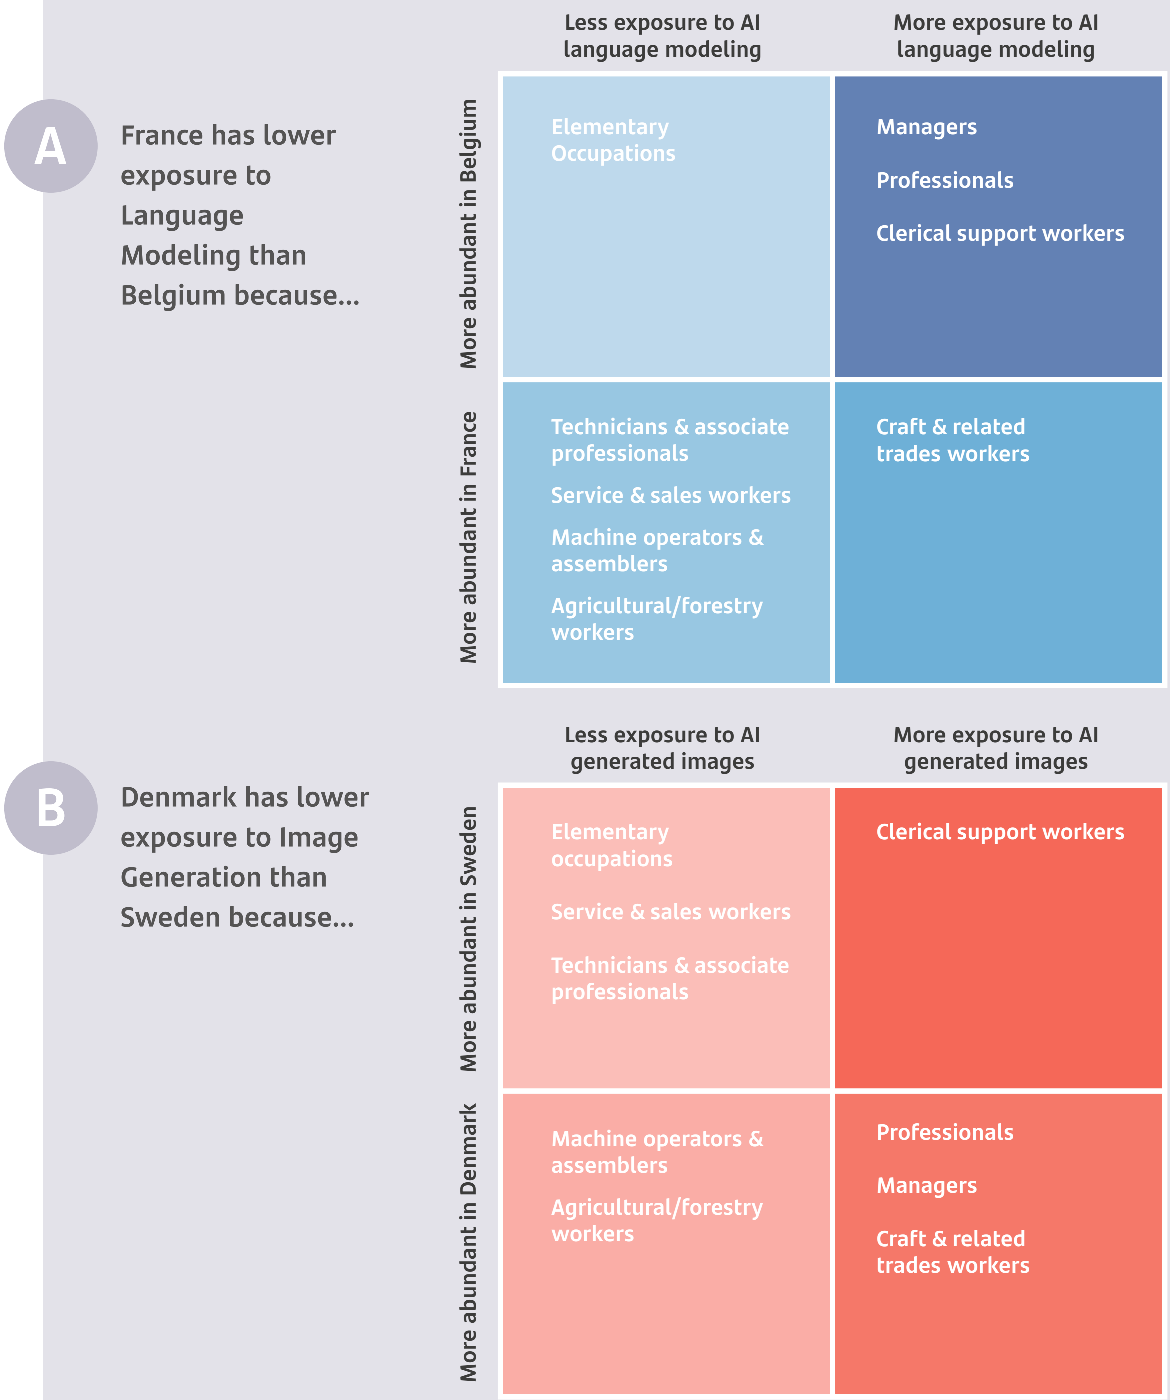

04

Demographics of exposure to generative AI

Since exposure to generative AI varies by occupations, exposure may also vary by demographic composition and location in European countries. Therefore, we explore GenAI exposure by gender and labor market type (i.e., urban or rural) using ILOSTAT employment statistics.

First, we expect urban employment exposure to account for a larger share of total exposed employment because cities have large populations and a larger share of white-collar workers. Figure 4A supports this hypothesis. However, when looking at exposure of workers to GenAI within urban or rural labor markets, the share of exposed employment is more comparable, especially for AI-generated images (see Figure 4B). These results indicate that exposure to generative AI is an equally urban and rural phenomenon, while exposure to AI language modeling is more abundant in urban settings.

Figure 4: generative AI exposure in urban and rural areas across European countries. (A) Exposure in urban labor markets make up a larger share of employment that is exposed to AI language modeling or AI image generation. (B) However, within urban or rural labor markets, relative exposure is comparable.

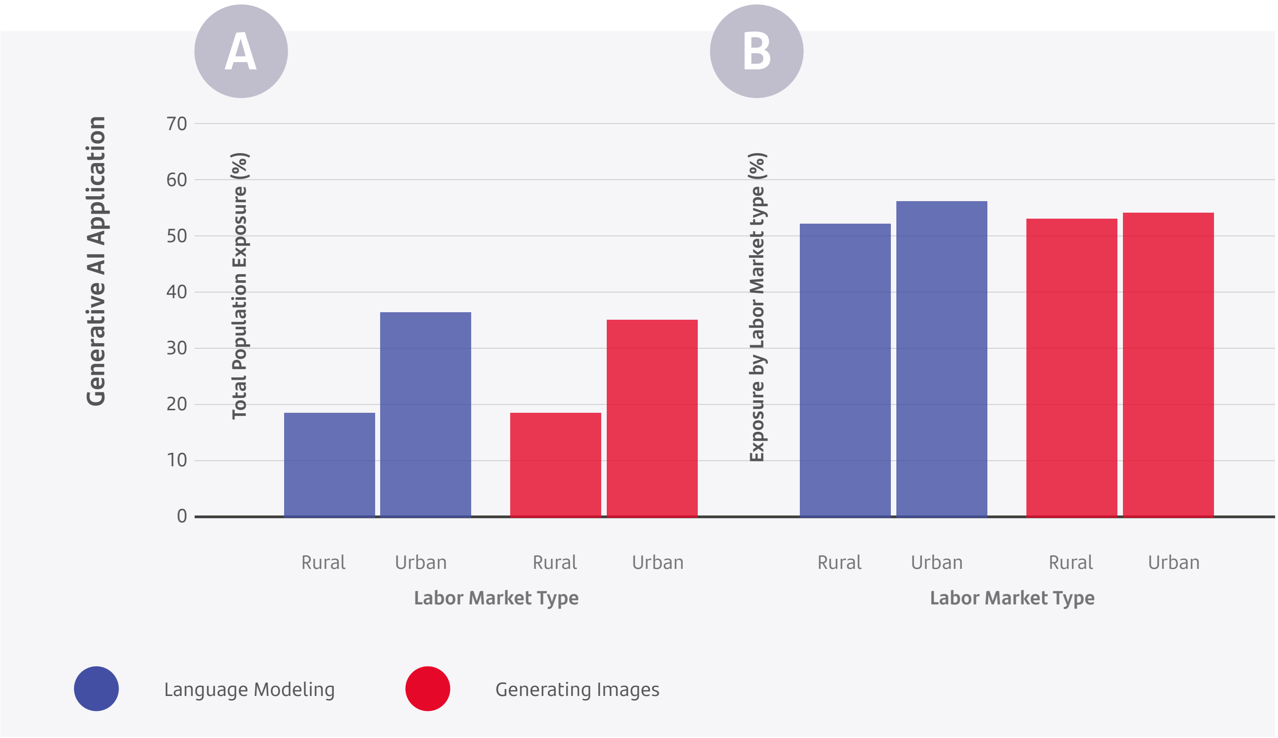

Second, we examine the exposure to generative AI by gender. Although males make up a larger share of overall exposed employment (see Figure 5A), the probability of exposure within each gender tells a different story. In Figure 5B, we see that exposure is higher for women than for men. This result is likely due to the fact that women are more likely to work in white-collar occupations, while men are more likely to work in blue-collar jobs. These results raise an interesting question regarding generative AI and equity. On one hand, generative AI may alter the ability requirements of white-collar occupations, lowering the barrier to entry for all potential workers. For example, potential software developers may need less programming knowledge with AI tools like ChatGPT. On the other hand, occupations that currently support a relatively large share of female employment are the most exposed to generative AI. Although AI exposure does not necessarily lead to employment loss, this suggests that workers in these occupations will have to adapt the most to generative AI. Disentangling these two possibilities requires further research into how occupations empirically adapt their workplace requirements.

Figure 5: generative AI exposure by gender across European countries.
(A) Across European countries, females make up a smaller share of employment that is exposed to AI language modeling or AI image generation. (B) However, given a worker's gender, females have greater exposure.

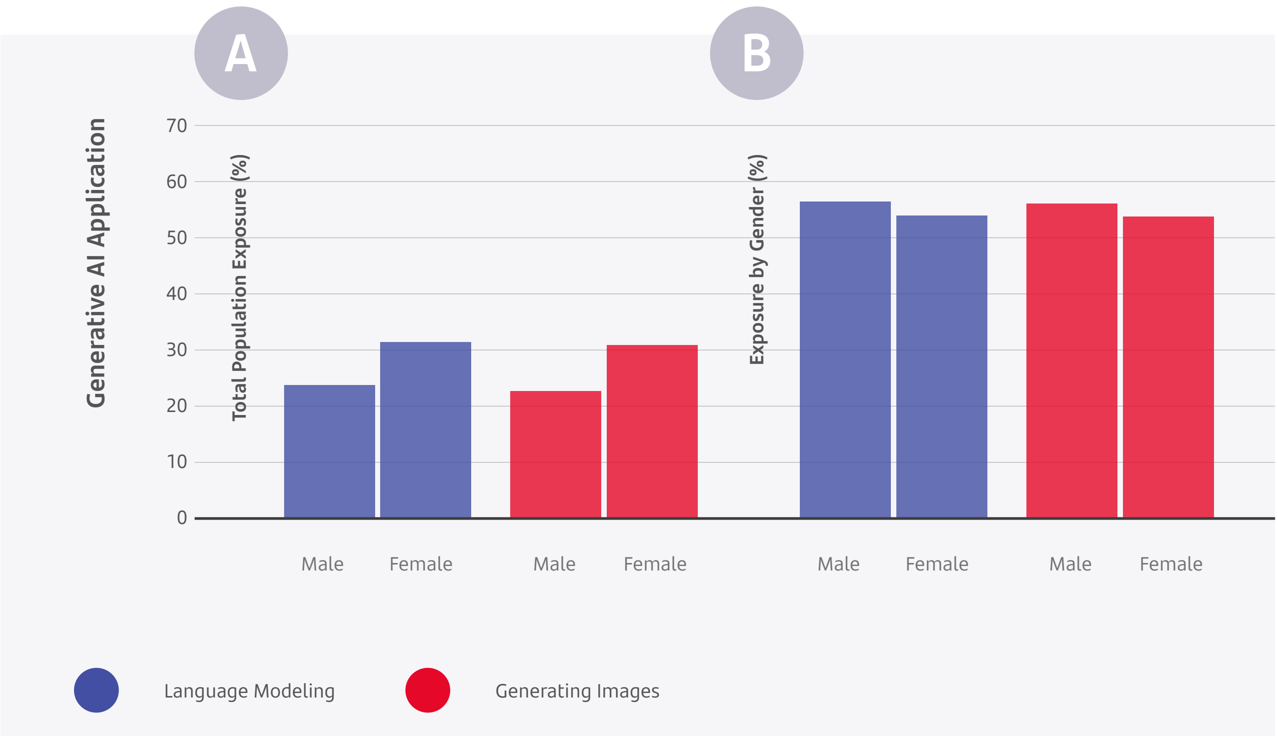

05

Substitution vs. complement

When generative AI automates a workplace activity, it typically only affects a small part of an occupation because occupations typically require many different abilities[6,7]. If similar abilities are bundled together across occupations and the technology alters the demand for those abilities, a larger impact may result[8]. Testing for correlated exposure within occupations helps determine the impact on workers. A negative correlation suggests that workers can adapt easily by shifting emphasis to the remaining workplace requirements while a positive correlation suggests that large portions of ability bundles of occupations may change, resulting in more dramatic implications for workers.

For example, consider the role of a customer service representative. This occupation requires a variety of abilities and activities such as communication, problem-solving, and empathy. When generative AI is used to automate certain tasks, such as data entry or basic inquiries, it affects only a small portion of the representative's overall job. However, if the use of AI leads to a larger shift in the way customer service is handled, such as the implementation of chatbots to handle a significant portion of inquiries, then the representative's bundle of abilities is more likely to change. This could lead to a positive correlation between ability bundles and the likelihood of being replaced by AI, suggesting more dramatic impacts on the workplace and its workers.

In Figure 6, we explore how abilities are grouped together across different occupations. For example, an occupation that requires the *Visualization* ability often also requires the *Written Expression* ability and so we draw a connection line between *Visualization* and *Written Expression.* In this way, we look at each pair of workplace abilities and construct a network that describes the bundling of abilities across all occupations. The network reveals a densely linked group of abilities that includes *Visualization* and *Written Expression.* These abilities are most common in cognitive white-collar occupations, including *Managing Directors and Chief Executives* and *Finance Managers.*

Ability bundles help contextualize the impact of AI on occupations by quantifying correlated AI exposure. For example, when both *Visualization* and *Written Expression* are exposed to AI, *Managing Directors and Chief Executives* and *Finance Managers* will need to adapt more to generative AI than in the scenario where only one ability is exposed. We quantify correlated exposure by comparing ability exposure scores, which are weighted by the probability that ability pairs are bundled together across occupations[ii]. Ability exposure to AI language modeling is more highly correlated (0.16) than exposure to AI image generation (0.05). Therefore, workers in occupations requiring abilities exposed to AI image generation may need

ii. Using a weighted Pearson correlation over network links.

to adapt only slightly to generative AI tools. However, workers in occupations exposed to AI language modeling tend to require several workplace abilities that are exposed to generative AI tools, suggesting that workers may need to undertake more drastic adaptations. For example, the ability group containing *Visualization* and *Written Expression* in Figure 6A shows a shade of blue that is darker than the red depicted in Figure 6B. This suggests that this ability bundle within white-collar occupations has similar exposure to AI language modeling, while only a few have high exposure to AI image generation.

Figure 6:

Mapping how workplace abilities are bundled together across occupations.

(A) Abilities are colored by their exposure to AI language modeling. Some example abilities are highlighted for context. (B) Abilities are colored by their exposure to AI generated images.

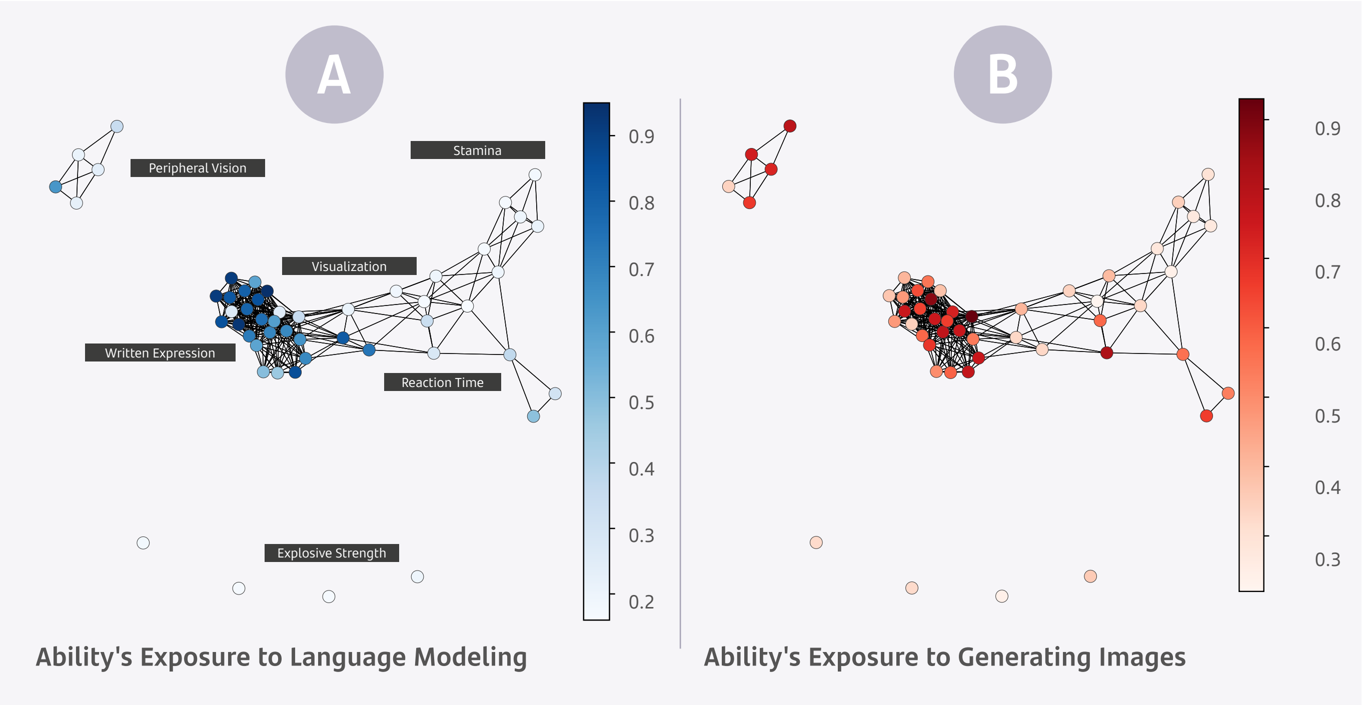

Although technology alters the demand for workplace skills and abilities, workers typically identify with their occupation. However, describing occupations as bundles of required workplace abilities yields critical insights into career mobility and the extent of adaptation required after a technological disruption. In Figure 7, we measure the similarity of abilities required by each pair of occupations[iii] and connect similar occupations with a link. Doing this for every pair of occupations, we produce an occupation network. Similar to Figure 6, we color occupations based on their exposure to AI language modeling or AI image generation and quantify their exposure homophily. White-collar occupations, including *Financial Analysts* and *Accounting Clerks,* are more exposed to generative AI than blue-collar occupations, including *Freight Handlers* and *Building Caretakers.* Exposure to language modeling is more correlated (0.15) than exposure to image generation (0.11). This suggests that workers in occupations exposed to generative AI tools may need to make larger career moves (e.g., changing sectors) to escape this exposure if they are unable to adapt their workplace abilities. For example, *Financial Analysts* are exposed to AI language modeling, and other occupations requiring similar abilities are similarly exposed, suggesting that workers who transition away from the *Financial Analysts* occupation may need to drastically reskill if they aim to escape exposure.

Figure 7: Mapping occupations (ISCO 4-digit) based on shared ability requirements.
(A) Occupations are colored by their exposure to AI language modeling. Some example occupations are highlighted for context. (B) Occupations are colored by their exposure to AI generated images.

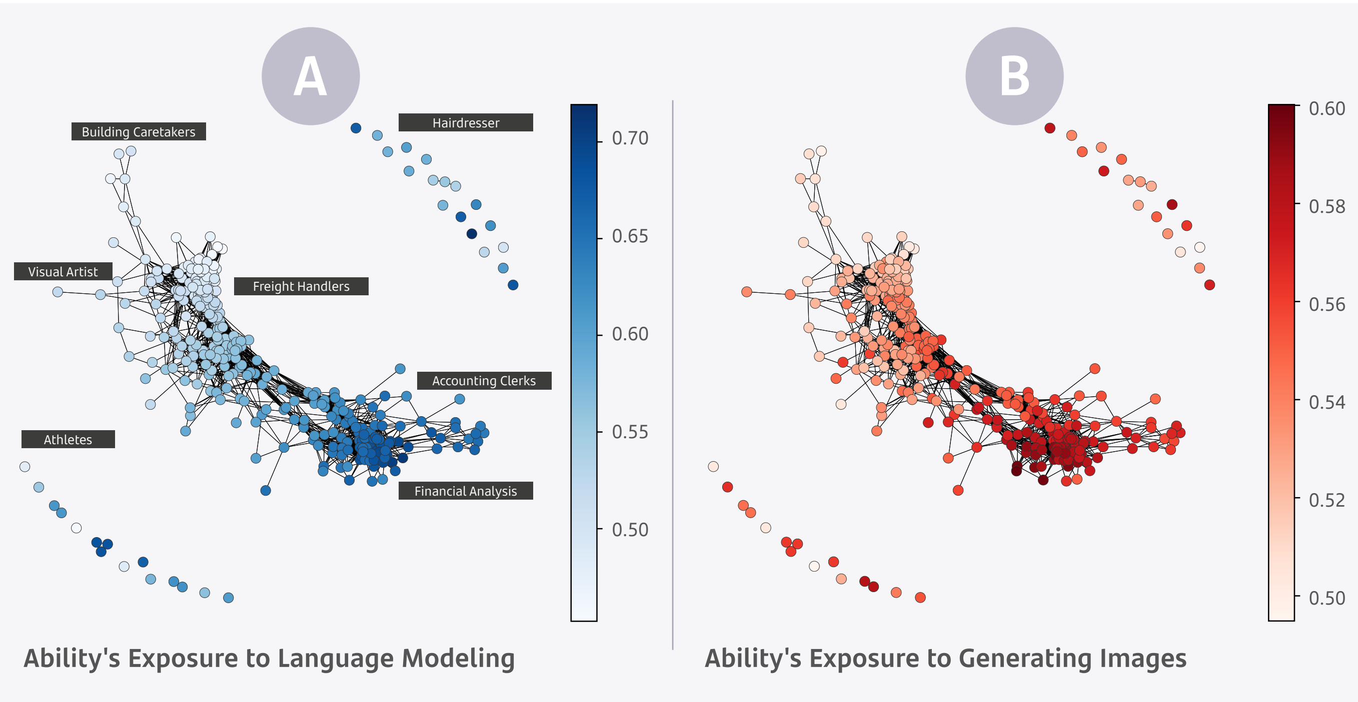

iii. Using Jaccard similarity, which measures the amount of shared ability requirements while controlling for the total amount of unique abilities required in each occupation.

Exposure to AI image generation is more correlated in the occupation space (0.11) compared to the workplace ability space (0.05). At first glance, this result may be confusing since both spaces are constructed from the same data relating occupations to their ability requirements. However, these different results arise because the ability space considers all observed bundles of abilities, whereas a single occupation represents only one realized bundle of ability requirements. Combining the two representations can therefore reveal how exposed occupations might shift their ability requirements in response to generative AI. The ability network's low exposure homophily for AI image generation demonstrates that a shift in ability requirements may occur following the connections between commonly bundled workplace abilities that are less exposed. For example, *Visualization* has relatively high exposure to AI image generation, but is often bundled with *Written Expression,* which has lower exposure; thus, workers who currently perform visual tasks may shift their efforts to written tasks in the presence of applications for AI image generation.

06

Roles, tasks, and sector implications

Traditionally, automation has had a significant impact on blue-collar jobs involving manual, repetitive tasks. However, as we have seen, generative AI has the potential to disrupt white-collar jobs that involve cognitive, analytical, and creative tasks. In the U.S., for example, the *Finance and Insurance* sector has the greatest exposure to AI language modeling when combining occupation-level exposure with sector employment distributions from the U.S. Bureau of Labor Statistics (see Table 1 in page 10). This differs from previous AI exposure studies. Generative "creativity" of generative AI is one of its most striking advancements over previous language models; all of these domains were considered safe from automation because they require creative ability. For example, a major Oxford study on the likelihood of computerization of occupations[9] asserted that:

> **"Because creativity, by definition, involves not only novelty but value, and because values are highly variable, it follows that many arguments about creativity are rooted in disagreements about value… In the absence of engineering solutions to overcome this problem, it seems unlikely that occupations requiring a high degree of creative intelligence will be automated in the next decades."**

As generative AI disrupts creative work (e.g., graphic design, musical composition, programming, and writing), we need to identify the creative workers whose work will adapt moving forward despite long periods of stable employment and wages.

Years of research on the impact of AI on our working places show that, rather than wiping out entire categories of jobs, automation can unlock human potential to do tasks differently and take on other higher-value tasks. A typical occupation consists of 20 to 30 distinct tasks, some of which are much easier for to automate through AI than others. Generative AI applications directly impact the need for human workers to perform certain tasks, which can lead to workers adapting their workplace activities in response to new technology. Workers who fail to adapt may experience job separations (i.e., from quitting or firing) or, more rarely, become unemployed if they fail to secure a new employment thereafter. Within an occupation, however, changes in ability demands and workers' skill adaption are the most likely outcome. For example, computer vision algorithms are sometimes better than humans at interpreting medical images. But the number of radiologists hired has increased in recent years because image interpretation is just one of the many job tasks performed by a typical radiologist. As another example, although the U.S. *Finance and Insurance* sector has high exposure to AI language modeling today (see Table 1), employment in occupations impacted by AI language modeling in this sector was expected to decline from 2020 to 2030 according to employment projections made by the U.S. Bureau of Labor Statistics in 2020 (see Table 3), two years before the release of ChatGPT.

Table 4: Expected change in exposure to AI language modeling and AI image generation US industries (two-digit NAICS codes) from 2020 to 2030 using employment statistics and projections from the US Bureau of Labor Statistics.

Industry	2020-2030 Language Modeling Exposure (%)	Change in Generating Images Exposure (%)
Construction	0.706	0.768
Utilities	0.217	0.387
Mining, quarrying and oil and gas extraction	0.143	0.410
Federal government	0.113	-0.085
Accommodation and food services	0.100	0.271
Government	-0.188	-0.103
Real estate and rental and leasing	-0.201	-0.022
Wholesale trade	-0.359	-0.160
Arts, entertainment, and recreation	-0.390	-0.422
Finance and insurance	-0.665	-0.759
Agriculture, forestry, fishing and hunting	-0.766	-0.122
Others services (except public administration)	-0.971	-0.595
Management of companies and enterprises	-1.022	-0.955
Administrative and support and waste management...	-1.210	-1.031
Educational services; state, local and private	-1.300	-1.263
Information	-1.879	-1.732
Professional, scientific, and technical services	-2.139	-2.157
Healthcare and social assistance	-3.018	-3.028

Unleashing the use of generative AI will offer workers and companies more creative ways to foster human-machine collaboration in those occupations and sectors with more exposure. A recent study by developers' collaborative platform Github found that developers using Copilot (an AI tool developed by GitHub and OpenAI to assist software developers using GPTs) were 11% more productive and 55% faster than developers not using Copilot[10]. The effect is so pronounced that the ban on ChatGPT in Italy early this year reduced the output of Italian GitHub users by around 50% in the first two business days after the initiation ban[11]. Another study found that the use of AI-based conversational assistants in customer support increased productivity, measured by issues resolved per hour, by an average of 14%[12]. In general, these studies align with predictions that generative AI will increase global productivity by 1.5% over a 10-year period[13]. The key idea is that these technologies will break down entry communication barriers between humans and machines, streamline business workflows, and automate routine creative tasks. Some analysis suggests that with access to LLM, around 15% of all work tasks in the U.S. could be completed significantly faster with the same level of quality[14].

These productivity gains are not equally shared by all employees. Preliminary research has shown that novice, low-skilled workers benefited the most from incorporating generative AI tools into their jobs. Newer workers can move up the experience curve more quickly. Also, low-skill workers seem more likely to adopt these tools than high-skill workers, as the latter may have less to gain from AI assistance precisely because AI recommendations capture their implicit job experience. Generative AI appears to level up years of experience across workers in the same company. This differs markedly from earlier waves of computer technology, in which low-skill workers did not benefit immediately from advances.

Because of their unique combination of creativity and high value proposition, professional services (e.g., legal), information, enterprise software, healthcare, and financial services industries are more exposed to the impact of generative AI. Jobs like telemarketers, proofreaders, bookkeepers, accountants, computer programmers, mathematicians, or web developers have tasks that are more likely to be automated by generative AI than more physical jobs like surgeons or construction workers. For this reason, large organizations, that are more likely to employ high-skill workers, will be more impacted by the generative AI revolution. According to some studies, the companies with the highest exposure to ChatGPT are IBM, Intuit, Qualcomm, NVIDIA, S&P Global Inc, and Microsoft[15]. With some major software companies poised to roll out their own generative AI tools to improve their products, business, and workforce, and with the lowering of entry costs that generative AI gives to new competitors, big companies in these sectors will have to act quickly to adopt generative AI. Big companies also have some advantages over small competitors. Rather than using generally trained tools like ChatGPT, they can train their own LLM models on their data, documents, and processes to develop better generative AI suited for their own purposes. For example, Bloomberg has developed BloombergGPT[16], an LLM specifically trained on a wide range of financial data to support a variety of NLP (Natural Language Processing) tasks within the financial industry. Over the next five years, we will see more companies adopt this framework to encode their large experience into generative AI tools without sacrificing the generality of these tools.

Despite the high exposure of jobs to generative AI and the uncertainty surrounding its potential, the sense of opportunity of generative AI tools for different businesses and sectors is reflected in the markets and forecasts. In general, following the release of ChatGPT, firms with higher exposure generated excess returns that were 0.4% higher on a daily basis than those of firms with lower exposure[17]. The financial sector, however, was a counterexample. Despite previous impressive improvements in AI and other technologies, productivity growth has actually slowed in recent years, from an average of over 2.4% per year between 1995 and 2005 to less than 1.3% per year since. The bottleneck is not the development of new technologies, but the lack of retraining of the workforce or business dynamism. Similar to past technological changes, the mere adoption of generative AI without making changes to the business organization, employee onboarding, reskilling, and merit evaluation will hinder the realization of the full potential of these tools in our workplace.

Undoubtedly, labor markets are beginning to experience this mix of optimism and caution. Job seekers and job descriptions now list ChatGPT as a qualification, especially for software developers and engineers, and a large fraction of companies are currently using ChatGPT even for the hiring process[18]. The rapid and widespread adoption of these technologies has suddenly created human-machine imbalances in our workplaces that may soon disappear. A very recent survey of 1,000 business leaders found that 1 in 4 companies in the U.S. have replaced workers with ChatGPT since November, which could lead to more layoffs. IBM is freezing hiring for roles that could be replaced with AI in the coming years, and tech giants are not just cutting jobs, but making them extinct because of the adoption of generative AI. However, workers who have some experience with these technologies are more likely to keep their jobs, showing the important role of reskilling, retraining, and job redefinition policies within companies and across sectors.

Generative AI has the potential to democratize access to today's white-collar jobs. As technology removes the need for workers to perform certain activities, the barriers to entry into an occupation may decrease. An occupation that requires fewer different workplace activities than before may become suitable for more job seekers with less reskilling. However, this outcome depends on how difficult is to acquire the skills appropriate for the tasks that are left over after automation. In addition, the tasks that remain will determine the value and quality of an occupation moving forward. History shows that new technologies often lead to increases in overall employment, allowing low-skilled workers to perform jobs previously reserved for skilled artisans. For example, manufacturing employment increased during the Industrial Revolution, as traditionally handcrafted goods were mass-produced in factories using steam engines and low-skilled labor. Despite the fact that net employment rose, the new jobs were of questionable quality: although these manufacturing workers outnumbered their artisan predecessors, they earned lower wages and worked in less desirable factory environments. The barrier to entry into the production of artisanal goods such as ceramics, textiles, and steel fell during the Industrial Revolution, when factory workers with little education or training replaced artisans with years of apprenticeship and experience. Generative AI could create similar shifts in labor markets by displacing temporally creative experienced workers by generative AI tools in many industries. However, in most labor markets, reskilling and retraining tools could facilitate the transition of these experienced workers into other occupations or tasks.

In focus

Policy implications

The localization of the potential impact of generative AI in specific sectors, jobs, and demographic groups, requires a holistic approach to realize the full potential of these tools while mitigating potential problems for our society and economy. Historically, automation technologies have been associated with increased economic inequality and labor disruption, mainly because the workers most affected were in low-skilled, low-wage occupations. The generative AI impact is different, primarily affecting white-collar workers in high-paying jobs. However, these jobs are in high-value, job-intensive industries such as finance, education, and professional services. Major disruptions associated with job losses or potential gains to the economy from the adoption of these new technologies in these industries can rapidly spread throughout the labor market. Therefore, as in previous labor market shocks, we must focus on easing the barriers in our economy to the successful adoption of these technologies in our classrooms and workplaces:

Implementing strategies to facilitate upskilling and reskilling in generative AI, especially for white-collar workers. Companies, governments, and educational institutions must embrace this new technology as a new tool. Investing in programs that train the workforce in generative AI will benefit both firms and workers. Early data show that workers who already have some experience with these technologies are more likely to keep their jobs. And also, that new workers are the ones more likely to embrace, use, and benefit from these technologies.

Easing barriers to job mobility. Generative AI will make some sectors to grow quickly while others may shrink. We have already seen some Big Tech companies take a hiring pause for certain roles. Upskilling and reskilling in other areas not related to generative AI could lower these barriers, as could changes in regulations or licensing that limit access to different occupations.

Reform of safety net programs. Generative AI labor market shocks might affect specific demographic groups, such as mid-career and highly experienced professionals. As we have seen, more women than men are exposed to AI automation. Since reskilling in generative AI could lead to long separations from the labor market for these groups, companies and governments may need to invest in safety net programs to help workers and their families find new employment opportunities.

Redefine education. We should redefine our education programs to promote the kinds of skills that machines cannot match, such as interpersonal skills. But new programs should also provide training in AI use. Educators and students should also explore ways to use AI in learning, such as using AI tools to improve teaching and learning outcomes.

07

Opportunities and future directions

To wrap up, we propose a roadmap indicating the areas of opportunity that arise from a combination of task substitution and complementarity. Here are two possible (non-excludable) avenues to explore:

a. Where and how does generative AI hold the greatest potential for creating new jobs and industries? | Vertical perspective

eCommerce, Marketing: creative AI is going to revolutionize the customer experience. The use of these tools will provide new ways to hyper personalize offers, products or services for customers. And it will do so in an interactive way, using natural language to iteratively search for a product and get recommendations, services, and content that is specifically tailored to our personality or interests. This could lead to new careers in AI-driven marketing strategy and data-driven advertising.

Cultural industries: the cultural industries will be the ones more affected by this hyper personalization. Amazon is already flooded with books generated by ChatGPT and Midjourney, and given recent advances in text-to-video tools, some directors anticipate that full-feature AI movies will be possible in two years[19]. But it could be also possible to evolve and transform storytelling by asking your streaming platform to curate a story (or create a new one) specifically for you and render a movie based on it. Similarly, music could be created on the fly or newspapers articles customised to your personality. These changes will not only change creative industries, but also create new avenues and new jobs like AI art consultants, AI music producers, etc., much like the technical and software revolutions in the video game industry.

Healthcare: AI is already being used in healthcare to recognize patterns and make predictions for disease diagnosis and prognosis or drug discovery. Because generative AI is better able to capture complex and contextual representations, it can be used to accelerate this process by exploring a broader space of drug candidate or personalizing diagnosis for specific patients. In addition, generative AI models can integrate many data sources, opening up new opportunities for biomedical applications. We can build bio-LLMs by augmenting them with our knowledge about genes, diseases, and drugs to develop more tailored tools for drug discovery. Or we can use data from millions of patients to answer questions about specific diseases and potential treatments, as well as patients likely to respond to those treatments. Generative AI models can learn from large datasets and provide insights tailored to individual patients. This could lead to more personalized treatment plans and improve patient outcomes. Or due to its ability to recognize patterns in different contexts in high-dimensional data, they could also help identify design potential drug candidates for rare disease treatments, where very little is known about their chemical structure.

Education: some of the jobs more exposed to generative AI are in the education sector. Generative AI is capable of answering math questions, writing essays, and even passing standardized university or professional tests. Despite initial concerns about the impact of these tools on testing what students have learned, generative AI holds immense potential to revolutionize the education sector by personalizing and enhancing learning experiences. It can generate dynamic, tailored educational content and assessments, act as AI tutor, recommend personalized learning paths, support social-emotional learning, and address special educational needs. However, many educators are overworked and under-resourced, which could hinder their ability to effectively leverage these new technologies. Careful consideration and planning is also needed to ensure that AI is properly integrated into the classroom. Some educators argue that the advent of AI tools like ChatGPT should prompt the revaluation of current pedagogical practices. Finally, generative AI can help make educational resources accessible to more and diverse people who have limited access to educational institutions. Who wouldn't be interested in learning from a potential tutor like MIT-GPT, an LLM trained in all courses at MIT? Over the next years, we expect to see steady but important changes in the way education is tailored and redefined for existing and new people.

Financial and legal services: the ability of generative AI to analyze large datasets, documents and enable human-like interactions can transform the way legal and financial business operate. Generative AI tools can also automate routine tasks, such as preparing legal and financial reports, handling customer inquiries, providing legal and financial advice, assisting with governance and regulatory compliance, etc. Some large financial companies are already developing sophisticated automated investment strategy advisors for their clients based on ChatGPT[20]. As we have seen, workers in financial and legal services are not only more exposed to generative AI, but in many countries and demographic groups they employ a large share of workers. As a result, these sectors are probably the ones that will witness more adaptation and transformation of their business in the coming years.

b. Human-AI collaboration | Horizontal perspective

Human-AI collaboration represents a promising opportunity to increase productivity and innovation, especially when the comparative strengths of humans and generative AI are effectively leveraged. Research in human-AI collaboration demostrates the complementarity and benefits of using AI tools in the workplace. In general, AI can sometimes be better than humans, but nothing beats a human using AI tools productively. For example, recent studies on the use of ChatGPT in customer support show that human agents can increase their productivity by 15% with ChatGPT. Companies that leverage human-machine collaboration have the best chance of success. Companies that already achieve significant financial benefits from AI and those that are developing systems where AI learns from humans and vice versa foster a symbiotic relationship that drives innovation . Humans play a crucial role in training AI systems, explaining their outputs, and ensuring their responsible use. Conversely, AI can assist humans in information gathering, data processing, and routine customer service, freeing them for higher-level tasks that require leadership, judgment, and other human skills. Optimizing this collaboration requires a redesign of business processes and a focus on developing employees who can work effectively at the human-AI interface.

This time, however, generative AI is approaching tasks traditionally reserved for humans, such as creative thinking. The ability to use ChatGPT or Midjourney to generate ideas quickly and efficiently can lead to a new form of interaction between humans and AI, where AI can help find more creative answers or even new questions in our organizations. Just as Google helps us find new information on the internet, generative AI could help us find new ideas.

On the other hand, while previous advances in AI require highly skilled knowledge to leverage their potential impact, the generality of generative AI tools like ChatGPT makes them accessible to a large fraction of people in an organization. This translates into a new kind of human-AI collaboration, in which entire departments, people within a project, or across the organization can simultaneously use these tools to find better answers, even if our organization has not officially implemented and adopted them.

Roadmap for future action

For businesses and organizations:

businesses should strive to incorporate AI into their daily processes and explore innovative ways to improve productivity and drive growth. This means investing in AI technology and training employees to use AI tools effectively. Generative AI should be used as a tool to stimulate creative thinking across all departments within an organization, leading to a new form of interaction between humans and AI.

For Employees and Individuals:

individuals should aim to upskill themselves and learn how to effectively use new AI technologies. They should be open to collaborating with AI in their work and recognize the potential benefits that can result from such collaboration. They should also be prepared for the changes AI will bring to some job tasks and be willing to adapt to those changes.

For policymakers and regulators:

policymakers should aim to create a regulatory environment that encourages the use of AI while protecting individuals and businesses. They should address issues such as data privacy, AI ethics, and the impact of AI on employment. They should also considerhow to promote AI education and training to ensure that individuals and businesses are able to take advantage of these new technologies.

For educational institutions:

educational institutions should strive to integrate AI into their curricula to prepare students for a future in which AI will be a fundamental part of many occupations. They should provide training on the use of AI and encourage students to think critically about the potential impact of AI on society. They should also look for opportunities to use AI in their own processes, such as using AI tools to improve teaching and learning outcomes.

a1

Annex 1 | Glossary

- **BARD:** artificial intelligence chatbot similar to ChatGPT.
- **DALL-E:** a variant of GPT-3 developed by OpenAI that generates images from text descriptions. DALL-E is capable of generating unique—and often surprising—images from even unusual or abstract prompts.
- **FINE-TUNING:** a process in machine learning in which a pre-trained model (a model trained using a large amount of data) is adapted or "fine-tuned" for a specific task.
- **GENERAL PURPOSE AI SYSTEMS:** these are AI systems that are not designed for a specific task but can be used for a variety of tasks. They are versatile and adaptable. Models like GPT-3 and ChatGPT fall into this category because they can generate text for a wide range of prompts, making them useful for a variety of applications.
- **NATURAL LANGUAGE PROCESSING (NLP):** an area of AI that focuses on the interaction between computers and humans through natural language. The ultimate objective of NLP is to read, decipher, understand, and make sense of human language in a valuable way.
- **MIDJOURNEY:** generative artificial intelligence program and service created and hosted by Midjourney Inc. that generates images from natural language descriptions called "prompts", similar to OpenAi's DALL-E.
- **PROMPT:** in the context of generative AI, the input given to the model that it uses as the basis for generating output.
- **PROMPT ENGINEERING:** this is the process of crafting effective prompts to obtain desired outputs from AI models. Prompt engineering is a crucial part of fine-tuning AI models like GPT-3 and requires a deep understanding of how the model responds to different inputs.
- **STABLE DIFFUSION:** deep learning, text-to-image model released in 2022. It was developed by the start-up Stability AI in collaboration with a number of academic researchers and non-profit organizations.

References

1. Goldman Sachs. (2023). ***Generative AI could raise global GDP by 7%.*** https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html

2. NVIDIA. (2023). ***NVIDIA: Reduce the cost of CPU-training an LLM from $10 million to just $400,000 USD by buying our GPUs***. https://wccftech.com/nvidia-reduce-cost-ai-training-llm-400000-usd-10-million/

3. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). ***GPTs are GPTs: An early look at the labor market impact potential of large language models.*** **arXiv preprint.** https://doi.org/10.48550/arXiv.2303.10130

4. Goldman Sachs. (2023). ***Generative AI could raise global GDP by 7%.*** https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html

5. Felten, E. W., Raj, M. & Seamans, R. (2023). **Occupational Heterogeneity in Exposure to Generative AI. SSRN.** http://dx.doi.org/10.2139/ssrn.4414065

6. Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., & Wang, D. (2019). **Toward understanding the impact of artificial intelligence on labor.** *Proceedings of the National Academy of Sciences, 116(14)*, 6531-6539. https://doi.org/10.1073/pnas.1900949116

7. Acemoglu, D., & Autor, D. (2011). **Skills, tasks and technologies: Implications for employment and earnings. In D. Card & O. Ashenfelter (Eds.),** Handbook of Labor Economics (Vol. 4, Part B, pp. 1043-1171). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

8. Alabdulkareem, A., Frank, M. R., Sun, L., AlShebli, B., Hidalgo, C., & Rahwan, I. (2018). **Unpacking the polarization of workplace skills.** *Science Advances,* 4(7), eaao6030. https://doi.org/10.1126/sciadv.aao6030

9. Frey, C. B., & Osborne, M. A. (2017). **The future of employment: How susceptible are jobs to computerisation?** *Technological Forecasting and Social Change,* 114, 254-280. https://doi.org/10.1016/j.techfore.2016.08.019

10. Kalliamvakou, E. (2022, September 7). **Research: Quantifying GitHub Copilot's impact on developer productivity and happiness.** *The GitHub Blog.* Retrieved from https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/

References

11. Kreitmeir, D. H., & Raschky, P. A. (2023). ***The unintended consequences of censoring digital technology: Evidence from Italy's ChatGPT ban.*** **arXiv preprint.** arXiv:2304.09339. https://doi.org/10.48550/arXiv.2304.09339

12. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). **Generative AI at work (No. w31161).** National Bureau of Economic Research. https://doi.org/10.48550/arXiv.2304.11771

13. Goldman Sachs. (2023). ***Generative AI could raise global GDP by 7%.*** https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html

14. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). ***GPTs are GPTs: An early look at the labor market impact potential of large language models.*** **arXiv preprint.** https://doi.org/10.48550/arXiv.2303.10130

15. Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). **Generative AI and firm values.** SSRN. https://doi.org/10.2139/ssrn.4436627

16. Wu, S., Irsoy, O., Lu, S., Dabravolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). ***BloombergGPT: A large language model for finance.*** **arXiv preprint.** arXiv:2303.17564. https://doi.org/10.48550/arXiv.2303.17564

17. Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). **Generative AI and firm values.** *SSRN.* https://doi.org/10.2139/ssrn.4436627

18. Resume Builder. (2023). ***1 in 4 companies have already replaced workers with ChatGPT.*** https://www.resumebuilder.com/1-in-4-companies-have-already-replaced-workers-with-chatgpt/

19. Variety. (2023). **'Avengers' Director Joe Russo predicts AI could be making movies in 'Two Years': It will 'engineer and change storytelling'.** https://variety.com/2023/film/news/joe-russo-artificial-intelligence-create-movies-two-years-1235593319/

20. CNBC News. (2023). **JPMorgan is developing a ChatGPT-like A.I. service that gives investment advice**. https://www.cnbc.com/2023/05/25/jpmorgan-develops-ai-investment-advisor.html

21. Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chu, M. & LaFountain, B. (October 2020). **Expanding AI's Impact With Organizational Learning.** *MIT Sloan Management Review and Boston Consulting Group* https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/

esade

Santander X Innovation Xperts

www.santander.com/santander-x-innovation-xperts-en

By Santander