Skill Demands in the Age of AI: Evidence from Continental Europe

Najada Feimi (LISER)
Christina Gathmann (LISER, University of Luxembourg, and CEPR)
Terry Gregory (LISER, ZEW, and IZA)
David Marguerit (LISER, and University of Luxembourg)

December 2, 2024

Abstract

Despite rapid developments in AI capabilities, our understanding how AI affects the demand for skills and jobs remains severely limited. We provide a comprehensive analysis of how AI has shifted the demand for skills in Continental Europe over the past years. Based on a comprehensive dataset of 31 million online job vacancies from Belgium, France, Germany, and Luxembourg, we use advanced ML models to extract and classify skills from multi-lingual job ads from 2018 to 2023. We document a significant growth in the demand for AI and data skills, a decline in the need for prediction skills and a stable demand for judgement skills. Linking the observed skills to measures of AI exposure, we find a strong positive relationship of AI with skills in AI and data management, in line with the ongoing adoption of AI in companies. AI also increases the demand for social skills, which points to a rising need to integrate AI and human-machine interactions into the production process.

Keywords: AI, automation, skill demand, decision-making, online job vacancies
JEL codes: O33, J23, J24

1 Introduction

Artificial intelligence (AI) has vastly expanded its capabilities, and the pace of change is even expected to increase in the near future (Korinek, 2024). AI has significantly lowered the costs of prediction, generating a surge in the demand for prediction tasks in the economy (Agrawal et al., 2018). More recently, GenAI has moved beyond that, being able to evaluate complex texts or even create new content. These rapid developments have spurred fears that many skills could become obsolete, leading to unemployment and rising inequality. A more optimistic view sees AI as a complement to human labor with a sizable potential to boost productivity and create new jobs.

We still have a limited understanding how AI has been transforming the demand for skills in the labor market. One reason is that the capabilities of AI have been expanded rapidly in different domains. Many tasks involving tacit knowledge that were traditionally reserved for humans are now within reach of algorithms. A second reason for our limited understanding originates in the fact that many firms have only started to adopt and implement some AI into their production processes. A third reason is that the impact of AI on skill demand and jobs is a-priori unclear and depends on whether AI is mostly used for automation or also creates sizable productivity gains and even entirely new skills or jobs. Hence, the impact of AI will crucially depend on how businesses decide to make use of it; whether they have the technical and managerial skills to make use of its many opportunities; and the rules guiding the diffusion and use of AI in the economy. As such, we expect that AI's impact will vary a lot across industries and countries (Agrawal et al., 2019).

This paper investigates how AI has changed the demand for skills in European firms over the past years. In particular, we study the following questions: How can we identify the demand for AI-related skills in the data? Which job skills are substitutes, and which complements to AI? How does the total demand for certain skills evolve with AI; and how does it evolve across industries, occupations and countries? To analyze these questions, we make use of Lightcast data on online job vacancies for several European countries. Our database contains more than 100 millions vacancies collected between 2018 and 2023, which were scraped from over 300 platforms across Belgium, France, Germany, and Luxembourg. After rigorous cleaning to remove duplicates, irrelevant data, and overly concise or incomplete job descriptions, we retain 31 million job ads for analysis. Skills are extracted using Mixtral 8x7B, a multilingual large language model (LLM) optimized for high-performance skill identification in English, French, German, and Dutch. Skill extraction follows a three-step process: identifying skills within job descriptions, translating them into English, and consolidating semantically similar skills through cosine similarity-based grouping. This process results in a clean set of 5.8 million unique skills. To evaluate Mixtral's performance, we compare its outputs against a human-labeled

benchmark, achieving an average F1-score of 0.78, which underscores its strong ability to extract relevant skills with high precision and recall.

We classify the 5.8 million unique job skills using a supervised machine learning approach into several predefined categories: AI skills, data skills, prediction skills, judgement skills, Decision-making, and social skills. AI skills like machine learning and image processing emphasize technical expertise, while Data skills such as data science focus on managing and interpreting datasets. Decision-making and Judgment highlight project management and customer orientation, reflecting strategic and interpersonal priorities. Prediction skills, including forecasting and cost management, show demand for analytical foresight, and Social skills like teamwork and customer service underscore the importance of collaboration. We use a Multi-layer Perceptron (MLP) model trained on a curated set of labeled skills. We train the model using 114 hand-labeled skills, with a test-train split ensuring robust evaluation. The classifier achieves an F1-score of 0.8, indicating strong predictive performance, with features such as frontier technology keywords and AI-related terms contributing most significantly to predictions. Our classification finds that 7% of skills are classified as AI-related, while 17% refer to data skills, for instance.

We have several preliminary findings. Our descriptive analysis reveals that the demand for AI-related skills has increased by 10% between 2018 and 2023 – likely reflecting growing adoption of AI technologies. Moreover, we see a 5% rise in the demand for data skills, which could indicate their role in supporting AI adoption and implementation. In contrast, we observe a decline in prediction skills by 5%, which is likely related to the fact that AI can perform routine forecasting tasks that were previously performed by humans. Interestingly, we find that the demand for judgment skills remained stable over time

We then turn to a more systematic investigation of how AI might influence the demand for different skill categories. To measure exposure to AI forces, we use AI exposure to automation and augmentation developed by Marguerit (2024a) (see also Autor et al., 2024a). AI automation measures how much AI substitutes labor in occupations, using AI-related Stack Overflow queries mapped to 2010 O*NET data. AI augmentation assesses AI's potential to enhance outputs, linking similar queries to the 2010 Census Alphabetical Index. The measures vary by occupation and year (automation) and industry and year (augmentation). We estimate how AI exposure shifts the probability of requiring a skill controlling for job characteristics, industry and time fixed effects. Our results show that exposure to automation significantly increases the demand for AI-related and data skills. Moreover, we also see an increase in judgment, social skills, and even prediction skills – indicating both complementarities and productivity gains. In turn, exposure to AI augmentation has mixed effects. Augmentation raises the demand for decision-making and social skills, but it reduces the demand for data skills.

Our paper contributes to at least three strands of the literature. Closest to our analysis are papers that use job vacancy data to learn about how skill requirements change in

response to the expanding capacities of AI. A few recent studies document an increasing demand for AI expertise in data science and machine learning in the United States and United Kingdom. They also report rising wage premiums for those skills, while the demand for other skills other than AI sometimes declines and sometimes increases in exposed firms (Alekseeva et al., 2021, Acemoglu et al., 2022, Babina et al., 2024, Samek et al., 2021). Samek et al. (2021) reports that demand for AI skills is linked to the demand for other cognitive and socio-emotional skills in AI-intensive professions. We contribute to this literature by providing a much more detailed analysis on skill categories that are potential substitutes and complements to AI based on a more sophisticated method to classify skills. We further provide novel results on the impact of AI in Contintental Europe, where AI and data protection are more stringently regulated than in the UK or US markets.

Second, we complement existing studies using job vacancies to measure the broad skills employers demand in the labor market. Deming and Kahn (2018) and Deming (2021) document rising returns to social and decision-making jobs. Ash and Hansen (2023) discuss the use of large language models (LLMs) for classifying unstructured text in economics, highlighting both their potential and limitations. Our paper contributes to this literature by using an automated method to extract and classify skills that goes beyond simple keyword search and provides a more detailed view on job skills.

We further contribute to studies focusing on the labor market impact of AI. An emerging literature has investigated the labor market effects of AI, focusing on employment and wages (e.g., Acemoglu et al., 2022, Alekseeva et al., 2021, Brynjolfsson et al., 2018, Bonfiglioli et al., 2024, Engberg et al., 2024, Felten et al., 2018, 2019, Mann and Püttmann, 2023, Webb, 2020). Most studies find no evidence for job displacement and a modest impact on wages or inequality. We contribute to this literature by providing evidence on the direct effect of how AI has shifted labor demand and skill requirements in Europe over the past years. Moreover, we can provide more detailed insights in which skills are more likely to be replaced and which seem to be complementary to AI.

The paper proceeds as follows. The next section spells out our understanding of different skill categories and how their demand could be affected by AI. Section 3 introduces the job vacancy data used in our analysis and discusses the ML-based extraction of skills. Section 4 outlines the machine learning approach to classify skills into predefined categories and highlights its accuracy and key results. Section 5 presents an in-depth analysis of trends in AI-related skill demand, examining their evolution over time and interaction with occupational exposure to AI. Finally, Section 6 discusses the implications of our findings and concludes.

2 Technological Change and Skills

2.1 Theoretical Mechanisms

How technological advances impact the labor market depends crucially on how technologies like AI shift the demand for job tasks. The task-based approach has stressed the importance of comparative advantage of human labor and the complementarities between different tasks as primary drivers of labor demand changes in response to technological change (Acemoglu and Autor, 2011, Acemoglu and Restrepo, 2019, Acemoglu et al., 2024, Autor et al., 2024a). In task models, output is produced by several complementary tasks, which can be produced by machines or workers. Technological advances like AI allows machines to take over some job tasks traditionally performed by human labor. As a result of automation, employers no longer require workers to perform those automated tasks. As a result, we should see their demand decline with the diffusion of AI.

Yet, the use of new technologies requires the use of complementary tasks to make productive use of its capacities. Earlier waves of technologies like robots have mostly replaced routine tasks that rely on codified rules and knowledge. In contrast, the demand for non-routine analytical tasks to train, monitor and improve the functioning of robots increased in response to the adoption of robots in manufacturing. The diffusion of AI allows automating prediction tasks that were previously performed by humans like inventory management, credit worthiness or screening applicants, to name a few examples. At the same time, the implementation of AI requires experts to make algorithms perform well for the specific application at hand. Such AI-related expert knowledge like programming or data science engineering are crucial for the adoption of AI in firms. Moreover, AI-based predictions require interpretation, validation and good judgement in order to capture hallucination effects, validate its output and increase its acceptance among employees, suppliers, customers and regulators. And while AI can do some tasks extremely well and fast, AI needs to be integrated into a work process that is populated by humans. Navigating and negotiating the interaction between humans and machines requires good managerial skills like transparency, communication and decision-making. Hence, the value of complementary skills, particularly data, judgment and decision-making skills should become more important (Deming, 2021). It is important to note that the adoption of AI may even see an increase in the demand for other tasks (not exposed to AI), if automation increases (total factor) productivity sufficiently to raise demand in non-exposed tasks and occupations (Acemoglu and Restrepo, 2019).

Finally, new technologies often create entirely new tasks and even new occupations. Using a long-term perspective, Lin (2011) and Autor et al. (2024a) show that the creation of new occupations plays an important role in the adjustment to technological change. Previous evidence indicates that some occupations that saw a notable increase in demand,

such as machine learning engineers or data scientists, are comparatively new (Acemoglu et al., 2022, Alekseeva et al., 2021). How labor demand for skills responds to the adoption of AI depends on whether AI is mostly used for automation or whether it is also used to enhance human capacities and creativity. The response will also depend on the speed of adoption and its successful implementation into the production process. We next explain how we measure the different types of skills whose demand might respond to AI.

2.2 Definition of Skill Categories

To capture the dynamic interplay between technological capacities and human skills, we define several different types of skill categories. Some skills might become automated with the diffusion of AI, others are required to adopt and implement AI, and yet, others might emerge as new skills in response to AI.

Table 1 defines several types of skills: AI and data skills, prediction skills, judgment skills, decision-making skills and social skills. We describe each of them in turn.

AI Skills: These skills involve the design, implementation, and operation of AI technologies. AI skills focus on leveraging advances in machine learning and other AI domains to automate processes and enhance decision-making systems. Examples include developing algorithms, training models, and deploying AI systems. These skills are crucial to adopt and implement prediction tasks Agrawal et al. (2018) and are likely to see growing demand with continued innovation in AI.

Data Skills: Data-related skills ensure that AI systems operate effectively by providing high-quality inputs. These skills include data collection, preparation, cleaning, labeling, and organization. The paper underscores the importance of prediction improvements arising from better data availability, making data skills crucial in supporting AI's predictive capabilities. As AI adoption grows, the demand for data skills is also expected to increase, given the reliance of AI on large, clean, and well-organized datasets.

Prediction Skills: Prediction skills focus on forecasting and classifying outcomes, tasks that AI systems increasingly automate. While AI excels in making predictions, the human ability to contextualize and interpret predictions remains vital in many settings. Examples include market analysis and trend forecasting. As noted by Agrawal et al. (2018), the improvements in AI prediction capabilities could reduce the demand for stand-alone human prediction skills while shifting focus toward integrating these predictions into broader decision-making processes.

Judgement Skills: Skills required to assess, interpret, and evaluate options, situations, or predictions, often under uncertainty, to make informed evaluations or recommendations. Agrawal et al. (2018) highlight judgment as a uniquely human skill critical for determining payoffs and guiding decisions in complex, uncertain environments. As AI

Table 1: Definitions of key skill types

Skill Type	**Definition**	**Example Skills**
AI	Skills related to the development, implementation, and use of artificial intelligence technologies.	Machine learning algorithm development, natural language processing, AI model training, AI system deployment, advanced data analytics.
Data	Skills associated with providing, preparing, and managing data necessary for the operation of AI systems.	Data preparation, data labeling, data collection, data cleaning, data organization.
Prediction	Skills that involve human-driven processes to forecast, classify, or predict outcomes, even when these tasks could be automated by AI.	Trend forecasting, market analysis, sales prediction, statistical modeling, risk assessment.
Judgment	Skills required to assess, interpret, and evaluate options, situations, or predictions, often under uncertainty, to make informed evaluations or recommendations.	Critical thinking, ethical reasoning, prioritization, assessing trade-offs, interpreting probabilistic outcomes, weighing risks and benefits.
Decision-Making	Skills involved in selecting the best course of action from alternatives based on evaluations, predictions, and goals.	Strategic planning, problem-solving, crisis management, resource allocation, operational decision-making, scenario analysis.
Social	Skills centered on interpersonal interaction, collaboration, and the ability to navigate social contexts effectively, which enhance teamwork and productivity in the workplace.	Teamwork, negotiation, persuasion, social perceptiveness, active listening, conflict resolution, leadership in group settings.

systems improve at prediction, the demand for judgment skills is expected to grow, as humans will be increasingly required to interpret AI outputs, assess potential actions, and make state-contingent decisions.

Decision-Making Skills: Decision-making skills focus on selecting the best course of action from alternatives based on evaluations, predictions, and goals. As highlighted by Deming (2021), the growing automation of routine tasks has elevated the importance of adaptive, context-sensitive decision-making in non-routine roles. These skills involve integrating AI predictions with human judgment to select optimal actions while considering ethical, strategic, and contextual factors. Decision-making complements AI by enabling humans to interpret and act on predictive outputs effectively in dynamic environments.

Social Skills: Social skills, such as teamwork, communication, and leadership, are essential for collaboration in increasingly complex, team-based work. Deming (2017) underscores their growing importance as jobs demand greater interpersonal interaction and adaptability. Unlike technical skills, social skills cannot be automated, making them critical for fostering trust, resolving conflicts, and integrating diverse perspectives, especially as AI transforms organizational processes.

We next discuss how we measure these skill concepts in our data.

3 Data Sources

3.1 Online Job Vacancy Data

To study the demand for skills that are related to AI diffusion, we use online job vacancy data from Lightcast (formerly Burning Glass). The data for Europe are collected by scraping job vacancies from over 300 websites. Websites include EURES, the job portal of EU employment agencies, company websties, recruitment agencies, but also large commercial job platforms such as Monster.de, Indeed etc. The algorithm also detects and removes clear duplicate job advertisements from the database. For our analysis, we use data for job vacancies located in Belgium, France, Germany, and Luxembourg from September 2018 to December 2023. In total, our dataset contains job advertisements over this period.

Recent studies using similar data for Europe have compared the distribution of online job vacancies by occupation and industry with data from job vacancy surveys, where firms report their near-future hiring needs (Beręsewicz and Pater, 2021, Cammeraat and Squicciarini, 2021, Fernandez Macias and Sostero, 2024, Vermeulen and Amaros, 2024). They show that online job vacancies are biased towards high-skilled professions and industries, while underrepresenting jobs in low-skilled services (Cammeraat and Squicciarini, 2021).[1] As a consequence, managerial positions are more likely to be captured by online job vacancies, while unskilled and helper professions are underrepresented (Fernandez Macias and Sostero, 2024, Beręsewicz and Pater, 2021). The selectivity with respect to higher-skilled professions does not pose a threat to our analysis, as we are explicitly interested in the impact of AI on those managerial positions and high-skilled professions.[2]

We first clean the data by removing redundant and irrelevant data. We eliminate all empty values and duplicate job advertisements. Then, we filter out the longest job

[1] Specifically, job vacancies in "Financial and Insurance Activities", "Professional, Scientific and Technical Activities", and "Information and Communication" tend to be overrepresented in online job vacancies, while vacancies in "Accommodation and Food Service Activities", "Agriculture, Forestry and Fishing", and "Construction" tend to be underrepresented (Beręsewicz and Pater, 2021, Vermeulen and Amaros, 2024).

[2] Cammeraat and Squicciarini (2021) propose weighting schemes to accurately reflect the entire job market.

descriptions since they result from web scraping the entire website and include unnecessary information such as HTML code or irrelevant text. To further refine the dataset and keep the core of the job adverts, we analyze the most commonly used phrases, symbols, and words in job descriptions. This analysis allows us to create a comprehensive removal list, starting with non-informative symbols, progressing to irrelevant single words, and finally targeting entire phrases that add no value.

The second step is to remove job descriptions that are not job adverts but statements that the position is no longer available, usually expressed by: "this job has already been filled" or "job offer has been removed." These patterns are used to flag and exclude the corresponding entries from the dataset. Additionally, we remove job adverts with descriptions including less than 210 characters. These typically include vague or overly concise statements that cannot be analyzed. By applying a character-length threshold, we ensure that these incomplete or uninformative entries are removed, leaving only descriptions that provide meaningful insights into the advertised positions.

This iterative cleaning process requires constant adjustments to address emerging issues as we refine the dataset. Each step is carefully reviewed to validate the changes and ensure no critical information is accidentally removed. The process involves going back and forth between filtering, analyzing results, and refining the criteria to adapt to new challenges. This iterative procedure ensures that the final dataset is clean and aligned with the project's goals of extracting valuable and accurate information on skills employers look for.

Finally, we drop job advertisements for which we have missing values for the covariates. In total, we are left with 31,014,925 job ads in our data, of which 43.9% are from France, 43% from Germany, 12.8% from Belgium and 0.3% from Luxembourg.

3.2 Extracting Skills from Job Ads

To extract the skills from the cleaned dataset on job ads, we use the large language model (LLM) Mixtral 8x7B, which belongs to the family of sparse mixture-of-experts models developed by Mistral AI (Jiang et al., 2024). Mixtral 8x7B has three main advantages over other LLMs and methods. First, it outperforms other LLMs in its category on most benchmarks, including Llama 2 70B and GPT-3.5. Second, it is a multilingual LLM that performs consistently well in multiple languages. This feature is important in our context where job ads are posted in English, French, or German.[3] Other popular methods for skill extraction, such as the Skills Extractor Library developed by NESTA (NESTA, 2024), are monolingual. Finally, Mixtral 8x7B is openly accessible to the developer community. Research has shown that prompts can significantly influence the output quality of LLMs.

[3]While Mixtral 8x7B has not been formally tested for Dutch, we observe that its performance in this language is comparable to that in French and German.

To ensure high-quality responses from Mixtral 8x7B, we adhere to best practices in prompt engineering by dividing the main task into subtasks. We also include numerous examples in the prompts to guide Mixtral 8x7B's expected behavior.

The skill extraction proceeds in three steps. In the first step, we extract raw skills from the job vacancy descriptions. To achieve this, we identify the language of each description and then query Mixtral 8x7B to identify the skills explicitly mentioned. We adapt the prompt to the language of the description because we have observed that it improves skill extraction.[4]

Box 1 presents the list of prompts used for skill extraction.[5] We begin by defining the role of the LLM and providing clear instructions for the task it must perform. Next, we supply a carefully chosen example that includes a mix of hard and soft skills, along with the expected response from Mixtral 8x7B. We complement this with another example that does not mention any skills, ensuring the model can recognize such cases and refrain from generating unnecessary skills because of hallucination. Finally, we input an actual job vacancy from our cleaned dataset.

InfoBox 1: Prompts for Extracting Skills from an English Job Ad

System: You are an expert human resource manager. Your task is to analyze the skills required in job offers. Follow these instructions carefully:

1. *You will be given a job offer in English.*
2. *Extract all the skills explicitly mentioned in the job offer.*
3. *Return a list of unique skills, with each skill on a new line starting with '- '.*
4. *Each skill should be a maximum of 5 words.*
5. *If no skills are found, return 'None'.*
6. *Do not include any elements other than skills.*
7. *Ensure the output is clear, precise, and follows these guidelines strictly.*

Table 2 gives an example for the skills extracted by Mixtral 8x7B for an English job ad in our dataset.

The second step in the skill extraction process involves translating the raw skills into English. We begin by selecting unique skills and then employ Mixtral 8x7B for translation. The prompts used are presented in Box 2. As with the skill extraction phase, we

[4]We adapt the prompt to the most predominant languages in the online job vacancies (Dutch, French, German, and English). For other languages (e.g. Luxembourgish), we use the prompt in English.

[5]To enhance the reproducibility of our approach, we set the temperature to 0.2, making the model more deterministic and consistent when running the same prompt multiple times. A low temperature also reduces the risk of hallucinations by Mixtral 8x7B.

Table 2: Example: job vacancy and extracted skills with Mixtral

Job ad	Skills extracted with Mixtral
We are looking for a Data Analyst specialized in labor economics, proficient in **Python** and **R** for advanced **econometric analyses** and able to create impactful **data visualizations**. The candidate should be familiar with **Machine Learning** techniques. The ideal candidate will also demonstrate excellent **communication** and **problem-solving skills**. Fluency in **English** is essential. The position requires a PhD in economics, with significant experience in data analysis. Send your CV and cover letter to recrutement@institut-economie.fr to join our dynamic team in Paris, with the possibility of partial telecommuting. Salary: 10 euros/hour.	Python R Econometric analysis Data visualization Machine Learning Communication skills Problem-solving skills English

define the role of the LLM, specify its task, and provide detailed instructions. These instructions specify that when multiple skills are provided together, they should be split, as we observed that Mixtral 8x7B sometimes outputs multiple skills in the initial extraction step. We also instruct the model to return "None" if the input does not represent a skill, to filter out false positives related to educational requirements, experience, or working conditions that were erroneously identified in the first step. Finally, we provide several examples in different languages to cover the common cases Mixtral 8x7B might encounter during translation. In total, we receive unique skills.

In the final step, we reduce the dimensionality of the skill set by removing skills that are essentially the same but expressed in different ways. For example, Mixtral 8x7B might generate "Microsoft Office skills," "MS Office skills," and "MS Office," all of which refer to the same skill. These variations, while phrased differently, are redundant and can create inconsistencies when calculating skill intensity measures in job advertisements. By consolidating such entries into a single, standardized form, we ensure a cleaner and more accurate analysis.

To identify and remove such redundant skills, we follow a four-step process. First, we query Mixtral 8x7B to generate detailed descriptions of each skill, using prompt engineering techniques that include defining the model's role, specifying the task, providing instructions, and offering examples (the prompts are presented in Box 3). Second, we compute the cosine similarity between each pair of skills using their descriptions and sentence embeddings (Reimers and Gurevych, 2019).[6] Cosine similarity measures the degree to which two descriptions share similar meanings, ranging from -1 (completely oppo-

[6]The sentence embeddings are created using the model `intfloat/e5-large-v2`

site meanings) to 1 (identical meanings). Third, we group skills with a cosine similarity equal to or greater than 0.8, merging them into the most frequently mentioned skill. For instance, "Microsoft Office skills," "MS Office skills", and "MS Office" have cosine similarities exceeding 0.8 and are thus combined. Since "MS Office skills" is the most frequently mentioned, we standardize the other variants to "MS Office skills". Finally, we remove the redundant skills from the dataset. After these steps, we are left with millions unique skills.

3.2.1 Comparison against Human Benchmark

To evaluate the skills extraction performed by Mixtral, we compare it to a benchmark coded by human evaluators. To create such a benchmark, we draw a random sample of 1,000 job ads from the entire database. We stratify the random sample by country (Germany, Luxembourg, France, Belgium), source (job search engine, public employment service, employment agency), and year (2018-2021). We ask ChatGPT to extract the skills. We use ChatGPT 4 using the prompt "Please provide a list of all skills required for the following job ad, ranked by importance. Consider both soft and hard skills. Do not add further information. Please adhere to the following response format: #. skill. Here is the job ad: {text}." We then ask 15 employees at LISER with different language backgrounds to read the job ad and assess whether the skills extracted by ChatGPT are correct or not. We organize the human evaluation in a way such that each job ad is evaluated by two evaluators. In total, we were able to evaluate 118 job ads.

We compare the Mixtral skill to the human-based benchmark in two main steps:

Step 1: Creating a Unified List of Evaluator Skills

The first step combines and deduplicates the skills provided by two evaluators into a single list of unique and semantically distinct entries. This process uses the SentenceTransformer `all-MiniLM-L6-v2` model[7] was used for evaluating skill similarities, which generates embeddings for each skill to measure their semantic similarity. Before this comparison, skills are normalized by removing prefixes such as "Ability to" or "Knowledge of," converting the text to lowercase, and trimming whitespace. A pairwise comparison is then conducted between normalized skills, with a similarity threshold set at 0.8. Skills exceeding this threshold in semantic similarity are grouped together, and only the original phrasing of the first occurrence is retained. This ensures that the final evaluator skill list is concise, non-redundant, and faithful to the original inputs, serving as a clean and standardized "ground truth" for subsequent evaluation.

Step 2: Comparing Skills of Evaluators with Extracted Skills by Mixtral

The second step evaluates how closely the skills extracted by Mixtral align with the consolidated skills list by human evaluators created in Step 1. For this, the SentenceTrans-

[7] `https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2`

former model `intfloat/e5-large-v2` [8], which is available on Hugging Face, is employed to encode both the evaluator and Mixtral skill lists into numerical embeddings. These embeddings represent each skill in a high-dimensional space, where their semantic similarity can be quantified using cosine similarity.

True positives (TP) are identified as matched skills that exceed the threshold, while false positives (FP) are the Mixtral-extracted skills with no match, and false negatives (FN) are the evaluator skills that Mixtral failed to capture. Using these metrics, precision, recall, and F1-scores are calculated for each row in the dataset. In order to optimize the performance, we handled threshold tuning to choose the most optimal threshold that balances precision and recall. At lower thresholds (e.g., 0.7), the system accepts more matches, resulting in very high recall since most relevant skills are captured. However, precision tends to drop as the system also includes more false positives. On the other hand, setting a threshold like 0.9 is counterproductive because it becomes so strict that it loses its ability to generalize. In our data, a threshold of 0.8 performs better because it balances recall and precision: it captures most relevant matches (high recall) while maintaining decent precision by avoiding too many false positives.

This process enables a detailed quantitative evaluation of Mixtral's skill extraction performance against the unified evaluator benchmark, with the results stored as an additional column in the dataset for further analysis or reporting.

Figure 1: Descriptive Statistics on Human vs. Machine Skills Extraction

	Accuracy score human vs. machine extraction (F1 score)
Mean	0.78
Median	0.79
Standard Deviation	0.08
Minimum	0.47
Maximum	0.92
25th Percentile	0.73
75th Percentile	0.84

Notes: Table shows the distribution of F1-scores resolving from comparisons between skills extracted by Mixtral and the consolidated evaluator skill list.

Table 1 summarizes the distribution of F1-scores, providing a comprehensive view of Mixtral's performance in extracting skills compared to the human-based benchmark. The mean F1 score of 0.78 suggests that Mixtral demonstrates a strong overall ability to identify relevant skills with high precision and recall. The median score, closely aligned with the mean at 0.79, indicates a balanced performance across most evaluated job ads in the evaluation data.

[8] `https://huggingface.co/intfloat/e5-large-v2`

The relatively narrow standard deviation of 0.08 highlights consistent performance, though some variability exists. This is further reflected in the minimum F1 score of 0.47, which points to certain outlier cases where Mixtral's extraction notably underperformed. Conversely, the maximum F1 score of 0.92 indicates that in some instances, Mixtral achieved near-perfect alignment with human evaluators. The interquartile range, with the 25th percentile at 0.73 and the 75th percentile at 0.84, shows that the majority of evaluations fall within a high-performance band.

These results underline Mixtral's reliability in skill extraction while also identifying areas for improvement. Overall, the descriptive statistics affirm that Mixtral provides a solid baseline for automated skill extraction.

3.3 Generating Skill-related Information

For the unique skills we extracted, we generate additional information to be used in the classification of skills. We first assess whether the skill is directly related to any digital technology (A). In case the technology was not related to any digital technology, we asked for a general description of the skill. If the skills is related to digital technologies, we ask for a concise description of the technology it is related to (B), as well an assessment of whether the skill is considered cutting-edge / frontier (C), directly related to artificial intelligence (D), data skills (e.g. data preparation) (E), self-controlled (F), IT-integrated (G), whether it is related to a prediction problem that AI can substitute (H), or whether it involves human judgment and decision-making based on insights from (AI-generated) data (I). The precise prompts are described in Appendix InfoBox 4. We create dummy variables from the responses of prompts (C)-(I), but also ask Mixtral to give us the reason for its subsequent assessments. As an example, Table 5 shows Mixtral's responses for selected skills. These columns will later serve as input features for the supervised machine learning approach, whereby descriptions will be transformed to vectors.

3.4 Keyword Search

We further enrich the data set by searching the skill descriptions of Mixtral for keywords directly related to AI (e.g. machine learning, artificial intelligence, recommender systems, copilot, chatbot) as well as keywords indirectly related to AI (e.g. data cleaning, data collection, data labeling, data organization, big data, etc.). In addition, we use further keyword lists for digital and frontier digital technologies. All keyword lists are collected manually based on our concept of skills.

After basic manipulation (lowercase, lemmatization), we loop through all descriptions and calculate the number of matched keywords. Table 6 shows the distribution of keywords found for a test sample of 10,000 randomly selected skills (from the total of 5.8

million skills that we observe). For instance, we found up to 6 keywords in the skills' descriptions directly related to AI for only 1.5% of skills, while 17% of skills descriptions entail keywords indirectly related to AI.

We then create two dummy variables for digital and frontier technologies, each capturing the number of keywords found. These will serve as further features in the supervised machine-learning approach.

4 Classification of Skills using ML

The aim is to classify our 5.8 million unique skills found across all 31 million job ads into AI, Data, Prediction, Judgement, Decision-Making, or Social, see Section 1, using a machine learning approach. We describe here how we classify AI and Data skills. By AI skills, we mean skills that involve the development, implementation, or use of AI technologies, such as machine learning algorithms, natural language processing, AI model training, AI system deployment, and advanced data analytics. Complementary Data skills include skills such as data preparation, data labeling, data collection, data cleaning, data organization, big data, and other supportive technologies that contribute to creating the ecosystem in which AI operates, but which are not directly related to AI as defined above. In both definitions, we assume what the technology can do today, not in the future. For the classification, we follow a supervised learning approach, i.e. we train a neural network classifier to predict the skill class based on labeled data together with a set of input features described below.

Neural Network Classifier. To predict the skill class of a job ad, we use a Multi-layer Perceptron (MLP) algorithm provided by pythons sklearn. MLP is a supervised learning algorithm that learns the non-linear relationship between numerical features X and the target y based on the training data using backpropagation. In our case, the target is the skill class taking either "AI", "Data", "Other" (in later versions also Prediction, Judgement, Decision-Making, and Social). The feature set X is discussed below. In contrast to classical logistic regressions, there can be one or more non-linear layers between the input and the output layer, called hidden layers.

Feature set. As input features, we use the skills' descriptions. The reason is that the descriptions extracted from the LLM Mixtral (see Section 3.3) contain more context to the meaning of the job skills. For instance, using word embeddings of the skill names is likely to assign all skills related to Microsoft products in one class due to the word "Microsoft". However, Microsoft Windows, Microsoft Azure and Microsoft Copilot are very different, especially w.r.t. to how they are related to AI. Whereas Microsoft operating system is a digital technology not directly related to AI, the AI tool Microsoft Copilot is directly

related to AI. Microsoft Azure in contrast is related to complementary Data skills, as it provides a computing environment that can, though not must be, used for AI. We therefore rather use the embeddings of the descriptions, instead of only the skills' name, to exploit this important detail. For this, we preprocess the descriptions (lemmatization, tokenization, stop word removal, removal of special characteristics, and stemming) before transforming the descriptions into a numerical vector representation using Word2Vec. In addition to the descriptions, we utilize dummy indicators to identify whether Mixtral classified the skill as related to AI, Data, Decision-Making, and Prediction (we will add Judgement, and Social in later iterations). In addition, we add generated assessments regarding whether a skill is related to a digital technology, frontier, self-controlled, and/or IT-integrated. Finally, we include separate dummy variables for whether keywords for AI and Data skills were found, as well as further variables capturing the number of keywords found for these dimensions.

Training. For the training data, we hand-label 114 out of 10,000 randomly selected skills into three classes: AI, Data, Other (we will add skills related to Prediction, Judgement, and Decision-Making). For this, we follow our concept of skills and use the descriptions as well as the assessments of Mixtral for further guidance. We only include cases in the training data, which well represent the subsequent groups. We call this our GroundTruth data. Based on the GroundTruth data, we then train the MLP classifier to predict the skill class based on the feature set. In the test-train split, we put 50 percent of the training data aside for the later evaluation. For the tuning of the parameters, we use GridSearch.

Evaluation. At this point, we can compare predictions of 57 job skills in the training data with the GroundTruth (i.e. hand-labeled) values. Table 3 shows the precision score, recall, and f1-score. Overall, we receive a relatively strong prediction power. For instance, according to the precision score, the model has made very few false positive predictions and has identified a large proportion of the positive instances in the dataset. The combined f1-score metric is around 0.8, which is close to perfect precision. The feature importance statistics in Figure 5 show that the number of frontier tech keywords, AI keywords found, and number of keywords found, have the highest feature importance score.

Prediction. For the final prediction, we retrain the model based on the entire training data, i.e. including the skills left by side for the evaluation and predict the skill class. In this process, we also store the underlying prediction probabilities. Overall, Figure **??** shows that 7% of skills are related to AI, while 17% are related to complementary Data skills. Note that in further iterations of the paper, we will label more skills based on the entire list of 5.8 million skills (i.e. increase the training data) and predict the skill group

Table 3: Evaluation report

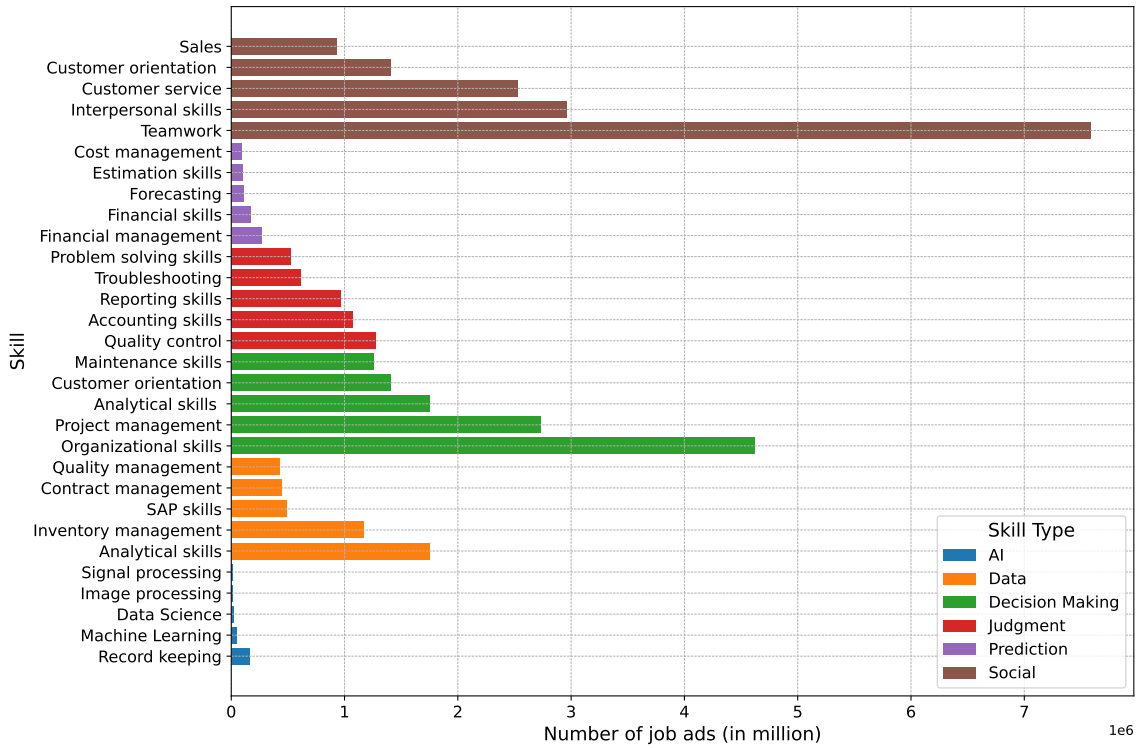
	Precision	Recall	F1-Score	Support
AI skill	0.71	0.67	0.69	15
Data skill	0.75	0.67	0.71	18
Other	0.85	0.96	0.90	24
Accuracy	0.79	0.79	0.79	1
Macro Avg	0.77	0.76	0.77	57
Weighted Avg	0.78	0.79	0.78	57

for all unique skills in the OJV data (i.e. not only on the sample of 10,000 as currently done). This will allow us to calculate the share of skills in by category for each job ad.

Top Five Skills by Type. As outlined in the previous section, we have not yet classified all 5.8 million skills using the ML-based classification method described in the previous section. For now, we thus use the keyword search only to classify job ads into our groups – AI, Data, Prediction, Judgement, Decision Making, and Social. In particular, we declare a job ad to fall into skill group k if at least one keyword appears in the job ad. Though we expect higher precision with the ML-based method and receive an intensity measure, we proceed with the keyword-based method for now.

Figure 2 shows the top 5 skills across AI, Data, Decision-making, Judgment, Prediction, and Social, based on the keyword-based method. Though preliminary, the figure reveals some insights into the specific skills demanded within each category. AI skills (blue) skills like machine learning and image processing dominate, reflecting a strong focus on developing advanced technologies, supported by data science and signal processing. Data skills (yellow) like record-keeping and data science are at the forefront, complemented by inventory management and SAP skills, which reflect the practical requirements of managing large-scale data environments. Decision-Making (green) and Judgment (red) categories emphasize project management, analytical skills, and quality control, illustrating the balance between data-driven insights and maintaining operational standards. Prediction (purple) skills such as forecasting and cost management remain critical for financial and operational foresight, but which may be conducted by AI as prediction costs fall. Meanwhile, Social skills (brown) like teamwork, customer service, and interpersonal communication highlight the necessity of collaboration and customer engagement. This diverse demand showcases the critical balance between advanced technical expertise and interpersonal capabilities, essential for addressing modern organizational challenges and leveraging technological advancements.

Figure 2: Top 5 most frequent skills mentioned in jobs ads by skill type



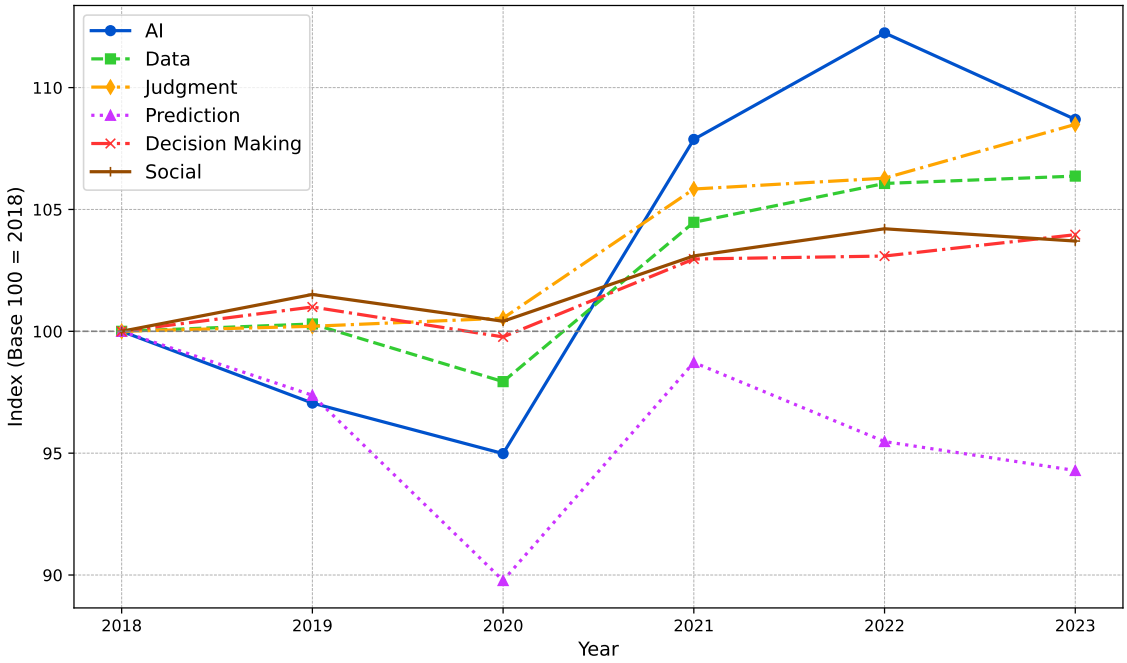
Notes: Figure shows the Top 5 most frequent skills identified based on keyword-based method (i.e. at least one keyword found in the skill requirements of the job ad) across all job ads for each skill type.

5 AI-related skill demand in job vacancies

The analysis of job advertisements from 2018 to 2023 in Figure 3 reveals significant shifts in the demand for different skill types, with a notable initial dip in demand in 2020, likely due to the economic disruptions caused by the COVID-19 pandemic. AI skills saw a substantial increase of approximately 10% over the five-year period, reflecting the growing importance of AI technologies as prediction costs fall. Data skills exhibited steady growth of about 5%, underscoring their critical role in supporting AI systems through the preparation, organization, and management of high-quality data. In contrast, prediction skills declined by roughly 5%, likely due to the automation of routine forecasting and classification tasks by AI systems, which reduces reliance on human expertise in this domain. Meanwhile, judgment skills remained stable, fluctuating minimally around the base level, highlighting their enduring importance in interpreting AI outputs, navigating uncertainty, and making context-sensitive decisions. These trends illustrate the complementary relationship between AI capabilities, human judgment, and data inputs, while prediction tasks are progressively delegated to AI systems.

Since there may be various reasons for the trends seen in Figure 3 beyond AI adoption in firms and industries, we next shed light into how much the trend can be explained

Figure 3: Evolution of skills demand (Base 100 = 2018)



Notes: Figure shows the share of job ads for skill type k (e.g. AI) where at least one keyword appears in skill requirements of the job ad. The time series are normalized to the starting year 2018.

by AI. For this, we regress the job-level skill demand on AI exposure measured at the occupational level. We use the measures of automation and augmentation AI exposure developed by Marguerit (2024b), as defined below:

- *Automation AI Exposure*: Refers to technologies that substitute labor inputs in occupations, potentially displacing workers by automating tasks.
- *Augmentation AI Exposure*: Represents technologies that enhance worker capabilities, improve output quality, and generate new demands for expertise and spec

The automation index is constructed by mapping AI-related questions posted on Stack Overflow—the leading Q&A forum for developers— onto occupational descriptions from the 2010 O*NET. This mapping results in a longitudinal measure of automation AI exposure at the occupational level from 2018 to 2022. To build the augmentation AI exposure index, AI-related questions from Stack Overflow are matched with the micro-titles listed in the 2010 Census Alphabetical Index (CAI). CAI provides thousands of micro-titles associated with occupations and industries, offering detailed descriptions of the outputs produced in these roles (Autor et al., 2024b). This measure is longitudinal and varies at the occupation-industry level for 2018-2022. Both measures are cross-walked to align with the industry NACE and occupational ISCO classifications. Appendix Figures 6 and

7 show the occupations with the highest and lowest automation and augmentation AI exposure.

In particular, we estimate the probability that the job ads contain skill k, with k={AI, data, prediction, judgement, decision-making and social}, at time t as follows:

$$\text{P(Y=1)}_{ijt} = \lambda_1 \text{Automation}_{jt} + \lambda_2 \text{Augmentation}_{ljt} + \zeta_i + \eta_t + \theta_{j\times l} + \varepsilon_{it} \qquad (1)$$

where Y is a dummy variable measuring whether a job advert requires skill k, and j is the occupation. ζ_i captures several job-specific fixed effects such as the job advertisement website, location (NUTS-1), industry (1-digit NACE), η_t reflects month and year fixed effects, and θ_{jxl} reflection occupation-industry fixed effects (l = industry). ε_i is the remaining error term.

The results are presented in Table 4. Column (1) shows that automation AI exposure positively correlates with the demand for AI skills (point estimate: 0.19; SE: 0.11), whereas augmentation AI exposure does not appear to be related to AI skills. In Column (2), automation AI exposure is much more strongly positively associated with the demand for data skills: a one standard deviation increase in automation AI exposure raises the probability of requiring data skills by 1.98 percentage points. Conversely, augmentation AI exposure negatively correlates with data skills (point estimate: -3.03; SE: 0.53). Column (3) documents the relationship between prediction skills –defined as human-driven processes to forecast, classify and predict outcomes– and AI exposure. In contrast to our expectations, Automation AI exposure is positively correlated with the demand for prediction skills (point estimate: 2.98; SE: 0.28), while augmentation AI exposure is not related to prediction skills. This result could suggest that it is not advanced technical AI skills per se that are demanded most from Automation AI, but rather the data skills provided to prepare and maintain data necessary for the operation of AI. The surprising fact that we observe a positive association between automation AI and humand-based prediction skills could reflect that most prediction tasks are still performed by humans, which might change when increasingly deligated to AI systems. In contrast to Automation AI, there seems to be no positive skill demand changes from Augmentation AI exposure in the domains AI, Data and Prediction.

We now focus on judgement, as other input to decision-making, besides data. For judgment skills (Column (4)), a one standard deviation increase in automation AI exposure significantly increases the probability of requiring judgment skills by 2.91 percentage points. Column (3) examines decision-making skills. The probability of requiring decision-making skills is not associated with automation AI exposure, while a positive association is found for augmentation AI exposure (point estimate: -1.02; SE: 0.34). Finally, Column (6) reports that the probability of requiring social skills increases by 7.68 percentage points with a one standard deviation increase in automation AI exposure. Similarly, aug-

Table 4: Occupational AI-Exposure and AI-related skill demands

Dependent variable: Probability that the job ad contains skill k:						
	AI	Data	Prediction	Judgment	Decision Making	Social
	(1)	(2)	(3)	(4)	(5)	(6)
Automation AI	0.191*	1.982**	2.983***	2.915***	0.245	7.683***
	(0.116)	(0.997)	(0.285)	(0.942)	(0.827)	(0.845)
Augmentation AI	0.123	-3.033***	0.070	-0.182	1.025***	1.125**
	(0.098)	(0.534)	(0.267)	(0.435)	(0.340)	(0.522)
R-squared	0.029	0.195	0.134	0.201	0.156	0.107
Obs (in mil.)	31.01	31.01	31.01	31.01	31.01	31.01
Mean of dep var (in %)	1.32	47.70	10.29	59.67	68.79	63.79

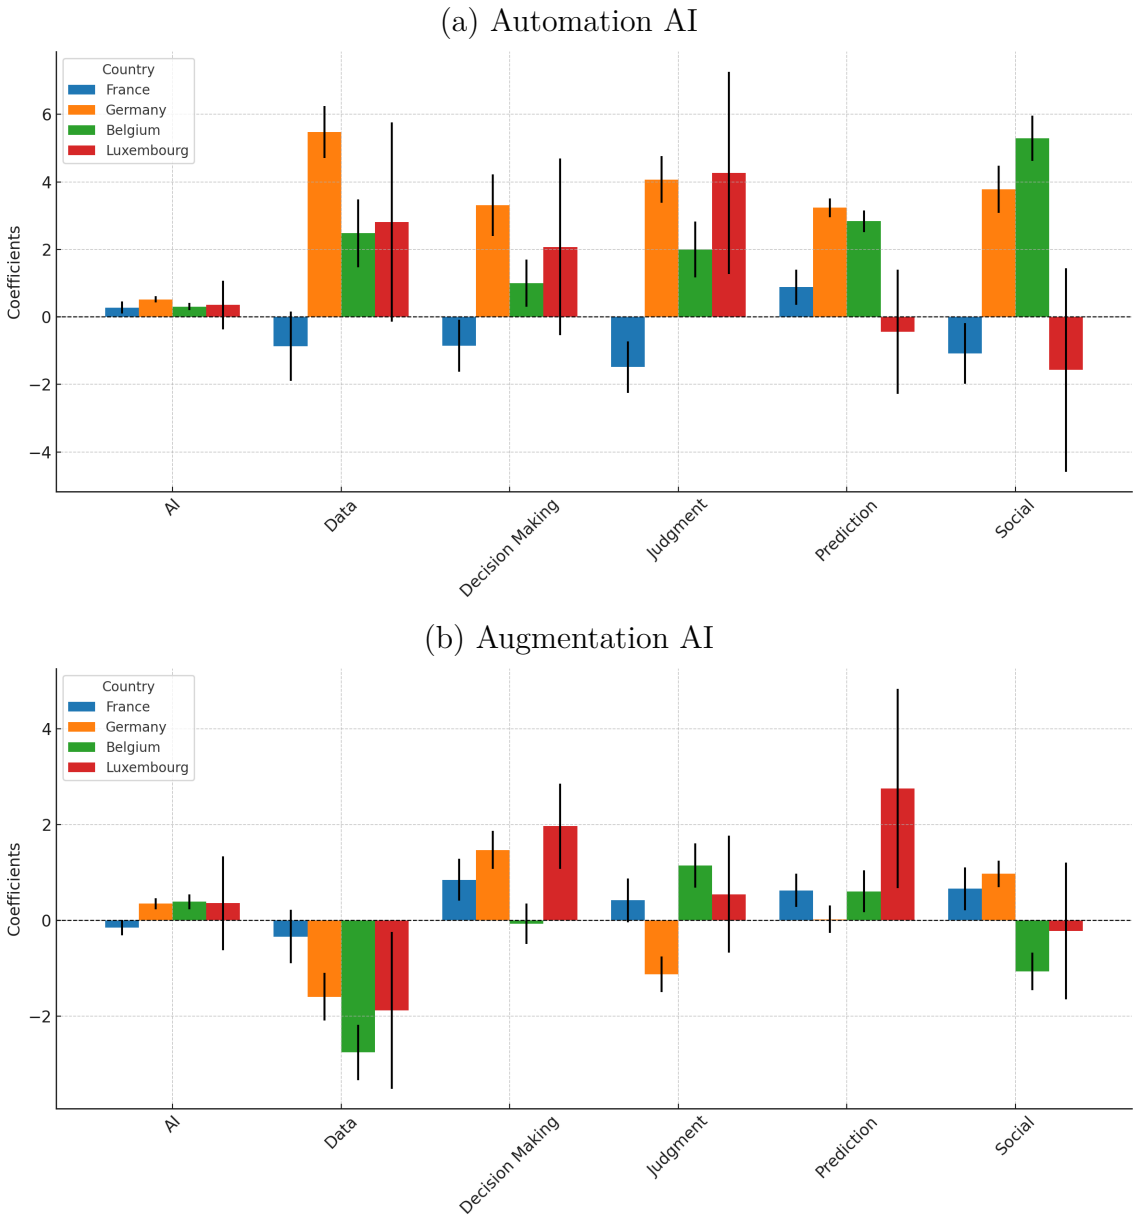
Notes: Table shows coefficients from regressions of Equation 1. Regressions control for educational requirements and include the following fixed effects: year, month, industry (1-digit) interacted with year and month, industry (2-digit) interacted with occupation (4-digit), job location (NUTS-1), and job advertisement website. Standard errors are in parentheses and clustered at the occupation-industry level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Coefficients are multiplied by 100.

mentation AI exposure is positively correlated with the demand for social skills (point estimate: 1.12; SE: 0.52). Though preliminary, the results suggest that Augmentation AI exposure has a stronger influence on decision-making and social skills, highlighting the need for workers.

In subsequent analyses, we will focus on the intensive margin of skill demand. Rather than estimating the probability that a job advertisement requires a particular skill category (e.g., AI, Data, Decision-making, Judgment, Prediction, or Social skills), we will assess the prominence of these skills within job advertisements. To do this, we will compute a skill intensity measure that quantifies the share of each skill category among all skills mentioned within job adverts.

By country. Figure 4a depicting the effects of Automation AI across countries and skill types reveals significant cross-country differences in its impact (see also Appendix Table 8. Germany stands out with strong positive effects across all skill types, particularly in AI, Data, and Judgment, suggesting widespread integration of automation into diverse sectors. Conversely, France and Luxembourg show a more nuanced picture, with positive impacts on AI and Prediction but notable negative effects on Judgment and Social, highlighting challenges in applying automation to tasks requiring nuanced human interaction. Belgium exhibits a mixed pattern, with strong positive effects on technical skills like Prediction but notable limitations in Social applications. Overall, Automation AI is shown to enhance technical and predictive capacities across most countries, but its application to collaborative and interpersonal tasks remains a challenge.

Figure 4: Occupational AI-exposure and AI-related skill demands by country



Notes: Figure shows coefficients from regressions of equation 1 by country. Each bar indicates the coefficient for Augmentation AI across different skill types. Regressions control for educational requirements and include the following fixed effects: year, month, industry (1-digit) interacted with year and month, industry (2-digit) interacted with occupation (4-digit), job location (NUTS-1), and job advertisement website. Error bars represent standard deviations. Observations: France (13616355), Germany (13336801), Belgium (3982736) and Luxembourg (79033).

Figure 4b illustrating the effects of Augmentation AI highlights its strengths in enhancing Prediction and Decision-making across all countries, particularly in Luxembourg, where it shows robust positive impacts on these skill types. However, it also underscores significant limitations in areas like Social and Data, with negative effects in several countries, including Germany and Belgium. France demonstrates moderate positive effects

across most skill types, suggesting that augmentation is being effectively leveraged to enhance analytical and cognitive processes. Despite its promise in forward-looking and decision-based skills, Augmentation AI appears to struggle with social and data-centric tasks, indicating potential barriers in collaborative and interpersonal applications. These findings suggest that while Augmentation AI excels in cognitive enhancement, further efforts are needed to improve its integration into areas requiring human interaction and data processing.

6 Conclusion

This paper investigates the impact of artificial intelligence (AI) on skill demands in the labor markets of Continental Europe, leveraging a dataset of over 31 million online job vacancies from Belgium, France, Germany, and Luxembourg spanning 2018 to 2023. Using advanced machine learning models, it classifies job skills into categories such as AI, data, prediction, judgment, decision-making, and social skills. The findings reveal a significant increase in the demand for AI-related skills (10%) and data skills (5%), reflecting the growing integration of AI technologies in business operations. Conversely, the need for prediction skills has declined, likely due to automation of routine forecasting tasks by AI, while judgment skills have remained stable, highlighting the enduring importance of human decision-making in AI-enhanced contexts.

The analysis also shows that occupations exposed to AI automation and augmentation experience varying shifts in skill demand. Automation exposure drives up the need for AI, data, and judgment skills, while augmentation positively influences the demand for decision-making and social skills. These trends underscore the complementary relationship between technical and human-centric skills in the context of AI adoption.

In conclusion, the study finds that while AI adoption accelerates the demand for advanced technical skills, it also amplifies the importance of uniquely human competencies such as judgment and decision-making. This dual shift highlights the necessity for targeted workforce development and upskilling strategies to ensure that labor markets can effectively adapt to the opportunities and challenges posed by AI integration. These findings have critical implications for policymakers, educators, and businesses aiming to prepare the workforce for an increasingly AI-driven future.

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, 4:1043–1171.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D., Kong, F., and Restrepo, P. (2024). Tasks At Work: Comparative Advantage, Technology and Labor Demand. *NBER Working Paper No. 32872*.

Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.

Agrawal, A., Gans, J., and Goldfarb, A. (2018). Prediction, judgment, and complexity: A theory of decision-making and artificial intelligence. In *The economics of artificial intelligence: An agenda*, pages 89–110. University of Chicago Press.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour economics*, 71:102002.

Ash, E. and Hansen, S. (2023). Text algorithms in economics. *Annual Review of Economics*, 15(1):659–688.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024a). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*, 139(3):1399–1465.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024b). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Beręsewicz, M. and Pater, R. (2021). *Inferring job vacancies from online job advertisements*. Publications Office of the European Union.

Bonfiglioli, A., Crinò, R., Gancia, G., and Papadakis, I. (2024). Artificial intelligence and jobs: Evidence from us commuting zones. *Economic Policy*, page eiae059.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? In *AEA Papers and Proceedings*, volume 108, pages 43–47.

Cammeraat, E. and Squicciarini, M. (2021). Burning Glass technologies' data use in policy-relevant analysis: An occupation-level assessment.

Deming, D. and Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. (2021). The growing importance of decision-making on the job. Technical report, National Bureau of Economic Research.

Engberg, E., Hellsten, M., Javed, F., Lodefalk, M., Sabolová, R., Schroeder, S., and Tang, A. (2024). Artificial intelligence, hiring and employment: Job postings evidence from sweden. Technical report.

Felten, E. W., Raj, M., and Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. In *AEA Papers and Proceedings*, volume 108, pages 54–57.

Felten, E. W., Raj, M., and Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *NYU Stern School of Business*.

Fernandez Macias, E. and Sostero, M. (2024). Skewed signals? confronting biases in online job ads data. Scientific analysis or review, Seville (Spain).

Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D. S., de las Casas, D., Hanna, E. B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L. R., Saulnier, L., Lachaux, M.-A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T. L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., and Sayed, W. E. (2024). Mixtral of experts.

Korinek, A. (2024). Economic policy challenges for the age of AI. Technical report, Paper for a conference organized by the Peterson Institute for International Economics and the International Monetary Fund on 'Rethinking Economic Policy: Steering Structural Change'.

Lin, J. (2011). Technological Adaptation, Cities, and New Work. *Review of Economics and Statistics*, 93(2):554–574.

Mann, K. and Püttmann, L. (2023). Benign Effects of Automation: New Evidence from Patent Texts. *Review of Economics and Statistics*, 105(3):562–579.

Marguerit, D. (2024a). Augmenting or automating labor? the effect of AI exposure on new work, employment, and wages. *mimeo*.

Marguerit, D. (2024b). Augmenting or automating labor? the effect of ai implementation on new work, employment, and wages. Technical report, LISER.

NESTA (2024). Nesta's skills extractor library.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Samek, L., Squicciarini, M., and Cammeraat, E. (2021). The human capital behind AI: Jobs and skills demand from online job postings.

Vermeulen, W. and Amaros, F. G. (2024). How well do online job postings match national sources in european countries?

Webb, M. (2020). The Impact of Artificial Intelligence on the Labor Market. Working Paper, Accessed at: https://www.michaelwebb.co/webb_ai.pdf.

InfoBox 2: Prompts for skills translation into English

System: You are a translator specialized in human resources. Your task is to translate a given skill from a job vacancy into English. Follow these instructions carefully:

1. *If the skill is already in English, leave it unchanged.*
2. *If the skill includes the name of a software, leave the software name unchanged.*
3. *Return the generic name of the skill instead of a literal translation.*
4. *Return only one translation for each skill, starting with '- '.*
5. *If multiple skills are provided, return a list of translated skills, each starting with '- ' on a new line.*
6. *Ensure each skill is a maximum of 5 words.*
7. *If the given item is not a skill, return 'None'.*
8. *Do not include any explanation or additional text beyond the translated skills.*
9. *Ensure the output is clear, precise, and strictly follows these guidelines.*

InfoBox 3: Prompts for describing skills

System: You are an expert in describing job skills. Your task is to provide a clear and precise description of a given skill from a job vacancy. Follow these instructions carefully:

1. *Describe the skill in detail.*
2. *Return only one description for each skill, starting with '- '.*
3. *If multiple skills are provided, return a list of descriptions, each starting with '- ' on a new line.*
4. *Ensure each description has a maximum of 50 words.*
5. *If the given item is not a skill, return 'None'.*
6. *Do not include any explanation or additional text beyond the description.*
7. *Ensure the output is clear, precise, and strictly follows these guidelines.*

InfoBox 4: Prompts for skills' assessments with Mixtral

1. *TEMPLATE_DIGITAL = "You are an AI assistant specialized in classifying job competencies based on my definitions. Digital technology refers to electronic devices or systems that create, store, and process data. Examples include computers, smartphones, the internet, software applications, or other devices that involve any digital components. Does the job competency {} directly involve any digital technology that exists today? Assume what this digital technology can do today, not what it will do in the future. Respond only with either 'yes' or 'no' followed by a concise explanation."*

2. *TEMPLATE_DESCRIPTION = "You are an AI assistant specialized in classifying job competencies based on my definitions. Provide a concise description, in running text, of the job competency {}. Answer in the format: {} is..."*

3. *TEMPLATE_DATA = "Data skills include skills such as data preparation, data labeling, data collection, data cleaning, data organization, big data, and other supportive technologies that contribute to creating the ecosystem in which AI operates. Assume what the technology can do today, not what it will do in the future. Is {} indirectly related to AI? Respond with 'yes' or 'no' followed by a concise explanation."*

4. *TEMPLATE_AI = "Does the skill {} belong to the field of artificial intelligence? Respond 'yes' if the skill mentions artificial intelligence, machine learning, or any related terms. Avoid overanalyzing; simply classify based on explicit mention of AI-related terms."*

5. *TEMPLATE_PREDICTION = "Some human tasks, such as forecasting, classifying, predicting outcomes, or organizing data, can be fully substituted by AI or prediction machines based on current capabilities. Is {} a task that could reasonably be completely substituted by an AI or prediction machine? Assume what the technology can do today, not what it will do in the future. Respond with 'yes' or 'no' followed by a concise explanation."*

6. *TEMPLATE_DECISION_MAKING = " Decision-making skills are those that involve making judgments, choices, or strategies informed by data insights, regardless of whether the data comes from AI models or other sources. Does {} relate to decision-making that could, in principle, be informed by data? Respond with 'yes' or 'no' followed by a concise explanation."*

InfoBox 5: Prompts for skills' assessments with Mixtral - cont'd

6. *TEMPLATE_FRONTIER = "Cutting-edge technology refers to advanced technological devices, techniques, or achievements that push the boundaries of innovation and represent significant progress in their field, even if they are not yet widely adopted or fully mature. Is {} related to a technology or process that is innovative or significantly advanced by today's standards? Respond only with either 'yes' or 'no' followed by a concise explanation."*

7. *TEMPLATE_SELF_CONTROLLED = "Self-controlled production technology performs work processes largely automatically and autonomously. Examples include smart factories, cyber-physical systems, and the Internet of Things (IoT). Is {} directly related to self-controlled production technology? Respond with 'yes' or 'no' followed by a concise explanation."*

8. *TEMPLATE_IT_INTEGRATED "IT-integrated office and communication technology automates office, communication, or data-processing tasks and performs them largely automatically. Examples include big data, cloud computing, and online platforms. Is {} directly related to IT-integrated office and communication technology? Respond with 'yes' or 'no' followed by a concise explanation."*

Table 5: Examples of skill-related information generated with Mixtral

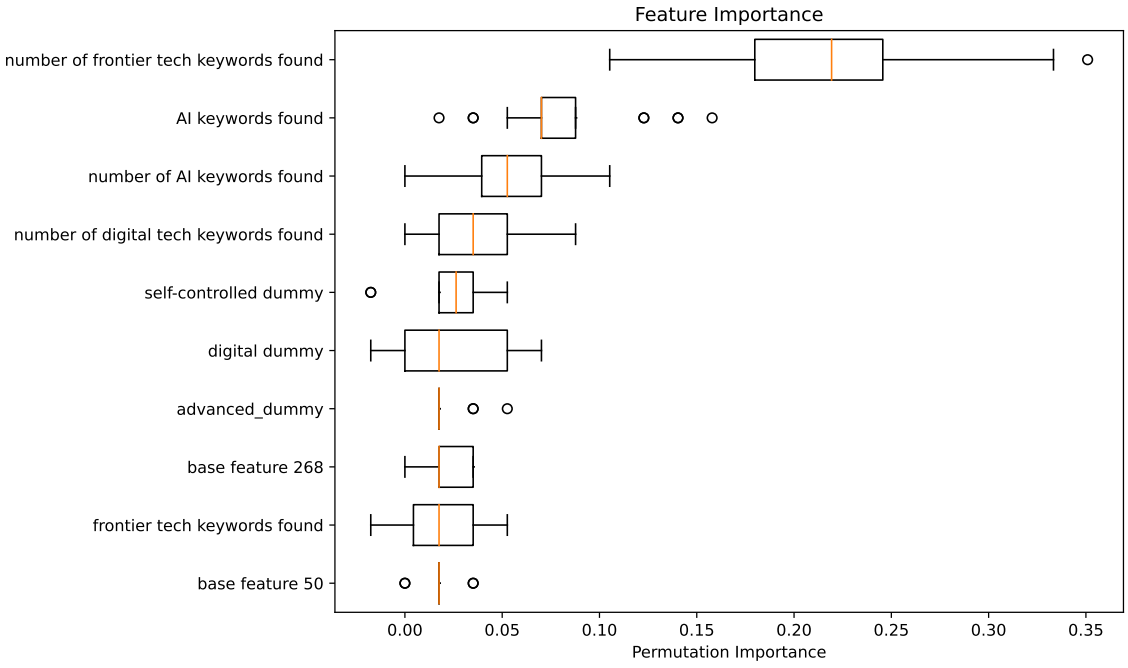
Skill	Description	Digital	Frontier	AI	Data	Self-Controlled	IT-Integrated	Prediction	Decision Making
ai-powered data visualization modules	AI-powered data visualization module be a job competency that involve the ability to use artificial intelligence tool and technique to create graphical representation of complex data set . This competency require a strong understanding of both data analysis and visual design principle , as well a the ability to effectively communicate insight and trend to stakeholder . Professionals with this competency be able to leverage AI algorithm to automate the data visualization process , make it fast and more efficient . They be also able to create interactive and dynamic visualization that allow user to explore and manipulate data in real-time . Overall , AI-powered data visualization module be a critical competency for data scientist , business analyst , and other professional who need to make data-driven decision in a fast-paced and complex environment .	1	1	1	0	0	0	0	1
big data architectures on hyperscalers	Big Data architecture on Hyperscalers be a job competency that involve design , building , and manage scalable and efficient Big Data architecture on cloud-based hyperscale platform , such a Amazon Web Services (AWS) , Microsoft Azure , and Google Cloud Platform (GCP) . This role require a deep understanding of Big Data technology , such a Hadoop , Spark , and Hive , as well a expertise in cloud service , include virtual machine , container , and serverless computing . A Big Data architect on Hyperscalers should be able to create data pipeline , optimize storage and processing , ensure data security and compliance , and provide technical leadership to data engineering team .	1	1	0	1	0	0	0	1
workshop teams construction	Workshop team construction be the ability to effectively design , assemble , and manage team responsible for conduct workshop . This involve identify the necessary skill and expertise require for the workshop , recruiting and select team member , set clear goal and expectation , and facilitate collaboration and communication among team member to ensure a successful workshop outcome . This competency also include the ability to adapt to change circumstance , resolve conflict , and provide constructive feedback to team member to enhance their performance and the overall workshop experience .	0	0	0	0	0	0	0	0

Table 6: Frequency of number of keywords found in skills' descriptions

# Keywords Found (1)	Digital (2)	Frontier (3)	Data (4)	AI (5)
0 (0.0%)	6068 (60.7%)	8126 (81.3%)	8313 (83.1%)	9846 (98.5%)
1 (1.8%)	2344 (23.4%)	1268 (12.7%)	1305 (13.1%)	60 (0.6%)
2 (3.6%)	980 (9.8%)	334 (3.3%)	274 (2.7%)	57 (0.6%)
3 (5.5%)	360 (3.6%)	130 (1.3%)	69 (0.7%)	20 (0.2%)
4 (7.3%)	152 (1.5%)	56 (0.6%)	22 (0.2%)	12 (0.1%)
5 (9.1%)	69 (0.7%)	31 (0.3%)	10 (0.1%)	3 (0.0%)
6 (10.9%)	19 (0.2%)	27 (0.3%)	3 (0.0%)	2 (0.0%)
7 (12.7%)	8 (0.1%)	16 (0.2%)	2 (0.0%)	0 (0.0%)
8 (14.5%)	0 (0.0%)	4 (0.0%)	1 (0.0%)	0 (0.0%)
9 (16.4%)	0 (0.0%)	6 (0.1%)	0 (0.0%)	0 (0.0%)
10 (18.2%)	0 (0.0%)	2 (0.0%)	1 (0.0%)	0 (0.0%)

Note: Tables based on random sample of 10,000 skills.

Figure 5: Feature importance statistics



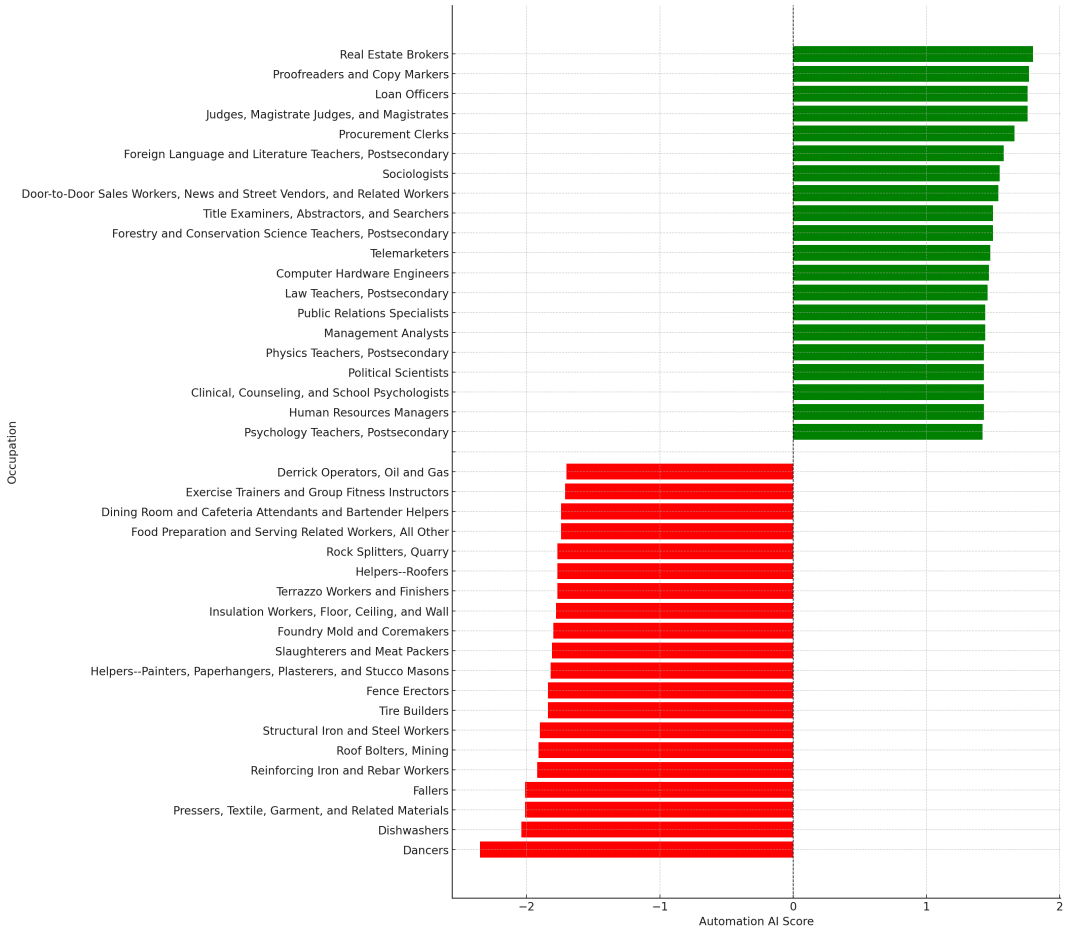
Notes: XXX

Table 7: Top 20 skills based on ML-based skill classification

Top 20 AI Skills	**Top 20 Data Skills**
Knowledge of Text Generation	SAP BI Development
AI for PET (Positron Emission Tomography) CT Scans	SAP BI and BW Reporting
AI-Based Development Methods	Tracking and Monitoring Dashboards
Modern Learning Architectures	Knowledge of Pentaho (Versions 5.2, 8.1, 9.0)
Background in E-Technology	Spark Business Intelligence (BI) Skills
Understanding of Dialogue Pipelines	UX (User Experience) Technologies
AI and Machine Learning Techniques in Life Sciences	ELT (Extract, Load, Transform) Pipelines
Machine Interface Technologies	Knowledge of Technical Constraints for Web Support
Mobile Solution Autonomy	Data Processing Automation for Gaming
Analyzing and Optimizing Machine Programs	Database Client Software Expertise
Transformer Optimization	HR QlikView Dashboards
Cost Engineering Digitization	IT Specialist Roles
Smart Logistics Systems	TCP/IP, Ethernet, HTTP, SMTP, NTP, SNMP Protocols
Secure AI Products	Self-Service BI&A (Business Intelligence and Analytics) Solutions
IEM (Integrated Enterprise Management) Development	Windows Server Products (SQL Server, RDS, WSUS, etc.)
Custom Profiling Skills	Verifying Server Compliance
Edge AI Solutions	Proficiency in W3C (World Wide Web Consortium) Technologies
E-Labeling	Expertise in Web Design
NMVE (Non-Volatile Memory Express) Technology	Developing Enhancements for Business Applications
Improvement of Optical Sorting Machines	Nomenclature Building in CMMS with SAP

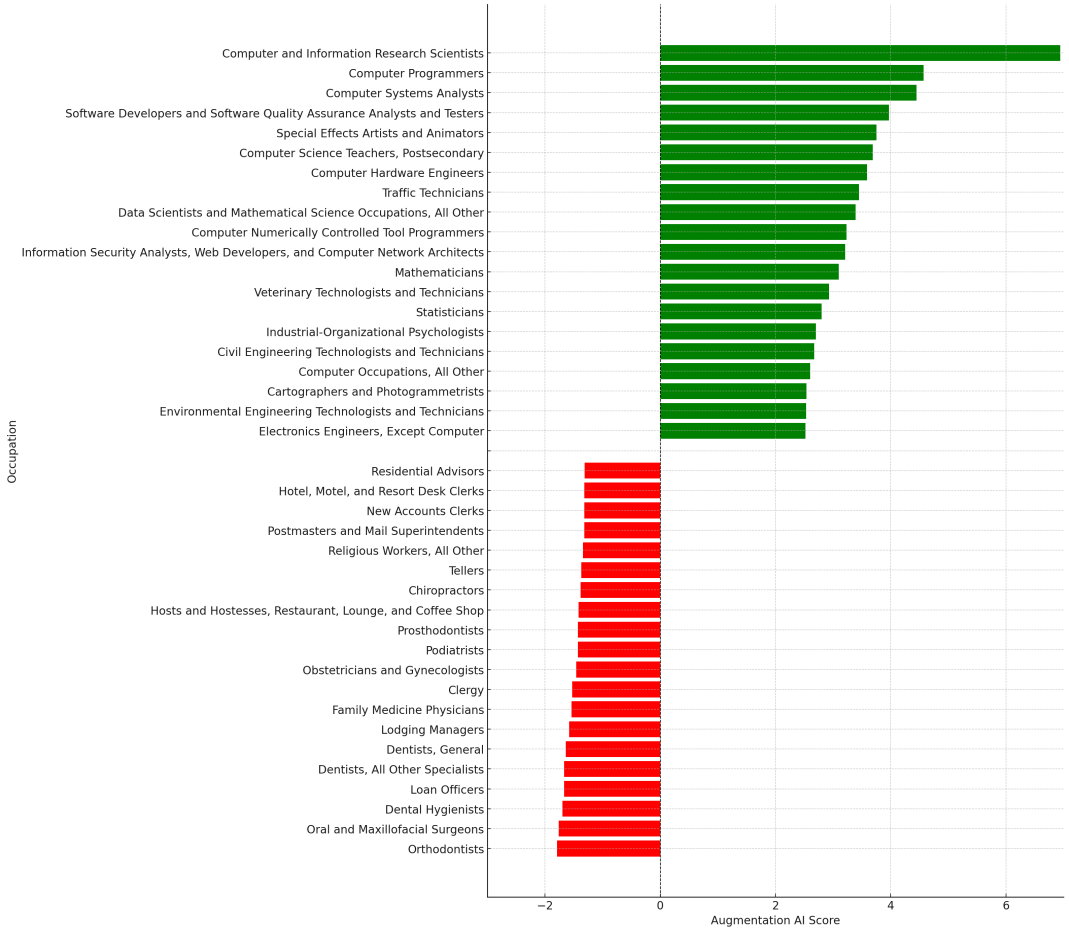
Notes: Table shows Top 20 skills ranked by the subsequent probability scores from the ML-bassed skill classification in Section 4 for the random sample of 10.000 skills.

Figure 6: Occupations with the highest and lowest automation AI exposure, 2022



Notes: Occupations are ranked by their score of automation AI exposure at 6-digit Standard Occupational Classification (SOC) level and computed following (Marguerit, 2024b). AI score is standardized to have a mean of 0 and a standard deviation of 1.

Figure 7: Occupations with the highest and lowest augmentation AI exposure, 2022



Notes: Note: The table shows the weighted average exposure at the occupational level, using employment size in industries as weight. Occupations are ranked by their score of augmentation AI exposure at 6-digit Standard Occupational Classification (SOC) level and computed following (Marguerit, 2024b). AI score is standardized to have a mean of 0 and a standard deviation of 1.

Table 8: Occupational AI-Exposure and AI-related skill demands by country

	Automation				Augmentation			
	France	Germany	Belgium	Luxembourg	France	Germany	Belgium	Luxembourg
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.28	0.52***	0.31***	0.35	-0.15	0.35***	0.39**	0.36
	(0.18)	(0.09)	(0.11)	(0.72)	(0.16)	(0.12)	(0.16)	(0.98)
Data	-0.87	5.48***	2.48**	2.81	-0.34	-1.60***	-2.76***	-1.88
	(1.03)	(0.78)	(1.00)	(2.96)	(0.56)	(0.50)	(0.58)	(1.64)
Decision Making	-0.86	3.31***	1.00	2.08	0.85*	1.47***	-0.07	1.97**
	(0.77)	(0.92)	(0.70)	(2.61)	(0.44)	(0.40)	(0.42)	(0.89)
Judgment	-1.48*	4.07***	2.00**	4.27	0.42	-1.13***	1.14**	0.55
	(0.76)	(0.69)	(0.82)	(3.00)	(0.46)	(0.37)	(0.46)	(1.22)
Prediction	0.88*	3.23***	2.84***	-0.44	0.63*	0.02	0.61	2.76
	(0.52)	(0.28)	(0.32)	(1.84)	(0.35)	(0.28)	(0.44)	(2.08)
Social	-1.08	3.78***	5.29***	-1.57	0.66	0.97***	-1.07***	-0.22
	(0.90)	(0.70)	(0.67)	(3.01)	(0.45)	(0.28)	(0.40)	(1.43)
R-squared	0.03	0.05	0.04	0.14	0.19	0.16	0.12	0.15
Observations (in mil.)	13.62	13.34	3.98	0.08	13.62	13.34	3.98	0.08

Notes: Table shows coefficients from regessions of Equation 1 by country. Regressions control for educational requirements and include the following fixed effects: year, month, industry (1-digit) interacted with year and month, industry (2-digit) interacted with occupation (4-digit), job location (NUTS-1), and job advertisement website. Error bars represent standard deviations.