

Guarascio, Dario; Reljic, Jelena; Stöllinger, Roman

Working Paper

Artificial Intelligence and Employment: A Look into the Crystal Ball

GLO Discussion Paper, No. 1333

Provided in Cooperation with:

Global Labor Organization (GLO)

Suggested Citation: Guarascio, Dario; Reljic, Jelena; Stöllinger, Roman (2023) : Artificial Intelligence and Employment: A Look into the Crystal Ball, GLO Discussion Paper, No. 1333, Global Labor Organization (GLO), Essen

This Version is available at:

<https://hdl.handle.net/10419/278106>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Artificial Intelligence and Employment: A Look into the Crystal Ball

Dario Guarascio^{1,2}, Jelena Reljic^{1,2*}, Roman Stöllinger³

¹ Sapienza University of Rome

² GLO

³ Delft University of Technology (TU Delft)

* Corresponding author: jelena.reljic@uniroma1.it

Abstract

This study provides evidence of the employment impact of AI exposure in European regions, addressing one of the many gaps in the emerging literature on AI's effects on employment in Europe. Building upon the occupation-based AI-exposure indicators proposed by Felten et al. (2018, 2019, 2021), which are mapped to the European occupational classification (ISCO), following Albanesi et al. (2023), we analyse the regional employment dynamics between 2011 and 2018. After controlling for a wide range of supply and demand factors, our findings indicate that, on average, AI exposure has a positive impact on regional employment. Put differently, European regions characterised by a relatively larger share of AI-exposed occupations display, all else being equal and once potential endogeneity concerns are mitigated, a more favourable employment tendency over the period 2011-2018. We also find evidence of a moderating effect of robot density on the AI-employment nexus, which however lacks a causal underpinning.

Keywords: Artificial intelligence, industrial robots, labour, regional employment, occupations
JEL codes: J21, J23, O33, R1

1. INTRODUCTION

On the 22 of March 2023, a group of more than 30 thousand IT CEOs, managers and scientists, including prominent business figures such as Elon Musk (CEO of SpaceX, Tesla & Twitter) and Steve Wozniak (Co-founder, Apple), signed an open letter¹ urging ‘AI labs to pause for at least 6 months the training of AI systems’. According to the petitioners, AI technologies with human-competitive intelligence ‘can pose profound risks to society and humanity’, thus ‘should be planned for and managed with commensurate care and resources. On the top, there are the risk that AI takes over existential activities, allowing no longer intelligible (and therefore hardly controllable) machines to be in charge of vital decisions in areas such as defence, security or health care. These concerns reflect a pessimistic view on technological progress, which is related to the risk that the diffusion of AI may turn into a new wave of technological unemployment, bringing back to the fore the same fears surrounding previous automation and digitalization waves (Frank et al., 2019; Autor, 2022).² Similar concerns are influencing policy making at the EU level. Speaking at the last G20 summit, the EU Commission President Ursula Von der Leyen argued that ‘Europe and its partners should develop a new global framework for AI risks...protecting against systemic societal risks and foster investments in safe and responsible AI systems’.³

When it comes to the potential labour disruptive effects of AI, there are at least three discontinuities that are worth mentioning. First, AI technologies are, for the first time, putting man and machine in competition over tasks so far unattainable for non-human devices, particularly in the service sector. Generative AIs – e.g., Open AI’s Chat-GPT or Google’s Bard – are already showing their potential replicating and, in some cases, outperforming humans in carrying out activities that are based on complex reasoning, value judgements and combination of highly heterogeneous elements, such as a legal or a clinical advice (Felten et al., 2023). For instance, consider the interpretation of x-ray images and other medical imaging diagnostics, which – when performed by humans – requires a lot of knowledge, experience and involves non-routine cognitive tasks. Put differently, human workers now face competition in their core competences which explains the ‘this time could be different’ attitude with regards to labour market effects of technology. Second, AI is expected to magnify the potential of key automation technologies (e.g., robots), opening the way for job disruption also in manufacturing. Third, as AI-related technologies, knowledge and capabilities are not homogeneously distributed, this technological wave may bring about polarisation across Global Value Chains (GVCs) as well as within-countries, with peripheral areas losing further grounds together with the emergence of new forms of techno-economic dependency (Korinek and Stiglitz, 2021). On the other hand, AI promises to complement and enhance workers in carrying out a number of tasks, in services and manufacturing alike (Gmyrek et al., 2023). Increasing productivity may spur labour demand, potentially compensating AI-induced job disruptions, should the latter take place. By the same token, organisational efficiency - including managerial tasks aimed at monitoring, coordinating and directing workers – takes a leap forward, with potential labour-saving effects as well as changes in the composition of occupational profiles and their relative position within organisations (Chowdhury et al., 2023). New sectors and market niches are also about to emerge. Thus, it is hard to predict how many

¹ The letter and the list of signatures is available at: <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>. Last access: September 28, 2023

² The ‘man vs machine’ race is already found in the writings of Keynes on the future of work (Keynes, 2010[1931]).

³ The full speech of the European Commission President is available at: https://ec.europa.eu/commission/presscorner/detail/en/statement_23_4424. Last access: September 28, 2023

jobs this will bring, as a significant share of disruptive AI applications are in their initial stage of development (Mondolo, 2022).

In other words, the ultimate employment effects of AI remain an open question (Acemoglu, 2021). What is sure is that exposure to AI technologies is highly heterogeneous across occupations, sectors, countries and regions (Georgieff and Hye, 2022). Therefore, structural change is a rather predictable outcome implying, as seen with the previous ICT waves (Autor, 2015), sectoral reshuffling, occupational transitions and reskilling processes. The magnitude and distribution of the associated socio-economics costs will depend on the speed of the AI diffusion process and the relative resilience of the agents involved.

In this context, empirical research is struggling to provide policy-makers robust evidence that could prove useful in tailoring AI-related policies and regulations (Acemoglu, 2021; Autor, 2022). As the diffusion of AI gains momentum and its potential applications flourish across sectors (e.g., education, finance, legal counselling, insurance), attempts to assess its employment impact multiply (among others, see Brynjolfsson et al. 2018; Felten et al. 2018, 2021; Webb, 2020; Acemoglu et al., 2022; Albanesi et al., 2023).⁴ Given its complex multi-purpose nature, measuring AI and, even more so, the impact it may have on employment is a challenging endeavour. To complicate things further, a large share of AI-related tasks is fundamentally defined by tacit knowledge, organisational skills and procedural routines, making occupational heterogeneity and related idiosyncratic characteristics a central element, to be analysed at the highest possible level of granularity (on this point, see, among others, Tolan et al., 2021). Moreover, AI technologies are often ‘embodied’ in artefacts and productive environments, with relevant effects on organisation, quantity and quality of employment but with little knowledge and control by the economic agents it may affect (Green et al., 2023).

The existing empirical evidence refers mainly to the US and largely converges towards a positive relationship between AI and employment, particularly concerning high-skill occupations. According to Felten et al. (2021), for instance, the top 5 AI-exposed occupations include ‘genetic counsellors’, ‘financial examiners’, ‘actuaries’, ‘purchasing agents’ and ‘budget analysts’. On the other hand, the least exposed occupations characterised by a prevalence of physical activities, such as ‘dancers’, ‘fitness trainers’, ‘helpers’ (including painters, paper hangers, plasterers), ‘iron and rebar workers’ and ‘textile/garment operators’. At the global level, heterogeneities emerge based on development levels, with high-income countries exhibiting greater exposure and women seemingly more affected than men (Gmyrek et al., 2023). Overall, no strong evidence of labour substitution seems to emerge. There can be at least three reasons to explain this result, though. First, virtually all the available evidence relies on ‘potential’ measures of AI exposure to assess employment dynamics of different occupations. Consequently, different outcomes may arise once data on actual adoption in sectors and firms are available, especially due to the early stage of the AI diffusion process. Second, earlier studies did not account for the joint action of AI and automation technologies, including robots. Indeed, this can be a relevant source of job disruption, as this kind of innovations are mostly aimed at increasing efficiency by reducing the amount of inputs employed in production (Barbieri et al., 2020). Third, no studies have explored the AI-employment nexus at the regional level, accounting for supply, demand and structural factors that may reshape the relationship at stake. All in all, the empirical evidence produced so far is rather US-centric while less is known on other advanced economies, including Europe.

This work aims at filling these literature gaps, providing causal evidence on the employment impact of AI in European regions. Following Albanesi et al. (2023), we crosswalk the occupation-based AI-exposure indicators proposed by Felten et al. (2018, 2019, 2021) to the

⁴ For an earlier review of the literature on the AI-employment nexus, see Barbieri et al. (2020) and Mondolo (2021).

European occupational (ISCO) classification, analysing the dynamics of EU regional employment between 2011 and 2018. We find that AI-exposure has a rather positive impact on employment, while controlling for a rich set of region-level variables, including labour market characteristics, the share of manufacturing and medium-large firms, R&D investment and gross value added. European regions characterised by a relatively larger share of AI-exposed occupations display stronger employment growth, all else being equal and after mitigating for potential endogeneity concerns. This points to the prevalence of complementarity effects on average, although potential cross-country and regional heterogeneities would require further investigation as differentiated patterns could be in order. The situation changes when the joint impact of AI technologies and robots is taken into account. Even if not confirmed across all model specifications, the relationship turns negative lending support to the hypothesis of a labour-saving impact of AI when associated with automation technologies.

The article is organised as follows. Section 2 provides a brief review of the literature, highlighting key issues and systematising the available empirical evidence by unit of analysis and AI-exposure measurement approach. Section 3 spells out our main research question, while data and descriptive evidence are provided in Section 4. Econometric strategy and results are illustrated in Section 5, while Section 6 concludes discussing policy implications, limitations of the analysis and avenues for future research.

2. THE AI-EMPLOYMENT NEXUS: KEY ISSUES AND BACKGROUND LITERATURE

Although relatively recent, the corpus of empirical literature focusing on the AI-employment nexus is large enough to be distinguished according to the adopted unit of analysis and related approach to measure the potential/actual penetration of such technologies.

A significant group of contributions rely on occupation and ability-based indicators to assess the relative ‘AI exposure’ (Felten et al., 2018), i.e., the likelihood that an occupation will come into contact with, be assisted or replaced by AI given the characteristics of the tasks performed and the underlying abilities. In line with the literature studying the employment impact of ICTs distinguishing occupations according to the degree of ‘routineness’ of their tasks (Autor et al., 2003) or automation probabilities (Frey and Osborne, 2017), this stream of works starts by ranking jobs considering the importance and prevalence of abilities that occupations ‘share’ with AI. Brynjolfsson et al. (2018) focus on advancements in machine learning (ML) technologies, which are at the basis of virtually all AI applications. Relying on the rubric evaluating task potential for ML proposed in Brynjolfsson and Mitchell (2017), the authors introduce a task-based measure of ‘Suitability for Machine Learning’ (SML) linking it to 18,156 tasks included in the O*NET⁵ database. Their key results are summarised as follows: i) most of the occupations included in O*NET display at least some SML tasks; ii) only few of them turn out to be fully replaceable by AI technologies; and iii) redesign of job task content is often required to employ such technologies. Another occupation-based measure is proposed by Felten et al. (2019, 2021), refining the indicator proposed in an earlier contribution (Felten et al., 2018). These authors link occupation-level O*NET abilities to the Electronic Frontier Foundation dataset (EFF) developed in the context of the ‘AI Progress Measurement initiative’. Scores are then assigned by matching 10 AI applications (abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modelling, translation, and speech recognition) to human abilities

⁵ The O*NET program is the US primary source of occupational information. Central to the project is the O*NET database, containing information on hundreds of standardized and occupation-specific descriptors. The database is continually updated by surveying a broad range of workers from each occupation.

included in O*NET. The matching is realized by administering a questionnaire to 2,000 individuals, reached with the Amazon's Mechanical Turk (mTurk) web service. Interviewees are asked whether AI applications are related to or could be used for each of the 52 abilities listed in the O*NET. Felten et al. (2021)'s AI occupational exposure scores (i.e. AIOE) point to white collars as the most exposed occupational group, being however silent regarding the likelihood of a complementary/substituting effect.

The only attempt to apply the Felten et al. (2021)'s occupation-based methodology to assess the employment impact of AI on the European economy is the one by Albanesi et al. (2023). Carrying out a crosswalk, analogous to the one upon which our analysis is based, to link the O*NET-based AIOE to European 3-digit occupations⁶, these authors find that, in Europe, employment shares tend to increase in occupations more exposed to AI. The evidence is particularly strong for those occupations characterised by a relatively higher proportion of younger and skilled workers. Focusing on 16 European countries over the period 2011- 2019, Albanesi et al. (2023) argue that, despite country-level heterogeneities do matter, particularly concerning differences in terms of pace of technology diffusion, education levels, product market regulation and employment protection laws, there is no EU country where the share of the most AI-exposed occupations tends to decline

In a more recent effort, Felten et al. (2023) updated their indicator to isolate only advances in Language Modelling (LM) – i.e., the AI technology which is key for the development of 'generative' applications as GPT-4 - and thus verifying if and to what extent such a specific technological development may have a peculiar impact on employment. To do so, the AIOE undergoes a weighting procedure allowing to order occupations according to the number of abilities that are related to LM disregarding the other AI-related abilities included in the original indicator. Although most of the top-exposed occupations can be also found in the list provided in Felten et al. (2021), some relevant 'new entries' are worth mentioning. Among the top occupations exposed to LM AI there are telemarketers and a variety of post-secondary teachers, such as English language and literature, foreign language and literature, and history teachers. Concerning the distribution of the LM AI indicator across industries, sectors displaying the highest values are, as expected, legal services and securities, commodities and investments.

Focusing on occupations but adopting a different approach, Gmyrek et al. (2023) assess the employment impact of Generative Pre-Trained Transformers (GPTs). Unlike Felten et al. (2021, 2023), these authors make use of Chat GPT-4 to estimate task-level scores of occupation exposure to AI. This ranking is then applied to quantify the impact of AI on employment and job quality, by country and income group. According to their estimations, only occupations related to clerical work are highly exposed to AI, with 24% of clerical tasks considered highly exposed and an additional 58% with medium-level exposure. Concerning other occupational groups, the greatest share of highly exposed tasks ranges between 1% and 4%, and medium exposed tasks do not exceed 25%. As a result, they discard the hypothesis of massive substitution, pointing instead to complementary effects that are concentrated among white collars and high-skilled occupations. A similar analysis is carried out by Elondou et al. (2023), who combine experts' opinion and GPT-4 classifications to quantify the impact of GPTs on the US labour market. Merging task-level information stemming from the O*NET and employment data drawn from the Bureau of Labour Statistics (BLS) referring to the years 2020 and 2021, these authors reveal that around 80% of the US workforce could have at least 10% of their work tasks affected by GPTs, while almost one fifth of occupations could have up to

⁶ The O*NET repository uses SOC occupational codes used in the US, while EU member states follow the International Standard Classification of Occupations (ISCO) to classify occupations. 3-digit ISCO codes are referred to as sub-minor groups. See also Section 4.

50% of their tasks impacted. Confirming previous evidence, the highest level of exposure is concentrated at the top of the occupational distribution, i.e., high-skilled and high-income workers.

Despite being very useful for characterising occupations according to their relative AI exposure, occupation and task-based measure do have limitations. First, this type of indicators provides a proxy of ‘potential’ AI exposure, remaining silent on whether such technologies are actually employed – along with the ‘how, when and where’ – by firms. Second, and relatedly, these indicators lack any information about industry- and firm-level technological heterogeneities. The latter, in turn, may be crucial both to more precisely characterise what AI technologies actually are, beyond linking AI applications and human abilities; as well to consider industry and firm level features (e.g., technological capabilities, distribution of complementary goods and competences) which may affect pace and distribution of the same technology. In an attempt to account, jointly, for technological and occupational heterogeneities, Webb (2020) developed an indicator tracking co-occurrence of verb-noun pairs in the title of AI patents and O*NET tasks. This way, he obtains a measure which considers, at the same time, technological choices of the firm (and fine-grained characteristics of specific AI technologies) as illustrated in patents; and task/abilities-related characteristics of occupations, as reported in O*NET. According to Webb (2020)’s results, AI is more likely to affect skilled and older workers than previous innovation waves, e.g., robots or software. The robustness of such a mixed patent-occupation based AI exposure measure, however, is partly undermined by the fact that the titles of patents do not contain a full description of the underlying technology. No less relevant, restricting co-occurrence to verb-noun pairs increases the risk of false positives.

Among the few firm-level studies analysing the employment impact of AI, there is a contribution by Damoli et al. (2023). These authors rely on a sample of 3,500 front-runner companies, stemming from the Orbis BvD database, which patented AI-related inventions over the period 2000–2016 (data are drawn from the PATSTAT database). The coefficient associated to AI patents is always positive and significant, despite relatively small in terms of size, pointing to a moderate positive employment impact of AI patenting (with a short-term elasticity of about 3-4%). Such a ‘labour-friendly’ effect is paralleled by an analogously positive and significant effect of other (non-AI) firm innovation activities. These findings confirm the employment-friendly nature of product innovation in general, provide novel firm-specific evidence on emerging AI technologies. However, it’s important to note that patents are not a proxy of actual AI adoption but rather a partial measure of product innovation, as not all innovation activities are patented. Additionally, as the study focuses solely on patenting companies, it remains silent on the net aggregate effect of actual AI adoption, as it fails to account for aspects such as ‘business stealing’ (see Calvino and Virgillito, 2018).

Another way to look at the relationship between AI technologies and employment is to use job-posting data. In a seminal work, Acemoglu et al. (2022) rely on Burning Glass Technologies data, which provide wide coverage of firm-level online job postings, linked to SOC occupational codes to assess the relative penetration of AI technologies at the establishment level in the US. To quantify the degree of firm-level AI exposure, three different definitions are employed, namely those proposed by Brynjolfsson et al. (2018), Felten et al. (2021) and Webb (2020). The authors do not find any clear employment effect of AI at the industry or occupation-level. Instead, some evidence of a re-composition effect towards more AI-intensive jobs seems to emerge. Overall, the lack of effects at the industry and occupation level is attributed to the relatively limited diffusion of AI technologies and, relatedly, of niche-level nature of adoption. In addition, no evidence of any direct complementarity between AI job posts and non-AI jobs arises, hinting at a prevalent substitution effect and workforce re-

composition, rather than productivity enhancement after AI adoption. While online job vacancies offer a rich data source, caution is necessary when it comes to their representativeness of the overall labour demand as they tend to be biased toward specific occupations, industries and countries.

A systematic overview of different AI proxies, along with the main findings and their limitations, is available in Table A1 in Appendix A.

It is not only about cognitive activities and services, however. AI is going to increase capabilities and scope of a number of automation technologies, including robots (Agrawal et al. 2019). This connection between AI and digital technologies more broadly is also in line with the claim by CIIP (2022) that digitalization, particularly in the industrial domain, is not so much about new tangible technologies but manifest itself in the integration of existing technologies stemming from the ‘physical’ and the ‘ICT’ world. As robots and other machine become ‘smarter’, namely capable to learn and adjust their ‘behaviour’ in ever more complex productive contexts, the opportunities for process automation and related efficiency gains grow too (Barbieri et al. 2020). If this is the case, a wave of AI-induced job destruction in manufacturing could be on the way (Autor et al., 2022). Unless, the same efficiency gains translate into compensation mechanisms capable to offset the not so remote possibilities of job destruction (for a discussion of these mechanisms, see Calvino and Virgillito, 2018). It is hard to say, at present, which is the most likely scenario, as no robust empirical evidence on the impact of AI on manufacturing seems to be available.

An important role could be also played by heterogeneously distributed supply, demand and structural factors that are likely to influence AI diffusion patterns (Reljic et al., 2021). Countries, sectors and regions characterised by a large share of knowledge intensive services are those where deployment of new technologies could be most rapid, especially with regard to AI applications capable of complementing human tasks. In contrast, where medium-technology sectors prevail and firms are predominantly adopters of technologies provided by external suppliers, it is more likely to observe labour substitution phenomena (Calvino and Fontanelli, 2023). On the other hand, the presence of labour market institutions aimed at protecting workers against dismissal could slow down an eventual process of AI-driven destruction of occupations (Reljic et al., 2023). The spread of AI could also produce a polarization in Global Value Chains (GVCs), rewarding the industries and firms that control the most strategic AI services and penalizing, in terms of the amount of value captured along the VC, the broad mass of those depending on it. These are all elements which may ultimately contribute to determine the net (direct and indirect) employment impact AI, as well as its distribution in space and time.

Four main take-home messages emerge from this brief literature review. First, much more needs to be said about the employment impact of AI, as the latter is in its initial stage of diffusion and novelties in terms of applications and potential impact on jobs quality and quantity are about to come. Second, albeit the available indicators represent a very useful base to assess exposure and (potential) employment impact of AI, further refinements, considering both the characteristics of occupations and actual business decisions, as well as the specifics of industries regarding the adoption process, would be of great advantage. Third, spatial specificities (e.g., characteristics of regions, provinces, cities or local labour markets) must be adequately taken into account, given the weight that these elements may have in determining the diffusion of AI and its impact on the labour market. Fourth, more evidence is needed regarding the intertwining of AI and automation technologies in manufacturing (e.g., robots). This is of particular relevance in the European case where manufacturing plays an important role and the diffusion of AI could be intertwined with phenomena having great transformational potential, such as the transition to electric in the energy and automotive sectors. We have now

reviewed the relevant literature and highlighted the key issues concerning the AI-employment nexus. In what follows, we spell out our research questions aiming to fill some of the abovementioned literature gaps.

3. THE IMPACT OF AI ON EUROPEAN REGIONS: POLICY CONTEXT, RESEARCH QUESTIONS AND CONTRIBUTION

As one of the largest markets in the world, characterised by a substantial share of knowledge intensive business services as well as of high-tech manufacturing industries, Europe is witnessing an increasing diffusion of AI technologies (European Commission, 2018). As alarms on the potential risks of AI started to spread globally, the European Commission (EC) reaction followed promptly. Building on its 2021 ‘European AI strategy’⁷, the EC published a ‘Regulatory framework proposal on AI’, aiming at addressing AI-related risks, including those connected to employment and labour markets. Despite not explicitly mentioning the threat of a massive labour disruption, the EC highlights a number of risks specifically created by AI applications pointing to the need of a European governance that, among other things, should carry out ‘a conformity assessment before the AI system is put into service or placed on the market’.⁸ On the other hand, Europe does not want to fall behind the AI technological frontier concerning the development of knowledge, competences and incremental innovation. Because of its general-purpose nature (Agrawal et al., 2019), AI is likely to trigger a myriad of transformative innovations, in manufacturing and services alike. As a result, economies lagging behind technological leaders risk increasing their dependence on external suppliers, with potentially negative consequences for growth and employment. European economic policy faces a double challenge, then. First, staying on the AI frontier: policy instruments are being developed to promote the diffusion of AI technologies and the development of the skills needed to seize their potential. Second, mitigating socio-economic risks and preventing the most serious adverse effects related to the diffusion of AI, including those that may affect the labour market.

Given the growing attention of EU policy makers, one would expect a blossoming of studies providing robust quantitative evidence concerning the effects of AI on the European labour markets. This not the case, though. While contributions on the US abound (see above), a certain darkness seems to characterize the European context. In particular, three are the main literature gaps worth underlining. First, while a number of studies provide evidence of the impact of AI exposure on the dynamics of US occupations considering, in a detailed manner, abilities and tasks that can be complemented/replaced by AI technologies, no study provides analogous evidence for the EU. A notable exception is the work by Albanesi et al. (2023), investigating the evolution of AI-exposed occupations/sectors. These authors, however, do not provide any causal evidence leaving ground for more exploration. Second, despite one of the key features of the European economy is its strong between and within-country structural heterogeneity concerning macroeconomic performance, innovativeness and labour force composition – it is not by chance that a large chunk of EU funds are destined to cohesion policies aimed at reducing macroeconomic and structural gaps across EU member states and regions (Darvas et al, 2019; Landesmann and Stöllinger, 2020) – no studies have so far analysed the distribution of AI exposed occupations and related employment dynamics taking into account such heterogeneity. Third, while the entanglement between AI and automation technologies is

⁷ Detail on the European AI strategy and on subsequent actions undertaken by the European Commission are documented at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2018:237:FIN>.

⁸ The Regulatory framework proposal on artificial intelligence of the European Commission can be accessed here: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>.

widely acknowledged, empirical evidence regarding their joint impact on employment is still missing.

In this work, we enrich the approach of Albanesi et al. (2023) studying the impact of AI on EU employment by adopting a regional perspective and relying on a novel IV strategy. AI exposure is measured by mapping Felten et al. (2021)'s AIOE indicator into 3-digit ISCO occupations and then to 2-digit NUTS2 regions (see Section 4 for a detailed description of the methodology). A large number of dimensions are considered, accounting for regional heterogeneities as well as for supply, demand and structural factors likely to affect the relationship at stake. Moreover, we include a regional measure of robot intensity, testing whether and to what extent the combined action of AI technologies and robots has a peculiar impact on European employment.

Our main research questions run as follows. First, we assess the relationship between occupational AI-exposure and employment dynamics in European regions. Theoretical expectations about the direction of the relationship are difficult to establish ex-ante. All other things being equal, both a complementarity as well as a substitution effect may prevail. If AI operates as a process innovation, increasing efficiency by reducing the amount of labour employed, a negative relationship can be expected. Conversely, if new tasks are required or if organizational innovation is stimulated enhancing the contribution of the most AI-exposed occupations, the effect may go in the opposite direction. Accordingly, the first research question (RQ1) reads as:

RQ1: Does AI increase or destroy European jobs?

Once tested RQ1, relying on different AI indicators – so to account whether different approaches to measure AI lead to heterogeneous results concerning its employment impact – and estimators, we move our focus to the linkage between AI and automation technologies, namely robots. As argued, AI technologies may improve the automation potential of robots positively affecting their performance and, hence, their labour-saving potential; as well as enlarging the set of tasks that such machines are capable to carry out. On the other hand, robotization and, even more so, ‘smart robotization’ is hardly a homogeneous process, as EU countries and regions are characterized by significantly heterogeneous productive and technological capabilities which are likely to influence pace, distribution and, ultimately, the employment impact of such process (Reljic et al., 2023). Therefore, we consider the joint effect of AI and robots on employment, considering all the region-specific structural factors that may shape the relationship. In principle, ex-ante expectations are relatively easier to establish in this case, given the fact that intelligent robots are in most cases aimed at increasing the efficiency of production processes, giving rise to labour-saving effects. Nonetheless, if intelligent robots help improving organisational efficiency this could have a larger and persistent impact on economic performance, potentially compensating the labour-saving effects. As a result, our second research question (RQ2) is openly worded as the previous one and reads as follows:

RQ2: What is the employment impact of the AI-robots combination on the employment dynamics of European regions?

The following sections illustrate data and indicators; offers a comprehensive assessment of AI occupational exposure across EU's occupations and regions; introduces the empirical strategy adopted to address RQ1 and RQ2 and provides the main results.

4. DATA AND DESCRIPTIVE EVIDENCE

4.1. Data description

To investigate the employment impact of AI in European regions, we combine data from several sources over the period 2011-2018⁹. Our sample includes 202 NUTS-2 regions from 23 European countries¹⁰. The NUTS-2 classification has undergone changes in some countries over time due to a combination of administrative and statistical-related factors (Eurostat, 2020). These amendments in NUTS-2 codes pose challenges for regional panel analysis. Thus, to ensure the consistency of our data throughout this period, we simply recode regions that have only changed the name of NUTS (e.g., French regions), while in other instances, we were compelled to aggregate NUTS-2 regions (e.g., LT00-02, IE01-06, DE40-42)¹¹.

Artificial Intelligence. Regarding AI, we draw on earlier works by Felten et al. (2018, 2019, 2021), who proposed and made available the indicator of AI occupational exposure. This indicator links various AI applications - abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modelling, translation, speech recognition and instrumental track recognition - to 52 workplace abilities (e.g., mathematical reasoning, speech recognition, written comprehension, originality, body coordination) using the mTurk web service survey.¹² Occupational exposure to AI (AIOE) is constructed by weighting the ability-level exposure to AI with their prevalence and importance within each occupation:

$$AIOE_k = \frac{\sum_{j=1}^{52} A_{ij} * L_{jk} * I_{jk}}{\sum_{j=1}^{52} L_{jk} * I_{jk}} \quad (1)$$

where A_{ij} stands for the ability-level AI exposure, calculated as a sum of relatedness scores across ten AI applications for each of the 52 abilities; L_{jk} and I_{jk} represent prevalence and importance of each ability (j) within each occupation (k).

This means that occupations characterised by a higher prevalence and importance of abilities classified as highly exposed to AI exhibit a relatively higher exposure to AI, and vice versa. Under the assumption that AI-related workplace abilities of US occupations are similar to those characterising their EU counterparts¹³ (Albanesi et al., 2023), we map Felten et al.'s AIOE available at the six-digit SOC occupations into the International Classification of occupations (ISCO-08) at the four-digit level, ultimately collapsing at the three-digit ISCO level (126) by calculating the mean exposure across occupations.

As focus of the analysis are European regions, we construct a regional AI exposure (AIRE) indicator following the approach suggested by Felten et al. (2021). To this end, we combine the occupational AI exposure (AIOE) with the occupational distribution (ISCO 3-digit) of employees within regions from the EU LFS, as follows:

⁹ Our sample starts in 2011 due to a major revision of ISCO (International Standard Classification of Occupations), when ISCO-88 was succeeded by ISCO-08, which makes comparisons before and after 2011 impossible.

¹⁰ Austria, Belgium, Bulgaria, Czech Republic, Germany, Denmark, Estonia, Spain, Finland, France, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Netherlands, Poland, Portugal, Romania, Sweden, Slovakia, United Kingdom.

¹¹ For example, in the case of Lithuania, the change in classification in 2013 resulted in the separation of Lithuania into two NUTS-2 regions.

¹² A detailed methodology is provided in Felten et al. (2021).


¹³ An analysis showing the US-EU within-occupation similarities in terms of digital task content has been recently provided by Gschwent et al. (2023).

$$AIRE_{ijt} = \sum_{k=1}^{126} \frac{EMP_{kijt}}{EMP_{ijt}} * AIOE_k \quad (2)$$

where EMP_{kijt} denotes the number of employees in occupation k in region j in country i and year t , while EMP_{ijt} stands for the total number of employees in region j in country i in year t . Thus, the first term denotes the employment share of each of the 126 ISCO 3-digit occupations in region r in year t , while $AIOE_k$ corresponds to occupational AI exposure, as defined in Equation 1¹⁴. The $AIRE$ indicator is normalised to have a zero mean and unit standard deviation in each year, representing relative AI exposure across regions.

The main limitation of Felten et al. (2021)'s indicator is that measures crowd-sourced opinions on relative exposure to AI technologies. It sheds light on which occupations, industries, countries and regions are most likely to be affected by advancements in AI rather than on actual adoption. Therefore, to check the robustness of our regional AI-exposure measure we rely on Eurostat ICT business survey data, recording the percentage of enterprises employing AI technologies. To allow for comparison, we calculate AI exposure at the country level by combining the AIOE data with the occupational distribution (ISCO 3-digit) of employees within each country. Figure 1 presents a scatter plot displaying the joint distribution of country level AI exposure and Eurostat's AI adoption indicator. Despite some noise, the positive correlation (correlation coefficient: 0.57) suggests that occupational AI exposure is likely related to AI adoption. Specifically, countries with higher degree of AI exposure tend to have a greater share of enterprises that adopt AI technologies.

Figure 1. AI exposure and share of enterprises using AI technologies by country



Source: Authors' elaboration based on Felten et al. and Eurostat's ICT business survey

Robots. We rely on International Federation of Robotics (IFR) database, which provides information on the robot stock and new instalments from 1994 to 2019 at the country-industry level in manufacturing sector. As extensively discussed in Fernández-Macías et al. (2021), the IFR data comes with an important caveat: they do not account for variations in robot quality across different industries, countries and time periods. Nevertheless, the IFR remains the most reliable source of data upon which empirical literature on the employment effects of robots has

¹⁴ Note that we introduce dynamics by allowing our indicator $AIRE$ to vary over time, reflecting the changes in the occupational distribution.

flourished (see Acemoglu and Restrepo, 2020; Fernandez-Macias et al., 2021; Petit et al., 2023; Reljic et al., 2023, among others).

In line with earlier studies (Dauth et al., 2021; Petit et al., 2023, among others), to construct an indicator of robot density at the regional level we assume that the distribution of robots within an industry is uniform across regions within a country, conditional on the industry-region employment shares. To this end, we combine industry-level data on robot stock with employment distribution across the 2-digit NACE Rev.2 industries within regions, as follows:

$$Robot\ stock_{rt} = \sum_{j=1}^J \frac{EMP_{jrt}}{EMP_{rt}} * Robot\ stock_{jt} \quad (3)$$

where EMP_{jrt} denotes the number of employees in industry j in region r in year t , EMP_{rt} stands for the total number of employees in region r in year t and J represents the complete set of industries for which robot stock ($Robot\ stock_{jt}$) is available. The remaining variables used in the empirical analysis and their sources are listed in Table A2 in Appendix A.

4.2. Descriptive evidence

In order to get a first impression of the data and its characteristics, we list the ten ISCO 3-digit occupations (sub-minor group) with the highest (Table 1) and the lowest AIOE scores (Table 2). The most striking feature of the top ten list is the dominance of high-skilled workers, predominantly stemming from the group of professionals (ISCO major group 2). This is fully in line with the ranking in Felten et al. (2021) for US occupations which also includes many professions such as financial examiners, budget analysts or mathematicians, which are all represented also in the list in Table 1. The relatively high AI exposure of high-skilled worker is also in line with the wide-spread perception that labour market effects of digitalisation – or industry 4.0 – will affect not only, and maybe even not most strongly, blue-collar workers as was the case with automation (Cirillo et al., 2021). In this context, it is noteworthy that the high score of professionals is not built in the construction of the AI index because the AI index is considered ‘neutral’. It means that a high exposure to AI does not necessarily imply that occupations will be substituted by AI technologies. Rather, they can also score high if they are complementary to AI, or as Felten et al. (2021) point out, the methodology for calculating the AI index is agnostic as to whether AI substitutes for or complements occupations (respectively the abilities needed in occupations).

Table 1. Top 10 most AI exposed occupations

Ranking	ISCO 3-digit	ISCO 3-digit label	AIOE
1	212	Mathematicians, Actuaries and Statisticians	1,66
2	241	Finance Professionals	1,63
3	261	Legal Professionals	1,63
4	242	Administration Professionals	1,46
5	431	Numerical Clerks	1,45
6	231	University and Higher Education Teachers	1,44
7	411	General Office Clerks	1,44
8	122	Sales, Marketing and Development Managers	1,42
9	251	Software and Applications Developers and Analysts	1,40
10	233	Secondary Education Teachers	1,38

Source: Authors’ elaboration based on Felten et al. (2021)

This characteristic of the AIOE presumably also explains why, along with various professionals, there are also some medium-skilled occupations present in the list of top-ranking

occupations such as numerical or general office clerks. In other words, the rationale of this AI-exposure index is partly different from the one characterising other well-known occupation-based indices, such as the routine-task intensity (RTI) index (Autor et al., 2003) or the offshoreability index (Acemoglu and Autor, 2011). The latter include a clear task-related occupational hierarchy concerning replacement risks vis-à-vis complementarity (i.e. occupations characterized by a relatively larger share of routinary task are considered more at risk of technology-driven substitution), while no such hypotheses are made to build the AIOE.

At the other end of the spectrum (Table 2), we find mostly low-skilled occupations, in particular elementary ones. Common traits include the relatively lower technological intensity of their tasks, which are mostly manual and physical but not necessarily repetitive (i.e. routinary). Several of these occupations have little or no relevance anymore in most EU member states, which is particularly true for subsistence farmers or subsistence fishers and hunters.

Table 2. Bottom 10 least exposed occupations


Ranking	ISCO 3-digit	ISCO 3-digit label	AIOE
126	931	Mining and Construction Labourers	-1,78
125	631	Subsistence Crop Farmers	-1,74
124	912	Vehicle, Window, Laundry and Other Hand Cleaning Workers	-1,69
123	911	Domestic, Hotel and Office Cleaners and Helpers	-1,64
122	713	Painters, Building Structure Cleaners and Related Trades Workers	-1,60
121	634	Subsistence Fishers, Hunters, Trappers and Gatherers	-1,53
120	932	Manufacturing Labourers	-1,51
119	921	Agricultural, Forestry and Fishery Labourers	-1,51
118	633	Subsistence Mixed Crop and Livestock Farmers	-1,47
117	712	Building Finishers and Related Trades Workers	-1,44

Source: Authors' elaboration based on Felten et al. (2021)

Looking at the distribution of the AI exposure at the level of European region (NUTS-2 regions), it is not surprising to find a high AI exposure in many high-income regions, including many capital regions. Examples include, Ile de France (Paris region), Vienna, Berlin, Warsaw metropolitan area, Prague and many more. In other cases, larger areas of the country are identified as having a high AI exposure, such as North Holland, South Holland and the Utrecht region in the Netherlands or Southern Sweden and Lower Bavaria in Germany. All these regions, however, are also high-income regions, both in an EU-wide comparison and a national comparison.

In contrast, in the Southern periphery (Spain, Italy) and the Eastern periphery (Romania, Bulgaria) there are numerous regions with very low levels of AI exposure. These patterns coincide well with other measures for implicit technological capabilities across Europe, such as, for example, functional specialisation patterns (Kordalska et al., 2022). These regional patterns of AI exposure and their development over time are crucial as they will be reflected in the econometric analysis.


Figure 2. AI exposure across European regions in 2018



Source: Authors' elaboration based on Felten et al. (2021)

It was already shown that the AI index is related to actual AI adoption (see Figure 1 above). However, it is important to see that the AI exposure index does not show any clear correlation with robot density, a measure that, following the seminal contribution by Graetz and Michaels (2018), has become popular for investigating the employment effects of automation. While regions with a low robot density may have very high or very low AI exposures, those with a very high robot density, for the most part, do not deviate too strong from the average AI exposure of region. This lack of correlation could indicate that the AI exposure and the installation of robots reflect different aspects of technological progress. More specifically, it could be argued that the former is directly related to digitalisation and the 'fourth industrial revolution' (or Industry 4.0), while robots are a manifestation of automation belonging to the 'third industrial revolution' (or Industry 3.0). Be that as it may, both 'generations' of technologies will be considered in the econometric analyses.

Figure 3. AI exposure and robot density by region, 2011-2018



Notes: Values refer to the averages over the period 2011-2018

Source: Authors' elaboration

5. ECONOMETRIC STRATEGY AND RESULTS

The impact of technology can vary significantly across countries, regions, industries, and occupations. Recent studies show that there are structural features likely to affect the capacity of countries and regions to leverage new technologies, including sectoral specialisation, technological capabilities, positioning in the GVCs and skills (Petit et al. 2023; Reljic et al., 2023). Industries and occupations tend to cluster geographically, not only across countries but also within them (Valentini et al., 2023). Hence, regions emerge as a pertinent unit of analysis in this context.

We estimate the effect of AI exposure on employment at the regional level using the following baseline specification:

$$\Delta Y_{ijt} = \alpha_0 + \beta_1 AIRE_{ijt-1} + \gamma * x'_{ijt-1} + \tau_j + \varepsilon_{ijt} \quad (4)$$

where the dependent variable is the annual employment change in region j in country i in year t , so that ΔY_{ijt} refers to the annual change in employment between year t and $t-1$ in region j and country i . $AIRE$ denotes the regional AI exposure, which is our main explanatory variable. The vector x' comprises a set of control variables, τ_j denotes region fixed effects and ε is the error term.

In line with the literature, the regional AI exposure indicator is lagged by one year because it may require some time to manifest its effects (Dosi et al., 2021), while other explanatory variables are also lagged in order to mitigate potential simultaneity bias (Blanas et al., 2019).

As discussed above, the capacity to embrace and reap the benefits of AI technologies is likely to vary across regions according to their characteristics, potentially complementing labour in some areas while substituting it in others.

To begin addressing these regional differences, we include controls for socio-demographic workforce characteristics within regions: the shares of female employees, non-standard

employment, youth workers (aged 15-24), senior workers (aged 55-64) and those with university degree within regions.

Furthermore, recent empirical findings highlight the role of firm size in AI adoption. Rammer et al. (2021), for instance, found that in Germany large firms with at least 1000 employees are nearly ten times more likely to adopt AI compared to the small business (5 to 9 employees). Similarly, AI adoption is unevenly distributed across sectors, with the latest Eurostat ICT business survey suggesting higher AI adoption rates in ICT services and professional business activities. To account for structural heterogeneities across regions that may affect the nexus between AI and employment, we control for the size of manufacturing sector and the share of firms with more than 50 employees. In addition, we also include regional levels of R&D investment and gross value added (proxying heterogeneities concerning aggregate demand).

While our AI indicator appears effective in capturing the occupational exposure to AI technologies (and to some extent AI adoption, see Figure 1), it falls short in accounting for AI's role in the realm of robotics. Indeed, Felten et al. (2019) explicitly acknowledge their focus on 'purely AI technologies,' intentionally omitting consideration of 'how the interaction between advanced AI and robotics technologies affects abilities or occupations.' Consequently, AI occupational exposure is inherently skewed toward cognitive abilities and tasks, as observed in the study by Felten et al. (2021). While this is not necessarily a limitation, it is important to note that this indicator does not account for the fact that AI is also enhancing automation potential in manufacturing industries by making industrial robots more flexible, autonomous and intelligent (IFR, 2022; Soori et al., 2023).

In line with this, we augment the model in Equation (4) by including an interaction between AI exposure and robot density in an attempt to measure AI embodied in robots:

$$\Delta Y_{ijt} = \alpha_0 + \beta_1 AIRE_{ijt-1} + \beta_2 AIRE_{ijt-1} * Robot\ density_{ijt-1} + \beta_3 Robot\ density_{ijt-1} + \gamma * x'_{ijt-1} + \tau_j + \varepsilon_{ijt} \quad (5)$$

Equation 4 and 5 are estimated using OLS and FE, and standard errors are robust to heteroscedasticity and clustered at the country-regional level to account for serial correlation. Results are available in Table 3.

The question that arises is whether we can consider AI exposure as a plausibly exogenous factor in driving employment dynamics in European regions. In their recent work on European sectors, Albanesi et al. (2023) employ Felten et al.'s indicator and argue that AI exposure is 'not that endogenous,' considering that it was originally constructed for the US economy.

To address endogeneity concerns and ideally isolate the component of AI exposure driven solely by advances in AI technologies, we adopt an instrumental variable approach using a new instrument. The approach consists of using the number of daily internet users relative to the total population as an instrument. This choice is grounded on the assumption that regions with a higher number of daily internet users are more likely to be exposed to and adopt AI technologies. This reconciles with the idea that AI technologies have the capacity to continuously 'learn' from large amount of unstructured data, and regions with a larger internet user base are likely to generate more data, thereby creating a conducive environment for fuelling advances in AI performance. At the same time, it is reasonable to assume that daily internet users are not directly influenced by changes in regional employment dynamics, and vice versa.

We attempt to estimate the causal impact of AI on employment, by employing a two-stage least squares, using the instrument discussed above. The results are presented in Table 4. We start off with discussing the Fixed Effects (FE) results (Table 3). Ignoring the pooled OLS estimates (specification 1), we find that the regional AI exposure, AIRE, as our main variable

of interest is positively associated with employment growth. This effect tends to get larger as more control variables are added to the model. Specification 4, for example, which includes structural and labour market controls as yields a coefficient of 2.3 for AIRE. This suggests that a 1 standard deviation increase in the AI exposure of a region is associated with an increase in the employment growth by 2.3 percentage points. The coefficient of AIRE tends to increase further as the model is augmented to include robot density, along with controls for R&D and demand. These results provide an answer to our first research question: on average, AI exposure increases regional employment, implying that the AI-job complementarities dominate the substitutive forces in the AI-job nexus. This outcome is different from the results reported by Felten et al. (2019) concerning the US. In their specification, these authors do not find a significant effect of AI on state-level employment growth for the full set of US occupations for the period 2010-2016. While the specification of Felten et al. (2019) differs from ours in many respects¹⁵, this comparison of results signals that the employment effects of AI across regions could be different in the EU and the US. This seems all the more likely as Albanesi et al. (2023) also find positive employment of AI in a sample of 16 EU countries, which is stronger for high-skilled occupations. Hence, the (very limited) existing evidence on labour market effects of AI in Europe seems to be in line with our regional results. It should be also mentioned that the US state-level analysis yield a positive employment effect of AI for the sub-set of high-skilled occupations.

Table 3. Employment change and AI exposure

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
AIRE	0.263 (0.252)	2.273* (1.201)	2.259* (1.193)	2.287* (1.193)	2.358** (1.196)	2.735** (1.213)	2.669** (1.164)
Robot density					0.378 (0.232)	0.383* (0.212)	0.471** (0.236)
AIRE x robot density						-0.314*** (0.114)	-0.276** (0.113)
Structural controls	No	No	Yes	Yes	Yes	Yes	Yes
Labour market controls	No	No	No	Yes	Yes	Yes	Yes
R&D	No	No	No	No	Yes	Yes	Yes
Demand	No	No	No	No	Yes	Yes	Yes
Time fixed effect	No	No	No	No	No	No	Yes
Constant	1.085*** (0.0838)	1.085*** (1.11e-)	6.047* (3.095)	9.442 (12.26)	7.115 (17.32)	7.130 (17.05)	4.924 (17.36)
Observations	1,413	1,413	1,413	1,413	1,413	1,413	1,413
R-squared	0.002	0.022	0.028	0.063	0.067	0.073	0.111
Number of ID		202	202	202	202	202	202

Notes: The dependent variable is annual employment growth. In all regressions except specification (1), region fixed effects are included. Controls consist of structural variables (average firm size and the share of manufacturing), labour market characteristics (the shares of female employees, non-standard employment, youth workers, senior workers, and those with a university degree), R&D and demand. Standard errors are robust against heteroscedasticity and serial correlation at the regional level (in parentheses *** p<0.01, ** p<0.05, * p<0.1).

Source: Authors' elaboration

¹⁵ A first important difference is that Felten et al. (2019) estimate longer term employment effects stretching over the period 2010-2016, while we investigate year-on-year employment changes. Second, Felten et al. (2019) use occupation-region specific employment changes as the dependent variable, while we use a regionalised version of the AI index (see Section 4.1).

The next step consists of taking into account regions' robot density. While robot density in itself does not seem to affect regional employment growth (specification 5), it does so in the interplay with AI exposure (Specifications 6 and 7). When the interaction term between AIRE and robot intensity is included, the positive coefficient of robot density becomes statistically significant (at least at the 10 per cent level), indicating that robot density, just like AI exposure, is also positively correlated with employment growth. Most importantly, in relation to our second research question, the results suggest an interesting relationship between AI exposure and robot density. While the coefficient of the AI exposure remains positive, the effect is counteracted by the robot density of regions, as indicated by the negative coefficient of the interaction term. This means that the joint AI-robots effect is heterogeneous across regions. More specifically, the combined effect of a 1 standard deviation increase in AIRE would be a 2.7 percentage points increase in the employment growth rate for a region with zero robot density (that is the main effect, $2.669 + (-0.276 \times 0)$). However, for a region with an average robot density (5.7 in our sample), this effect would shrink to 0.96 percentage points ($2.669 + (-0.276 \times 5.7)$). The 'break-even' point for an overall positive employment effect of the joint AI exposure and robot implementation, that is the robot density which would imply no employment effect for a region, is reached at a robot density of 8.7. Given the large dispersion of robot intensities at the regional level (see Figure 3), there are numerous regions for which a negative overall employment effect must be expected. This would concern above all the automotive-oriented regions, including for example Wolfsburg (DE91), headquarter of German carmaker Volkswagen and the region with the highest robot density in the sample.

The economic interpretation of this moderating impact of robot density on the nexus between AIRE and employment growth is that the implementation of industrial robots is also associated with labour shedding in EU regions, which have high exposure to AI (that is positive values of AIRE). This would correspond to a pattern where both Industry 3.0 technologies (industrial robots) and Industry 4.0 technologies (digital technologies) are associated with labour growth. However, in regions with a high robot density, AI exposure is expected to lead to less benign employment outcomes. In such an environment, a high regional AI exposure is expected to be detrimental to job creation, which are presumably routine-task intensive jobs. Hence, there is not a unique answer to RQ2, dealing with the combined employment impact of AI exposure and robot density, as these effects can be very heterogeneous across EU regions.

In order to make sure that our results are not driven by reverse causality or some other confounding factors, we repeat the estimations with the IV approach described above (Table 4). The IV results question the effects of the robot density¹⁶, and that of the interaction terms but confirm the relevance of the regional AI exposure, which is the main variable of interest. In fact, the coefficient of AIRE remains positive and statistically significant across all specifications. And since neither robot density, nor the interaction term seems to matter, it is preferable to rely on a simpler model such as the one in specification 4, which yields a coefficient of 2.1 for the AIRE variable, which is slightly lower than in the fixed effect regression.

Hence, based on the IV results, we could conclude that high AI exposure is beneficial for regions in terms of employment creation. That is, the answer to RQ1 remains the same, while RQ2, in light of the IV results, becomes less relevant because the AI-robots-nexus for regional employment development cannot be confirmed. Overall, these findings are confirmed even when we employ different versions of AI exposure indicator, as presented in Appendix B.

¹⁶ However, these results should be interpreted with caution, as instrumenting the interaction raises concerns about weak identification in specification (6).

Table 4. Employment change and AI exposure - IV estimates

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>First stage</i>							
Daily internet	0.0396*** (10.41)	0.0267*** (7.00)	0.0244*** (6.08)	0.0193*** (5.40)	0.0175*** (3.45)	0.0166** (2.79)	0.0106** (2.58)
Daily internet						0.000118 (0.26)	0.00108*** (4.05)
<i>Second stage</i>							
AIRE	0.529* (2.19)	1.105* (2.35)	1.584** (2.84)	2.104** (2.62)	6.747* (1.97)	6.273 (1.92)	3.022* (2.18)
Robot density				0.0525 (1.44)	0.141 (1.54)	0.0943 (1.33)	0.0964 (1.93)
AIRE x robot						-0.267 (-1.44)	-0.144 (-1.42)
Structural controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Labour market	No	No	Yes	Yes	Yes	Yes	Yes
R&D	No	No	No	Yes	Yes	No	No
Demand	No	No	No	Yes	Yes	No	No
Time fixed effects	No	No	No	No	Yes	Yes	Yes
Country fixed	No	No	No	No	Yes	Yes	Yes
Kleibergen-Paap	108.307	48.948	36.955	29.200	11.895	5.462	18.623
Constant	1.101*** (12.33)	2.110* (2.21)	4.449 (1.44)	8.142 (1.80)			
N	1322	1322	1322	1322	1322	1322	1322

Notes: The dependent variable is annual employment growth. Controls include structural variables (average firm size and the share of manufacturing), labour market characteristics (the shares of female employees, non-standard employment, youth workers, senior workers, and those with a university degree), R&D and demand. The AI exposure index (AIRE) is instrumented with the share of daily internet users relative to the population in the same regions. First-stage results are reported in the first two rows. In columns 5-7, the equation is estimated using deviations from means, and column 7 reports regression results weighted by the average number of employees for each region. Standard errors are robust against heteroscedasticity and serial correlation at the regional level (t-statistics in parentheses, * p < 0.05, ** p < 0.01, *** p < 0.001). Source: Authors' elaboration

6. CONCLUSION AND POLICY IMPLICATIONS

This paper adds to the nascent literature on the employment effects of the accelerating diffusion of AI technologies throughout the European economy. Mirroring the situation in the technological domain itself, Europe is also trailing behind the US when it comes to the literature discussing the consequences of AI for labour market outcomes. This lag is owed primarily to better data availability for the US economy, where researchers can draw on publicly available data on occupational profiles as well as employment and wage data at a very granular level.

Country and region-specific empirical analysis are indispensable though because the effects of technology adaption and diffusion are likely to differ significantly from one location to the other. This is related to a more general aspect of technological change which is 'Kranzberg's law' which holds that "*Technology is neither good nor bad; nor is it neutral*" (Kranzberg, 1986, p. 545). The point the technological historian put forward in his writings is that the consequences of technology depend not only on its technical features but on the societal and temporal context. Hence, chances are that there is not one answer to the question of what AI means for jobs that researchers are so eager to answer. Rather, answers can only be partial, specific to the locations and period in time analysed.

Applied to the research questions put forward in this paper this means that there is no universal answer to the question, if – and in the affirmative – how AI will affect labour markets in general. We do not know whether the ‘future of work’ will resemble the rosy world envisaged by Keynes (2010[1931]) in his essay on the *Economic Possibilities for Our Grandchildren*, where new technologies – in our context AI – leads to such massive increases in productivity, essentially freeing society from scarcities allowing people to indulge in science, arts and philosophy. We would see this as the positive or ‘Star Trek’ scenario. Things could play out very differently, though. As sketched by Leontief (1983), humans may face the same destiny as horses in their function as ‘workforce’, meaning they will just not be needed anymore, apart for some curious nostalgic purposes such as tourist entertainment showing the ways of the past, sports, or the circus. This ‘Death of the Workhorse’ scenario in which men lose the race against machines has become popular among economists with many facets of it. This prevalence of the pessimistic view is the results of two characteristics and their interactions, one related to AI, the other to the current economic paradigm. Leontief’s point that the horse as a production input became obsolete because the steam engine outcompeted the workhorse in its core competences – physical strength and stamina. Likewise, AI in many work contexts now outcompetes humans in cognitive tasks and also seriously challenges it in (simple) social interactions. This feeds into the ‘this time is different’ narrative which often comes with a Luddite undertone but in principle could be counteracted by the fact that, unlike the horse, humans themselves can decide whether, to what extent and for which purposes, they want to introduce the new technologies which become available. However, this safety valve for meaningful labour risk being undermined by the current economics paradigm, which induces firms to use AI not primarily to find valuable solutions to societal challenges but to maximise shareholder value. The latter typically involves the replacement of labour with AI algorithms (see Acemoglu and Johnson, 2023)¹⁷ – typically in combination with robots and other machines. This is what we find in our quantitative analysis of regional employment growth across EU regions. Regions with a high robot density, in interplay with a high AI exposure, may end up losing jobs. This underlines the fact that labour market effects emanating from AI may be very heterogenous across EU regions. However, overall, the results for EU regions indicate positive employment effects for the average EU region stemming from AI exposure. This is encouraging news and could suggest that the EU is still more on heading towards a Star Trek scenario. However, this is only one possible interpretation of the positive AI-employment growth nexus in EU regions.

In fact, our results – reassuring as they are – must be interpreted with great care for a number of very different reasons. An important caveat is that by focusing on employment, only one aspect of labour markets is captured, while neglecting other relevant dimensions of work, in particular working conditions. It could very well be that while new jobs are created, these jobs are of a poor quality, which are low in terms of skill requirements but above all lack a sense of meaning. The working conditions of many these newly created jobs could be described as underpaid, isolated, where workers are stuck at home in front of their computers with work and leisure time getting increasingly blurred. While the results of this first regional analysis have little to say on this, the fact that AI exposure is skewed towards high-skilled jobs and that other studies (Albanesi, et al., 2023; Felten et al., 2019) found that AI leads to employment growth primarily for high-skilled labour may question this prediction. At the same time, it cannot be ruled out that even jobs of high-skilled workers are getting increasingly monotonous and meaningless.

¹⁷ The situation is aggravated by the fact that in the US, a major innovator in this domain, AI technologies are under the control of a few powerful firms.

In addition to the omission of job quality, there are a number of important methodological limitations which have to be kept mind. First, quantitative work of the kind undertaken here is bound to make inferences from the past onto the future. While this is legitimate, the predictions emerging from such an undertaking may be less accurate and reliable when they deal with a potentially disruptive technology, i.e, AI, which by definitions marks a break in the technological trajectory. Secondly, the diffusion of AI in the economy may still be too limited so that its macroeconomic consequences (such as employment growth) are hard to identify in the data. A related factor, which implies a downward bias in the effect of AI exposure on employment as well, is that new general-purpose technologies have the potential to create entirely new markets and associated labour demand that cannot be captured by historical data. Furthermore, there are institutional factors, notably the existence of labour unions, which are likely to influence labour market outcomes. More specifically, labour unions may to some extent be able to soften AI-related labour shedding, so that ‘redundant’ labour is employed elsewhere in the company. This is a somewhat different aspect as it could potentially be handled in a quantitative analysis, provided the necessary information were available. As such this constitutes an interesting avenue for further research as does a more differentiated analyses of employment effects by skill groups.

Finally, it is imperative to be aware of what the AI exposure measure can capture. The AI exposure index by Felten et al. (2018) and consequent work on the index merely captures technological aspects of occupations. More specifically, they capture the overlap in abilities and tasks which can be performed by AI algorithms with the abilities and tasks characterising occupations. In other words, it is a measure of technological feasibility of ‘interaction’ between AI and human workers, where this interaction could mean labour substitution, task sharing or task expansion. The measure does not capture any form of actual implementation of AI technologies in the economy. This is in contrast to the industrial robot measure and identifying or developing some sort of AI investment indicator, which to the best of our knowledge is currently not available yet, would be a great leap forward for the possibilities to analyse the labour market effects of AI related technologies.

REFERENCES

- Acemoglu, D., D. Autor, J. Hazell and P. Restrepo (2022). ‘Artificial intelligence and jobs: Evidence from online vacancies’, *Journal of Labor Economics*, 40(1), 29-40.
- Acemoglu, D. (2021). ‘Harms of AI’ (No. w29247), National Bureau of Economic Research.
- Acemoglu, D. and S., Johnson (2023). ‘What’s Wrong with ChatGPT?’, Project Syndicate Commentary, February. Available at: <https://www.project-syndicate.org/commentary/chatgpt-ai-big-tech-corporate-america-investing-in-eliminating-workers-by-daron-acemoglu-and-simon-johnson-2023-02>
- Acemoglu, D. and P. Restrepo (2020). ‘Robots and jobs: Evidence from US labor markets’, *Journal of Political Economy*, 128(6), 2188-2244.
- Acemoglu, D. and D. Autor (2011). ‘Skills, Tasks and Technologies: Implications for Employment and Earnings’ In O. Ashenfelter and D. Card (edited by), *Handbook of Labor Economics*, 1(4), 1043-1171.
- Agrawal, A., J. Gans and A. Goldfarb (2019). ‘Economic Policy for Artificial Intelligence’, *Innovation Policy and the Economy*, 19(1), 139-159.
- Albanesi, S., A. D. da Silva, J. F. Jimeno, A. Lamo and A. Wabitsch (2023). ‘New Technologies and Jobs in Europe’ (No. w31357), National Bureau of Economic Research.
- Autor, David H., Frank Levy, R. J. Murnane (2003). ‘The Skill Content of Recent Technological Change: An Empirical Investigation’, *Quarterly Journal of Economics*, 118(3), 1279-1333.
- Autor, D. H. (2015). ‘Why are there still so many jobs? The history and future of workplace automation’, *Journal of Economic Perspectives*, 29(3), 3-30.
- Autor, D. (2022). ‘The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty’ (No. w30074), National Bureau of Economic Research.
- Barbieri, L., C. Mussida, M. Piva and M. Vivarelli (2020). ‘Testing the employment and skill impact of new technologies’, *Handbook of labor, human resources and population economics*, 1-27.
- Blanas, S., G. Gancia and S. Y. Lee (2019). ‘Who is afraid of machines?’ *Economic Policy*, 34(100), 627-690.
- Bogliacino, F. and M. Pianta (2010). ‘Innovation and employment: A reinvestigation using revised Pavitt classes’, *Research Policy*, 39(6), 799–809.
- Brynjolfsson, E., Tom Mitchell (2017). ‘What Can Machine Learning Do? Workforce Implications’, *Science*, 358 (6370), 1530–1534.
- Brynjolfsson, E., T. Mitchell and D. Rock (2018). ‘What can machines learn and what does it mean for occupations and the economy?’ *AEA papers and proceedings*, 108, 43-47.
- Calvino, F. and M. E. Virgillito (2018). ‘The innovation-employment nexus: a critical survey of theory and empirics’, *Journal of Economic Surveys*, 32(1), 83-117.
- Calvino, F. and L. Fontanelli (2023). ‘A portrait of AI adopters across countries: Firm characteristics, assets’ complementarities and productivity’. OECD.

- Cambridge Industrial Innovation Policy (CIIP) (2022). ‘Policies and institutions for industrial digitalisation’, IfM Engage. Institute for Manufacturing, University of Cambridge.
- Cetrulo, A., V. Cirillo and D. Guarascio (2019). ‘Weaker jobs, weaker innovation. Exploring the effects of temporary employment on new products’, *Applied Economics*, 51(59), 6350-6375.
- Cetrulo, A., D. Guarascio and M. E. Virgillito (2020). ‘Anatomy of the Italian occupational structure: concentrated power and distributed knowledge’, *Industrial and Corporate Change*, 29(6), 1345-1379.
- Chowdhury, S., P. Dey, S. Joel-Edgar, S. Bhattacharya, O. Rodriguez-Espindola, A. Abadie and L. Truong (2023). ‘Unlocking the value of artificial intelligence in human resource management through AI capability framework’, *Human Resource Management Review*, 33(1), 100899.
- Cirillo, V., R. Evangelista, D. Guarascio and M. Sostero (2021). ‘Digitalization, routineness and employment: An exploration on Italian task-based data.’ *Research Policy*, 50(7), 104079.
- Damioli, G., V. Van Roy, D. Vertesy and M. Vivarelli (2023). ‘AI technologies and employment: micro evidence from the supply side’, *Applied Economics Letters*, 30(6), 816-821.
- Darvas, Z., J. Mazza and C. Midões (2019). ‘How to improve European Union cohesion policy for the next decade’, Bruegel Policy Brief, 8, May.
- Dauth, W., S. Findeisen, J. Suedekum and N. Woessner (2021). ‘The adjustment of labor markets to robots’, *Journal of the European Economic Association*, 19(6), 3104-3153.
- Dosi, G., M. Piva, M. E. Virgillito and M. Vivarelli (2021). ‘Embodied and disembodied technological change: The sectoral patterns of job-creation and job-destruction’, *Research Policy*, 50(4), 104199.
- Eloundou, T., S. Manning, P. Mishkin and D. Rock (2023). ‘Gpts are gpts: An early look at the labor market impact potential of large language models’, arXiv preprint arXiv:2303.10130.
- Ernst, E., R. Merola and D. Samaan (2019). ‘Economics of artificial intelligence: Implications for the future of work’, *IZA Journal of Labor Policy*, 9(1).
- Eurostat-European Commission (2020). ‘Statistical regions in the European Union and partner countries: NUTS and statistical regions 2021’. Luxembourg: Publications Office of the European Union.
- European Commission (2018). ‘Artificial intelligence: a European perspective. Luxembourg: Publications Office of the European Union.
- Felten, E. W., M. Raj and R. Seamans (2018). ‘A method to link advances in artificial intelligence to occupational abilities’, *AEA Papers and Proceedings*, 108, 54-57. American Economic Association.
- Felten, E. W., M. Raj and R. Seamans (2019). ‘The occupational impact of artificial intelligence: Labor, skills and polarization’, NYU Stern School of Business.

- Felten, E., M. Raj and R. Seamans (2021). ‘Occupational, industry and geographic exposure to artificial intelligence: A novel dataset and its potential uses’, *Strategic Management Journal*, 42(12), 2195-2217.
- Felten, E., M. Raj and R. Seamans (2023). ‘How will Language Modelers like ChatGPT Affect Occupations and Industries?’ arXiv preprint arXiv:2303.01157.
- Frank, M. R., D. Autor, J. E. Bessen, E. Brynjolfsson, M. Cebrian, D. J. Deming, M. Feldman, M. Groh, J. Lobo, E. Moro, and D. Wang (2019). ‘Toward understanding the impact of artificial intelligence on labor.’ *Proceedings of the National Academy of Sciences*, 116(14), 6531-6539.
- Frey, C. B. and M. A. Osborne (2017). ‘The future of employment: How susceptible are jobs to computerisation?’ *Technological Forecasting and Social Change*, 114, 254–280.
- Georgieff, A. and R. Hyee (2022). ‘Artificial Intelligence and Employment: New Cross-Country Evidence’, *Frontiers in Artificial Intelligence*, 5.
- Graetz, G. and G. Michaels (2018). ‘Robots at Work’, *The Review of Economics and Statistics*; 100(5), 753–768.
- Green, A., A. S. del Pero and A. Verhagen (2023). ‘Artificial intelligence, job quality and inclusiveness.’ Paris: OECD.
- Gschwent, L., D. Guarascio, S. Jestl, A. Sabouniha and R. Stöllinger (2023). ‘Digital Tasks and ICT Capital: Methodologies and Data (No. 11).’ Vienna Institute for International Economic Studies, wiiw.
- IFR (2022). ‘Artificial Intelligence in Robotics, Position paper.’
- Keynes, J.M. (2010[1931]). ‘Economic Possibilities for Our Grandchildren’, In: *Essays in Persuasion*. Palgrave Macmillan, London.
- Kordalska, A., M. Olczyk, R. Stöllinger and Z. Zavorská (2022). ‘Functional Specialisation in EU Value Chains: Methods for Identifying EU Countries’ Roles in International Production Networks’, wiiw Research Report 461.
- Korinek, A. and J. E. Stiglitz (2021). ‘Artificial intelligence, globalization, and strategies for economic development’ (No. w28453), National Bureau of Economic Research.
- Kranzberg, M. (1986). ‘Technology and History: “Kranzberg’s Laws”’, *Technology and Culture*, 27(3), 544–560.
- Landesmann, M.A. and R. Stöllinger (2020). ‘The European Union’s Industrial Policy’ In Oqubay, A., Chang, H.-J., Cramer, C., Kozul-Wright, R. (edited by) *The Oxford Handbook of Industrial Policy*, Oxford University Press. Oxford.
- Leontief, W. (1983). ‘Technological Advance, Economic Growth, and the Distribution of Income’, *Population and Development Review*, 9(3), 403-410.
- Mondolo, J. (2022). ‘The composite link between technological change and employment: A survey of the literature’, *Journal of Economic Surveys*, 36(4), 1027-1068.
- Petit, F., F. Jaccoud and T. Ciarli (2023). ‘Heterogeneous Adjustments of Labor Markets to Automation Technologies’, CESifo Working Paper No. 10237.

- Raj, M. and R. Seamans (2019). 'Primer on artificial intelligence and robotics', *Journal of Organization Design*, 8, 1-14.
- Rammer, C., D. Czarnitzki and G. P. Fernández (2021). 'Artificial intelligence and industrial innovation: Evidence from firm-level data', ZEW-Centre for European Economic Research Discussion Paper, (21-036).
- Reljic, J., R. Evangelista and M. Pianta (2021). 'Digital technologies, employment and skills', *Industrial and Corporate Change*.
- Reljic, J., V. Cirillo, D. Guarascio (2023). 'Regimes of robotization in Europe', *Economics Letters*, 232, 111320.
- Reljic, J., A. Cetrulo, V. Cirillo and A. Coveri (2023). 'Non-standard work and innovation: evidence from European industries', *Economics of Innovation and New Technology*, 32(1), 136-164.
- Soori, M., A. B. and R. Dastres (2023). 'Artificial intelligence, machine learning and deep learning in advanced robotics, A review', *Cognitive Robotics*.
- Tolan, S., A. Pesole, F. Martínez-Plumed, E. Fernández-Macías, J. Hernández-Orallo and E. Gómez (2021). 'Measuring the occupational impact of AI: tasks, cognitive abilities and AI benchmarks', *Journal of Artificial Intelligence Research*, 71, 191-236.
- Valentini, E., F. Compagnucci, M. Gallegati and A. Gentili (2023). 'Robotization, employment and income: regional asymmetries and long-run policies in the Euro area', *Journal of Evolutionary Economics*, 1-35.
- Webb, M. (2020). 'The impact of artificial intelligence on the labor market', Available at SSRN 3482150.

APPENDIX A

Table A1. AI indicators: Findings and Limitations

Data Source	Findings	Limitations
AI-related Online Job Vacancies (e.g., Burning Glass)	Acemoglu et al. (2020) found no significant effects of AI at the occupation and industry level in the US	Online job vacancies are not representative of overall labour demand; occupation- industry- and country-biased
Patents	Damioli et al. (2022) studied 3,500 leading companies with AI-related patents and found a moderate positive employment impact of AI	Not an indicator of adoption, but innovation (partial, as not all innovations are patented); sample includes only patenting companies, silent on net effects (i.e., business stealing)
AI Patents and O*NET Tasks	Web (2020) used verb-noun pairs in AI patent titles and O*NET tasks to measure automation. AI more likely to affect skilled and older workers compared to previous innovation waves (ICT and robots)	Focuses on exposure rather than adoption; patent titles may not fully describe the underlying technology; selection of keywords is arbitrary
Occupation-based Indicators	Felten et al. (2018, 2019, 2021) found positive effects on wages but no impact on employment in the US; Gmyrek et al. (2023), focusing on GPTs, reveal that 24% of clerical tasks are highly exposed	Focuses on exposure rather than adoption; silent on industry and firm-level technological differences

Source: Authors' elaboration

Table A2. List of variables

Variable	Definition	Source
<i>AI exposure</i>	Standardised with 0 mean and unit standard deviation	Felten et
<i>Robot density</i>	Robot stock in manufacturing industries per 1000 employees	IFR
Total employment	Annual employment growth	EU LFS
Women (%)	Percentage of female employees over the total	EU LFS
Young workers (%)	Percentage of young employees (15-24 years old) over the total	EU LFS
Senior workers (%)	Percentage of senior employees (55-64 years old) over the total	EU LFS
University degree (%)	Percentage of employees with tertiary degree over the total	EU LFS
Non-standard	Percentage of employees with either temporary, part-time contract,	
<i>Economic variables</i>		
Manufacturing firms	Percentage of manufacturing firms out of the total	EU LFS
Large firms (%)	Percentage of large firms (50+ employees) out of the total	EU LFS
R&D investments	R&D investment as percentage of GDP	Eurostat
Demand	Gross value added in constant (2010) prices	Eurostat
<i>Instrumental variable</i>		
Daily internet users (%)	Percentage of individuals who used internet daily	Eurostat

Source: Authors' elaboration

APPENDIX B. ROBUSTNESS ANALYSIS USING DIFFERENT INDICATORS

In this section, we conduct a series of robustness checks aimed at examining the robustness of our findings based on AIOE (forward), when using alternative definitions and calculations of AI exposure, summarised in Table B1. We delve into various dimensions, including:

- (i) Construction of the matrix for matching AI applications and workplace abilities: we examine differences between crowd-source survey conducted on Amazon Mechanical Turk (MTurk) and insights gathered from PhD students.
- (ii) Different O*NET Waves: we examine whether the choice of the O*NET dataset's wave affect results by comparing the indicators using the 2009 vis-à-vis 2019 O*NET waves.
- (iii) Backward vs. Forward-Looking Indicator: we investigate whether our findings are robust to the choice between backward and forward-looking indicators.

Table B2 presents a pairwise correlation between different AI exposure indicators at the occupational level. It suggests that all indicators are highly correlated with our preferred measure (AIOE forward), with correlation coefficients ranging from 0.937 to 0.987.

Table B1. Alternative AI indicators

Abbreviation	Reference	Description	Matching
AIOE forward	Felten et al. 2021	Match between forward-looking AI applications and workplace abilities (O*NET 2019)	Matrix constructed with inputs from mTurk survey (n=2000)
AIOE backward	Felten et al. 2018	Match between backward-looking (2010-2015) AI progress (Electronic Frontier Foundation - EFF) and workplace abilities (O*NET 2009)	Matrix constructed using inputs from computer science PhD students. We modify it to account for importance and prevalence of each ability within occupations
AIOE backward mTurk	Felten et al. 2018	Match between backward-looking (2010-2015) AI progress (Electronic Frontier Foundation - EFF) and workplace abilities (O*NET 2009)	Similar to “AIOE backward” but modified by employing a matrix that matches abilities with AI applications using mTurk survey (n=2000)
AIOE hybrid	Felten et al. 2019	Match between backward-looking (2010-2015) AI progress (Electronic Frontier Foundation - EFF) and workplace abilities (O*NET 2019)	Matrix constructed with inputs from mTurk survey (n=2000)

Source: Authors' elaboration based on Felten et al. (2018, 2019, 2021)

Table B2. Correlation between Felten et al.'s (2018, 2019, 2021) indicators at the occupational level

Variables	(1)	(2)	(3)	(4)
(1) AIOE (f)	1.000			
(2) AIOE hybrid	0.987*	1.000		
(3) AIOE (b & mturk)	0.937*	0.900*	1.000	
(4) AIOE (b)	0.975*	0.981*	0.902*	1.000

Source: Authors' elaboration based on Felten et al. (2018, 2019, 2021)

The high correlation between the alternative AI indicators is reassuring in particular with respect to the robustness of the employment effects. Indeed, the results in the main text are fully supported by a series of robustness checks (Table B3). Irrespective of the AI indicators chosen, they all suggest a positive employment effect of AI exposure for the average EU region. Hence, our results are not sensitive to the choice of the AI measure.

Table B3. Employment change and AI exposure, alternative AI indicators (*IV estimates*)

	(1)	(2)	(3)	(4)
<i>First stage</i>				
Daily internet users	0.0106** (2.58)	0.01000* (2.36)	0.00976* (2.56)	0.00954* (2.33)
Daily internet users x robot density	0.00108*** (4.05)	0.00106*** (4.11)	0.00142*** (5.10)	0.00106*** (4.07)
<i>Second stage</i>				
AIRE (f) our indicator	3.022* (2.18)			
AIRE (f) x Robot density	-0.144 (-1.42)			
AIRE (b & O*NET 2009)		3.046* (2.17)		
AIRE (b & O*NET 2009) x Robot		-0.140 (-1.46)		
AIRE (b & mturk)			3.132* (2.27)	
AIRE (b & mturk) x Robot density			-0.147 (-1.63)	
AIRE (b & O*NET 2019)				3.242* (2.16)
AIRE (b & O*NET 2019) x Robot				-0.146
Robot density	0.0964 (1.93)	0.0973 (1.90)	0.115* (2.04)	0.107* (1.96)
Structural controls	Yes	Yes	Yes	Yes
Labour market controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes
Kleibergen-Paap rk Wald F statistic	18.623	17.775	20.406	18.465
N	1322	1322	1322	1322

Notes: Alternative indicators as explained in Table B1. Specification refers to those of specification 7 in Table 4 of the main text using alternative AI indicators. The dependent variable is annual employment growth. Controls include structural variables (average firm size and the share of manufacturing), labour market characteristics (the shares of female employees, non-standard employment, youth workers, senior workers, and those with a university degree), country and time fixed effects. Regression results weighted by the average number of employees for each region. The AI exposure index (AIRE) is instrumented with the share of daily internet users relative to the population in the same regions. Standard errors are robust against heteroscedasticity and serial correlation at the regional level (t-statistics in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

Source: Authors' elaboration