The Diffusion of Digital Technologies and its Consequences in the Labor Market

Christina Gathmann, LISER, University of Luxembourg and CEPR
Felix Grimm, University of Heidelberg

April 28, 2022

Abstract

To analyze the impact of digital technologies on the labor market, we develop new measures for AI and robotics technologies in Europe using Natural Language Processing techniques on patent data from the European Patent Office. Robotics technology has taken off earlier but is concentrated in a few manufacturing industries. AI technologies, in turn, are still in its infancy but have started to diffuse into a broader range of industries very recently. Combining our patent-based measures with administrative data on establishments in Germany, we investigate the labor market consequences of both technologies. Based on a shift-share design, we estimate overall employment and wage effects in local labor markets. For AI, we find positive effects on employment, esp. in the service sector. Automation plays a more important role for robotics, as we find negative effects on employment. Our results suggest few effects on local or sectoral wages. Finally, we investigate who benefits from the two technologies and who loses by investigating differences by skill level.

1 Introduction

The Fourth Industrial Revolution has dramatically improved the technical capabilities of artificial intelligence (AI) enabling machines to perform and learn tasks at human-like levels of capability in domains including translation and visual image recognition (Pratt, 2015; Schwab, 2016). Improvements in underlying techniques such as machine- or deep learning open up new possibilities for applications in AI, which may be used in a wide variety of industries. Similarly, robots have been diffusing in the economy and further advances in AI could act as a catalyst for robots to become smarter, less dependent on human guidance and thereby more efficient. There is widespread belief that both technologies will be reshaping the way we work and live.

How robots and AI affect the labor market and individual workers is far from clear, however. From a theoretical perspective, automation technologies may have three main effects on labor demand. There is a direct displacement effect as machines (or algorithms) take over some tasks previously performed by humans, which reduces the demand for some labor. The second one is a productivity effect on remaining workers that tends to increase labor demand for their services. Finally, new technologies may also create entirely new tasks and hence, job profiles, which need to be filled by workers (Acemoglu and Restrepo, 2018 2019). The net effect could be positive, zero or negative depending on the technology and the incentives to adopt including the costs of different types of labor.

Empirical studies have mostly focused on the diffusion of robots in the manufacturing sector. The results differ widely ranging from negative (Acemoglu and Restrepo, 2020), close to zero (Graetz and Michaels, 2018) or even positive employment effects (Dauth et al., 2021). Firm-level evidence, in contrast, indicates that adopters of robotics technology are not only more productive, but also grow after adoption and outperform their competitors within the same industry (Acemoglu et al., 2020; Alderucci et al., 2021; Benmelech and Zator, 2022; Koch et al., 2019). The empirical evidence for AI technologies is very scarce and shows few links between AI technology measures, employment or wages by industry or occupation (Acemoglu et al., 2021).

Assessing the impact of new digital technologies on the labor market faces substantial methodological challenges. A key issue is how to capture their advancement and diffusion in the economy. Existing studies on digitalisation have followed different paths to generate proxy variables for the importance of new digital technologies in the workplace. A first approach uses broad measures such as firms' R&D expenditure or investments in information and communication technologies (ICT)

(Bloom et al., 2014; Bresnahan et al., 2002; Caroli and van Reenen, 2001). Other studies use direct measures on specific technologies such as the number of robots installed in broad industries (Graetz and Michaels, 2018; Dauth et al., 2021; Acemoglu and Restrepo, 2020) or by firms (Koch et al., 2019; Acemoglu et al., 2020; Dixon et al., 2020; Bonfiglioli et al., 2020). One drawback is that these measures are available only for a limited set of technologies and, in the case of robots, for a small number of broad industries only.

A second approach uses occupation-level measures constructed from information on tasks performed on the job. Earlier research relied on experts (Frey and Osborne, 2017) or crowd workers (Brynjolfsson et al., 2018) to assess task automation. These measures provide a snapshot of how automation could replace occupations in the future; because they are cross-sectional, they tell us little about the dynamics of the digital transformation, however. Furthermore, Arntz et al. (2017) demonstrate that expert assessments overstate automation potentials as they do not account for the heterogeneity and shifts in task usage within occupations. Firms might reshuffle the set of tasks performed in a job or add new tasks in response to automation of some tasks. Likewise, workers may specialize in tasks that cannot be easily automated to avoid displacement, for instance.

A third approach has relied on patents, which have long been used in the innovation literature to proxy innovation (Griliches, 1990) and technology diffusion (Jaffe et al., 1993). Patents are exclusive rights of use for novel solutions to technical problems. In exchange for these exclusive rights, all patent applications are published, revealing technical details of the invention. Patent databases therefore contain the latest technical information which can be used for research. Hence, patents are a natural candidate for measuring technological progress and frequently serve as proxies for innovation. Especially in emerging technologies such as AI, new patents can be seen as a shift in the technological frontier which increases the possibilities for firms to adopt this new technology in their production processes. The number of patents and patent meta-data, such as citation counts or the location and identity of inventors have been used frequently (Hall et al., 2001; Acemoglu et al., 2014; Bell et al., 2018) in innovation research. More recent approaches have gone beyond the sheer number of patents and analyzed the actual text of patent documents (see Bessen and Hunt, 2007, for an early example). A few studies have generated patent-based measures of automation potentials (Mann and Püttmann, 2017; Dechezleprêtre et al., 2020; Danzer et al., 2020; Montobbio et al., 2020). Very recently, authors have combined patent data with information on tasks performed on the job to quantify the automation potential of digital technologies (Felten et al., 2019; Webb, 2020). Like other occupation-level measures, these are cross-sectional and focus on the replacement of labor through automation - abstracting from the evolution of new tasks, for instance.

In this paper, we develop new patent-based measures of AI and robotics for Europe covering the period from 1990 to 2018. These measures allow tracking the advances in robotics and AI technologies and their diffusion in detailed industries and over time. Our measures make use of the European Patent Office's PATSTAT database, which includes patent applications and grants for all EPO member countries. Applying text mining and natural language processing, we extract patents related to AI and robotics based on patent codes and keyword searches in the patent titles and abstracts. In the next step, we link patents to industries of use by applying a probabilistic concordance scheme developed by Lybbert and Zolas (2014) [1].

The measures proxy for the adoption and diffusion of AI and robotics technologies at the industry level. Our measures differ from existing measures along a number of dimensions. Some earlier studies have focused on the actual installation of industrial robots (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021) or on automation technologies more broadly (Frey and Osborne, 2017; Arntz et al., 2017; Mann and Püttmann, 2017). Other studies cover the industries of invention but not the industry where the patent is likely to be used in production (Dechezleprêtre et al., 2020; Montobbio et al., 2020). Our measures characterize the evolution and diffusion of two major technologies, AI and robotics, that will shape the economy for decades to come. Moreover, these technologies may replace some workers, but also raise the productivity of other workers or even result in the reorganization of work and the emergence of new tasks. Hence, our evidence is not limited to identifying the automation effect of digital technologies. In addition, our measures go beyond the actual usage of new technologies in industries, but capture the evolving technological frontier. And finally, our measures reflect the diffusion of digital technologies in Europe, which has experienced another dynamic in digital technologies than the United States. Europe, and in particular France and Germany, are leaders in the adoption of robotics technology; at the same time, they lag behind the United States and China in AI development and provision.

We complement previous studies that focus on the effects of technologies such as AI on tasks in occupations (Brynjolfsson et al., 2018; Felten et al., 2018; Gregory et al., 2019; Webb, 2020). Recently, Webb (2020) proposed new measures of time-invariant exposure of occupations to three different technologies: information technology, robots and AI. However, the measure is based on US patent and occupational data and therefore not easily transferable to European data. Also, it is a

[1] Lybbert and Zolas (2014) match keywords from the description of patents to keywords extracted from the definition of industries according to SITC and ISIC codes. Then, they construct a probability match of IPC/CPC code classes to industries based on the amount of keyword matches obtained.

static measure that does not provide variation of time.

The paper proceeds as follows. In the next section, we present our patent-based measures and explain how we identify AI and robotics patents in the PATSTAT database. We describe in detail how we link patents to industries and how we measure industry-level exposure to AI and robotics technology. Section 3 shows descriptive evidence on our patent measures and compares them to existing proxy variables such as data on the installation of robots and on investment in ICT capital. In Section 4, we introduce our administrative labor market data in Germany and discuss the empirical strategy to identify the employment and wage effects of exposure to digital technologies as measured by our patent data. In Section 5, we present our results on overall employment and wage effects as well as on the effects in manufacturing and services; finally, Section 6 discusses the implications of our findings and concludes.

2 Patent-based Measures of AI and Robotics for Europe

We construct measures of the technological opportunities of digitalisation, specifically artificial intelligence and robotics technologies, using patent data from the European Patent Office (EPO). Our approach proceeds in three steps. First, we prepare the PATSTAT data for applying text analysis to the content of each patent. In the second step, we use natural language processing techniques and the IPC/CPC codes to identify patents in robotics and AI technologies. In the final step, we match those patents to the industries that are most likely to use of them.

2.1 European Patent Data

We use data from the World Patent Statistical Database (PATSTAT) as of 2019, which contains detailed bibliographical and technical information on all patents filed in 86 countries. We focus on patents granted by the European Patent Office (EPO) between 1990 and 2018. Important innovations are patented in all major patent offices and any invention a firm wants to have protected in the European market will be patented at the EPO even if the innovation occurred abroad. Our data contain a total of about 7 million patent documents, which are identified by 3.5m unique application ids [2]. Of the 7 million documents, 5 million are patent applications and about 2 million are patent grants. The patent documents include the title and abstract of each patent,

[2] The smaller number of unique applications reflects the fact that most patents have multiple entries in the PATSTAT database, one for the patent application, others for revisions and yet another for the patent grant if the patenting process was successful.

the name, company and location of the inventor, the dates of application and grant of the patent. The technical content of a patent is characterized by its IPC or CPC codes, which are assigned by highly specialized experts, the patent examiners. The older IPC and newer CPC frameworks are very detailed with several thousand entries.

We analyze the titles and abstract of patents to determine whether the patent describes an innovation in the field of AI or robotics technologies [3]. Though each patent document includes a title, abstracts are missing in about 30% of the patent grants (670,000 cases) we extracted. Rather than dropping patents with missing abstract, we impute the abstract specifying the technical content of a patent by using the concept of patent families. Patents belong to a narrow patent family, which contains all documents of patents covering the same technical content or to an extended patent family, which combines patents covering a similar technical content. A patent family is defined based on patents with the same (detailed or slightly broader) IPC/CPC codes. As the patent classification of technologies is very detailed, the technical content is very similar even within an extended patent family. We first assign missing abstracts an abstract from the narrow patent family; if that was not successful, we use an abstract from the extended patent family instead. Overall, we are able to impute about 450,000 abstracts with these steps of which only 45,262 abstracts are based on the extended patent family. We drop the remaining patents with missing abstracts after imputation.

A patent can be filed at the EPO in one of the three official languages English, French and German. Patents filed in another language need to provide a translation into one of the official ones. While patent claims are published in all three languages, abstract and patent description are published in the official language the patent was filed in. For our purpose, we restrict attention to documents with an abstract in English as other languages are not compatible with our keyword search.[4] To perform the text search on the sample of patents, we convert all patent abstracts and titles to text corpora. We then pre-process the text as follows: we convert all text to lower cases; then remove numbers, special characters, punctuation and stop words. We then strip the text of any blanks and white spaces. Finally, we extract word stems and divide the text into *tokens*.

[3]Following the patent literature, we do not use the full text of the patent description or claims. These texts are written by patent lawyers in generic language to increase the protective scope of a patent.

[4]PATSTAT typically records the language of the abstract but this information is missing for about 250,000 patent documents. We use natural language processing to identify the language of the abstract for documents missing that information. We then drop all documents that do not contain any information in English, which reduces our sample by only 7%.

2.2 Identifying Patents in Robotics and AI

To identify patents related to AI technologies or robotics, we use a combination of patent classification codes (IPC/CPC) and keyword searches of the patent titles and abstracts.

Robotics Patents In comparison to the broad and emerging field of AI technologies, the technology of robots is defined more precisely. According to the ISO 8373 definition, a robot is an "actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks". Robots are further grouped into industrial or service robots based on their intended application. We identify robotics patents if they belong to the CPC code B25J9: "Programme-controlled manipulators" or if they match a keyword search conducted over the titles and abstracts of all patents. A patent is then classified as a match for robotics if one or more keyword tokens match with tokens of the text corpora of titles and abstracts.

AI Patents Artificial intelligence is a very broad concept and there are several important underlying technologies. We can identify AI patents in our data by a few IPC/CPC codes that are directly connected to an AI technology or sub-field such as machine learning, neural networks or fuzzy logic. Examples include the code G06N7/046 'Computer systems based on specific mathematical models - implementation by means of a neural network' or the code H04N21/4662 'Selective content distribution, e.g. interactive television or video on demand - characterized by learning algorithms'. In the first step, we select all patents with a IPC/CPC code that is directly related to AI technology using a list of AI-specific IPC/CPC codes from the World Intellectual Property Organization (2019).

Yet, there are only a small number of IPC/CPC codes for software and algorithms. Most AI-related inventions are not identified by these codes, however. Many AI innovations are patented if their purpose is to solve a specific technical problem. As a result, AI innovations are often embedded in patent applications for innovations in many different technology fields. Examples are speech and image recognition, two of the most important applications of AI technologies.[5] To identify such patents, we conduct a keyword search over the titles and abstracts of all remaining patents. The list of keywords we use is a synthesis from the World Intellectual Property Organization (2019) and Baruffaldi et al. (2020).[6] Examples of keywords include machine learning, natural language processing, fuzzy logic or decision tree. The keyword list is pre-processed using the same steps

[5] One example for such an embedded innovation is the case of level 4 and 5 autonomous driving, which relies heavily on image recognition through artificial intelligence.

[6] Our keyword list is substantially shorter than the list used in World Intellectual Property Organization (2019) in order to reduce false positives. Their keyword list includes keywords like network, algorithm, logic or boost, which can potentially be found in many patents that are unrelated to AI technologies.

as for the patent documents. A patent is classified as a match for AI technologies if one or more keyword tokens match with tokens of the text corpora of titles and abstracts. A general pattern is that patent titles often include specific techniques such as neural networks while more general technological concepts such as artificial intelligence or machine learning are more likely to appear in the abstract instead.

Results For technologies in robotics, our search yields 14,235 patent documents of which 92% contain one or more of the keywords and 8% are included based on the CPC code 'B25J9'. Among these documents, around 11,000 are actual applications or grants; the remainder contain supplementary information to existing applications [7] For AI technologies, the combined approach of codes and keyword search yields 10,311 AI patent documents of which 90% contain on or more of the AI-specific keywords and 10% are included purely on their CPC codes. After excluding supplementary documents, we are left with around 7,000 applications and grants.

Panel (A) of Figure 1 shows the evolution of robotics patent grants and applications between 1990 and 2018. Robot patents show a first peak in the mid-1990s and then again in the late 2000s. Patent applications for robotics continue to grow throughout the whole time period. Panel (B) in Figure 1 shows that AI patent grants start to emerge in the mid-1990s but remain at low levels until 2018. In contrast, patent applications for AI technologies start to grow strongly after 2015 and especially in 2017 and 2018.

To see in which broad sectors of the economy AI and robotics patents play a role, we aggregate patents to broader technology classes. We use a mapping of the more recent IPC codes at the 4-digit level to thirty-five technology classes developed by Schmoch (2008).[8] We then aggregate the thirty-five technology classes into five broad sectors: Electrical engineering, Mechanical engineering, Instruments, Chemistry and Other fields. Instruments include optical instruments, control technology and medical technology. Chemicals include pharmaceuticals, biotechnology, food and materials. Other includes many consumption goods like furniture, games but also civil engineering.

Appendix figure A1 shows that robotics technology is heavily concentrated in mechanical engineering (see Panel (A)). Since 2010, robotics patents have become more prevalent in the sector "Instruments and other fields", which points to new applications beyond mechanical engineering and industrial robots. The picture looks very different for AI patents (shown in Panel (B)): AI tech-

[7] Such supplementary documents can be corrections to existing applications or supporting material such as search reports.

[8] We prefer this classification over the one in Hall et al. (2001) because the latter is much older and thus less accurate in capturing recent developments in AI and robotics technologies.

nologies are most prominent in electrical engineering, but have recently become more important in instruments.

2.3 Technological Opportunities in Robotics and AI at the Industry Level

The final step is to match the patents to the industries that adopt the evolving technological opportunities of AI and robotics. This matching step is crucial as it will define the exposure of particular industries to the technological possibilities of AI and robotics innovations. We rely on a walkover between CPC codes and 4-digit ISIC industry codes from Lybbert and Zolas (2014) [9]. A key advantage is that the walkover has been developed more recently than concordance schemes previously used in the literature such as Kortum and Putnam (1997) or Silverman (2002) and thus is more likely to capture developments in AI and robotics technologies.

The basic idea is to use the description of an industry from the official classification. The keywords describing the activity in a certain industry are then used to search for patents that contain these keywords. The result is a list of patents with their IPC/CPC codes and the industries matched on keywords. The match frequency is then used to calculate a probabilistic weight for each industry. The weight is based on Bayes rule taking into account the number of possible codes and how often a code is matched to an industry.[10]

Appendix table A1 shows in the top panel the 4-digit industries with the highest number of AI patents. AI related patents are heavily used in the manufacturing of ICT, but also in machinery and measuring equipment. Interestingly, AI patents are also important in the music and film industry (see ISIC codes 5912 and 5920 in appendix table A1). The table further reports for each industry the share of patent grants in AI. The average share of AI patents across all industries is 0.21% implying that only a small share of the patents used in an industry has so far occurred in AI. Yet, all industries with a high total number of AI patents, with the exception of medical and dental practice activities, also have a higher share of AI patents than other industries. Hence, AI plays an important role in these industries not only in absolute terms but also relative to all other innovations. Among the industries without any AI patents, we find many services like real estate, care, tourism but also early childhood education and farming. Similarly to the case of AI patent grants, the bottom panel of appendix table A1 shows the industries with the highest

[9] see also Goldschlag et al. (2019) for further walkovers to industry and trade classification schemes.

[10] The hybrid probability weight we use adjusts the weights upward for specific codes that are strongly linked to particular industries. The rationale for this reweighting is that matches with high specificity in a certain technology class may indicate particularly important linkages to certain industries compared to broad technologies that match to many different industries.

number of robotics patent grants. As we would expect, the industries with most robotic patents are in manufacturing with the exception of "2592 Treatment and coating of metals". Further, all industries except "2620 Manufacture of computers and peripheral equipment" are robot intensive in absolute and relative terms as their share of robotic grants is higher than the average of 0.32% across all industries. Based on the number of patents used, we construct measures reflecting the evolution of AI and robotics technology at the industry level. We interpret patents as shifting the technological frontier, which improves existing or opens up entirely new possibilities in production. As such, industries with a strong growth in patents are treated as relatively more exposed to these new technological opportunities than industries where the new digital technologies are not used as intensively.

3 Descriptive Evidence on Patent Measures in Europe

Our patent-based measures describe technological opportunities of robotics and AI that may or may not be heavily used by firms. To check whether the technological opportunities laid down in patents are used by firms in the industries assigned, we relate them to industry-level proxies for technological change previously used in the literature. More specifically, we compare robotics patents to data on industry-level robot installations and AI patents to industry-level investments in ICT capital. As new technologies embodied in patents take time to be adopted and diffuse in the economy, we allow for a time gap of three years between the timing of the patent grant and an industry's exposure and potential adoption to it. Hence, a patent granted in 1990 can have a measurable impact on innovation at the earliest in 1993.

3.1 Robotics Patents and Robot Installations

To validate our patent measures on robotics technology, we use information on robot installations from the International Federation of Robotics (2020). The IFR data track the installation and stock of robots in around fifty countries. The data contains a count of annual robot installations by broad industries mostly in manufacturing and an estimate of the annual stock of robots. We use the information on robots installed in European countries to ensure that the data cover the same set of countries as contained in the European Patent Convention. Our patent data is available at the 4-digit level and hence, much more detailed than the industry classification used in the IFR data, which contains information on thirteen industries in manufacturing (food and beverages; textiles including apparel; wood and furniture; paper and printing; plastics and chemicals; minerals; basic metals; metal products; industrial machinery; electronics; automotive; shipbuilding and aerospace; and miscellaneous manufacturing including production of jewelry and toys) and on six other broad sectors (agriculture, forestry and fishing; mining; utilities; construction; education, research and development; services). We therefore aggregate our patent data to the 2-digit level and match these to the IFR classification. About 30 percent of robots are not classified into one of the nineteen IFR industries though this share declines over time. Following Acemoglu and Restrepo (2020), we allocate unclassified robots to industries in the same proportions as in the classified data. Panel (A) of figure 2 shows a strong positive correlation between patents in robotics technology and robot installations three years later (both measured in logs). The grey area represents the 95% confidence intervals. The correlation between patent applications in robotics and robot installations is much more pronounced than for patents grants in robotics and installations.

3.2 AI Patents and ICT Capital

We next compare our measure of AI patents to data on investments in ICT equipment by country and industry between 1995 and 2015 from the EU-KLEMS database [11]. We again focus on European countries, which includes the European Union member states and the UK. All countries in this sample are also members of the European Patent Convention and are hence provided patent protection by the EPO. Industries in the EU-KLEMS database are available at the 2-digit level though some industries are combined. We define overall ICT investments as the sum of investments in information technology equipment, in communication technology equipment and in software and databases, all measured in 2010 prices. The link between AI and ICT is likely weaker than for patents in robotics and robot installations as not all investments in ICT capital like a monitor need to be AI-related; and not all AI innovations are part of ICT equipment and software – a smart washing machine being one example. Panel (B) of figure 2 shows that there is a weak positive correlation between AI patent applications and ICT investments, while there is strong correlation between AI grants and ICT investments.

[11] See Stehrer et al. (2019) for details on the methodology of KLEMS and Stehrer and Adarov (2019) for empirical results on ICT and productivity growth. For an earlier release of KLEMS see Timmer et al. (2007).

4 Data and Empirical Strategy for Local Labor Market Analysis

With our measures of AI and robotics technologies in hand, we now turn to an analysis of their consequences in the labor market. In particular, we merge our industry-level measures of technological opportunities to plant-level data from Germany. We first describe the data followed by a discussion of our estimation strategy.

4.1 Administrative Establishment-level Data

We use administrative data from the German Establishment History Panel (BHP), a 50% random sample of all establishments with at least one employee covered by the social security system in Germany (see Ganzer et al., 2020, for more details). The social security data cover around 80% of the German labor force excluding civil servants, military personnel and the self-employed. Our plant sample spans the years from 1990 to 2018. We match our measures of AI and robot technologies, which vary by detailed industry and year, to the establishment data by the detailed (3-digit) industry and year allowing for a three-year time lag.[12]

The data include detailed information on the socio-economic composition of the workforce by age, gender and skill in each establishment. We distinguish three skill groups - low, medium and high skilled - based on the highest qualification obtained. High-skilled workers are workers who have graduated from a college or university. Medium-skilled workers have completed a vocational training program or obtained the university entrance certificate after high school (*Abitur*). Low-skilled workers have lower qualifications or no qualifications at all. In the raw data, the education variable is missing for about 9% to 37% of the observations depending on the year. We use imputation procedures to fill in missing education information, which reduces missings to less than 1% (see Fitzenberger et al., 2005). We further distinguish three broad age groups (20-34, 35-49 and 50-64). We also expect that digital technologies affect some occupational groups more than others. To analyze who might benefit and who might lose, we use information on the occupational structure in the plants (e.g. the share of simple manual jobs, clerks, technicians, skilled manual labor, engineers, professionals and managers) and on the type of employment contracts used (e.g. fixed-term contract, temporary worker).

Finally, we also observe establishment wages. As is common in social security records, wages

[12] We use a crosswalk provided by EU RAMON to convert the ISIC rev. 4 classification of our patent measures to the European NACE rev. 2 classification, which is equivalent to the German industry classification (WZ08) at the 3-digit level.

are right-censored at the highest level of earnings that are subject to social security contributions. Wages are imputed based on the imputation procedure of Gartner (2005). We observe the wage distribution characterized by the 25th, 50th and 75th percentile wage. The wage information is available for full-time workers, by the three skill groups and by gender.

We aggregate the plant-level data to the district level and explore how advances in robotics and AI technology affect local labor markets. In total, there are 401 districts (*Kreise*) in Germany. The main outcome variables are changes in employment (defined as Δ *log employment* or % Δ*employment*) and wages (measured as Δ *log wages*) at the district level between 1990 and 2018 as well as for several sub-periods (1990-1997, 1998-2004, 2005-2011 and 2012-2018). Table 1 reports summary statistics.

4.2 Estimation Strategy

To identify the impact of robotics and AI technologies in the labor market, we rely on a shift-share design (Bartik, 1993). Shift-share designs have become popular to study the impact of trade and technology shocks on local labor markets. In our context, the shift-share design combines a 'shift' variable to represent the overall technological advancement potentially used in an industry with a 'share' variable to proxy for how much a local labor market is possibly affected by the new technology.

Cumulative Knowledge Creation in AI and Robotics: We define the 'shift variable' by exploiting our patents in robotics or AI. More specifically, we sum all patents used in an industry over our sample period to reflect the cumulative nature of knowledge creation, which is codified in patents. The sum of patents reflects the additional knowledge of emerging technologies created relative to 1990, the beginning of our sample period. Hence, to trace the long-run development related to robotics and AI technology, we use the following measure:

$$TotPat_i^c = \sum_{1990}^{2018} Log(1 + Pat_{i,t}^c) \tag{1}$$

where $c = (Robotics, AI)$ denotes the technology considered, i stands for the industry and t for year. The measure $TotPat_i^c$ reflects the total accumulation of new knowledge in technology c to be used in industry i; it accounts for the relative importance of patents in different industries since the logarithm of patents puts less weight on industries with many patents.

To investigate the influence of digital technologies on the labor market, we characterize each local labor market by its industry structure in the base year. We choose 1993 as our base year as this is the first year when reliable labor market data is available for East Germany. Our key independent variable, i.e. the local exposure to the technological innovations in robotics or AI, is then defined as the interaction between initial employment shares in industry i and region r ('shares') and the evolution in AI and robotics technologies in industry i over time ('shift'). The exposure measure using the long-run accumulation of knowledge in AI and robotics is thus calculated as:

$$Exposure_r^c = \sum_{i=1}^{I} \left(\frac{Emp_{i,r}^{1993}}{Emp_r^{1993}} * TotPat_i^c \right) \tag{2}$$

where r denotes the local labor market and i the 3-digit industry. The first term ($\frac{Emp_{ir}^{1993}}{Emp_r^{1993}}$) measures initial employment shares in industry i in the base year (1993). The second term, $TotPat_i^c$, captures the growth in AI or robotics patents as a proxy of the technological opportunities in industry i.[13].

Figure 3 shows the geographic variation in exposure to robotics and AI patents where exposure is constructed as the combination of initial industry shares and the overall growth in patent grants between 1990 and 2018 (according to equation 2). Most notably, there is a marked difference between East and West Germany as districts in West Germany are much more likely to be exposed to both AI and robotics than districts in East Germany. More districts are exposed to robotics technology than AI technologies, which might in part reflect the wider diffusion of robotics technology in industries esp. in the manufacturing sector. AI on the other hand has diffused less broadly geographically and is concentrated on districts in the South and West of Germany.

To provide a first assessment of the link between exposure to AI and robotics technologies and labor market outcomes at the local level, we plot scatterplots between changes in employment and wages and our long-run exposure measure to AI and robotics. In figure 4, each district represents a dot. We also add a linear regression line and the 95% confidence interval. There is a positive relationship between exposure to AI and robotics and changes in district employment whereas the effects of both technologies on wage changes are negative.

The correlations in Figure 4 might be spurious if there are regional labor demand or supply shocks that affect employment or wages. To control for these other influences, we estimate models of the following form:

$$\Delta Y_{r,t} = \beta Exposure_r + \gamma_1 \Delta Trade_{r,t} + \gamma_2 \Delta ICT_{r,t} + \delta' X_r + \theta_I + \alpha_R + \epsilon_{r,t} \tag{3}$$

where $\Delta Y_{r,t}$ denotes the local employment or wage changes between 1993 and 2018 and $Exposure_r^c$ characterizes the district's overall exposure to AI or robot technologies over this period. We control

[13]We find very similar results if we use the sum over three years (1990-1992 and 2016-2018) to calculate the long-run growth in patents instead.

for the local structure of the workforce by including controls for employment shares by age, skill group and gender (X_r) as well as employment shares by broad (1-digit) industry (θ_I); all are measured in the base year. We further include dummies for the broad region (North, East, South, West), α_R. To control for potential confounding effects of international trade, we adjust for changes in net exports per worker ($\Delta Trade_{r,t}$). To control for other technology-driven changes in labor demand, we further include a variable for general investments in ICT capital per worker ($ICT_{r,t}$).

Our main parameter of interest is β, which measures how employment or wages in districts exposed to AI and robotics technologies change relative to districts less exposed to the new technologies. Note that β combines any direct effect on plants in industries highly exposed to the new technologies, adjustments in wages, and potential local spillover effects on plants linked to exposed industries through input-output linkages or local multiplier effects in the same region. We return to this issue below.

Exploiting the Panel Dimension of Knowledge Creation: The measure in equation (1) varies across industries only. To incorporate the dynamic development of the two technologies over time, we define the following panel measures:

$$Pat_{i,P}^{c} = \sum_{s \in P} Log(1 + Pat_{i,s}^{c}) \tag{4}$$

where P denotes the sub-periods 1990-1997, 1990-2004, 1990-2011, 1990-2018. The measure in equation (4) reflects the cumulative knowledge creation in each period. Hence, the shift variable is now the cumulative sum of patent grants between 1990 and the end of the period e.g. the sum from 1990 to 1997 for the first period and likewise for the other periods. The measures in equation (4) take on four values for each industry.

The panel exposure of a local labor market is then defined accordingly:

$$Exposure_{r,t}^{c} = \sum_{i=1}^{I} \left(\frac{Emp_{i,r}^{t}}{Emp_{r}^{t}} * Pat_{i,t}^{c} \right) \tag{5}$$

where r and t now indicate that exposure varies by both district and time. The time variation comes from two sources: first, the accumulation of patents now varies not only across industries but also over time. Second, the exposure measure uses the industry shares of the first year in each period to account for the fact that districts might adjust their industry structure over time.

Exploiting our panel dimension, we then estimate panel models of the following form:

$$\Delta Y_{r,t} = \beta Exposure_{r,t} + +\gamma_1 \Delta Trade_{r,t} + \gamma_2 \Delta ICT_{r,t} + \delta' X_r + \theta_I + \alpha_r + e_{r,t} \tag{6}$$

Here, $\Delta Y_{r,t}$ are changes in employment and wages in each sub-period.[14] As before, our main

[14]Hence, for the first sub-period, for instance, the dependent variables are changes in employment or wages between

parameter of interest is β, which captures the impact of exposure to AI and robotics on the local labor market. α_r denote district fixed effects, which control for a district-specific linear trend in employment or wages. All other variables are measured as before. Here, we cluster standard errors at the district level.

For the shift-share design to be valid requires that either the employment shares or the shift (here, the growth in patents) have to be exogenous (Goldsmith-Pinkham et al., 2018; Borusyak et al., 2018). It is important to stress that, in our setting, the growth in knowledge as codified in patents is measured at the European level. Hence, we consider how AI patents produced and patented in e.g. Finland impact local labor markets in Germany. In addition, we estimate the effect for firms *using* patents in the production of goods and services, not for firms producing the patents. It is highly unlikely that the employment conditions and wage levels of firms using the knowledge codified in a patent have an impact on the likelihood or timing of patenting an invention in AI or robotics technology. Both conditions suggest it is reasonable to assume that the shift variable is exogenous to local labor market conditions of using firms. We will provide evidence on the employment shares below after presenting our main results.

A remaining concern of our estimation approach is that there could be differential labor market shocks in regions with industries that are exposed to greater advances in robotics and AI than other regions. A possible concern is that industry-specific demand shocks might lead to higher usage of AI and robotics in some industries than in others. A carmaker who is exposed to smart driving technology might implement electric vehicles faster if there is negative shock to the production of fuel cars or some problem in the supply of parts, for instance. To mitigate that concern, we control in our estimation for trade flows and investments in ICT. In addition, we also control in our panel estimation in equation (6) for district-specific trends to capture any differential trajectories on the labor demand or supply side.

Adão et al. (2019) point out that regions with similar industry structures (and hence, similar initial industry shares) are likely to be subject to similar technology- or demand-driven shocks. Hence, their error terms are likely to be correlated. To account for this spatial correlation, we calculate standard errors by clustering by regional industry shares.

1990 and 1997.

5 AI and Robotics Technologies and Local Labor Markets

5.1 Employment and Wages in the Local Economy

We start out with estimating the long-term influence of AI and robotics technologies on local labor markets according to equation (3). The dependent variables are long-term changes in employment and wages between 1993 and 2018; the key independent variables are the local exposure to AI and robot technologies as defined by equation (2).

Table 2 shows the results for log employment changes in Panel A. The first two columns report estimates for exposure to AI technologies, while columns (3) to (4) report the results for robotics technologies separately. In column (5), we include both exposure measures to AI and robotics technologies simultaneously to see how a set of digital technologies like AI affects labor market outcomes conditional on exposure to robotics technology. The first specification (in columns (1) and (3)) controls for the skill, age, gender and industry structure of the local labor market as well as for broad regions. The second specification (in columns (2), (4) and (5)) adds changes in net exports and changes in ICT investments to the specification in order to control for other demand-side forces through international trade and other technological advancement.

The long-run cumulative effect of exposure to AI technologies has increased local employment irrespective of the set of control variables included. Based on the specification in column (2), a one standard deviation change in the local exposure to AI raises local employment by around 3 percentage points over the 1993-2018 period. Exposure to robots, in contrast, seems to have had little long-run impact on local employment as the estimates fail to reach statistical significance.

Local economies that employ robotics technology might also be very active in the use and diffusion of AI technologies in production processes. An example is the car industry where robots have been heavily used in the actual production of vehicles, while AI technologies play an important role in the development of smart and self-driving vehicles. We could have regional economies with a weak industrial base but a strong, prosperous service sector. As robots are mostly used in manufacturing, these regions could be little exposed to robotics technology but at the forefront of using AI technologies. The correlation between the exposure measures in AI and robotics is only around 0.54; as such, we have a lot of independent variation in each local exposure measure.

Including both exposure measures simultaneously, we still find that AI technologies encourage employment growth over the long-run, while robotics has little effect (see column (5) in table 2). We find few effects on average wages in the region. Panel B of Table 2 shows that exposure to AI and robotics has no statistically significant impact on local wages over the 1993-2018 period. Our long-run results are in line with existing evidence on robot installations in Germany (Dauth et al., 2021). This congruence is important as our local exposure measure is based on patents rather than the number of robots installed. The advantage of our patent-based measure is that exposure is available for a much broader set of industries, including many outside manufacturing.

One concern with our estimation results in table 2 is that employment and wage changes are identified from variation across local labor markets only (see equation 3). Yet, local economies that are very exposed to one or both technologies might have different wage and employment trajectories than regions that are less exposed. If the adoption of technologies is in part motivated by competitive pressure from international markets and trade, labor demand changes are likely to differ between more and less exposed districts, for instance. Moreover, the dynamics of the two technologies might vary over time: robotics technologies diffused much earlier into the manufacturing sector, while AI technologies have spread more recently but diffused into more sectors in the economy.

To address these issues, we exploit the panel structure of our data. In particular, we split our time period into four sub-periods: 1993-1998, 1999-2005, 2006-2011 and 2012-2018. The dependent variables are now log employment or wage changes within each sub-period. The shift variable now measures the accumulation of knowledge in the two technologies within each sub-period (according to equation (5)). For a given district, the measure for exposure is now higher in a period where a lot of knowledge has been created and diffused into the economy; and lower if there was little knowledge accumulation in that period.

Table 3 shows the panel results for our baseline specifications (in columns (1)-(3)). To address the concerns that labor markets with more exposure have a different employment or wage trajectory than local labor markets with little exposure, we further include district fixed effects, which allow for district-specific trends in employment and wages (in column (4) to (6)). All specifications include the full set of demographic, regional and industry controls as well as demand-side changes as in previous tables.

The panel estimates without district fixed effects show the same pattern: AI technologies seem to increase local employment, while robotics technologies seem to replace labor and reduce local employment. The fact that AI and robotics technologies have opposing signs for employment changes over the long-run is important for the discussion on the digital transition. Given the scarcity of empirical results on AI, the public but also academic debate is often based on existing results for robots technology that are available for several advanced countries (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021).

The picture changes if we allow for district-specific time trends. Columns (4)-(6) suggest that trend breaks of employment tend to be negative in local labor markets highly exposed to AI or robotics technologies. To gauge the size of these effects, we use the estimates including both technologies in column (6) and consider a one standard deviation change in local exposure. These calculations indicate a modest reduction in employment: A one standard deviation increase in AI exposure would reduce employment by 1.1 percentage points, while an increase in robot exposure by a standard deviation reduces local employment by 1.8 percentage points. For wages, the results shown in Panel B of Table 3 are very similar to the observed impact over the total period: the economic effect is small and most do not reach statistical significance.

These overall effects on local labor markets suggest that a flexible specification that controls for district-specific labor market trajectories shows that both technologies have an automation potential that is not compensated by productivity or innovation effects, at least not at the local labor market level. We next turn to the question how the two technologies have impacted jobs and workers in manufacturing and the service sector.

5.2 Manufacturing versus Services

Beyond the overall effects on the local labor market, we would expect that robotics technology has a stronger effect in the manufacturing sector where robots have been most heavily used. In contrast, robots are not much used in the service sector; any impact we see on service sector employment would thus emerge from an indirect effect through sectoral mobility or sector-specific job creation or destruction by firms. In contrast, AI technologies might have diffused into both sectors though the direction of their effect is a-priori unclear. The impact in each sector depends on at least three factors: how much tasks in each sector are susceptible to automation through AI technologies; how strong the offsetting forces of increased productivity and creation of new tasks are; and how attractive the adoption of AI technologies is in each sector, which depends, among others, on the price of labor.

To investigate this empirically, we re-estimate equation 5 where the dependent variables are now employment or wages changes in the manufacturing sector and outside of manufacturing, which mostly covers the service sector. The panel estimates for employment changes are shown in table 4. Columns (1) to (3) do not include district fixed effects, while columns (4) to (6) add them.

If we analyze AI and robotics technologies in isolation, we see that both technologies replace labor in the manufacturing sector (see Panel A of table 4) reflecting their substantial automation potential. Robot exposure decreases local employment by roughly 6.6 percentage points based on a standard deviation change in exposure. Once we condition on AI exposure (see column (6)), exposure to robotics technologies reduces local employment by around 3.5 points. AI exposure reduces local employment in manufacturing by around 4 percentage points. The advancement of both technologies is therefore likely to hit areas with a strong industrial base especially hard.

Outside of manufacturing, the situation is different (see Panel B of table 4). AI technologies create employment in the service sector irrespective of whether we include district fixed effects or not (see column (3) and (6)) once we control for robot exposure in the local labor market. Robot exposure, in turn, reduces employment in the service sector across all specifications. The effect is about half the negative employment effect in manufacturing. One explanation for the employment decline outside of manufacturing is that robots, by destroying jobs in the local economy, but not raising wages, reduce the demand for local goods and services.

Do we see any differential effect on wages in manufacturing and outside of it – despite the fact that we saw no impact on average local wages? We report the results for wage changes based on the same panel specifications as for employment changes in appendix table A3. As for average wages, we find few effects of the two technologies on manufacturing wages. Outside of manufacturing, AI exposure tends to reduce wages. Together with the employment growth of AI outside of manufacturing, our results indicate that AI technologies create mostly low-wage jobs. For robotics technologies, wages tend to increase outside of manufacturing, though the effect vanishes if we allow for district-specific wage trends.

We next investigate which workers might be most affected by the employment shifts in manufacturing and outside of it. Do robots mostly replace low-skilled jobs, but create new jobs further up the skill distribution with few net effects on the total number of jobs? And do AI technologies replace jobs further up the skill distribution as some have claimed? To shed light on these questions, we study employment changes in manufacturing and services for three different skill groups: low-skilled (without a vocational degree), medium-skilled (with a vocational degree) and high-skilled workers (with a college or university degree).

We again use our panel specification from equation (6) using both the baseline and allowing for district-specific effects in employment trajectories. The results in Panel A of Table 5 show robotics technology primarily automates manufacturing jobs of low-skilled workers. Based on the estimates in column (4), an increase of robot exposure by one standard deviation leads to a decline in low skill manufacturing employment of around 8.1 percentage points. Interestingly, AI technologies destroy manufacturing jobs further up the skill distribution: employment of medium-skilled workers decreased by almost 10 percentage points if AI exposure would increase by a standard deviation.

Outside of manufacturing, the pattern looks quite different. Here, AI technologies create more jobs up the skill ladder: a one standard deviation increase raises high-skilled employment outside of manufacturing by 2.9 percentage points. Effects for medium- and low-skilled workers are much smaller and typically do not reach statistical significance. Exposure to robotics reduces employment outside of manufacturing for all skill groups but esp. for the low- and medium-skilled (with employment declines of around 4 and 1.5 percentage points, respectively).

Overall, these results indicate that high-skilled workers seem to benefit from AI technologies. Medium-skilled workers actually see their job prospects endangered by both AI and robotics technologies. Low-skilled workers, in turn, are mostly harmed by the diffusion of robotics technologies as their job prospects decline strongly both within and outside of manufacturing. These results, which have not been documented so far, show that the diffusion of digital technologies has sizable distributional consequences and hence, potential implications for training and other policies to support the transition into the digital era.

Our findings paint a nuanced picture of the impact of different digital technologies on the labor market. Robots with their strong automation potential of routinized jobs have different effects than AI technologies, which seem to replace workers higher up the skill distribution. When discussing the economic consequences and policy implications of the digital transition, it is therefore important not to draw conclusions on evidence from just one technology and extrapolate it to another.

6 Conclusion

We develop new measures for the advancement of robotics and AI technologies in Europe applying natural language processing on patent data from Europe. Our measures for robotics are strongly correlated with robot installations but are available for many more industries in manufacturing as well as outside manufacturing compared to existing data on industrial robots. Our measure for AI technologies show a positive, but weaker correlation with ICT capital. Overall, knowledge in robotics technology has been more prominent over the 1990-2018 period, but has diffused into a small set of industries in Germany. The patenting of knowledge in AI technologies, in turn, has only picked up since 2015 but has started to diffuse into more industries.

We then use our new measures to explore the labor market consequences of the new technologies. Using a shift-share approach and data on local labor markets in Germany, we find that exposure to AI and robotics technologies reduce local employment with few effects on local wages. Most importantly, the small average effects mask considerable heterogeneities across sectors of the economy: employment declines are much more pronounced in manufacturing than in the service sector. We also investigate what happens if we control for both types of digital technologies simultaneously. These conditional estimates indicate that AI technologies destroy more jobs in manufacturing, while robotics reduces employment more outside of it.

Finally, we investigate whether how different skill groups are affected by the new technologies. We find that high-skilled workers tend to benefit from the diffusion of AI technologies, while medium-skilled workers see their job prospects decline through both technologies. Low-skilled workers are hit hardest by the diffusion of robotics technology, which replaces low-skilled jobs both within and outside the manufacturing sector.

Our results for robotics are consistent with earlier evidence using installations as direct measure of robot diffusion in manufacturing (Dauth et al., 2021). The consistency of results for the two measures provides additional support for our approach to proxy the advancement of digital technologies using patent data. For AI, our approach provides a novel measure at the industry level over three decades, which complements recent attempts to quantify the future automation potential at the occupation level.

Unlike previous studies that consider AI and robotics as automation technologies (see Mann and Püttmann (2017) for example), we find considerable differences in the labor market effects of the two technologies, especially if we consider differences across sectors and skill groups. The most likely explanation is that they are used differently in production and thereby vary in the way they substitute for or enhance human labor. In this context, the impact of robots is stronger labor-replacing than that of AI, which tends to increase employment, especially non-manufacturing employment.

References

Acemoglu, D., U. Akcigit, and M. Celik (2014): "Young, Restless and Creative: Openness to Disruption and Creative Innovations," *NBER Working Paper No. 19894*.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2021): "AI and Jobs: Evidence from Online Vacancies," Tech. Rep. No. 28257, NBER Working Paper.

Acemoglu, D., C. Lelarge, and P. Restrepo (2020): "Competing with Robots: Firm-Level Evidence from France," *AEA Papers and Proceedings*, 110, 383–88.

Acemoglu, D. and P. Restrepo (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.

——— (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.

——— (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

Adão, R., M. Kolesár, and E. Morales (2019): "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

Alderucci, D., L. Branstetter, E. Hovy, A. Runge, and N. Zolas (2021): "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata," .

Arntz, M., T. Gregory, and U. Zierahn (2017): "Revisiting the Risk of Automation," *Economics Letters*, 159, 157 – 160.

Bartik, T. (1993): "Who Benefits from Local Job Growth: Migrants or the Original Residents?" *Regional Studies*, 27, 297–311.

Baruffaldi, S., B. van Beuzekom, H. Dernis, D. Harhoff, N. Rao, D. Rosenfeld, and M. Squicciarini (2020): "Identifying and Measuring Developments in Artificial Intelligence," *OECD Science, Technology and Industry Working Papers*.

Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2018): "Who Becomes an Inventor in America? The Importance of Exposure to Innovation," *The Quarterly Journal of Economics*, 134, 647–713.

Benmelech, E. and M. Zator (2022): "Robots and Firm Investment," Tech. rep., National Bureau of Economic Research.

Bessen, J. and R. Hunt (2007): "An Empirical Look at Software Patents," *Journal of Economics and Management Strategy*, 16, 157–189.

Bloom, N., L. Garicano, R. Sadun, and J. van Reenen (2014): "The Distinct Effects of Information Technology and Communication Technology on Firm Organization," *Management Science*, 60, 2859–2885.

Bonfiglioli, A., R. Crinò, H. Fadinger, and G. Gancia (2020): "Robot Imports and Firm-Level Outcomes," *CESifo Working Paper No. 8741*.

Borusyak, K., P. Hull, and X. Jaravel (2018): "Quasi-Experimental Shift-Share Research Designs," Working Paper 24997, National Bureau of Economic Research.

Bresnahan, T., E. Brynjolfsson, and L. Hitt (2002): "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics*, 117, 339–376.

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.

Caroli, E. and J. van Reenen (2001): "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments," *The Quarterly Journal of Economics*, 116, 1449–1492.

Danzer, A., C. Feuerbaum, and F. Gaessler (2020): "Labor Supply and Automation Innovation," *IZA Discussion Papers No. 13429*.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021): "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 19, 3104–3153.

Dechezleprêtre, A., D. Hémous, M. Olsen, and C. Zanella (2020): "Automating Labor: Evidence from Firm-Level Patent Data," *CEPR Discussion Paper No. 14249*.

Dixon, J., B. Hong, and L. Wu (2020): "The Robot Revolution: Managerial and Employment Consequences for Firms," *NYU Stern School of Business*.

Felten, E., M. Raj, and R. Seamans (2018): "A Method to Link Advances in Artificial Intelligence to Occupational Abilities," *AEA Papers and Proceedings*, 108, 54–57.

——— (2019): "The Variable Impact of Artificial Intelligence on Labor: The Role of Complementary Skills and Technologies," Tech. rep., Princeton University.

Fitzenberger, B., A. Osikominu, and R. Völter (2005): " Imputation Rules to Improve the Education Variable in the IAB Employment Subsample," *FDZ-Methodenreport 03*.

Frey, C. and M. Osborne (2017): "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change*, 114, 254–280.

Ganzer, A., L. Schmidtlein, J. Stegmaier, and S. Wolter (2020): "Establishment History Panel 1975-2019," *FDZ-Datenreport 16*.

Gartner, H. (2005): "The Imputation of Wages above the Contribution Limit with the German IAB Employment Sample," *FDZ-Methodenreport 02*.

Goldschlag, N., T. Lybbert, and N. Zolas (2019): "Tracking the Technological Composition of Industries with Algorithmic Patent Concordances," *Economics of Innovation and New Technology*.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2018): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

Graetz, G. and G. Michaels (2018): "Robots at Work," *The Review of Economics and Statistics*, 100, 753–768.

Gregory, T., A. Salomons, and U. Zierahn (2019): "Racing With or Against the Machine? Evidence from Europe," IZA Discussion Papers 12063, Institute of Labor Economics (IZA).

Griliches, Z. (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 1661–1707.

Hall, B., A. Jaffe, and M. Trajtenberg (2001): "The NBER Patent Citation Data File: Lessons, Insights and Methodologial Tools," *NBER Working Paper No. 8498*.

International Federation of Robotics (2020): "World Robotics Industrial Robots and Service Robots," Tech. rep.

Jaffe, A., M. Trajtenberg, and R. Henderson (1993): "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 108, 577–598.

Koch, M., I. Manuylov, and M. Smolka (2019): "Robots and Firms," CESifo Working Paper Series 7608.

Kortum, S. and J. Putnam (1997): "Assigning Patents to Industries: Tests of the Yale Technology Concordance," *Economic Systems Research*, 9, 161–176.

Lybbert, T. and N. Zolas (2014): "Getting Patents and Economic Data to Speak to each other: An 'Algorithmic Links with Probabilities' Approach for joint Analyses of Patenting and Economic Activity," *Research Policy*, 43, 530–542.

Mann, K. and L. Püttmann (2017): "Benign Effects of Automation: New Evidence from Patent Texts," *SSRN Electronic Journal*.

Montobbio, F., J. Staccioli, M. Virgillito, and M. Vivarelli (2020): "Robots and the Origin of their Labour-Saving Impact," IZA Discussion Papers 12967, Institute of Labor Economics (IZA).

Pratt, G. (2015): "Is a Cambrian Explosion Coming for Robotics?" *Journal of Economic Perspectives*, 29, 51–60.

Schmoch, U. (2008): "Concept of a Technology Classification for Country Comparisons," *World Intellectual Property Organization*.

Schwab, K. (2016): "The Fourth Industrial Revolution: What It Means, How To Respond," *Foreign Affairs*, 94.

Silverman, B. S. (2002): *Technological Resources and the Logic of Corporate Diversification*, Routledge.

Stehrer, R. and A. Adarov (2019): "Tangiblen and Intangible Assets in the Growth Performance of the EU, Japan and the US," *wiiw Research Report No. 442*.

Stehrer, R., A. Bykova, K. Jäger, O. Reiter, and M. Schwarzhappel (2019): "Industry Level Growth and Productivity Data with Special Focus on Intangible Assets," *The Vienna Institute for International Economic Studies*.

Timmer, M., M. O'Mahony, and B. van Ark (2007): "EU KLEMS Growth and Productivity Accounts: An Overview," *International Productivity Monitor*, 14, 71–85.

Webb, M. (2020): "The Impact of Artificial Intelligence on the Labor Market," *Stanford University*.

World Intellectual Property Organization (2019): "WIPO Technology Trends 2019: Artificial Intelligence," Geneva.

Figure 1: Number of Patents in AI and Robotics, 1990-2018

(a) Robotics Patents

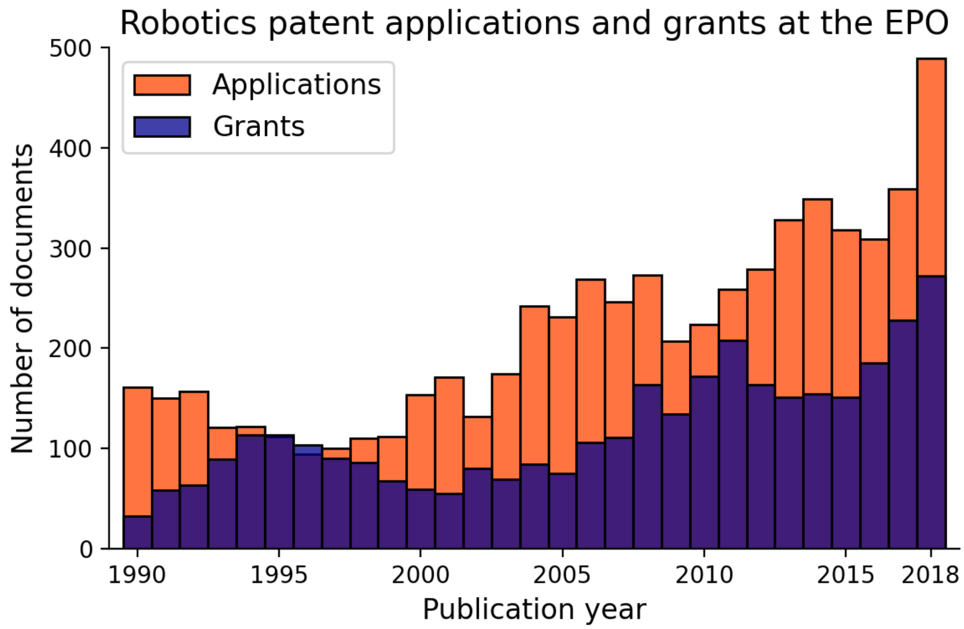

(b) AI Patents

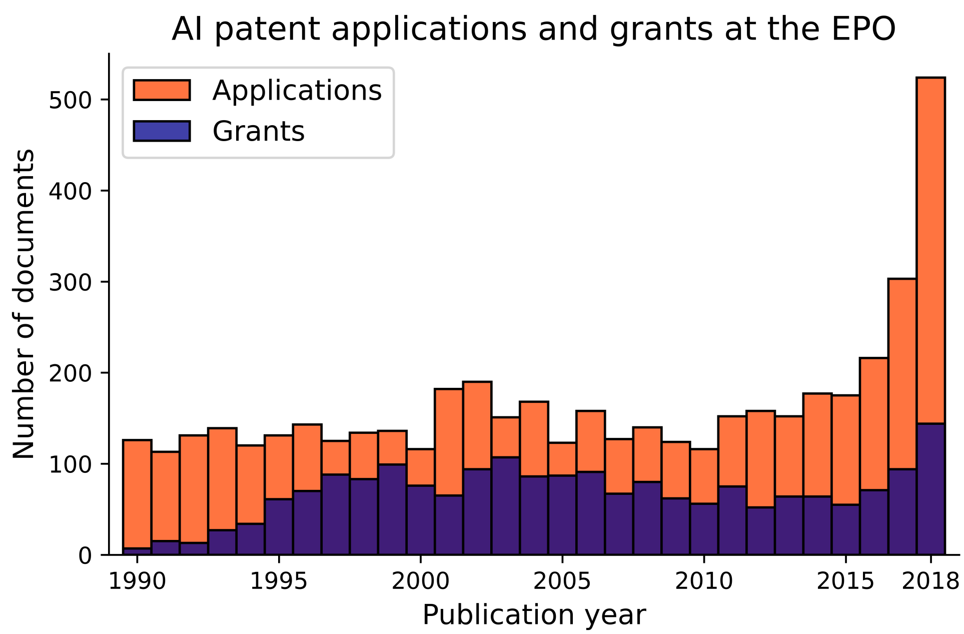

Figure 2: Correlation of Patents with ICT Investments and Robot Installations, 1990-2018

(a) Robotics patents and Robot installations

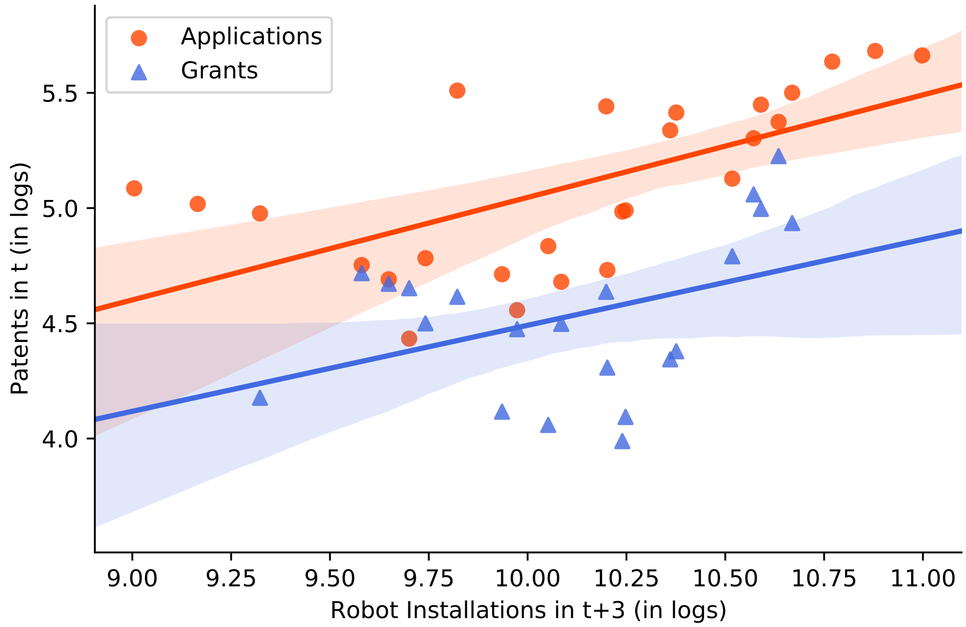

(b) AI patents and ICT Investments

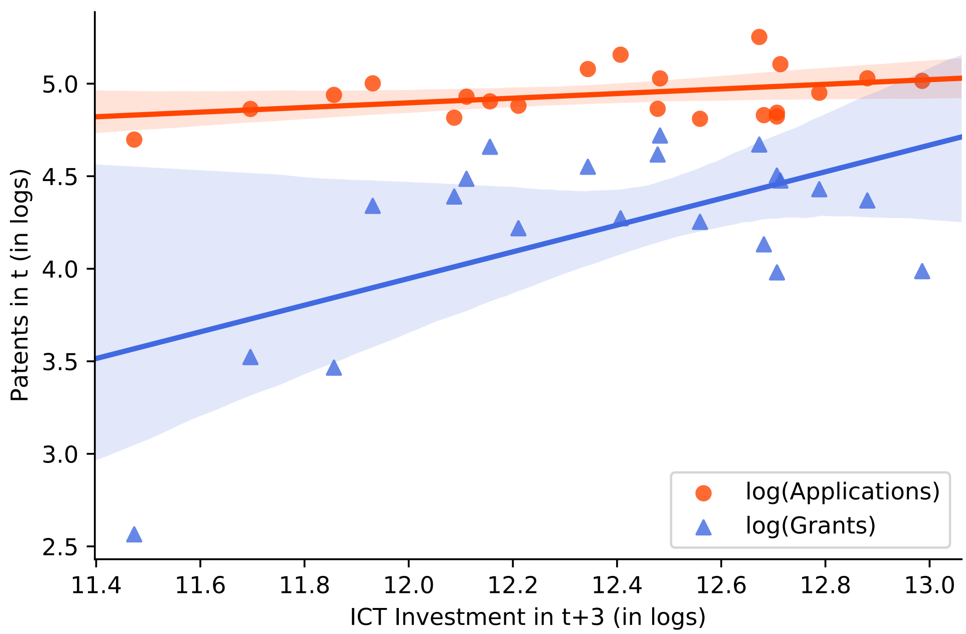

Figure 3: Regional Exposure to AI and Robotics in 2018

(a) AI exposure

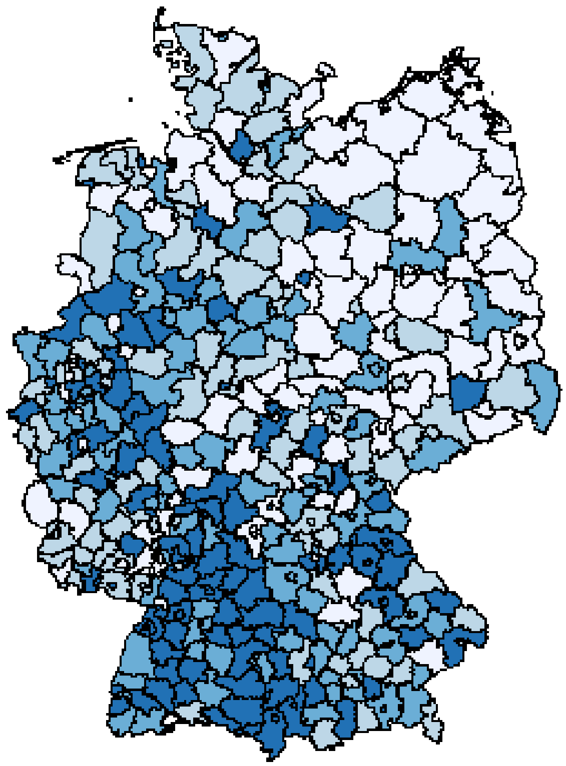

(b) Robot exposure

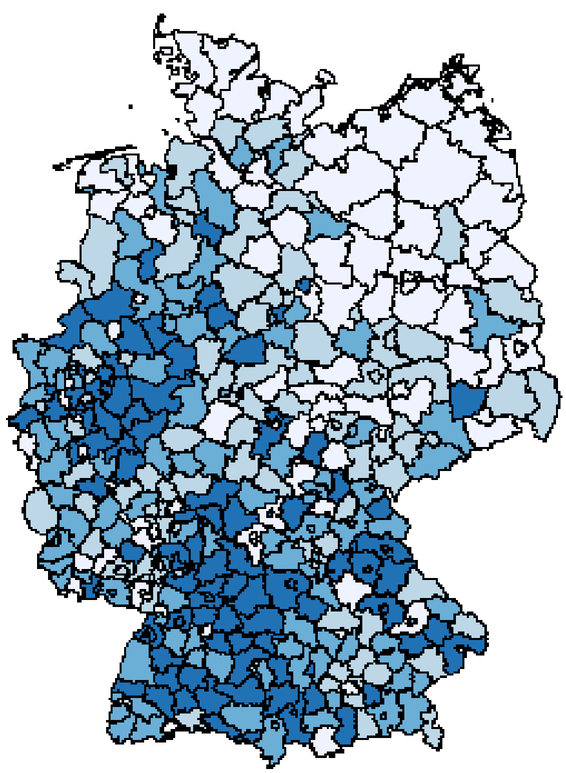

Figure 4: Correlation between Patents and Local Labor Market Outcomes

(a) Robot Exposure and Local Employment

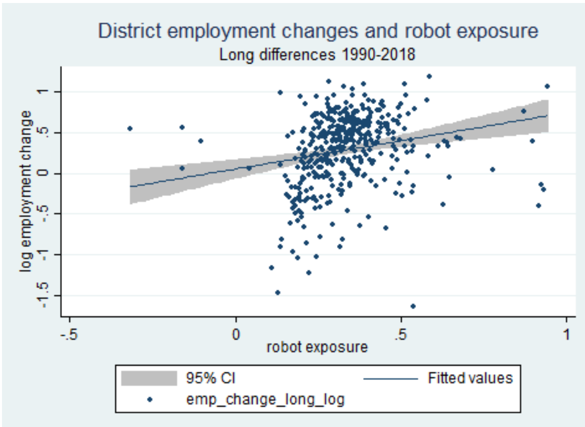

(b) AI Exposure and Local Employment

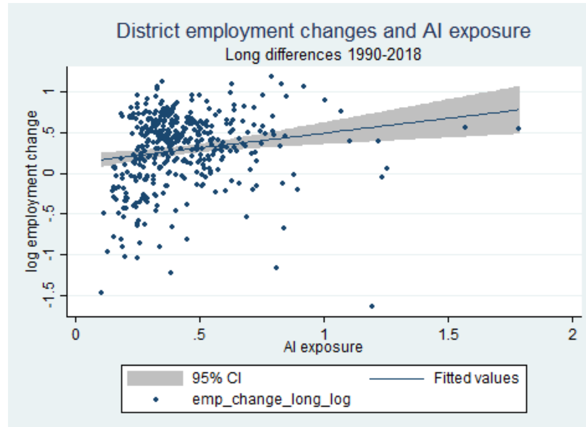

(c) Robot Exposure and Local Wages

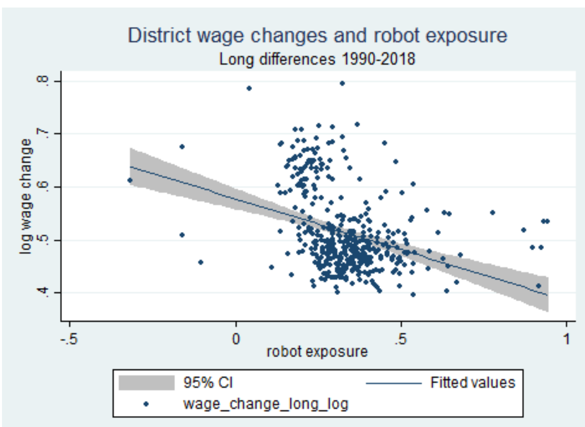

(d) AI Exposure and Local Wages

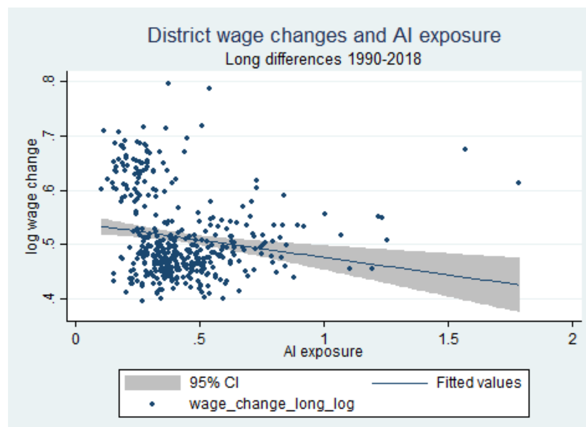

Table 1: Local Labor Market Characteristics and Exposure Measures

	Obs	Mean	Std. Dev.	Min	Max
Outcome variables:					
Δ Log Employment	401	-0.5317	0.8632	-3-4593	1.5946
Δ Log Daily Wages	401	0.4853	0.0970	0.2594	0.8871
Exposure Measures:					
Δ AI Exposure (log)	401	3.1013	2.5634	.3372	20.3104
Δ Robot Exposure (log)	401	8.6630	7.8316	.3149	54.7126
Control variables:					
Manufacturing Employment (1993)	401	0.2797	0.1686	0.0185	0.8896
High-skilled Workers (1993)	401	0.0773	0.0535	0.0155	0.4126
Medium-skilled Workers (1993)	401	0.7332	0.0507	0.4703	0.8773
Low-skilled Workers (1993)	401	0.1792	0.0537	0.0617	0.3569
% Female Employment (1993)	401	0.4220	0.0778	0.1282	0.6377
% Young Workers (1993)	401	0.3975	0.0392	0.2846	0.5346
% Prime-Aged Workers (1993)	401	0.3488	0.0301	0.2770	0.4381
% Older Workers 50-64 (1993)	401	0.2033	0.0286	0.1094	0.2744
Δ Net exports	401	0.7584	0.2411	0.7308	0.8768
Δ ICT investment	401	0.2527	0.0268	0.1854	0.3950

Notes: The outcome and exposure measures are measured at the district level and calculated as log changes between 1993 and 2018. The exposure measures are a shift-share variable consisting of two components: the shift variable denotes the cumulative number of patents in robotics and AI technologies used in a certain industry. The information on patents is extracted from EPO data using natural language processing techniques. The share variable is the employment share of an industry in the district in 1993. All control variables are measured at the district level and refer to 1993. Young workers are between 25 and 24 years of age; prime-aged workers between 35 and 49 years of age; and older workers between 50 and 64 years of age.

Table 2: Employment and Wage Effects of Exposure to AI and Robotics Technologies

Panel A: Employment Changes	(1)	(2)	(3)	(4)	(5)
ΔAI Exposure	0.0104**	0.0105**			0.0155***
	(0.00452)	(0.00452)			(0.00572)
ΔRobot Exposure			0.000929	0.000905	-0.00351
			(0.00205)	(0.00204)	(0.00258)
Δ Net exports	No	Yes	No	Yes	Yes
Δ ICT investment	No	Yes	No	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes
Industry shares	Yes	Yes	Yes	Yes	Yes
Region dummies	Yes	Yes	Yes	Yes	Yes
Observations	401	401	401	401	401

Panel B: Wage Changes	(1)	(2)	(3)	(4)	(5)
ΔAI Exposure	0.000162	0.000175			-0.000837
	(0.00169)	(0.00168)			(0.00232)
ΔRobot Exposure			0.000482	0.000471	0.000709
			(0.000811)	(0.000806)	(0.00108)
Δ Net exports	No	Yes	No	Yes	Yes
Δ ICT investment	No	Yes	No	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes
Industry shares	Yes	Yes	Yes	Yes	Yes
Region dummies	Yes	Yes	Yes	Yes	Yes
Observations	401	401	401	401	401

Notes: The table reports estimates from equation (3) where the dependent variables are log employment changes between 1993 and 2018 in Panel A and log wage changes between 1993 and 2018 in Panel B. The exposure measures are shift share variables as defined in equation (2). Demographic controls include the share of female workers, the share of high-, medium- and low-skilled workers, the share of young, prime-aged and older workers and the industry shares at the one-digit level. All demographic control variables refer to 1993. Δ Net exports is the change in net exports between 1993 and 2018 at the one-digit industry level, adjusted by the total wage bill. Δ ICT is the change in ICT investment per worker between 1993 and 2018 at the one-digit industry level. Region dummies refer to broad regions (North, South, East, West). Standard errors are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Table 3: Panel Estimates of Employment and Wage Effects

Panel A: Employment Changes						
	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.000389		0.00309	-0.0131**		-0.00623*
	(0.00163)		(0.00211)	(0.00295)		(0.00356)
ΔRobot Exposure		-0.00137	-0.00233**		-0.00634***	-0.00425**
		(0.000863)	(0.00114)		(0.00143)	(0.00174)
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Panel B: Wage Changes						
	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.00132		-0.00267*	-0.00239		-0.000648
	(0.00108)		(0.00138)	(0.00223)		(0.00288)
ΔRobot Exposure		0.0000798	0.000904		-0.00129	-0.00107
		(0.000527)	(0.000667)		(0.000955)	(0.00123)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Notes: The table reports estimates from equation (6) where the dependent variables are log employment changes in four sub-periods in Panel A and log wage changes in Panel B. The exposure measures are shift share variables as defined in equation (5). Demographic controls include the share of female workers, the share of high-, medium- and low-skilled workers, the share of young, prime-aged and older workers and the industry shares at the one-digit level. All demographic control variables refer to the first year of the respective sub-period. Δ Net exports is the change in net exports at the one-digit industry level, adjusted by the total wage bill. Δ ICT is the change in ICT investment per worker at the one-digit industry level. Region dummies refer to broad regions (North, South, East, West). Standard errors are clustered at the district level and are reported in parentheses. Significance levels: * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table 4: Panel Estimates of Employment Effects in and outside of Manufacturing

Panel A: Manufacturing Employment						
	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.00350		-0.00380	-0.0361***		-0.0224**
	(0.00519)		(0.00466)	(0.00105)		(0.0110)
ΔRobot Exposure		-0.000973	0.000199		-0.0160***	-0.00845
		(0.00282)	(0.00298)		(0.00526)	(0.00607)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Panel B: Non-Manufacturing Employment						
	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.000808		0.00382*	-0.00231		0.00459
	(0.00150)		(0.00205)	(0.00292)		(0.00432)
ΔRobot Exposure		-0.00192***	-0.00309***		-0.00273**	-0.00426**
		(0.00106)	(0.00163)		(0.00243)	(0.00403)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Notes: The table reports estimates from equation (6) where the dependent variables are log employment changes in four sub-periods in Manufacturing in Panel A and outside of Manufacturing in Panel B. The exposure measures are shift share variables as defined in equation (5). Demographic controls include the share of female workers, the share of high-, medium- and low-skilled workers, the share of young, prime-aged and older workers and the industry shares at the one-digit level. All demographic control variables refer to the first year of the respective sub-period. Δ Net exports is the change in net exports at the one-digit industry level, adjusted by the total wage bill. Δ ICT is the change in ICT investment per worker at the one-digit industry level. Region dummies refer to broad regions (North, South, East, West). Standard errors are clustered at the district level and are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Table 5: Panel Estimates of Employment Effects by Skill Level in and outside of Manufacturing

Panel A: Employment Changes in Manufacturing	Low-skilled Workers (1)	Medium-skilled Workers (2)	High-skilled Workers (3)	Low-skilled Workers (4)	Medium-skilled Workers (5)	High-skilled Workers (6)
ΔAI Exposure	0.00513	-0.00465	-0.00734	-0.00226	-0.0240**	-0.0165
	(0.00631)	(0.00460)	(0.00120)	(0.0153)	(0.0105)	(0.0199)
ΔRobot Exposure	-0.00401	0.000357	-0.00120	-0.0194***	-0.00641	-0.0171
	(0.00360)	(0.00283)	(0.00465)	(0.00674)	(0.00553)	(0.0111)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Panel B: Employment Changes outside of Manufacturing	Low-skilled Workers (1)	Medium-skilled Workers (2)	High-skilled Workers (3)	Low-skilled Workers (4)	Medium-skilled Workers (5)	High-skilled Workers (6)
ΔAI Exposure	0.000412	0.00352*	0.00867**	0.00202	0.00213	0.0162*
	(0.00295)	(0.00196)	(0.00366)	(0.00725)	(0.00390)	(0.00914)
ΔRobot Exposure	-0.000981	-0.00367***	-0.00323**	-0.00959***	-0.00360*	-0.00496
	(0.00130)	(0.000810)	(0.00146)	(0.00341)	(0.00184)	(0.00389)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Notes: The table reports estimates from equation (6) where the dependent variables are log employment changes by skill level in Manufacturing in Panel A and outside of Manufacturing in Panel B. The exposure measures are shift share variables as defined in equation (5). Demographic controls include the share of female workers, the share of high-, medium- and low-skilled workers, the share of young, prime-aged and older workers and the industry shares at the one-digit level. All demographic control variables refer to the first year of the respective sub-period. Δ Net exports is the change in net exports at the one-digit industry level, adjusted by the total wage bill. Δ ICT is the change in ICT investment per worker at the one-digit industry level. Region dummies refer to broad regions (North, South, East, West). Standard errors are clustered at the district level and are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

A Additional Results

Table A1: Industries Using the Most AI and Robotics Patents

ISIC	Industry	AI Patents	% AI Grants in Industry
2620	Computers and peripheral equipment	582	1.49
2640	Consumer electronics	316	4.45
2630	Communication equipment	139	0.55
2817	Office machinery and equipment	138	2.00
2822	Metal-forming machinery and machine tools	74	0.86
2670	Optical instruments and photographic equipment	72	0.25
5912	Motion picture, video and television programme post-production activities	69	0.47
5920	Sound recording and music publishing activities	65	0.24
2651	Measuring, testing, navigating and control equipment	52	0.35
8620	Medical and dental practice activities	40	0.18
ISIC	Industry	Robotics Patent	% Robotics Patents in Industry
2814	Bearings, gears, gearing and driving elements	413	1.80
2822	Metal-forming machinery and machine tools	228	2.63
2651	Measuring, testing, navigating and control equipment	169	1.13
2750	Domestic appliances	148	1.15
2592	Treatment and coating of metals	138	1.89
2811	Engines and turbines	137	0.41
1050	Dairy products	137	4.93
2816	Lifting and handling equipment	117	0.82
2670	Optical instruments and photographic equipment	101	0.36
2620	Computers and peripheral equipment	97	0.24

Notes: The table reports the top ten four-digit industries using AI (top panel) and robotics patents (bottom panel). The second column reports the total number of patent grants used in the industry during the 1990-2018 period, while the last column reports the share of AI resp. robotics patents to all patents used in the industry.

Table A2: Industries with Strongest Growth in AI and Robotics Patents

ISIC	Industry	Growth in AI Patents
262	Manufacture of computers and peripheral equipment	5.76
264	Manufacture of consumer electronics	3.71
862	Medical and dental practice activities	3.70
263	Manufacture of communication equipment	3.53
267	Manufacture of optical instruments and photographic equipment	3.36
265	Manufacture of measuring, testing, navigating and control equipment; watches and clocks	3.26
592	Sound recording and music publishing activities	3.01
281	Manufacture of general-purpose machinery	2.76
282	Manufacture of special-purpose machinery	2.52
749	Other professional, scientific and technical activities	2.43
ISIC	Industry	Growth in Robotics Patents
862	Medical and dental practice activities	4.35
267	Manufacture of optical instruments and photographic equipment	3.36
262	Manufacture of computers and peripheral equipment	3.32
325	Manufacture of medical and dental instruments and supplies	3.00
105	Manufacture of dairy products	2.74
360	Water collection, treatment and supply	2.52
310	Manufacture of furniture	2.40
202	Manufacture of other chemical products	2.17
960	Other personal service activities	2.11
201	Manufacture of basic chemicals, fertilizers and nitrogen compounds, plastics and synthetic rubber	2.11

Notes: The table reports the top ten four-digit industries using AI (top panel) and robotics patents (bottom panel). The second column reports the total number of patent grants used in the industry during the 1990-2018 period, while the last column reports the share of AI resp. robotics patents to all patents used in the industry.

Table A3: Panel Estimates of Wage Effects in Manufacturing vs. Non-Manufacturing

Panel A: Manufacturing Wages

	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.0000814		-0.000509	-0.00586*		-0.00401
	(0.00141)		(0.00172)	(0.00326)		(0.00422)
ΔRobot Exposure		0.000129	0.000285		-0.00249	-0.00115
		(0.000674)	(0.000839)		(0.000170)	(0.00224)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Panel B: Non-Manufacturing Wages

	(1)	(2)	(3)	(4)	(5)	(6)
ΔAI Exposure	-0.000926		-0.00268**	-0.00220		-0.00188
	(0.000817)		(0.00105)	(0.00153)		(0.00215)
ΔRobot Exposure		0.000348	0.00118***		-0.000829	-0.000201
		(0.000335)	(0.000436)		(0.000731)	(0.00103)
District FE	No	No	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry employment shares	Yes	Yes	Yes	Yes	Yes	Yes
Δ Net exports	Yes	Yes	Yes	Yes	Yes	Yes
Δ ICT investment	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1604	1604	1604	1604	1604	1604

Notes: The table reports estimates from equation (6) where the dependent variables are log employment changes in four sub-periods in Panel A and log wage changes in Panel B. The exposure measures are shift share variables as defined in equation (5). Demographic controls include the share of female workers, the share of high-, medium- and low-skilled workers, the share of young, prime-aged and older workers and the industry shares at the one-digit level. All demographic control variables refer to the first year of the respective sub-period. Δ Net exports is the change in net exports at the one-digit industry level, adjusted by the total wage bill. Δ ICT is the change in ICT investment per worker at the one-digit industry level. Region dummies refer to broad regions (North, South, East, West). Standard errors are clustered at the district level and are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Figure A1: Evolution of Patents by Broad Technology Class

(a) Robotics Patents

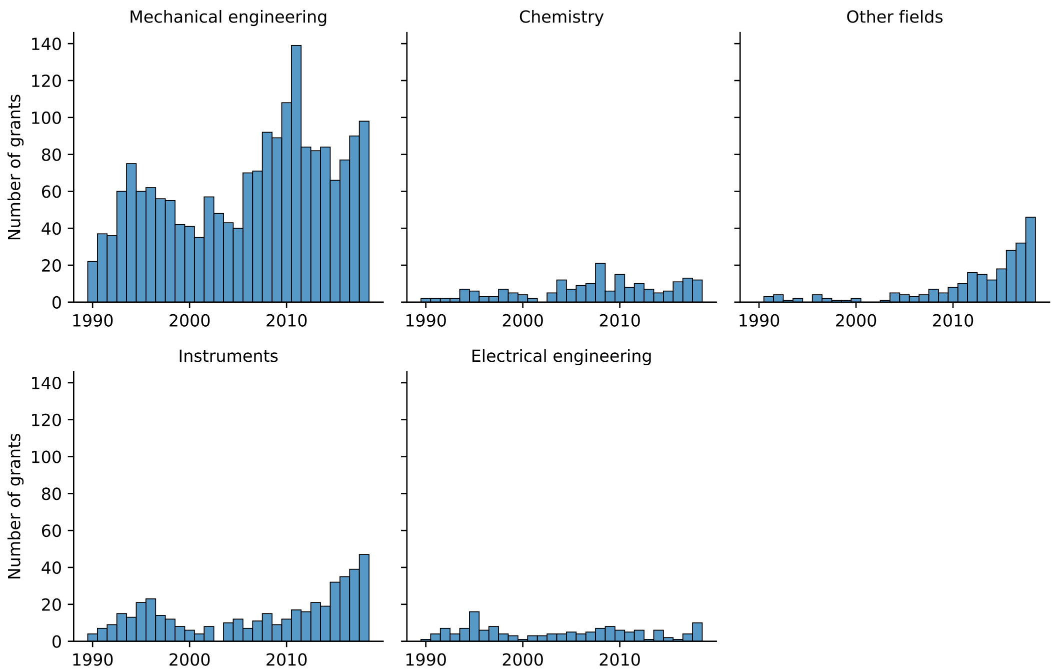

(b) AI Patents

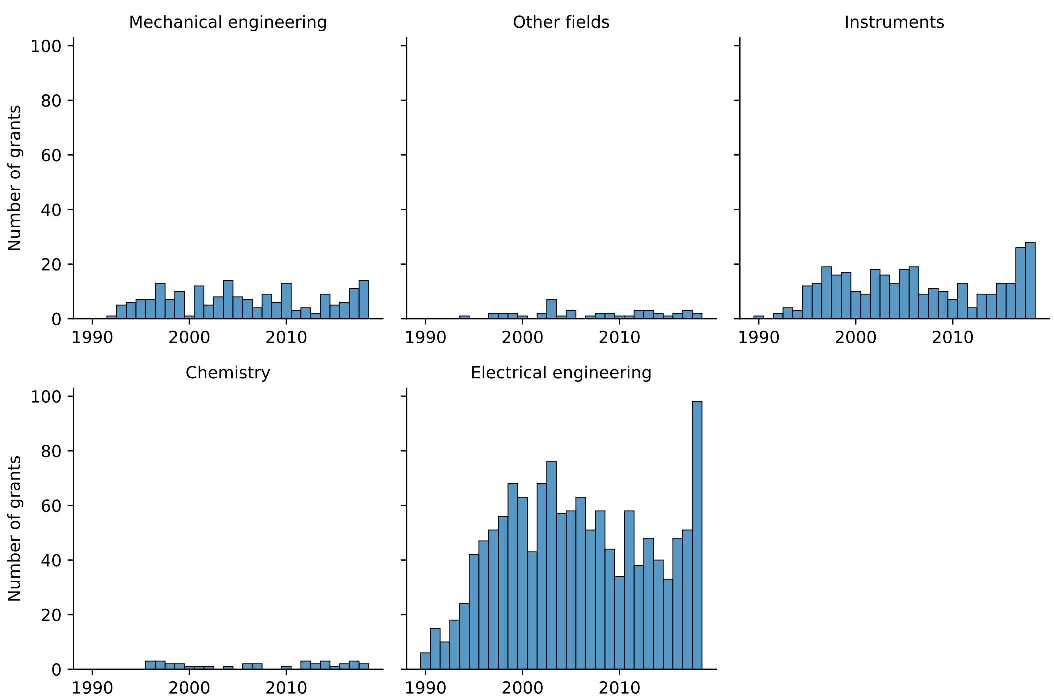

Notes: The figures show the number of patent grants in Robotics (Panel (A)) and AI (Panel (B)) in broad technology classes over time.