

Omar GUENNICH, PhD Candidate

omar.guennich@um5r.ac.ma

Faculty of Law, Economics and Social Sciences, Mohammed V University in Rabat,
Souissi-Rabat, Morocco

Abdellah YOUSFI, PhD

a.yousfi@um5r.ac.ma

Faculty of Law, Economics and Social Sciences, Mohammed V University in Rabat,
Souissi-Rabat, Morocco

Evaluating the Potential Economic Impact of Artificial Intelligence on Professions: A Refined Approach Using Ensemble Embeddings

Abstract. *This study explores the potential impact of artificial intelligence (AI) on the labour market by leveraging a deep learning-based natural language processing framework. We apply an ensemble embeddings approach to assess the alignment between AI capabilities and the skill requirements of occupations. Our findings indicate that cognitive-intensive skills, such as reading comprehension, operations analysis, and writing, exhibit the highest exposure, at the occupational level, professions like technical writers and mathematicians emerge as among the most exposed, while at the sector level, the results reveal that Education, Information and Communication Technology (ICT), and Finance/Insurance demonstrate the highest levels of AI exposure. The results also indicate a significant positive correlation between AI exposure and salary levels, with occupations requiring higher education degrees showing greater AI exposure, suggesting that such jobs are more likely to be influenced by AI capabilities.*

These findings challenge the conventional view that technology primarily automates routine tasks, instead revealing significant AI exposure in cognitively demanding, high-wage occupations. these results underscore the need for policy interventions focused on workforce adaptation and reskilling programs for safely integrating AI to work.

Keywords: *jobs, Artificial Intelligence, professions, technological unemployment, AI and Economy.*

JEL Classification: J23, J24, O33.

Received: 7 March 2026	Revised: 11 September 2026	Accepted: 20 September 2026
------------------------	----------------------------	-----------------------------

1. Introduction

Throughout history, many inventions like the steam engine, railways, electricity, and other innovations increased the productivity during the 19th and 20th centuries, these innovations led to significant increases in total factor productivity and gross domestic product. However, this technological progress not only increased productivity but also resulted in an increase in unemployment due to mechanization,

DOI: 10.24818/18423264/60.3.26.15

© 2026 The Authors. Published by Editura ASE. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

this phenomenon is known as "technological unemployment". A prime example is the onset of massive industrialization in Europe, where factories replaced artisanal workshops and machines supplanted human labourers. Moreover, this trend intensified during the Second Industrial Revolution, which transformed the economic structures and the nature of work.

In recent years, a similar trend has been observed with the rise of artificial intelligence, the advances in this field have progressed exponentially and many studies state that AI has caught or even surpassed human performance in some tasks like image classification (He et al., 2015).

This development raises a fundamental question: which types of human work are most exposed to AI, and what determines the degree of this exposure? What is the nature of this exposure? The economic consequences of AI depend on whether it substitutes for human labour, complements workers by increasing their productivity, or generates entirely new tasks and forms of employment.

2. Problem Statement and Methodology

This paper focuses on the analysis of AI exposure. Specifically, it applies a sophisticated textual analysis to identify the skills, professions, and economic sectors most susceptible to be exposed to artificial intelligence and to uncover the underlying factors that drive this exposure. First It characterizes professions using a streamlined set of 35 base skills derived from the O*NET (Occupational Information Network) database. Second, it incorporates 28 distinct AI application categories based on Stanford University AI Index Report (2024). Third, it introduces an ensemble embeddings approach by utilizing two specialized sentence embedding models to calculate AI exposure for individual skills, occupations and sectors, ensuring a robust and comprehensive analysis. and Finally, it checks the validity and robustness of the resulting metric against external criteria, and benchmarks a range of alternative encoders to justify the choice of these two.

3. Literature Review and Previous Work

Since the Industrial Revolution, the relationship between technological advancement and employment has been a focal point of debate. Over the past few decades, researchers have examined the nuanced effects of technology on labour markets.

Feldmann (2013), in his empirical analysis based on data from 21 countries, found that technology initially led to a significant increase in unemployment during the first three years.

In contrast, Autor (2015) presented a more optimistic perspective, suggesting that automation often generates new tasks and complementary roles.

Frey and Osborne (2017) addressed the question: To what extent are jobs likely to be computerized? They calculated the probability of computerization for 702

professions using a Gaussian model, they found that up to 47% of jobs in the US are at risk of automation.

A study (Manyika et al., 2017) from the McKinsey Global Institute indicated that half of tasks carried out by workers could be automated using the current technology.

Gregory et al. (2022) examined routine-replacing technological change in 238 European regions and found that the benefits of product demand more than offset the substitution effect, increasing net labour demand.

In recent years, this discussion has evolved with the emergence of Artificial Intelligence (AI), raising specific questions about the impact of this specific technology on employment, the professions most affected, and how to measure these effects.

Brynjolfsson et al. (2018) addressed the same issue by constructing a measure termed “Suitability for Machine Learning” (SML) to evaluate the potential impact of AI on labour inputs in the US economy. Building upon (Brynjolfsson and Mitchell’s, 2017) rubric, they applied this framework to a comprehensive dataset from the O*NET database, containing 2,069 work activities, 18,156 tasks, and 964 occupations. Their analysis revealed that while most occupations in most industries have at least some tasks highly suitable for machine learning, only a very few are fully automatable. Finally, the study highlighted the importance of job redesign, calling for greater focus on it from researchers, managers, and entrepreneurs.

Webb (2020) utilized a text-based analysis method, quantifying the overlap between job task descriptions (from O*NET) and technology-related patent texts to measure the potential impacts of various technologies, including AI, on occupations. Based on US labour market data, his analysis revealed that, unlike software and robots, AI targets high-skilled tasks, making occupations requiring higher education more susceptible to automation. Webb's research also indicated significant gender disparities, with traditionally "male" occupations facing much higher AI exposure than "female" occupations. Finally, his analysis predicted that AI implementation would reduce wage inequality between workers at the 90th and 10th percentiles, while leaving the top 1% earnings largely unchanged.

Felten et al. (2021) proposed a methodology to measure the potential impact of AI on occupations through their AI Occupational Exposure (AIOE) metric. This approach quantifies the overlap between occupational abilities (from O*NET) and AI applications, using a crowd-sourced survey to measure the degree of relatedness. Their research reveals that AI applications are predominantly associated with cognitive and sensory abilities, leaving certain jobs more exposed than others. Their results also highlighted that occupations requiring high levels of problem-solving, logical reasoning, and perception are significantly more exposed than those relying on physical abilities. Finally, their results indicate that exposure also varies significantly across sectors and regional contexts.

Similarly, Septiandri et al. (2024) developed a methodology to quantify the potential impact of AI innovations on occupational exposure by using deep learning to analyse which occupational tasks are influenced by AI advancements, their

approach relied on a comprehensive dataset of 17,879 task descriptions and quantifies AI's potential impact through the analysis of 24,758 AI patents filed with the United States Patent and Trademark Office between 2015 and 2022. Similar to prior research, their findings reveal that exposure to AI varies significantly across occupations in the United States, with certain jobs facing higher levels of exposure. Importantly, the results demonstrate that AI affects not only routine tasks but also non-routine cognitive tasks, such as diagnosing health conditions, programming computers, and tracking flight routes.

While previous studies have focused on measuring the impact of AI on the labour market, their findings often diverge due to fundamentally different methodologies. For example, when considering sectoral exposure in the Finance and Insurance sector, Felten et al. (2021) identified it as among the most exposed, with high scores for Securities and Investments and Insurance Carriers.

In contrast, Septiandri et al. (2024) find Finance and Insurance to be one of the least exposed sectors, these contrasting results illustrate how methodological differences can lead to different conclusions and therefore highlight the sensitivity of results to methodological choices.

Despite these divergent findings, the studies converge on certain points. For example, both studies identify Scientific and Technical Services as highly exposed to AI. Additionally, they agree on the low exposure of sectors grounded in physical labour, such as Construction and Accommodation and Food Services.

Each methodological approach carries inherent limitations that fundamentally shape its analytical scope and findings. Felten et al. (2021) adopt a micro-level approach that links AI capabilities to specific occupational abilities Occupational, yet their framework encounters significant challenges when attempting to map AI applications to certain capabilities such as balance, stamina, or finger dexterity, abilities that remain difficult to quantify through their crowdsourced assessment methodology. In contrast, patent-based approaches employed by Webb (2020) and Septiandri et al. (2024) face a different set of constraints. These methodologies analyse the overlap between patent text descriptions and occupational task descriptions to predict technological impacts on labour markets. However, this approach inherently captures only crystallized technological capabilities at specific temporal snapshots. For instance, Septiandri et al. (2024) analysis, which reflects AI capabilities documented in patents only through 2022, exemplifies a fundamental limitation of patent-based methodologies.

This paper investigates the potential impact of AI on the labour market from a different angle from the studies discussed above: instead of relying on tasks or on abilities, it maps occupational skills to AI applications, drawing on the Stanford University AI Index Report as its primary source, through an ensemble embeddings approach. A comprehensive description of the methodology is provided in the subsequent section.

4. Methodology Presentation

In this paper we quantify the potential impact of AI on employment through semantic exposure. It allows to measures the overlap between AI capabilities and the skills required by occupations. Exposure in this sense is a statement about overlap, not about outcomes. It indicates that AI can affect a skill, occupation, or sector, but it does not say what the outcome will be.

4.1 Used Databases

In this study, we incorporated several key databases that served as the backbone of our analysis. These resources provided essential data for evaluating the potential impact of AI across different dimensions. Specifically, we utilized the following databases:

- The Artificial Intelligence Index Report 2024 by Stanford University: serves as the primary source for identifying AI applications. In their study, these AI applications were referred to as task categories.
- O*NET (Occupational Information Network): A database developed by the U.S. Department of Labor that provides comprehensive information about professions in the United States, including:
 - o Job Descriptions: Details about the responsibilities and tasks associated with each profession.
 - o Skills and Qualifications: Information on the skills, abilities, knowledge, and qualifications required for each occupation.
 - o Working Conditions and Contexts: Information on typical work environments, working conditions, and other contextual characteristics of professions.
- The Occupational Employment and Wage Statistics (OEWS): conducted by the U.S. Bureau of Labor Statistics (BLS), collects comprehensive employment and wage data across industries in the United States. Formerly known as the Occupational Employment Statistics (OES) survey, this program provides detailed occupational information for wage and salary workers, making the data publicly available for research and analysis.
- European Union Labour Force 2022 Survey (EU LFS): A comprehensive resource that compiles data from surveys conducted in the EU member states. this database includes detailed statistical information that enables comparative analyses and trend monitoring over time, providing a deeper understanding of the socio-economic landscape within the EU.
- The Anthropic Economic Index: reports the share of approximately one million anonymised Claude conversations classified against O*NET task statements, which we aggregate to occupations and sectors to obtain a behavioural measure of where AI is actually used.

- Eurostat's survey on ICT usage in enterprises: reports the share of enterprises using at least one AI technology by NACE activity, which provides an official measure of adoption for the European Union.

4.2 Representation of AI Applications as Embedding

To quantify the potential impact of AI, we utilized AI applications from the Stanford AI Index Report 2024 (e.g., Text Classification, Image Generation, Reading Comprehension, and Text Summarization. The complete list of these AI applications is detailed in Table 1). We transformed the text labels into vector representations, known as embeddings, that encode their semantic meanings for use in AI applications.

Table 1. List of AI Task Categories from the Stanford AI Index Report 2024

Text Classification	General language	Image Generation	Medical image segmentation	Moral Reasoning
Text Summarization	Activity Recognition	Image Segmentation	Code Generation	Mathematical Reasoning
Sentiment Analysis	Speech Recognition	Semantic Segmentation	Text Generation	General Reasoning
Reading Comprehension	Factuality	Image Classification	Audio Generation	Causal Reasoning
Natural Language Inference	Agent-based Behavior	Image Instruction Following	Planning	
English language understanding	Human pose estimation	Image Editing	Visual Reasoning	

Source: Authors' processing.

For this purpose, we utilized an ensemble embedding by combining the representations generated by two complementary models: PatentSBERTa (Bekamiri et al., 2021), specialized in patent analysis, and SciBERT (Beltagy et al., 2019), trained on scientific literature. This combination is motivated by the dual register of the texts being compared. They are drawn from a research-oriented taxonomy and therefore largely reflect the vocabulary of the scientific literature and technical analysis. Since the two models produce vectors with different magnitudes, each embedding is first normalized to unit length. Let $M = \{m_1, m_2\}$ denote the set of encoders, $|M| = 2$, and let $e_{m_i}(t)$ be the embedding of text t generated by model $m_i \in M$:

$$\bar{e}_{m_i}(t) = \frac{e_{m_i}(t)}{\|e_{m_i}(t)\|}$$

The two normalized embeddings are then combined with equal weights:

$$E(t) = \frac{1}{|M|} \cdot \sum_{m_i \in M} \bar{e}_{m_i}(t)$$

This ensemble representation $E(t)$ therefore gives equal importance to the semantic information captured by the two models while placing them on the same scale. It provides the common vector representation used in the subsequent analysis to compare AI applications with professional skills.

4.3 The Potential AI Exposure Metric for Professional Skills

Using the ensemble embeddings described above, we measured how closely each AI application is related to the professional skills listed in the O*NET database (e.g., 'programming,' 'writing,' 'oral expression,' 'mathematics,' 'sciences,' 'critical thinking,' and 'active learning'). For each professional skill j , we considered its similarity with the AI applications, to measure the similarity, several metrics can be used, including Euclidean distance, cosine similarity, and Jaccard similarity.

In this study, we opted for cosine similarity, a widely used metric for assessing the alignment between two vectors. For two vectors u and v , cosine similarity is defined as:

$$\text{Sim}(u, v) = \frac{\langle u, v \rangle}{\|u\| \cdot \|v\|}$$




Figure 1. SBERT architecture at inference, for example, to compute similarity scores

Source: Using contextual sentence analysis models to recognize ESG concepts
VL - DO - 10.48550/arXiv.2207.01402 ER.

Following the calculation of relatedness between AI applications and professional skills, we defined the Potential Degree of Exposure of a Skill j to Artificial Intelligence $PDESk(j)$. This metric is used to quantify the overall exposure of a specific skill to the AI applications. For a given skill j , it is calculated as follows:

$$PDESk(j) = \sum_{k=1}^{28} Sim(AI_k, Skill_j)$$

Where $Sim(AI_k, Skill_j)$ represents the similarity between the vector representation of the AI_k and the O*NET title and description of the $Skill_j$.

4.4 The Potential AI Exposure Metric for Professions

Building on the measures proposed by Webb (2020) and Felten et al. (2021), we introduced an enhanced metric utilizing $PEDO(o)$. Consequently, our Potential Exposure Degree of an Occupation o to Artificial Intelligence, denoted as $PEDO(o)$, is defined by the following equation:

$$PEDO(o) = \frac{\sum_{j=1}^{35} PDESk_j \cdot L_{j,o} \cdot I_{j,o}}{\sum_{j=1}^{35} L_{j,o} \cdot I_{j,o}}$$

We weighted the Potential Degree of Exposure of a Skill to AI by the prevalence $L_{j,o}$ and the importance $I_{j,o}$ of each skill within a given occupation. This metric is based on the methodology of Felten et al. (2021). This weighting approach accounts for the varying presence of skills across different professions. Essential skills were assigned higher prevalence and importance scores, while less critical or less frequently used skills were given lower scores. This approach ensures that the potential impact of AI on an occupation accurately reflected the significance of the skills required for that profession. Finally, because the resulting index varies within a narrow band, all results are reported in standardised form across the occupations, so that a negative value indicates below-average exposure rather than its absence.

4.5 The Potential AI Exposure Metric for Sectors

To evaluate the potential impact of AI on a sector s , we developed a metric, denoted as $PEDSe(s)$, that quantifies the degree of AI's potential impact on the sector. This metric integrates the impact degrees of various professions influenced by AI, weighted by each profession's relative share of employment within the sector's total workforce. By aggregating these weighted impact values across all relevant professions within the sector, it is defined as follows:

$$PEDSe(s) = \sum_{o \in occupation(s)} PEDO(o) \times \frac{EMP_o}{\sum_{o \in occupation(s)} EMP_o}$$

Where EMP_o represents the employment in occupation o , and the term $\frac{EMP_o}{\sum_{o \in occupation(s)} EMP_o}$ represents the proportion of employment for occupation o within sector s .

By weighting each occupation's potential impact according to its relative share of employment within the sector, this method assigns greater significance to occupations that make up a larger portion of the sector's workforce, while diminishing the influence of occupations with smaller employment shares. Similarly, to PEDO, Sector scores are also standardised within each economy, across the 19 US and the 16 EU sectors respectively.

4.6 Implementation Methodology

During our analysis, we utilized the Python programming language, taking advantage of its comprehensive libraries and tools for data processing and analysis. For generating embedding vectors, we utilized the sentence-transformers package, leveraging pre-trained models PatentSBERTa and SciBERT from Hugging Face to map data into a 768-dimensional dense vector space. These embeddings were created using the SentenceTransformer class and its encode function. To measure similarity, we employed the cosine_similarity function provided by scikit-learn library to compute the semantic similarity between vectors.

5. Results

5.1 The Potential AI Exposure Metric for Professional Skills

We first represented the AI application labels and professional skills listed in the O*NET database as embeddings. These embeddings were created using an ensemble of the pre-trained models PatentSBERTa and SciBERT, as detailed previously. We then computed the semantic similarity between AI applications and professional skills.

Using the similarity matrix, we computed the AI exposure metric $PDES_k$ across the skills listed in the O*NET database to quantify the Potential AI Exposure to Professional Skills, providing insights into how AI technology may impact the workforce depending on the skills involved. The calculated exposure scores for the skills are presented in Table 2.

Table 2. Potential Degree of Exposure Scores of Skills to Artificial Intelligence (PDESk)

Skill	PDESk Score	Skill	PDESk Score
Negotiation	4.014	Repairing	5.942
Management of Financial Resources	4.016	Judgment and Decision Making	6.007
Time Management	4.711	Speaking	6.008
Equipment Maintenance	4.753	Complex Problem Solving	6.397
Service Orientation	4.920	Technology Design	6.466

Skill	<i>PDESk</i> Score	Skill	<i>PDESk</i> Score
Operations Monitoring	5.003	Social Perceptiveness	6.574
Monitoring	5.080	Active Learning	6.630
Installation	5.122	Instructing	6.648
Quality Control Analysis	5.201	Critical Thinking	6.711
Management of Personnel Resources	5.295	Coordination	6.722
Management of Material Resources	5.349	Learning Strategies	6.728
Troubleshooting	5.514	Mathematics	6.808
Systems Evaluation	5.723	Science	6.856
Equipment Selection	5.766	Programming	7.323
Persuasion	5.767	Operations Analysis	7.421
Operation and Control	5.826	Writing	8.079
Active Listening	5.886	Reading Comprehension	8.394
Systems Analysis	5.921		

Source: Authors' processing.

5.2 The Potential AI Exposure Metric for Professions

Moving from skills to broader professional impacts, we used the *PEDO* metric to assess how entire professions might be exposed to AI technologies. The calculated scores are presented in Table 3, while their distribution is illustrated in the histogram shown in Figure 2.




Figure 2. Histogram of the Potential Impact of AI on O*NET Professions

Source: Authors' processing.

Table 3. Potential Degree of Exposure Scores of a Sample of Professions to Artificial Intelligence (PEDO)

Profession	PEDO Score	Profession	PEDO Score
Chief Executives	-0.248	Operations Research Analysts	2.097
General and Operations Managers	-0.458	Statisticians	2.690
Advertising and Promotions Managers	0.127	Directors, Religious Activities and Education	0.361
Credit Analysts	1.212	Title Examiners, Abstractors, and Searchers	1.430
Personal Financial Advisors	0.660	Business Teachers, Postsecondary	1.911
Insurance Underwriters	1.137	Computer Science Teachers, Postsecondary	1.666
Financial Examiners	1.068	Mathematical Science Teachers, Postsecondary	2.122
Computer Programmers	1.436	Technical Writers	2.734
Web Developers	1.427	Writers and Authors	2.369
Actuaries	1.123	Interpreters and Translators	1.152
Mathematicians	3.088	Telemarketers	0.059

Source: Authors' processing.

We further aggregated these professional-level results across major occupational groups using data from the EU Labour Force Survey (EU LFS). This aggregation revealed systemic patterns in AI exposure across different professional categories, as illustrated in Figure 3.




Figure 3. Potential Degree of Exposure of Professions by 1-digit ISCO Occupational Groups to Artificial Intelligence – European Union (PEDO)

Source: Authors' processing.

The following presents key findings from the O*NET database, identifying the occupations that are most and least susceptible to AI-driven impacts. These results provide valuable insights into the differential influence of AI across the occupational landscape. The results are presented in Table 4.

**Table 4. Potential Degree of Exposure Scores
of Professions to Artificial Intelligence – Highest and Lowest Values (PEDO)**

Highest scoring	Lowest scoring
Mathematicians	Agricultural Inspectors
Editors	Cement Masons and Concrete Finishers
Statisticians	Fishers and Related Fishing Workers
Political Scientists	Outdoor Power Equipment and Other Small Engine...
Phlebotomists	Home Appliance Repairers
Technical Writers	Earth Drillers, Except Oil and Gas
Social Scientists and Related Workers, All Other	First-Line Supervisors of Housekeeping and Jan...
Nursing Instructors and Teachers, Postsecondary	Farmworkers, Farm, Ranch, and Aquacultural Ani...
Agricultural Sciences Teachers, Postsecondary	Service Unit Operators, Oil, Gas, and Mining
Psychology Teachers, Postsecondary	Mine Cutting and Channeling Machine Operators
Paralegals and Legal Assistants	Solar Photovoltaic Installers
Recreation and Fitness Studies Teachers, Posts...	Heating, Air Conditioning, and Refrigeration M...
Historians	Tile and Marble Setters
Operations Research Analysts	Locksmiths and Safe Repairers

Source: Authors' processing.

5.3 The Potential AI Exposure Metric for Sectors

At the sector level, we employed the *PEDSe* metric to evaluate the potential impact of AI technologies across industries in both the United States and the European Union and uncovered systemic patterns of AI exposure across diverse advanced economies.

We first calculated *PEDSe* scores for U.S. economic sectors. The results are shown in Figure 4 and detailed in Table 5.




Figure 4. Potential Degree of Exposure of Sectors to Artificial Intelligence – United States (PEDSe)

Note: *, ** and *** denote for Professional, Scientific, and Technical Services; Administrative and Support and Waste Management and Remediation Services; and Other Services (except Public Administration) respectively.

Source: Authors' processing.

Table 5. Potential Degree of Exposure Scores of Sectors to Artificial Intelligence – United States (PEDSe)

Sector	PEDSe Score	Sector	PEDSe Score
Transportation and Warehousing	-0.811	Administrative and Support and Waste Management	-0.287
Manufacturing	-0.573	Management of Companies and Enterprises	1.115
Finance and Insurance	1.440	Retail Trade	-0.381
Professional, Scientific, and Technical Services	1.638	Wholesale Trade	-0.009
Arts, Entertainment, and Recreation	-0.096	Information	0.676
Construction	-1.065	Agriculture, Forestry, Fishing and Hunting	-1.513
Other Services (except Public Administration)	-0.302	Utilities	-0.310
Health Care and Social Assistance	1.074	Mining, Quarrying, and Oil and Gas Extraction	-1.194

Sector	PEDSe Score	Sector	PEDSe Score
Educational Services	1.797	Real Estate and Rental and Leasing	-0.299
Accommodation and Food Services	-0.901		

Source: Authors' processing.

A corresponding analysis was conducted for EU sectors. The results are presented in Figure 5 and detailed in Table 6.




Figure 5. Potential Degree of Exposure of Sectors to Artificial Intelligence – European Union (PEDSe)

Note: * denotes Arts and recreation and other services

Source: Authors' processing.

Table 6. Potential Degree of Exposure Scores of Sectors to Artificial Intelligence – European Union (PEDSe)

Sector	PEDSe Score	Sector	PEDSe Score
Accommodation Food	-0.571	Health and Social Care	0.866
Administrative services	-0.379	ICT services	1.411
Agriculture Fishing Hunting	-1.434	Manufacturing	-0.599
Arts and recreation and other services	0.266	Mining and Quarrying	-1.108
Construction	-1.133	Professional Services	1.257

Sector	PEDSe Score	Sector	PEDSe Score
Education	1.649	Transport Storage	-0.807
Energy Supply Service	0.367	Water and Waste Treatement	-0.698
Finance and Insurance	1.089	Wholesale Retail Trade	-0.176

Source: Authors' processing.

6. Exposure to AI by Education Level and Salary

6.1 Potential AI Exposure by Education Level

To further explore the impact of AI, we examined how exposure varies across professions based on the education level required for each job, categorizing professions accordingly to highlight variations in AI exposure. The results are presented in the next Figure.




Figure 6. Potential Degree of Exposure by Education Level to Artificial Intelligence (PEDO)

Note: This figure plots the PEDO metric (y-axis) categorized by levels of education (x-axis) for each profession.

Source: Authors' processing.

6.2 Potential AI Exposure by Salary

Similarly, we analysed the relationship between AI exposure degrees and salaries, investigating the linear correlation between the AI exposure levels and annual wages for each profession. The results are presented below.




Figure 7. Relationship Between the Potential Degree of Exposure to Artificial Intelligence and Annual Mean Wage (PEDO)

Note: This figure plots the regression between the PEDO metric(x-axis) and mean salary (y-axis) for each profession, where professions are grouped into 100 equal sized bins based on PEDO metric.

Source: Authors' processing.

Table 7. Correlation between Salary and Potential Degree of Exposure (PEDO) scores

	Pearson	Spearman
PEDO vs annual mean wage (levels)	0.3915***	0.5038
PEDO vs log annual mean wage	0.4732***	0.5038

Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Source: Authors' processing.

We find a correlation of 0.392 between PEDO and annual mean wage, rising to 0.473 in logs. We also estimate least-squares specifications with progressively richer controls, to ensure robustness. Standard errors are heteroskedasticity-robust. It is also, quantile regressions of log annual mean wage on PEDO at five points of the conditional distribution were also used. The control variables included are log occupational employment and, in the full models, fixed effects for the 22 major SOC groups. results are shown in table 8 and 9.

Table 8. Regression of Log Annual Mean Wage on the Potential Degree of Exposure (PEDO)

Ordinary least squares	β (PEDO)	Robust s.e.	R²
(1) Bivariate	0.215***	(0.013)	0.224
(2) + log employment	0.215***	(0.013)	0.230
(3) + major occupational group fixed effects	0.165***	(0.022)	0.562
(4) Specification (3), employment-weighted	0.093*	(0.048)	0.722

Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Source: Authors' processing.

In the bivariate regression, a one-standard-deviation increase in PEDO is associated with a 0.215 log-point higher annual mean wage, and exposure alone accounts for roughly 22 per cent of the cross-occupational variance in log wages. Adding log employment leaves the coefficient unchanged, ruling out an occupation-size artifact. The relationship is not merely a reflection of the broad occupational hierarchy: within the 22 major SOC groups the coefficient remains at 0.165, about three-quarters of its unconditional value, even as the fixed effects raise the explained variance to 0.562. Weighting by employment reduces the coefficient to 0.093, indicating that the gradient is steepest among small, specialised, highly paid occupations and considerably flatter for the typical worker.

Table 9. Quantile Regressions of Log Annual Mean Wage on the Potential Degree of Exposure (PEDO)

Quantile regression	B	s. e.
$\tau = 0.10$	0.122***	0.014
$\tau = 0.25$	0.183***	0.017
$\tau = 0.50$	0.216***	0.016
$\tau = 0.75$	0.223***	0.018
$\tau = 0.90$	0.278***	0.027

Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Source: Authors' processing.

The quantile estimates show where in the wage distribution this association is concentrated. the coefficient rises monotonically from 0.122 at the tenth percentile to 0.278 at the ninetieth, so that exposure is most tightly linked to pay at the top of the distribution.

7. Validation and Robustness

Moreover, it is important to find whether the metric agrees with independently constructed measures and with observed behaviour, and whether it survives the choice of encoder.

7.1 Criterion Validity at the Occupational Level

The behavioural criterion is the Anthropic Economic Index (Handa et al., 2025), which classifies roughly one million anonymised Claude conversations against O*NET task statements; aggregating those shares gives an observed measure of where generative AI is actually used. Each row below is computed on a single common sample, so the three columns are directly comparable.

Table 10. Criterion validity: rank correlations with external criteria

Criterion	This paper	Felten et al. (2021)	Webb (2020)
Observed AI use, occupations	0.594	0.679	0.081
Observed AI use, US sectors	0.893	0.730	-0.296
Enterprise AI adoption, EU sectors	0.952	—	—

Source: Authors' processing.

At the occupational level our index trails the AIOE, which is what a measure built without survey input should do. Aggregation to sectors reverses the ordering.

7.2 The Choice of Encoder

Finally, it is worth justifying the choice of these two particular models. Table 11 re-estimates the whole chain with each encoder alone and with five widely used alternatives, holding everything else fixed: only the embedding changes, and every specification passes through the identical weighting, aggregation and standardisation.

Table 11. Embedding-model Ablation

Embedding specification	AI use (occupations)	US sectors	EU sectors	Overall (Mean)
Ensemble: PatentSBERTa + SciBERT (this paper)	0.594	0.893	0.952	0.813
PatentSBERTa only	0.598	0.872	0.952	0.807
SciBERT only	0.596	0.828	0.976	0.800
all-MiniLM-L6-v2	0.587	0.884	0.903	0.791
all-mpnet-base-v2	0.546	0.888	0.903	0.779
BGE-large-en-v1.5	0.477	0.796	0.976	0.750
Qwen3-Embedding-0.6B	0.332	0.842	0.988	0.721
JobBERT-v2	0.499	0.828	0.782	0.703

Source: Authors' processing.

Taken on the overall criterion, the ensemble embedding performs best and is the most balanced between the three tests, where every alternative gains on one only by losing on another. PatentSBERTa alone is marginally ahead against observed AI use and Qwen3 against European adoption, but the latter is the weakest specification in the table against observed use, at 0.332. Register matters more than scale: two 768-dimensional domain encoders outperform general-purpose models several times their size, while an encoder trained on job postings ranks last overall.

8. Discussion

8.1 The Potential AI Exposure Metric for Professional Skills

The analysis of professional skills reveals substantial variation in their susceptibility to AI impact, as measured by the Potential Degree of Exposure Scores (PDES_k) presented in Table 2. These scores are based on the semantic similarity of each skill to AI applications, with lower scores indicating skills that are less likely to be affected by AI capabilities, while higher scores suggest greater exposure and, therefore, a higher likelihood of being influenced by AI technologies.

The results show low exposure for skills such as Negotiation (4.014), Management of Financial Resources (4.016), and Service Orientation (4.920), which

rely heavily on human judgment, and interpersonal interaction. Additionally, physical and manual skills like Repairing (5.942) and Troubleshooting (5.514) also have relatively low exposure scores, as these are more affected by robotics and mechanical automation (Webb, 2020), which typically automate manual, repetitive tasks rather than the cognitive tasks targeted by AI. In contrast, our findings suggest high exposure for skills like Writing (8.079), Reading Comprehension (8.394), Programming (7.323), and Operations Analysis (7.421), indicating that AI is more likely to impact skills that involve pattern recognition, data processing, and the generation or interpretation of structured information.

8.2 The Potential AI Exposure Metric for Professions

The distribution of AI exposure scores across professions illustrated in the histogram in Figure 2 reveals some important patterns in the potential impact of AI technologies on the labour market. As shown in the histogram, the majority of professions cluster between -1 and 1, indicating that AI is likely to have a moderate impact on most occupations, with fewer professions experiencing either very low or very high degrees of exposure. This pattern shows that while AI's influence may be widespread, its intensity varies substantially across different occupational categories.

Tables 3 and 4 further show that the most significantly affected professions are concentrated in sectors such as Education, Finance, and Research, highlighting a heightened susceptibility to transformation in knowledge-based work as AI adoption accelerates. Conversely, professions involving routine manual tasks, such as those in Construction and Manufacturing, exhibit lower exposure to AI, reflecting its greater influence on cognitive rather than repetitive physical work and underscoring that AI's impact extends far beyond purely mechanical automation.

The substantial impact on education observed in our analysis warrants particular attention given its societal significance, with 86% of educational organizations now utilizing generative AI technologies, the highest adoption rate across all sectors. Microsoft Special Report (2025). And also, the global AI in education market is projected to expand from \$5.18 billion in 2024 to \$112.3 billion by 2034. World Economic Forum (2025). Moreover, according to the U.S. Department of Education (2023), AI has the potential to enhance educational experiences by facilitating novel teacher-student interactions and enabling personalized learning pathways. Teachers are evolving from primary knowledge sources to learning facilitators and mentors, emphasizing critical thinking development and independent research skills. Our results further reveal that this impact extends beyond traditional classroom settings into technical and professional domains, highlighting AI's transformative role in reshaping how knowledge is acquired and applied in real-world, skill-intensive fields, the 2024 Stack Overflow Developer Survey indicates accelerating AI adoption among developers, with usage rising from 70% in 2023 to 76% in 2024, demonstrating how AI tools are rapidly becoming integral to professional practice in technology-intensive domains.

8.3 The Potential AI Exposure Metric for Sectors

Passing to the sectorial exposure analysis, the results in Table 5 reveal distinct patterns of AI potential impact across sectors, with Education and Professional, Scientific, and Technical Services demonstrating the highest exposure levels, followed by Finance, Healthcare which underscore the transformative potential of AI in cognitive tasks and information synthesis, with notable variation across sectors.

On the other hand, sectors such as Construction, Manufacturing, Mining, and Agriculture demonstrate the lowest exposure levels because they rely heavily on physical skills, such as equipment repair and troubleshooting, which are found to be less susceptible to being impacted by AI capabilities, according to our PDES_k metric. However, these sectors are not immune to technological disruption, having already experienced significant impacts from previous advancements, particularly robotics. Robotics have led to substantial employment reductions in manufacturing, agriculture, and mining (Kapetanidou & Pissarides, 2020; Zhao et al., 2022) by increasingly taking over routine manual tasks such as handling, welding, assembling, and dispensing (Kapetanidou & Pissarides, 2020).

While the capacities of AI remain limited in such sectors, they are not exempt from further transformation. The integration of advanced robotics with AI-enhanced systems could drive significant changes in employment within those sectors.

Comparing the results between the United States (Figure 4) and European Union (Figure 5) reveals remarkably similar patterns of potential AI impact, despite some variations in the precise ordering of sectors. This consistency across developed economies suggests that AI's transformative effects may follow comparable trajectories in advanced industrial nations, transcending regional differences.

In terms of comparison with existing literature, our analysis aligns with Felten et al. (2021) in identifying a strong AI influence on the Education, Research, and Financial sectors, though with notable distinctions. While our results suggest a stronger impact on Education and Research, they indicate only moderate effects on Marketing and professions requiring social interaction and ethical judgment, such as legal and religious occupations. In contrast, Septiandri et al. (2024) identify Financial Services, Real Estate, and Education as some of the least impacted sectors, diverging from our findings. Notably, both studies align with our conclusions regarding AI's significant impact on the Information Technology and Healthcare sectors.

8.4 Exposure to AI by Education Level and Salary

And finally, in term of Education Level and Salary, the findings in Figure 6 indicate that individuals with lower levels of education experience less impact from AI, whereas those with intermediate qualifications (such as an associate's or bachelor's degree) and higher education levels are more significantly affected. This suggests that AI predominantly influences medium and high-skilled labour. This trend underscores the uneven distribution of AI's impact across different segments

of the labour market, with potential implications for workforce training and education strategies.

Similarly, the findings in Figure 7 indicate a positive linear relationship between the degree of AI exposure and wages. This suggests that higher-paid professions are more likely to be exposed to AI. One possible explanation is that higher-paying roles often involve tasks that align with AI capabilities, such as data analysis, decision-making, and complex problem-solving. Conversely, lower-paid professions might include tasks that are less structured or require physical labour, which appear to be less exposed to AI.

9. Conclusions

This study presents a novel approach to measuring and analysing the potential impact of AI on the labour market, leveraging a deep learning-based natural language processing framework specifically, using an ensemble embeddings approach.

The analysis reveals a clear pattern of exposure across different skill sets. Cognitive-intensive skills such as reading comprehension, operations analysis, programming, and writing show the highest exposure to AI capabilities, suggesting that AI is particularly well-suited to tasks involving pattern recognition, data processing, and the generation or interpretation of structured information. In contrast, skills such as negotiation, management of financial resources, service orientation, as well as certain manual skills like repairing and troubleshooting, exhibit lower levels of exposure, indicating these areas may be less likely to be impacted by AI technologies.

Our findings suggest that AI is likely to have a moderate impact on most professions, with the majority clustering between -1 and 1 on the exposure scale. However, the impact is not uniform, occupations requiring high cognitive skills, such as technical writers and mathematicians, show significantly higher exposure. Sectors like Education, Finance, Research, and Healthcare are particularly susceptible, reflecting AI's stronger influence on knowledge-based work. In contrast, jobs involving routine manual labour, such as those in Construction and Manufacturing, tend to be less affected by AI technologies. Overall, the findings highlight that while AI's reach is broad, its effects vary significantly across occupations.

Additionally, a particularly compelling finding emerges from our analysis of education and salary correlations: professionals with higher educational attainment and salary levels face substantially greater AI exposure. This unexpected pattern challenges conventional assumptions about technological disruption and suggests a paradigm shift in how we conceptualize workforce vulnerability to AI advancement.

Overall, these findings carry some very important implications for policymakers, business leaders, and educational institutions, highlighting the need for targeted interventions and adaptive strategies. Navigating the ongoing evolution of AI requires a sophisticated understanding of its diverse exposure patterns across the professional landscape to ensure developing effective policies and support systems that promote a balanced and equitable technological transition while

minimizing risks related to AI's impact on the labour market. Therefore, future research should prioritize the development of dynamic methodologies to measure AI's multifaceted impacts across the professional spectrum, extending this analysis to economies beyond advanced ones. Additionally, more empirical work should be devoted to distinguishing between AI's substitutive and complementary effects on human labour, as well as to in-depth research into effective policy interventions and adoption strategies.

References

- [1] 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.
- [2] Bekamiri, H., Hain, D.S., Jurowetzki, R. (2024), *PatentSBERTa: A deep NLP based hybrid model for patent distance and classification using augmented SBERT*. *Technological Forecasting and Social Change*, 206, 123536, <https://doi.org/10.1016/j.techfore.2024.123536>.
- [3] Beltagy, I., Lo, K., Cohan, A. (2019), *SciBERT: A pretrained language model for scientific text*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (3615–3620), Hong Kong, China, November 3-7, 2019, Association for Computational Linguistics, Stroudsburg, PA, USA, <https://doi.org/10.18653/v1/D19-1371>.
- [4] Brynjolfsson, E., Mitchell, T. (2017), *What can machine learning do? Workforce implications*. *Science*, 358(6370), 1530-1534.
- [5] Brynjolfsson, E., Mitchell, T., Rock, D. (2018), *What can machines learn, and what does it mean for occupations and the economy? AEA Papers and Proceedings*, 108, 43-47.
- [6] Feldmann, H. (2013), *Technological unemployment in industrial countries*. *Journal of Evolutionary Economics*, 23(5), 1099–1126.
- [7] Felten, E., Raj, M., Seamans, R. (2021), *Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses*. *Strategic Management Journal*, 42(12), 2195-2217.
- [8] Frey, C.B., Osborne, M.A. (2017), *The future of employment: How susceptible are jobs to computerization? Technological Forecasting & Social Change*, 114, 254-280.
- [9] He, K., Zhang, X., Ren, S., Sun, J. (2015), *Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification*. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)* (1026-1034), Santiago, Chile, December 7-13, 2015, IEEE, <https://doi.org/10.1109/ICCV.2015.1233>.
- [10] Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., Dewhurst, M. (2017), *A future that works: Automation, employment and productivity (Report)*. McKinsey Global Institute, McKinsey & Company, New York, NY, USA, <https://www.mckinsey.com/featured-insights/digital-disruption/harnessing-automation-for-a-future-that-works>.
- [11] Gregory, T., Salomons, A., Zierahn, U. (2022), *Racing with or against the machine? Evidence on the role of trade in Europe*. *Journal of the European Economic Association*, 20(2), 869-906, <https://doi.org/10.1093/jeea/jvab040>.

- [12] Microsoft Education Team. (2025), *AI in education report: Insights to support teaching and learning*. Microsoft. Available at <https://www.microsoft.com/en-us/education/blog/2025/08/ai-in-education-report-insights-to-support-teaching-and-learning/> (accessed 18 October 2025).
- [13] Reimers, N., Gurevych, I. (2019), *Sentence-BERT: Sentence embeddings using Siamese BERT-networks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (3982-3992), Hong Kong, China, November 3-7, 2019. Association for Computational Linguistics, Stroudsburg, PA, USA, <https://doi.org/10.18653/v1/D19-1410>.
- [14] Septiandri, A.A., Constantinides, M., Quercia, D. (2024), *The potential impact of AI innovations on U.S. occupations*. *PNAS Nexus*, 3(9), page 32.
- [15] Stack Overflow. (2024), *Stack Overflow developer survey 2024*. Available at <https://survey.stackoverflow.co/2024/> (accessed 18 October 2025).
- [16] U.S. Department of Education, Office of Educational Technology. (2023), *Artificial intelligence and the future of teaching and learning: Insights and recommendations* (Report, 64 pp.). U.S. Department of Education. Washington, DC, USA, <https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>.
- [17] Webb, M. (2020), *The impact of artificial intelligence on the labor market*. *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.3482150>.
- [18] World Economic Forum. (2025), *How AI and human teachers can collaborate to transform education*, <https://www.weforum.org/stories/2025/01/how-ai-and-human-teachers-can-collaborate-to-transform-education/> (accessed 18 October 2025).