



ARTICLE

Generative AI and the Evolution of Skill Requirements in Job Postings across Labor Markets

Diana Maria Popa, Simona-Vasilica Oprea^{*} and Adela Bâra

Department of Economic Informatics and Cybernetics, Bucharest University of Economic Studies, Bucharest, Romania

^{*}Corresponding Author: Simona-Vasilica Oprea. Email: simona.oprea@csie.ase.ro

Received: 29 April 2026; Accepted: 24 June 2026; Published: 13 August 2026

ABSTRACT: This paper investigates how generative-artificial intelligence (AI) is influencing job requirements, skill compositions and sectoral dynamics across global labor markets. It examines the evolving frequency and framing of AI-related competencies in job postings, exploring whether generative-AI functions primarily as an augmentative or substitutive component in the workplace. A large-scale, multi-source corpus of over 150,000 English-language job postings (2018–2025) is compiled from twelve open-access datasets and one public API. The analytical framework integrates lexical skill extraction, semantic framing, topic modeling and time-series forecasting. Skill mentions are categorized into five dimensions: AI_Data, Routine, Soft_Meta, Domain_Specific and Leadership, while cross-sectoral analyses and correlation matrices quantify interdependencies between competencies. Sentence-transformer embeddings and cosine similarity are used to compute a Framing Index, distinguishing augmentation- vs. automation-oriented discourse. Results reveal a statistically significant increase in AI-related skill mentions after 2021 ($t = -4.70, p = 0.016$), alongside a decline in routine-task skills such as data entry and manual coding. Forecasts indicate potential continued growth in AI_Data and Soft_Meta skills through 2025, signaling a structural convergence toward hybrid human-AI expertise as a new foundation of employability. However, the results do not imply direct evidence of workforce transformation beyond what job-posting data can support. Investigating job postings, our research contributes a replicable, data-driven methodology for mapping the diffusion of AI-related skills across industries and time.

KEYWORDS: Generative-AI; labor market evolution; skill evolution; semantic framing; topic modeling; workforce adaptation

1 Introduction

The rise of generative artificial intelligence (AI) has triggered deep and multifaceted changes in labor market structures and the skillsets required across occupations. Unlike traditional automation, which primarily targeted repetitive physical or computational tasks, generative-AI systems such as ChatGPT, Copilot, and Claude can produce text, images, code, abstraction, and analytical insights with human-like fluency and machine-like rigor [1,2]. This capability fundamentally redefines the nature of work, creativity, and professional identity across some industries. Consequently, job requirements are undergoing rapid changes, demanding new technical and cognitive competencies. Traditional job structures, once centered around the execution of defined tasks, are evolving into collaborative ecosystems where humans and AI systems co-create value. This change affects the skills required for specific professions and also the broader dynamics of organizational culture, education, and regulations [3]. Thus, understanding how generative AI

influences job requirements and what forms of workforce adaptation are necessary has become an important question for researchers, policymakers, and business leaders alike [4].

The changing nature of job requirements can be observed in several interrelated dimensions [5,6]. First, *AI literacy* and *human-AI collaboration* are emerging as core competencies across all sectors. Workers are expected to know how to collaborate with generative models rather than merely operate them. Skills such as prompt engineering, model interpretation, fine-tuning and evaluation of AI outputs are increasingly valued. Second, *creative* and *analytical roles* are being redefined; professionals in fields like journalism, marketing and software development are shifting from producing content to curating, supervising and strategically integrating AI-generated content [7]. Finally, generative-AI has led to the emergence of entirely *new job roles*, from prompt engineers to AI ethicists, reshaping the labor market around hybrid human-machine expertise [8].

One of the most pressing challenges accompanying the expansion of generative-AI is the widening gap between existing educational frameworks and the evolving demands of the labor market [9]. As technological innovation accelerates, academic curricula, vocational training and certification standards struggle to keep pace, resulting in a growing skills adaptation deficit. In many cases, educational systems continue to prepare graduates for roles that no longer reflect the realities of an AI-driven economy [10]. From a technical and institutional perspective, this misalignment manifests in two main ways. First, conventional training pathways remain anchored in pre-AI paradigms that emphasize procedural knowledge and repetitive execution. Second, the lack of unified frameworks for AI literacy and competency assessment slows the diffusion of interdisciplinary expertise across sectors such as healthcare [11], law and engineering.

Given this evolution, *workforce adaptation* has become an economic and social necessity. The rapid pace of AI innovation means that skills have a shorter lifespan, making continuous reskilling and lifelong learning essential for employability. Educational institutions and organizations have to redesign training programs to incorporate AI-related content, while individuals cultivate adaptability and self-directed learning habits. Simultaneously, human-centric competencies such as reasoning, creativity and interdisciplinary collaboration are becoming more valuable, as they represent areas where human intelligence complements AI capabilities rather than competes with them [12]. At the organizational level, adaptation involves restructuring workflows, redefining job roles and fostering an innovation-oriented culture that embraces AI as a collaborative partner. By automating routine processes and augmenting creative and analytical functions, it enables professionals to focus on higher-value, strategic and innovative tasks. The future of work will therefore be characterized by symbiosis, where humans and AI systems learn from one another, co-evolve and jointly drive progress [13,14].

Despite the growing literature on generative-AI and labor market evolution, several important gaps remain. Existing studies primarily focus on productivity effects, occupation-level AI exposure or isolated sector-specific analyses, while relatively few examine large-scale longitudinal changes in skill demand using job-posting data. Moreover, prior research rarely combines skill extraction, semantic framing and cross-sectoral analysis within a unified framework, limiting the understanding of how organizations position AI as either an augmentative or substitutive technology.

Our study addresses these limitations by integrating NLP-based skill extraction, semantic framing, topic modeling and temporal forecasting on a multi-source corpus of over 150,000 job postings collected between 2018 and 2025. The proposed framework enables a longitudinal and cross-sectoral analysis of how generative-AI reshapes workforce requirements, organizational discourse and emerging skill dynamics.

Our research explores how generative-AI is influencing job requirements, identifying emerging skills and professions, and outlining the pathways for effective workforce adaptation in the era of intelligent

automation. We investigate how the emergence and diffusion of generative-AI technologies influence the composition of required skills in job postings. Job advertisements provide a high-resolution and near-real time representation of employers' expectations, revealing both technical and transversal competencies. By systematically processing these texts, we detect sectoral and temporal changes in labor demand, aiming to contribute to the literature through an analysis of job-posting discourse. Finally, we aim to answer the following research questions (RQ):

RQ1: How has mention-frequency of AI-related skill tokens (like “prompt engineering”, “model validation”) changed across occupations and sectors in time?

RQ2: Are generative-AI skill mentions correlated with reductions in mentions of routine tasks (like “data entry”, “rote coding”) within the same job titles?

RQ3: How do the semantic frames surrounding AI (augmentation vs. automation) evolve across sectors, and what do these shifts indicate about the perceived role of AI in the workplace?

2 Literature Review

From a theoretical perspective, the labor-market effects of generative-AI can be interpreted through the lens of skill-biased technological change and task-based labor model theories. Skill-biased technological change suggests that technological innovation disproportionately increases demand for highly skilled cognitive labor while reducing the relevance of routine activities. Complementarily, task-based labor models argue that automation primarily substitutes repetitive and codifiable tasks, while augmenting non-routine analytical, creative and interpersonal work. The two frameworks provide a conceptual foundation for interpreting the observed rise in AI-related and transversal competencies alongside the gradual decline in routine skill mentions across sectors.

The diffusion of generative-AI and digital technologies has intensified changes in global labor markets, influencing both the structure of work and the demand for skills. A growing body of empirical research highlights how globalization, digitization and automation jointly drive shifts in occupational profiles, productivity and workforce adaptability. Recent studies employ diverse methodologies, from large-scale analyses of job postings to firm-level field experiments, to capture the evolving interplay between human capabilities and AI systems. This emerging literature provides a multidimensional understanding of how AI adoption influences skill demand, job design and productivity across sectors, forming the foundation upon which the current research builds.

For instance, Ref. [15] examines how globalization, digitization and rapid technological change are influencing labor market skill requirements, emphasizing the need for continuous workforce adaptation. Using over 500,000 Lithuanian job postings, the authors applied natural language processing (NLP) and clustering techniques to automatically identify emerging skill profiles and job categories. Job requirements were extracted using regex-based parsing, while BERT sentence transformers were employed for feature vectorization. The resulting clusters were used to generate job descriptions with generative-AI, which experts confirmed as accurate and practical. Overall, the findings demonstrate that NLP and clustering can enable real-time monitoring of labor market skill demand, supporting data-driven workforce and education policy planning. Also, Ref. [16] shows in a large firm-scale deployment that a generative-AI assistant increased customer-support productivity by ~15%, with the biggest gains for lower-tenure agents, evidence for complementarity and skill-gap compression. Ref. [17] likewise finds faster task completion and higher writing quality with ChatGPT in randomized experiments. Further, Ref. [18] meta-analyzes >100 studies and clarifies when human-AI teams outperform: content-creation benefits are stronger than for decision tasks. Together, these studies support the “augmentation” side as other researchers [19]. Using administrative claims

and occupation exposure metrics, Ref. [20] shows AI exposure predicts unemployment risk for specific occupation-region cells, highlighting uneven impacts. At the vacancy level, Ref. [21] documents rapid growth in AI-related postings and reconfiguration of skills within adopting establishments, consistent with our postings-based approach.

Analyzing millions of postings, Ref. [22] finds a marked rise in AI skill terms and a drift toward skill-based hiring (with reduced degree requirements), empirical backing for our RQ1 on AI-token frequencies. Thus, our lexicon/taxonomy pipeline aligns with this stream by operationalizing category-level counts over time. A management-practice lens shows nuance: Ref. [23] synthesizes six paradoxes of generative-AI customer service (e.g., efficiency vs. empathy), reinforcing why augmentation narratives dominate in service contexts even as automation pressures rise. In research workflows, Ref. [24] documents algorithmic management in science, again indicating hybrid human-AI orchestration rather than wholesale substitution. In marketing, Ref. [25] outlines how generative-AI restructures content creation, product development and customer interaction, consistent with our sectoral spikes post-2021 and other researchers in different sectoral anchor [26]. For creative production, Ref. [27] shows text-to-image generative-AI boosts creative productivity and value signals at scale. In healthcare, Ref. [28] (systematic review) and Ref. [29] (meta-analysis) temper enthusiasm: governance, bias and mixed diagnostic accuracy argue for cautious augmentation, echoing our legal/health findings.

Additionally, Ref. [30] shows “double-edged” generative-AI roles in creative digital work, gains with new coordination costs, while Ref. [31] theorizes generative-AI’s transformative mechanisms on digital platforms (e.g., recombination, autonomy), offering constructs one can reuse for “augmentation vs. automation” framing at the workflow level.

Across industry contexts, Ref. [32] reviews Large Language Models (LLMs) industrial applications; Ref. [33] links AI adoption to worker well-being dynamics, reminding that how AI is integrated matters for job quality; Ref. [34] synthesizes organizational productivity effects and sustainability implications; Ref. [35] quantifies skill-bias patterns in labor demand under AI. These aspects further complement our cross-sector trend analysis.

Existing NLP-based approaches for labor-market analysis vary substantially in methodological depth and interpretability. Lexicon- and regex-based methods provide high transparency, reproducibility and computational efficiency, making them suitable for large-scale skill extraction and longitudinal monitoring. However, such approaches may fail to capture semantic nuance, contextual variation and emerging terminology. In contrast, embedding-based models using transformer architectures (BERT, SentenceTransformers) capture contextual semantic relationships between terms and documents, enabling more accurate identification of conceptual similarities and framing patterns. Nevertheless, embedding approaches are often less interpretable and computationally more intensive. Combining lexical and embedding-based techniques therefore offers a complementary strategy that balances interpretability with semantic sensitivity.

The analytical pipeline proposed in this study adopts a complementary multi-method approach in order to overcome the limitations of individual techniques. Lexicon-based extraction ensures interpretability and consistency in longitudinal skill tracking, while embedding-based semantic analysis captures contextual relationships and framing dynamics that cannot be identified through keyword matching alone. Similarly, BERTopic, LDA and KMeans provide complementary perspectives on thematic structure, contextual, probabilistic and geometric, allowing cross-validation of emerging labor-market patterns. The integration of forecasting and sectoral analysis further extends the framework beyond descriptive text mining, enabling the investigation of temporal dynamics and cross-industry diffusion of generative-AI competencies.

While the aforementioned research delivers important insights, three gaps remain: (1) existing work seldom uses job-posting texts at large scale to track skill-mention frequencies across time and sectors; (2) few studies operationalize both skill-demand evolution and semantic-framing shifts (augment vs. automate) in the same dataset; (3) there is limited longitudinal, multi-sector evidence covering the diffusion of generative (rather than generic) AI skills from ~2018 onward. Our research addresses these gaps by combining a multi-source job-posting corpus over 2018–2025, lexicon-based skill taxonomy, semantic-framing analysis and forecasting of skill trajectories across sectors.

3 Methodology

3.1 Input Data

The empirical foundation of this research consists of a large-scale, multi-source corpus of job postings, aggregated from 12 publicly available open-access datasets and one public API (presented in [Table 1](#)). The datasets are selected for their transparent licensing, reproducibility and global representativeness across industries and regions. All data sources are open, either hosted on Kaggle or accessible through the ReliefWeb public API, making the entire corpus suitable for academic and non-commercial use. The data collection strategy is designed to achieve both longitudinal and cross-sectoral coverage. To this end, job postings are gathered from major employment platforms such as Indeed, CareerBuilder, SimplyHired, Reed UK and JobStreet, complemented by LinkedIn job datasets focusing on AI, data science and business analytics roles. In addition, domain-specific sources, such as *eMedCareers* (healthcare) and *ReliefWeb* (digital development and humanitarian work) are integrated to ensure balanced representation of both traditional and emerging fields.

Table 1: Input data sources.

No.	Dataset Name	Description	Records (Approx.)	Format	Access Link	Platform	Open Access
1	ReliefWeb API (2018–2025)	International postings in digital development, public administration, and humanitarian work	50,000+	JSON	ReliefWeb API	ReliefWeb	Yes
2	eMedCareers EU Job Board	Healthcare postings from EU markets	40,000	XML	Healthcare Jobs EU	Kaggle	Yes
3	Indeed Job Postings Dataset	General global job listings	100,000+	CSV	Indeed Dataset	Kaggle	Yes
4	CareerBuilder Job Listing (USA)	U.S.-based job postings	20,000	LDJSON	CareerBuilder Dataset	Kaggle	Yes

(Continued)

Table 1 (continued)

No.	Dataset Name	Description	Records (Approx.)	Format	Access Link	Platform	Open Access
5	SimplyHired Job Listing (2019)	General U.S. job market sample	10,000	XML	SimplyHired Jobs	Kaggle	Yes
6	Reed UK Job Board	UK employment data	50,000	CSV	Reed UK Dataset	Kaggle	Yes
7	LinkedIn Data Analyst Jobs (USA, Canada, Africa)	Regional data analytics and AI-related roles	30,000	CSV	LinkedIn Analyst Jobs	Kaggle	Yes
8	JobStreet All Job Dataset	Southeast Asian job postings	60,000	CSV	JobStreet Dataset	Kaggle	Yes
9	AI Job Dataset (Global)	AI/ML-related job listings with required skills	25,000	CSV	Global AI Job Dataset	Kaggle	Yes
10	AI & ML Job Listings (LinkedIn)	U.S. LinkedIn postings related to AI and ML	15,000	CSV	AI & ML Listings	Kaggle	Yes
11	Business Analyst Job Listings (LinkedIn)	Business and analytics roles, 2024	10,000	CSV	Business Analyst Jobs	Kaggle	Yes
12	AI Job Dataset	Alternative version including additional skill annotations	15,000	CSV	AI Job Dataset v2	Kaggle	Yes

All datasets include structured metadata (like job title, description, publication date, location), but only two fields are retained for the semantic and temporal analyses: the job description text and the publication date. This ensures analytical consistency across diverse file formats (CSV, XML, JSON, LDJSON). The preprocessing pipeline harmonizes the input from multiple sources into a unified schema, removing duplicates, standardizing encodings and cleaning boilerplate or repetitive content. For the Kaggle datasets, each source was downloaded manually and stored locally in standardized formats (CSV, XML, LDJSON). For the ReliefWeb API, a Python script was developed to collect international job postings between 2018 and 2025 directly from the official REST endpoint.

The obtained dataset contained job postings, each structured into two primary variables: Date-representing the posting's publication date, enabling longitudinal trend analysis from 2018 to 2025; and Description-the complete text of each job advertisement, used as input for NLP and semantic modeling tasks.

3.2 Framework

The proposed analytical pipeline developed for this research operationalizes the conceptual framework outlined through a sequence of Python-based modules (as presented in Fig. 1), designed to capture the evolving dynamics of skill requirements across multiple industries. Each stage corresponds to a distinct methodological layer, from data acquisition and preprocessing to semantic modeling and forecasting, ensuring both reproducibility and interpretability.




Figure 1: Full analytical pipeline.

3.2.1 Text Preprocessing and Structuring

The analytical framework is implemented as a modular pipeline designed to transform the heterogeneous corpus of job descriptions into a structured, analyzable format suitable for semantic modeling, skill extraction and temporal analysis. The records in the final dataset contain two essential variables: a publication date, ranging from 2018 to 2025, and a textual description representing the body of the job posting.

Data Cleaning and Standardization

The initial phase of the framework consists of cleaning and standardizing the raw corpus of job descriptions. This pipeline combines language detection, text normalization and heuristic filtering to retain only meaningful English-language content.

Textual cleaning relies entirely on rule-based regular expressions to remove control characters, redundant whitespace and recurring non-informative patterns typically found in recruitment templates. Standardized expressions such as “Equal Opportunity Employer”, “About the Company” and “Benefits include” are eliminated to preserve only the substantive portions of the text related to job content and required competencies.

A language detection module is integrated to automatically exclude non-English records, thereby maintaining linguistic consistency across the dataset. The publication year is derived from the date field to enable longitudinal comparisons, while extremely short or incomplete entries are discarded.

Following this multistage cleaning and filtering process, the dataset is reduced to 157,552 (out of an initial 230,922) valid English-language postings, each containing relevant descriptions suitable for semantic and temporal analysis. This procedure ensures that the subsequent modeling stages are grounded in a reliable and methodologically transparent textual foundation.

Although the preprocessing stage relies primarily on rule-based regular expressions, this approach was selected to maximize transparency, reproducibility and scalability across heterogeneous datasets. However, purely lexical filtering may overlook semantic variations or context-dependent expressions and may introduce certain classification biases. To reduce this limitation, the preprocessing pipeline is complemented in subsequent stages by embedding-based semantic analysis using SentenceTransformer models, which capture contextual similarities beyond explicit keyword matching.

A potential limitation of job-posting analyses is the phenomenon commonly referred to as “AI washing”, whereby organizations include AI-related terminology primarily for signaling purposes rather than reflecting substantive changes in job requirements. To reduce this risk, the preprocessing pipeline removed repetitive boilerplate recruitment language and promotional content. Furthermore, AI-related skills were identified through manually curated competency taxonomies and contextual semantic analyses rather than simple keyword frequency counts. Representative postings were iteratively reviewed across sectors and years to ensure that extracted AI-related mentions corresponded to concrete competencies, tools or tasks rather than generic references to AI trends. Nevertheless, some degree of signaling behavior may remain present and should be considered when interpreting longitudinal trends.

Semantic Anchors and Skill Taxonomy Construction

A central methodological element of the analytical framework is the creation of two complementary lexical resources: (a) a set of semantic anchors for framing analysis, and (b) a hierarchical skill taxonomy for classification. Both are designed to capture the linguistic and conceptual evolution of generative-AI discourse across professional sectors.

The semantic anchors are structured into three main vectors: *AI*, *Augment* and *Automate*, representing distinct orientations in the human-machine narrative. The *AI* anchors include tokens such as generative-AI, LLMs, prompt engineering and model monitoring, reflecting the core technical lexicon of the field. The *Augment* group comprises expressions like assist, co-create, hybrid intelligence and human-in-the-loop, emphasizing human-AI collaboration. In contrast, the *Automate* anchors capture substitutional language, including automated, replace and robotic process automation. This structuring is performed manually, following a qualitative analysis of the job descriptions and their linguistic patterns. Each anchor is selected based on its frequency, contextual meaning and relevance to one of the three conceptual categories (AI, Augment, Automate). To ensure contextual precision, domain-specific subsets are also defined for Legal, Healthcare, Software Engineering and Design, enabling sector-based filtering of job postings via lexical triggers such as contract review, diagnosis or creative direction. These subsets are also manually curated

after analyzing the typical terminology and task descriptions specific to each field, so that only the most representative and context-appropriate expressions are retained for sector-level filtering.

Complementing this, the skill taxonomy groups extracted terms into *five standardized macro-categories*: AI_Data, Routine, Soft_Meta, Domain_Specific and Leadership (as in [Table 2](#)). Each category is constructed based on conceptual reasoning and iterative inspection of representative job descriptions. For example, AI_Data aggregates explicit AI and machine learning skills (GPT, fine-tuning, MLOps), Routine refers to procedural tasks (data entry, manual coding), while Soft_Meta and Leadership capture transversal and managerial competences (strategic planning, communication).

Table 2: Skills classified into 5 dimensions.

Category	Description
AI_Data	AI, ML and data-centric technical skills
Routine	Repetitive, procedural or clerical tasks
Soft_Meta	Transversal analytical, ethical and communication skills
Domain_Specific	Sector-specific knowledge and applied expertise
Leadership	Strategic, managerial and decision-making abilities

Sectoral classification and skill extraction relied on manually curated keyword dictionaries and representative regex-based matching rules (as in [Table A1](#)). The taxonomy was designed to capture recurring occupational terminology across major professional domains, including IT, healthcare, legal, education, finance, logistics, sales, design, and management. To improve transparency and reproducibility, representative keyword groups and classification examples are provided in [Appendix A Table A1](#). Also, to facilitate reproducibility and future methodological extensions, the skill taxonomy, semantic anchor dictionaries, keyword lists used to construct regular-expression patterns and implementation examples have been made publicly available in an online repository.

3.2.2 Skill Extraction and Categorization

Skill detection is implemented through a lexicon-based approach operationalized in the next Python module. This procedure relies on a curated hierarchical dictionary stored in JSON format, integrating canonical skill taxonomies with additional generative-AI terms manually compiled (for instance, prompt engineering, fine-tuning, model monitoring). The detection process used regular expressions to identify skill mentions and applied pattern normalization to capture linguistic variations such as “Python programming” or “proficiency in Python”. Each skill category, AI_Data, Routine, Soft_Meta, Domain_Specific and Leadership, is defined by a set of regular expression patterns automatically compiled and applied to all cleaned job descriptions.

To improve consistency, the skill taxonomy and regex patterns were iteratively refined through manual inspection of representative job postings sampled across sectors and years. Ambiguous or context-sensitive expressions were reviewed to minimize semantic misclassification and overlapping category assignments. Although the framework prioritizes interpretability and reproducibility, the possibility of lexical drift and evolving terminology over time remains a methodological limitation.

For every posting, Boolean indicators are generated to mark the presence or absence of terms belonging to each skill category. Annual aggregates are then computed, yielding normalized frequencies expressed as mentions per 1000 postings. These results are exported in a file and visualized as a graphic, representing temporal dynamics of core skill categories across 2018–2025. This lexicon-driven framework offers high

interpretability and consistency across heterogeneous datasets, while remaining adaptable to emerging terminology. The proposed method captures the growing diffusion of AI and data-related competencies in multiple sectors, including e-commerce, where AI-driven personalization, marketing automation and data analytics have become dominant skill clusters.

3.2.3 Comparative Semantic and Framing Analysis

Semantic-Shift and Framing Analysis

To examine how AI is linguistically and conceptually represented within job postings, a semantic framing analysis is conducted. The objective is to determine whether AI is portrayed primarily as a tool for augmentation (supporting and enhancing human work) or automation (replacing human labor). Three semantic anchor groups are defined (as in Table 3) and stored in a JSON configuration file to guide the analysis.

Table 3: Semantic anchor groups.

Semantic Anchor	Description
AI anchors	Capturing technological entities such as GPT, deep learning and transformer model
Augment anchors	Emphasizing collaboration and assistance, including assist, decision support and human-in-the-loop
Automate anchors	Associated with full replacement of human input, such as automation, autonomous and replace

Each job description is converted into dense numerical embeddings using the SentenceTransformer model *all-MiniLM-L6-v2*. The script then computes cosine similarities between each job vector and the three anchor vectors, obtained by averaging embeddings of the respective term lists. The Framing Index (FI) is defined as:

$$FI = \text{similarity (Augment)} - \text{similarity (Automate)} \quad (1)$$

This continuous score quantifies whether the discourse surrounding AI framed it as a collaborative enabler (positive FI) or as a replacement technology (negative FI). Since the semantic anchors are manually defined, a degree of subjectivity in term selection may influence framing scores. To reduce this limitation, anchor vocabularies were iteratively refined through corpus-level frequency inspection, contextual validation and semantic consistency analysis across sectors. Moreover, the use of SentenceTransformer embeddings reduces dependence on exact lexical matches by capturing contextual semantic proximity between job descriptions and anchor groups. Additional robustness considerations were introduced by examining augmentation and automation similarity distributions across sectors and temporal intervals. To estimate uncertainty around yearly Framing Index values, bootstrap resampling was additionally applied to yearly document-level framing scores, generating 95% confidence intervals for the aggregated yearly estimates. The process is fully implemented in Python using the *sentence-transformers* and *pandas* libraries, with outputs including yearly averages and graphical visualizations of similarity trends.

Cross-Sectoral Framing Trends

Job descriptions are further categorized into nine macro-sectors (IT, Healthcare, Legal, Education, Finance, Design, Logistics, Management and Sales) using keyword-based classification rules embedded in the analysis scripts. Aggregating results by year and sector yields a longitudinal dataset covering the 2018–2025 interval. The classification rules are defined manually through an iterative process of lexical inspection, in which sector-specific terminology is identified and validated based on its frequency and contextual relevance within job descriptions. Each category includes a controlled vocabulary of professional

and technical terms (“developer”, “nurse”, “lawyer”, “designer”, “manager”) that serve as regular expression triggers for sector detection. These rules are implemented programmatically in Python, using precompiled regex patterns to ensure consistent and reproducible tagging of over 150,000 job postings.

The resulting classification enables descriptive statistics and also temporal and cross-sectoral analyses of skill dynamics, providing the empirical foundation for the longitudinal trends discussed in [Section 4](#).

3.2.4 Semantic Topic Modeling with BERTopic, LDA and KMeans

Three topic modeling techniques are applied in order to extract relevant topics from job postings. A short description of these techniques is provided in [Table 4](#).

Table 4: Brief comparison of the three topic modeling techniques.

Aspect	BERTopic	LDA	KMeans
Core method	Transformer-based contextual embeddings + UMAP + density clustering	Probabilistic topic modeling based on word co-occurrence	Vector-based clustering using semantic embeddings
Data input	Sentence embeddings (SentenceTransformer all-MiniLM-L6-v2)	CountVectorizer (lexical frequencies)	Same BERT embeddings as BERTopic
No. of topics/clusters	6 major topics (with temporal tracking)	6 latent topics	6 semantic clusters
Interpretability	High-semantic, contextual and temporal	Moderate—lexical but clear word patterns	High—semantically coherent, less lexical
Representative themes	AI Governance, Healthcare, E-commerce Automation, Creative Design	Customer Service, Healthcare/Education, Data Infrastructure, Marketing/Sales	E-commerce Marketing, Data Analytics, Leadership/Innovation
E-commerce focus	Strongly present; AI-driven personalization and automation	Present; customer/service/marketing/sales topic	Prominent—campaign, performance, personalization

A small-scale inter-rater validation exercise was additionally conducted for representative topic labels derived from BERTopic, LDA, and KMeans outputs. Cohen’s Kappa indicated moderate agreement between annotators, suggesting reasonable consistency in topic interpretation despite the qualitative nature of semantic labeling.

BERTopic Analysis

The first stage employs BERTopic, a transformer-based topic modeling technique designed to capture contextual relationships between words beyond surface lexical patterns. The implementation combines UMAP dimensionality reduction and density-based clustering, allowing the discovery of semantically coherent topic groups directly from the embeddings. Parameters are optimized for interpretability, setting a minimum topic size of 200 and cosine similarity as the distance metric.

Formally, each document d_i is encoded into a contextual embedding vector e_i using a pre-trained transformer model:

$$e_i = f_\theta(d_i) \in \mathbb{R}^n \quad (2)$$

Pairwise similarity between two documents is computed as the cosine of their embedding vectors:

$$\text{sim}(d_i, d_j) = \frac{e_i \cdot e_j}{\|e_i\| \|e_j\|} \quad (3)$$

These embeddings are projected into a lower-dimensional manifold using Uniform Manifold Approximation and Projection (UMAP), which preserves both global and local semantic structures:

$$z_i = \text{UMAP}(e_i) \in \mathbb{R}^k, k < n \quad (4)$$

Finally, the reduced embeddings are clustered using the HDBSCAN algorithm, which groups documents according to their local density in semantic space:

$$C = \text{HDBSCAN}(\{z_i\}, \text{min_cluster_size} = 200) \quad (5)$$

This method operates by grouping documents according to the proximity of their contextual embeddings rather than by word frequency. As a result, BERTopic captures subtle conceptual distinctions, such as differences between “AI-assisted design” and “AI automation”, that traditional lexical models overlook. The output included document-level topic assignments, topic summaries and temporal dynamics.

The longitudinal data produced by BERTopic are exported, and interactive visualizations are generated, allowing the observation of how specific topics (healthcare, finance, e-commerce) evolved in prominence over time. The resulting temporal matrices express the relative topic weight in year topic, normalized by the total number of documents in that interval:

$$T_{y,t} = \frac{N_{y,t}}{\sum_t N_{y,t}} \quad (6)$$

Latent Dirichlet Allocation (LDA) Analysis

To complement the transformer-based modeling, an LDA analysis is conducted to extract probabilistic topics based on lexical co-occurrence patterns. Using a CountVectorizer, the text corpus is converted into a document-term matrix, with English stopwords removed and thresholds applied to retain terms with appropriate frequency. The model is configured to extract six optimal latent topics, aligning with the number of clusters used in the other models.

Formally, LDA assumes that each document is generated as a mixture of topics, and each topic is a distribution over words. The generative process can be summarized as follows:

For each topic $k \in \{1, \dots, K\}$, draw a word distribution:

$$\phi_k \sim \text{Dirichlet}(\beta) \quad (7)$$

Topic-word distribution defines the probability of each word belonging to topic k .

For each document d_i , draw a topic distribution:

$$\theta_i \sim \text{Dirichlet}(\alpha) \quad (8)$$

Document-topic distribution represents how much each topic contributes to document i .

For each word in document:

- Draw a topic assignment:

$$z_{i,j} \sim \text{Multinomial}(\theta_i) \quad (9)$$

- Draw a word from the corresponding topic:

$$w_{i,j} \sim \text{Multinomial}(\phi_k z_{i,j}) \quad (10)$$

The joint probability of the corpus is then given by:

$$p(W, Z, \Theta, \Phi \mid \alpha, \beta) = \prod_{k=1}^K p(\phi_k \mid \beta) \prod_{i=1}^D p(\theta_i \mid \alpha) \prod_{j=1}^{N_i} p(z_{i,j} \mid \theta_i) p(w_{i,j} \mid \phi_k z_{i,j}) \quad (11)$$

Joint generative probability represents how all documents, topics and words jointly contribute to the model's probabilistic structure.

LDA conceptualizes each job posting as a mixture of underlying topics, each represented by a probability distribution over words. This method's advantage lies in its ability to reveal dominant lexical structures and shared terminology across documents. Although LDA lacks contextual sensitivity compared to embedding-based models, it offers high interpretability through direct inspection of the top-weighted terms per topic. These terms can be visualized as word clouds for qualitative validation of topic coherence. The outputs, including topic-term distributions and document-topic probabilities, are stored for comparative analysis against BERTopic and KMeans. Methodologically, the LDA model served as a lexical benchmark, revealing how frequently co-occurring terms organize thematic domains in traditional probabilistic space.

KMeans Clustering Analysis

The third analytical approach employs KMeans clustering applied directly to the semantic embeddings generated by the transformer model. Unlike LDA, which depends on word frequency, KMeans operates in a continuous embedding space, grouping job descriptions based on their semantic similarity. The algorithm partitioned the embedding space into six clusters, minimizing intra-cluster variance and maximizing inter-cluster separation.

Each cluster is subsequently analyzed using vectorization to extract the most representative terms, providing a bridge between geometric (embedding-based) and lexical (word-based) perspectives. This procedure enables the identification of conceptually cohesive groups such as marketing and personalization, data analytics, management and innovation. Furthermore, word clouds are generated for each cluster to visualize their linguistic markers, paralleling the visualization approach used in LDA.

KMeans proves effective for detecting broad thematic structures and conceptual proximities across job roles, offering an interpretable geometric model that complements the probabilistic and transformer-based methods.

Formally, given a set of n document embeddings $\{e_1, e_2, \dots, e_n\} \in \mathbb{R}^d$, K-Means aims to find K centroids $\{\mu_1, \dots, \mu_K\}$, that minimize the within-cluster sum of squared distances (WCSS):

$$J = \sum_{i=1}^K \sum_{e_j \in C_i} \|e_j - \mu_i\|^2 \quad (12)$$

The centroid of each cluster C_i is iteratively updated as the mean of all points assigned to it:

$$\mu_i = \frac{1}{|C_i|} \sum_{e_j \in C_i} e_j \quad (13)$$

Each embedding vector is reassigned to the nearest cluster according to the Euclidean distance:

$$\text{assign}(e_j) = \arg \min_i \|e_j - \mu_i\|^2 \quad (14)$$

The algorithm iterates between Eqs. (13) and (14) until convergence, defined as no further change in cluster assignments or a minimal improvement in the objective function J .

After clustering, the content of each group is vectorized using TF-IDF weighting to extract the most representative lexical units. For each cluster C_i , the average TF-IDF weight of term t is given by:

$$w_{i,t} = \frac{1}{|C_i|} \sum_{d \in C_i} \text{tfidf}_{d,t} \quad (15)$$

3.2.5 Temporal Forecasting, Correlation and Sectoral Analysis

Following the extraction and normalization of skill mentions from the cleaned corpus, the next analytical stage focuses on modeling temporal evolution, inter-category relationships and sectoral diffusion patterns between 2018 and 2025. Three complementary modules are implemented to address these dimensions: temporal forecasting, correlation analysis and sectoral decomposition.

Temporal Forecasting Analysis

Forecasting is implemented using autoregressive integrated moving average (ARIMA) models combined with exponential smoothing, reflecting both long-term trends and short-term oscillations. Two configurations are applied for robustness: a smoothed ARIMA(1,1,1) model for general trend estimation, and an ARIMA(2,0,2) specification to explore potential short-term oscillatory patterns. ARIMA models were selected due to their interpretability and suitability for short annual time series characterized by limited longitudinal observations. Prior to model estimation, Augmented Dickey-Fuller (ADF) tests were conducted to evaluate stationarity conditions. Results indicated that most skill-series exhibited non-stationary behavior, justifying the use of differencing procedures integrated within the ARIMA framework. Forecast confidence intervals were additionally computed to quantify uncertainty associated with projected trajectories. To further improve robustness, ARIMA forecasts were compared with a Prophet-based forecasting baseline, which produced directionally consistent trend estimations across the major skill categories. Given the relatively short annual time series, the forecasting component should be interpreted cautiously. The ARIMA models are intended primarily to illustrate potential directional tendencies rather than provide precise long-term labor market predictions.

Correlation Analysis

The correlation module extended the temporal modeling by investigating interdependencies among skill categories. Using the normalized annual frequencies generated in the previous step, Pearson correlation coefficients are computed pairwise across all skill families. It produces a symmetric correlation matrix, which is visualized through a heatmap, facilitating the identification of co-evolution patterns between categories such as technical, routine and managerial skills. The analysis enables the detection of both positive and inverse relationships, offering a statistical perspective on complementarity and substitution effects among skills. By linking skill trajectories through correlation coefficients, this module provides an empirical measure of convergence and divergence between human-centric and AI-centric competencies.

Sectoral Analysis

This Python module introduces an additional dimension of granularity through sectoral decomposition. Job postings are classified into nine macro-sectors: IT, Healthcare, Legal, Education, Design, Finance, Logistics, Sales and Management, based on regular expression keyword matching within job descriptions. This rule-based classification system is optimized for recall and validated on sampled subsets to ensure accurate sector identification. Following classification, the skill detection pipeline is re-applied within each sector to compute normalized mention rates, both cumulatively and over time. It illustrates the diffusion of AI-related terminology across industries. Methodologically, the sectoral analysis integrated the temporal and correlation outputs, revealing how skill categories evolve heterogeneously across domains. By comparing normalized trajectories per sector, the module enabled the quantification of domain-specific adoption patterns and the identification of temporal asymmetries in AI diffusion.

3.3 Temporal Alignment and Normalization

To analyze longitudinal dynamics, all postings are aligned into yearly intervals covering 2018–2025. For each category and sector, skill mentions are normalized by the total number of postings in that year, producing standardized metrics expressed as mentions per 1000 job postings. The temporal modeling and forecasting stages are implemented using several Python modules which applied time-series models such as exponential smoothing and ARIMA to capture both smooth and oscillatory skill trajectories.

3.4 Validation of the Classification Framework

The rule-based classification framework was validated through manual inspection of a representative sample of job postings collected from the ReliefWeb humanitarian dataset and supplementary academic and professional sources. The validation process focused on evaluating the consistency of the regex-based extraction rules and the semantic categorization of competencies across the predefined taxonomy categories, including AI_Data, Leadership, Routine, Domain_Specific and Soft_Meta skills.

The evaluation demonstrated that the proposed approach performs effectively in identifying explicitly stated competencies, technical tools, and sector-specific requirements. Keywords associated with data analysis, monitoring and evaluation, stakeholder coordination, humanitarian operations, financial reporting, and digital tools were generally classified correctly within their corresponding categories. The system also showed good consistency in recognizing recurring terminology across humanitarian and development-oriented job descriptions.

However, several limitations were identified during the validation phase. Some competencies were expressed indirectly through contextual descriptions rather than explicit terminology, reducing extraction accuracy. In addition, multilingual postings and semantically ambiguous expressions occasionally produced incomplete or overlapping classifications. Hybrid positions combining managerial, operational, and analytical responsibilities also introduced challenges for strictly rule-based categorization.

To assess overall reliability, the automatically generated labels were compared against manual semantic interpretation of the selected postings. The highest agreement was observed for technical and domain-specific competencies, while lower consistency appeared in the identification of soft skills and implicit leadership requirements.

Overall, the validation confirmed that the proposed methodology provides a transparent and computationally efficient solution for large-scale labor market analysis in humanitarian and development contexts. Future improvements may include the integration of NLP techniques, semantic embeddings, or large language models to improve contextual understanding, multilingual adaptability, and classification precision.

4 Results

The results of the analytical pipeline highlight several trends. They do not constitute direct evidence of workforce transformation, given the inherent limitations of job-posting data. Since 2021, there has been a sharp increase in AI-related skill mentions, coinciding with the diffusion of LLMs and generative tools. This growth is strongest in technology-intensive fields, particularly IT, Finance and e-commerce, where automation and personalization drive performance. At the same time, the decline in Routine skills reflects a transition toward automation, while framing patterns indicate a hybrid dynamic in which AI is predominantly framed as an augmentative partner rather than a fully substitutive technology. Positive correlations between AI_Data and Leadership suggest that strategic roles are becoming more data-driven, whereas Routine and AI_Data are inversely related, suggesting a gradual reduction in repetitive-task demand. In e-commerce, this evolution is most evident: new hybrid roles such as *AI Marketing Analyst* or *Product Data Strategist* illustrate how generative-AI enhances analytical and creative processes. Employers frame AI as a productivity amplifier, not a replacement for human expertise. Forecasting models further indicate a possible continuation of current trends in AI_Data and Soft_Meta competencies through 2025, while Routine skills will keep declining, signaling a lasting shift toward data-centric, ethically guided and strategically integrated forms of digital work.

The 2018–2021 period primarily reflects broader AI, automation and machine-learning discourse, whereas the post-2022 interval more strongly captures the diffusion of generative-AI terminology associated with large language models and content-generation systems. A sensitivity analysis was conducted by comparing the pre-GenAI period (2018–2021) with the GenAI era (2022–2025). The results indicate that the increase in AI_Data skills and the shift toward augmentation-oriented framing remain observable under this temporal partition, supporting the robustness of the main findings. The robustness check is provided in Table 5.

Table 5: Robustness check using 2022 as the GenAI breakpoint.

Variable	2018–2021	2022–2025	% Change	<i>p</i> -Value
AI_Data	81.3	146.5	+80.2%	0.016
Routine	58.9	51.4	–12.7%	0.044
Framing Index	0.12	0.19	+58.3%	0.001

4.1 Skill Extraction and Categorization

Fig. 2 illustrates the longitudinal evolution of skill mentions per 1000 job postings, based on the standardized lexical categories. A visible drop occurs in 2021 across all categories, potentially reflecting pandemic-related disruptions in recruitment activity and fluctuations in job-posting availability. After this temporary dip, AI_Data, Soft_Meta and Domain_Specific skills show a sharp recovery in 2022, indicating renewed emphasis on digital and analytical competencies following the global acceleration of AI adoption.

Domain_Specific and Soft_Meta skills consistently maintain the highest frequencies, reflecting the sustained importance of specialized and transversal abilities. Leadership follows a similar but more moderate trend, while Routine skills remain relatively stable at lower levels. Fig. 2 captures a structural shift in labor demand: post-pandemic recovery coincides with the rise of AI-driven, cognitive and domain-specific competencies, signaling a transition toward more analytical and technology-integrated job profiles.




Figure 2: Skill mentions per 1000 job postings.

To further evaluate the significance of post-2021 labor-market changes, independent-sample *t*-tests were conducted comparing pre-2021 and post-2021 skill frequencies. Results indicate a statistically significant increase in AI_Data skill mentions after 2021 ($t = -4.70$, $p = 0.016$), with a very large effect size (Cohen's $d = 3.43$), supporting the observed acceleration in demand for AI-related competencies following the diffusion of generative-AI technologies. In contrast, changes in the remaining skill categories were not statistically significant at the 5% level, although directional increases were observed for Soft_Meta, Domain_Specific and Leadership skills.

4.2 Semantic-Shift and Framing Analysis

As shown in the left plot in Fig. 3, the FI displays a sharp peak in 2021, reflecting an intensified perception of AI as an augmentative rather than substitutive force. This peak coincides with the post-pandemic recovery period, when organizations accelerated digital evolution and emphasized human-AI collaboration to sustain hybrid and remote work structures. The subsequent decline between 2022 and 2023 suggests a temporary shift toward automation-focused narratives, likely driven by efficiency pressures and the rapid diffusion of generative models such as GPT-3 and early ChatGPT prototypes. However, the second right plot in Fig. 3 indicates that, despite short-term oscillations, the semantic proximity to *AI* and *Augment* anchors remains consistently higher than to *Automate* terms. This suggests a relatively stable discursive orientation toward augmentation, AI as a supportive, knowledge-enhancing instrument rather than a labor-replacing mechanism. The mild decline in the augmentative framing after 2023 may be interpreted as a phase of normalization and critical reassessment of generative-AI systems, reflecting both ethical concerns and uncertainty about long-term integration. By 2024–2025, the modest recovery in the FI corresponds to the institutional stabilization of AI adoption across industries. Organizations increasingly frame AI as a co-creative and decision-support partner, signaling a maturing perception of hybrid intelligence rather than technological substitution. This transition suggests that the discourse surrounding AI in job postings has evolved from experimental enthusiasm toward pragmatic, governance-oriented integration.




Figure 3: Temporal dynamics of AI-related semantic proximity and framing index (2018–2025).

The framing analysis relies on predefined semantic anchors associated with augmentation and automation narratives. While the anchor selection was informed by iterative corpus exploration and prior literature, alternative lexical anchors may produce variations in the resulting Framing Index values. Therefore, the index should be interpreted as an exploratory longitudinal indicator rather than a definitive measure of organizational intent.

To improve robustness and uncertainty estimation, bootstrap-based 95% confidence intervals were computed for the yearly Framing Index values. As illustrated in Fig. 4, the general augmentative orientation remains observable despite temporal fluctuations, although the confidence intervals indicate moderate variability across years. The pronounced increase observed in 2021 should therefore be interpreted cautiously as an exploratory semantic trend rather than definitive evidence of structural labor-market evolution.




Figure 4: Bootstrap-based 95% confidence intervals computed for the yearly Framing Index values.

4.3 Sectoral Analysis

The exceptionally high concentration of domain-specific terminology in *Healthcare* indicates that automation and AI systems have not yet penetrated deeply into clinical or diagnostic roles (as in Fig. 5), where regulatory, ethical and safety constraints continue to favor human oversight. In contrast, the relatively lower weight of AI_Data and Leadership skills in Healthcare points to a slow digital transition, limited primarily to data management and administrative optimization rather than core medical decision-making. In *IT*, the pattern is almost the opposite. Notably, AI_Data and Soft_Meta skills dominate, marking this sector as the epicenter of generative-AI diffusion. The strong presence of transversal competencies such as

communication, problem-solving and adaptability reflects the hybridization of technical and managerial profiles, professionals are expected not only to build AI systems but also to integrate them ethically and strategically into production environments. This combination signals a maturing AI ecosystem, where technical innovation is increasingly paired with governance and leadership functions. The *Legal* sector exhibits a similar dual structure, with high *Domain_Specific* and *Soft_Meta* scores, suggesting early-stage digital evolution mediated by language-based automation tools (like document summarization, compliance screening). However, the relatively modest presence of *AI_Data* skills implies that legal professionals remain consumers rather than developers of AI solutions, relying on external technologies rather than in-house model expertise. In *Finance*, the balance between *Soft_Meta*, *Domain_Specific* and *Leadership* categories suggest an environment of controlled experimentation. Financial institutions have embraced AI for analytics and risk modeling but continue to emphasize governance and interpretability, resulting in a skill mix that values analytical literacy and ethical oversight rather than deep model engineering. The *Education* sector shows an atypical trajectory. While *AI_Data* mentions are generally modest compared to IT or Finance, there is a clear spike around 2022, indicating a short-lived but intense phase of experimentation with AI-related competencies. This surge likely corresponds to the post-pandemic acceleration of digital evolution in learning environments, when educational institutions rapidly adopted AI-powered tools for online assessment, adaptive learning and content generation. *Design and Sales* occupy an intermediate zone where AI serves primarily as an augmentation tool. In these creative and customer-oriented domains, the relatively balanced mix of *Soft_Meta*, *Routine* and *Leadership* competencies indicates the integration of AI-assisted design, marketing automation and personalization technologies. Notably, AI acts as a force multiplier for human creativity rather than a substitute, amplifying productivity while preserving the centrality of human judgment.




Figure 5: Evolution of AI mentions across sectors and corresponding skill profiles.

These models reveal a pronounced upward trend in *AI_Data* skills and a corresponding decline in *Routine* tasks, consistent with increasing automation tendencies affecting repetitive roles through cognitive systems and generative-AI technologies. Particularly in e-commerce, the frequency of AI-related skill mentions rose sharply after 2022, aligning with the widespread commercial adoption of recommendation engines, generative marketing tools and automated analytics platforms.

4.4 Correlation Analysis

The correlation matrix in Fig. 6 illustrates the interdependencies between the five skill dimensions: *AI_Data*, *Routine*, *Soft_Meta*, *Domain_Specific* and *Leadership*. Pearson coefficients range between 0.49 and 1.00, indicating varying degrees of association among categories. The highest correlations appear between *Soft_Meta*, *Domain_Specific* and *Leadership* (above 0.9), forming a cohesive cluster of cognitive, professional

and managerial competencies. These categories frequently co-occur within job postings, suggesting that analytical reasoning, domain expertise and strategic oversight are increasingly integrated within the same roles. In contrast, *AI_Data* exhibits the weakest correlations with other skill groups, ranging from 0.49 to 0.63. This relative independence indicates that AI- and data-related skills often appear in more specialized or technically focused postings, distinct from those emphasizing human-centric or managerial capacities. *Routine* skills occupy an intermediate position, maintaining moderate correlations with both *AI_Data* and the human-managerial cluster, reflecting a transitional space where automation and oversight functions coexist.




Figure 6: Correlation matrix.

Pearson correlation analysis revealed moderate positive associations between *AI_Data* and *Leadership* skills ($r = 0.49$), as well as between *AI_Data* and *Routine* skills ($r = 0.50$). However, due to the limited number of annual observations, these relationships should be interpreted cautiously and are better viewed as indicative trend patterns rather than strong causal dependencies.

Some observed correlations may partially reflect broader post-pandemic temporal trends shared across multiple skill categories rather than direct functional interdependencies between competencies.

Additionally, some observed correlations may partially reflect common post-pandemic temporal trends shared across multiple skill categories rather than direct structural interdependencies.

4.5 Semantic Topic Modeling with BERTopic, LDA and KMeans

4.5.1 BERTopic Analysis

Fig. 7 depicts the evolution of significant themes over time extracted from the BERTopic output. It captures the gradual evolution of the digital landscape and the shifting demand for competencies across sectors associated with e-commerce, finance and technology-driven business models. Between 2018–2020, the trend shows an early phase dominated by themes such as Corporate/Manufacturing and Accounting/Receivables, reflecting the first wave of process digitalization and the integration of data-based management systems. The COVID-19 period coincided with a sharp rise in Finance and Banking topics, potentially reflecting

the accelerated migration of financial operations toward digital environments and increased reliance on automated infrastructures.




Figure 7: Evolution of themes over time with BERTopic.

After 2020, dynamics reveal a deeper structural shift, from basic digitalization to human-machine augmentation through predictive and analytical AI. The Finance topic exhibits the steepest increase between 2021–2024, suggesting that the financial sector, particularly within e-commerce ecosystems, exhibited one of the strongest concentrations of AI-related themes associated with data analytics, decision support and commercial optimization. In parallel, the rise of Tech Business points to the globalization of this evolution: The rise of Tech Business topics suggests an increasing strategic integration of AI into business operations, logistics and market personalization. Meanwhile, Corporate/Manufacturing follows a more oscillating path, partly due to the pandemic’s impact on supply chains and production but also because of a gradual transition toward hybrid forms of automation and augmentation. This trend is consistent with the FI analysis, which documented a semantic shift from viewing AI as a substitute for human labor (“automate”) to positioning it as a collaborative and supportive partner (“augment”). Similarly, Accounting/Receivables and Banking reflect the consolidation of digital maturity within financial operations. This pattern is consistent with Skill-Biased Technological Change perspectives, according to which technological innovation increases demand for analytical, supervisory and high-skill cognitive functions, while repetitive activities become progressively automated.

A particularly notable moment occurs between 2023–2024, coinciding with the mainstream diffusion of generative-AI models, such as ChatGPT. This period appears associated with a convergence between technological and commercial domains, accompanied by increased emphasis on AI literacy and interpretative competencies. Employers begin to emphasize AI literacy, the ability to understand, interpret and responsibly apply model outputs, as a defining competency in digital commerce. By 2025, the relative decline in topic frequencies indicates a phase of market consolidation rather than regression. The relative stabilization of topic frequencies after 2024 may indicate a transition from exploratory AI adoption toward more standardized and operational integration across organizational workflows.

Each topic in Fig. 8 encapsulates a distinct cluster of labor-market discourse, organized around core skill and activity lexicons, thereby illustrating the latent semantic segmentation of the job landscape. Topic 0 represents the marketing and commercial communication domain, centered on terms such as *marketing*, *sales* and *customer*. This lexical field captures the language of e-commerce, CRM and digital promotion,

where AI is increasingly integrated into automation tools for campaigns, behavioral analytics and customer personalization. The topic reflects an augmentation-oriented framing in which AI appears associated with enhanced creative and operational efficiency rather than full labor substitution. Topic 1 revolves around *care* and *patient*, delineating the healthcare and caregiving sector. The vocabulary emphasizes interpersonal and ethical dimensions, reflecting the human-centered nature of medical practice. In this area, AI is framed as a supportive technology for diagnostics, monitoring and patient management, functioning as an analytical assistant while preserving professional judgment. The relatively limited automation-oriented framing in healthcare may reflect regulatory, ethical and professional constraints associated with medical decision-making. Topic 2 corresponds to the AI and data science domain, structured around highly technical terms such as *python*, *tensorflow*, *pytorch*, *mathematics*, *kubernetes*, *hadoop* and *statistics*. This topic encapsulates the linguistic core of computational infrastructure, data engineering and machine learning specialization. The strong co-occurrence of programming frameworks and orchestration tools suggests that model training pipelines and cloud-based automation form the backbone of AI-related professional discourse, underscoring persistent technical specialization. This thematic concentration aligns with Task-Based Labor Model perspectives emphasizing the growing importance of non-routine technical and analytical competencies in technologically intensive occupations.




Figure 8: BERTopic generated topics.

Topic 3 focuses on *accounting*, *financial*, *accounts*, *finance* and *tax*, representing the financial and administrative sector. Its formal and standardized vocabulary signals a mature discourse shaped by precision, compliance and auditability. Within this field, AI adoption is procedural and explainable, used primarily for risk modeling, anomaly detection and reporting automation, emphasizing reliability over disruption. Topic 4 brings together *clinical*, *medical*, *pharmaceutical* and *regulatory*, corresponding to the medical-pharmaceutical research ecosystem. Compared to Topic 1, the language is more institutional and compliance-oriented, reflecting the role of AI in regulated experimentation and data-driven medicine. AI is conceptualized as an augmentative analytical instrument, enhancing precision, discovery and efficiency in data interpretation rather than displacing human expertise. Topic 5 features terms such as *hr*, *payroll*, *employee*, *recruitment* and *resources*, mapping to the human resources and organizational management

domain. The lexical focus on personnel management and evaluation points to the diffusion of AI into workplace analytics. Automation streamlines repetitive and administrative tasks, while interpretive and social dimensions remain human-led, maintaining a supportive, augmentative role for AI in people management.

4.5.2 LDA Analysis

The first topic obtained with LDA (Fig. 9) is characterized by words such as *customer*, *apply*, *role*, *skills*, *work* and *company*, which indicate a focus on customer-facing and service-oriented occupations. This topic reflects the operational and interpersonal dimension of AI adoption, where roles emphasize adaptability, collaboration and skill diversification. The second topic, dominated by terms like *technical*, *support*, *development*, *data* and *project*, reflects the technical infrastructure underlying AI deployment. It emphasizes the demand for professionals capable of integrating data-driven systems, maintaining automated workflows and providing technical support.




Figure 9: Wordclouds-LDA.

In the third topic, words such as *consultant*, *client*, *marketing*, *media* and *team* reveal an orientation toward digital marketing, creative consultancy and AI-assisted communication strategies. This cluster highlights how AI tools are embedded in marketing ecosystems, influencing recruitment, branding and customer engagement through personalized and data-driven media practices. The fourth topic, represented by *patient*, *program*, *staff*, *education* and *health*, indicates an association with healthcare and educational evolution. The fifth topic revolves around *financial*, *accounting*, *report*, *requirements* and *management*, pointing to the financial and analytical dimension of AI applications. Finally, the sixth topic, featuring words such as *team*, *hours*, *duties*, *service* and *skills*, captures the operational and organizational routines associated with modern work structures.

4.5.3 KMeans Analysis

The first cluster obtained with KMeans (as in Fig. 10) corresponds to the financial and accounting sector, where the dominant words, *accounting*, *financial*, *experience* and *sales*, indicate roles focused on business operations, reporting and corporate management. The association between *sales* and *financial* terms shows the growing overlap between financial analysis and commercial strategy, suggesting that economic roles are increasingly data-driven and connected to e-commerce performance indicators. The second cluster

reflects the healthcare and clinical field, characterized by the frequent appearance of *patient*, *medical*, *clinical* and *care*. It represents the ongoing digitalization of healthcare, where data processing, monitoring systems and AI-based diagnostics support medical decision-making. The third cluster captures the technical and programming area, defined by *python*, *tensorflow*, *pytorch*, *kubernetes* and *sql*. These terms point to positions related to software development, data infrastructure and machine learning model deployment. The fourth cluster includes terms like *management*, *project*, *team*, *data* and *experience*, describing project coordination and organizational management. This cluster emphasizes the growing need for professionals who can translate technical innovation into operational efficiency. The fifth cluster is centered on recruitment, communication and client relations. Words such as *recruitment*, *client*, *apply* and *role* reveal a focus on human resource management and personalized engagement. The sixth cluster groups together the most service-oriented terms, *customer*, *service*, *skills*, *job* and *work*, illustrating the direct interface between businesses and consumers. This vocabulary is closely associated with e-commerce, where companies rely on AI systems to manage customer interactions, predict behavior and tailor user experiences. The repetition of *customer* and *service* emphasizes the centrality of personalization and responsiveness in the digital economy, as companies strive to balance automation with human-centered design.



Figure 10: Wordclouds KMeans.

The number of topics/clusters ($K = 6$) was selected following iterative experimentation across BERTopic, LDA and KMeans models in order to balance semantic coherence, interpretability and cross-method comparability. Lower values generated overly broad thematic structures that merged distinct labor-market domains, while higher values produced fragmented clusters with substantial semantic overlap and reduced interpretability. The $K = 6$ configuration provided the most stable separation between major thematic areas such as AI engineering, healthcare, finance, marketing and organizational management, while maintaining consistency across probabilistic, embedding-based and transformer-based topic modeling approaches.

4.5.4 Comparison of the Results

A comparison of the results obtained by the three techniques is offered in [Table 6](#).

Overall, the thematic convergence observed across BERTopic, LDA and KMeans suggests a broader organizational transition toward hybrid human-AI work structures. The increasing prominence of AI_Data, analytical and communication-oriented themes, combined with the relative decline of routine-oriented discourse, is consistent with both Skill-Biased Technological Change and Task-Based Labor Model approaches.

These findings indicate that generative-AI adoption is associated less with complete labor substitution and more with the reconfiguration of organizational workflows toward augmentation, supervision and data-driven decision-making.

Table 6: Thematic correspondence across techniques.

Core Theme	BERTopic (Transformer-Based)	LDA (Probabilistic)	KMeans (Embedding-Based)
Organizational and workforce evolution	AI Governance, Workforce AI Adoption	Customer, company, service, role, skills	Operations, management, leadership, client
Education and workforce preparation	Learning Adaptation, AI Training	Patient, staff, program, education	Training, support, communication, learning
Technical and data foundations	Model Development, LLM Engineering, MLOps	Technical, data, engineer, project, development	Analytics, python, automation, systems
E-commerce and sales innovation	E-commerce Automation, Generative Content Marketing	Marketing, content, media, sales	Campaign, customer, digital, performance, recommendation
Societal and creative implications	Creative Design Augmentation, Ethical AI	Consultant, creativity, ethics, innovation	Communication, design, ethics, creativity

4.6 Cross-Sectoral Framing Trends

The sectoral framing analysis in [Fig. 11](#) illustrates the heterogeneous ways in which AI and automation are being linguistically positioned across professional domains. By comparing the trajectories of AI, Augment and Automate framings between 2018–2025, a divergence emerges in how sectors conceptualize the integration of intelligent technologies, ranging from cautious experimentation to structural adoption and normalization. In the *IT* sector, the framing of AI shows a consistent upward trend, consolidating after 2021. The discourse is characterized by a pragmatic balance between innovation and control. Mentions of augmentation rise sharply in 2022, corresponding to the diffusion of generative-AI and automation frameworks into software development, data engineering and DevOps environments. The language of augmentation dominates over that of automation, reflecting a shift from viewing AI as a threat to technical roles toward understanding it as an efficiency multiplier that enhances coding, testing and deployment workflows.

The *management* sector presents a similar evolution but with a stronger and earlier adoption of augmentation-oriented framing. This indicates that organizational discourse increasingly centers on AI-assisted decision-making, strategic foresight and leadership in digital evolution. In *finance*, the narrative is shaped by a dual emphasis on automation and augmentation. Early fluctuations give way to a steady increase in augmentation framing after 2022, coinciding with the expansion of AI in risk modeling, fraud detection and predictive analytics. However, automation retains a secondary presence, revealing the sector's reliance on rule-based systems and its regulatory caution. The *education* sector displays a distinctive pattern, with a pronounced spike in AI-related framing in 2022. This anomaly corresponds with the mass

adoption of EdTech and learning analytics tools during the post-pandemic digital shift. However, the trend stabilizes afterward, suggesting that AI integration in education remains situational rather than systemic. The dominance of augmentation language underscores the framing of AI as a pedagogical assistant rather than a substitute for educators, enhancing assessment, adaptive learning and content personalization while preserving human oversight and empathy as core values of the teaching profession.




Figure 11: Semantic framing across sectors.

In *healthcare*, the discourse remains grounded in caution. Domain-specific and ethical constraints limit the prevalence of automation framing, while augmentation dominates. This reflects the sector's emphasis on AI-assisted diagnostics, patient monitoring and administrative optimization, rather than fully autonomous systems. The pattern points to a narrative of collaborative intelligence, where AI augments clinical judgment but does not replace it. The temporary rise in AI framing in 2022–2023 may correspond to the pandemic's acceleration of telemedicine and medical imaging analytics. The *legal* sector mirrors healthcare in its conservatism, showing modest growth in AI and augmentation framings but minimal automation language. The *logistics* sector exhibits volatility, with short-lived peaks in AI and augmentation framings around 2022. This likely corresponds to post-pandemic restructuring of global supply chains and the integration of predictive analytics in routing and inventory systems.

In *sales* and particularly in *e-commerce*, the discourse undergoes a striking evolution. After 2021, augmentation framing rises dramatically, accompanied by sustained growth in AI mentions. This linguistic pattern reflects the mainstream adoption of AI-driven recommendation engines, marketing automation and content generation. E-commerce stands out as the sector where AI is most tightly coupled with creativity and personalization. The coexistence of high augmentation and moderate automation framing signals a redefinition of human-machine interaction: automation handles transactional efficiency, while humans leverage AI to craft narratives, optimize engagement and innovate in customer experience. The language reflects co-creativity rather than replacement, making e-commerce a prototype for future hybrid digital labor models.

Findings showed a consistent trend toward augmentation-oriented framing, especially after 2021, when the adoption of generative-AI technologies accelerated. In e-commerce, this tendency was particularly pronounced, with frequent references such as AI-assisted marketing, decision-support analytics, and human-in-the-loop optimization. These expressions framed AI as a co-creative and productivity-enhancing force, reinforcing human roles in strategic and creative functions. Conversely, the IT and Finance sectors leaned more toward automation-oriented language, reflecting widespread process optimization. In contrast, Healthcare, Design, and Education maintained predominantly collaborative framings.

4.7 Forecast

The comparative ARIMA forecasts in Fig. 12 provide complementary perspectives on the temporal dynamics of skill demand across the five macro-categories. Given the limited number of annual observations, these forecasts should be interpreted as exploratory projections intended to illustrate potential directional tendencies rather than precise long-term labor market predictions. The smoothed ARIMA(1,1,1) projection models a stable, gradually converging trend after 2025, indicating that the market for digital and cognitive competencies is approaching a structural equilibrium. The trajectories of Domain_Specific and Soft_Meta skills remain dominant, showing steady growth that plateaus at high frequency levels, reflecting their consolidation as fundamental components of employability in AI-mediated labor markets. Leadership follows a parallel but lower trajectory, suggesting a sustained managerial demand oriented toward coordination and oversight rather than expansion. Conversely, Routine and AI_Data exhibit possible mild downward tendencies in the smoothed scenario, consistent with a gradual automation of repetitive tasks and the standardization of technical AI competencies. The smoothing effect implies a stabilization of demand rather than volatility, suggesting that the diffusion of AI tools has reached maturity and is now embedded in broader professional frameworks.




Figure 12: Forecast.

The oscillatory ARIMA(2,0,2) model explores potential short-term oscillatory patterns that the smoothed specification attenuates. Notably, Soft_Meta and Leadership continue to expand moderately, but with visible short-term oscillations, indicating alternating phases of emphasis on adaptive and managerial capabilities. Domain_Specific remains robust yet shows slight periodic corrections, which may reflect the saturation of specialized roles.

Augmented Dickey-Fuller (ADF) tests indicated that most annual skill-series exhibited non-stationary behavior ($p > 0.05$), supporting the use of differencing procedures within the ARIMA framework. Only the Domain_Specific series satisfied conventional stationarity criteria (Table 7).

Table 7: Results of Augmented Dickey-Fuller (ADF) stationarity tests for annual skill-series.

Category	ADF Statistic	p -Value	Stationarity
AI_Data	-1.408	0.579	Non-stationary
Routine	-2.750	0.066	Non-stationary
Soft_Meta	-2.396	0.143	Non-stationary
Domain_Specific	-3.139	0.024	Stationary
Leadership	-2.853	0.051	Non-stationary

To provide an additional comparative reference, exploratory Prophet-based forecasts were also generated for the AI_Data category. As illustrated in Fig. 13, both ARIMA and Prophet models exhibit broadly similar post-2022 upward trajectories, although the wide confidence intervals reinforce the uncertainty associated with forecasting based on short annual time series.




Figure 13: ARIMA-based and Prophet-based exploratory forecasting of AI_Data skill trends (2018–2030).

5 Conclusions

The analysis reveals an accelerating integration of AI-related competencies across professional domains, particularly after 2021, when generative-AI entered mainstream application through LLMs and automated content systems. Mentions of skills such as prompt engineering, data validation and model monitoring have expanded rapidly, especially within IT, Finance and e-commerce, signaling a structural redefinition of what constitutes digital expertise.

At the same time, the steady decline of Routine skill mentions points to an ongoing process of task automation, yet the data suggests that this is not a simple substitution dynamic. Rather than replacing human

labor, AI is increasingly portrayed as an augmentative partner, a tool that enhances analytical reasoning, creativity and strategic decision-making. Employers appear to value professionals who can integrate AI tools into existing workflows, interpret their outputs and ensure ethical oversight.

Sectoral variation reinforces this interpretation. Technology and financial fields exhibit the strongest concentration of AI-related skill mentions, while e-commerce demonstrates a hybrid pattern that bridges automation and human creativity. Job descriptions in this sector increasingly reference AI-assisted marketing, personalized recommendation systems and data-driven content optimization, reflecting a form of applied human-AI collaboration where automation enhances productivity without displacing creative roles. Education shows a temporary spike in AI-related terminology in 2022. By contrast, healthcare and legal sectors remain grounded in domain-specific terminology, underscoring the cautious and regulated nature of these environments. The next phase of workforce development will thus hinge on cultivating AI fluency, the capacity to guide, question and strategically employ AI in support of human judgment and creativity.

The framing perspective provides insight into how organizations conceptualize the role of AI in their workflows: as a collaborative tool that amplifies human capability (augmentative framing) or as an autonomous system that substitutes human labor (automotive framing). Two hypotheses guide this analysis. The first one posits that firms in high-adoption sectors such as IT and e-commerce will shift earlier toward augmentative framing, emphasizing human-AI collaboration and oversight. The second one suggests that, across industries, increases in generative-AI skill mentions will coincide with a decline in automation-oriented framing, reflecting a broader cultural and organizational transition toward viewing AI as an assistive rather than a replacement technology.

Nevertheless, the results reflect changes in employer demand as captured in job postings and should not be construed as direct evidence of broader workforce transformation. While the study adopts a robust and reproducible computational framework, several methodological limitations must be acknowledged. First, job postings primarily reflect stated requirements, which may not accurately correspond to the actual skills used in daily professional activities or the final hiring outcomes. Despite de-duplication procedures, some postings may be reposted or cross-listed across multiple platforms, potentially inflating certain categories. Furthermore, the presence of AI-related terminology can occasionally serve marketing or branding purposes rather than denote genuine technical needs. Second, sectoral imbalances exist due to uneven data availability: digital industries such as IT and e-commerce are heavily represented in online job boards, whereas traditional sectors like education and healthcare appear less frequently. In some cases, AI-related terminology may reflect recruitment signaling or perceived technological alignment rather than substantive organizational AI integration.

Further, changes in token frequencies and semantic framing should not be interpreted as direct evidence of workplace evolution or technological displacement. The analysis primarily captures shifts in recruitment discourse and employer communication strategies, which may only partially reflect underlying organizational change.

The augmentation-vs.-automation framework represents a simplified semantic abstraction of AI discourse and does not fully capture the multidimensional nature of organizational AI adoption narratives. Accordingly, the Framing Index should be interpreted as an interpretable semantic orientation indicator rather than a comprehensive representation of organizational AI behavior. In addition, topic interpretation involved qualitative inspection of representative keywords and document clusters, introducing an element of researcher judgment despite the use of multiple topic-modeling approaches. The rule-based classification framework also prioritizes interpretability and scalability over deep contextual understanding, which may limit classification precision for hybrid or multi-domain occupations.

Additionally, the forecasting results should be interpreted as exploratory projections intended to illustrate potential directional tendencies rather than precise long-term labor market predictions, particularly given the limited number of annual observations available in the dataset.

From a policy perspective, the findings highlight the urgent need for governments, educational institutions and organizations to develop coordinated strategies for workforce adaptation in the era of generative-AI. Public policies should support continuous reskilling and lifelong learning initiatives focused on AI literacy, data analytics, critical thinking and human-AI collaboration competencies. Universities and vocational training providers should update curricula to integrate interdisciplinary AI-related skills, ethical governance and practical exposure to generative technologies across both technical and non-technical domains. At the organizational level, firms should promote augmentation-oriented AI adoption models that combine automation with human oversight, creativity and strategic decision-making, particularly in sectors such as healthcare, education and legal services where ethical and regulatory considerations remain essential. Furthermore, policymakers should encourage transparent AI governance frameworks and inclusive labor-market policies aimed at reducing digital skill gaps and minimizing inequalities generated by uneven technological diffusion.

Acknowledgement: This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI—UEFISCDI, project number ERANET-ERAMIN-3-ValorWaste-1, within PNCDI IV.

Funding Statement: This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI—UEFISCDI, project number ERANET-ERAMIN-3-ValorWaste-1, within PNCDI IV.

Author Contributions: Diana Maria Popa and Simona-Vasilica Oprea: Conceptualization, Methodology, Formal Analysis, Investigation, Writing—Original Draft, Writing—Review and Editing, Visualization, Project Administration. Adela Băra and Diana Maria Popa: Validation, Formal Analysis, Investigation, Resources, Data Curation, Writing—Original Draft, Writing—Review and Editing, Visualization, Supervision. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: Data is available here: https://github.com/diana181103/job_postings_analysis.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

Table A1: Representative skill taxonomy and classification rules.

Sector	Representative Keywords/Patterns
IT	Developer, software engineer, python, machine learning, cloud, cybersecurity
Healthcare	Nurse, hospital, clinical, patient care, pharmaceutical
Legal	Lawyer, compliance, regulation, contracts, litigation
Education	Teacher, university, curriculum, pedagogy, e-learning
Finance	Accounting, banking, investment, audit, portfolio management
Sales & Marketing	CRM, digital marketing, ecommerce, customer success
Management	Leadership, project manager, operations, strategy

The dictionaries were iteratively refined through exploratory inspection of job advertisements and domain-specific terminology. Regex matching was primarily used for interpretable large-scale categorization rather than fine-grained semantic disambiguation.

References

1. Zhang Y. The impact of ChatGPT on HR recruitment. *J Educ Humanit Soc Sci.* 2023;19:40–4. doi:10.54097/ehss.v19i.10949.
2. Skondras P, Zotos N, Lagios D, Zervas P, Giotopoulos KC, Tzimas G. Deep learning approaches for big data-driven metadata extraction in online job postings. *Information.* 2023;14(11):585. doi:10.3390/info14110585.
3. Makridis CA, Mishra S. Artificial intelligence as a service, economic growth, and well-being. *J Serv Res.* 2022;25(4):505–20. doi:10.1177/10946705221120218.
4. Gutiérrez E, Lourie B, Nekrasov A, Shevlin T. Are online job postings informative to investors? *Manag Sci.* 2020;66(7):3133–41. doi:10.1287/mnsc.2019.3450.
5. Ge C, Shi H, Jiang J, Xu X. Investigating the demand for blockchain talents in the recruitment market: evidence from topic modeling analysis on job postings. *Inf Manag.* 2022;59(7):103513. doi:10.1016/j.im.2021.103513.
6. Sugar W, Hoard B, Brown A, Daniels L. Identifying multimedia production competencies and skills of instructional design and technology professionals: an analysis of recent job postings. *J Educ Technol Syst.* 2012;40(3):227–49. doi:10.2190/et.40.3.b.
7. Dixon N, Goggins M, Ho E, Howison M, Long J, Northcott E, et al. Occupational models from 42 million unstructured job postings. *Patterns.* 2023;4(7):100757. doi:10.1016/j.patter.2023.100757.
8. Liu R, Li M, Sin S, Tan M. Ethical requirements in job advertisements: a deep learning approach. *Eur Manag Rev.* 2024;21(1):134–49. doi:10.1111/emre.12569.
9. Pojednic R, O'Neill DP, Flanagan MG, Bartlett A, Carter BL, Kennedy MA. Exercise professional education, qualifications, and certifications: a content analysis of job postings in the United States. *Front Sports Act Living.* 2024;6:1338658. doi:10.3389/fspor.2024.1338658.
10. Martin F, Chen Y, Oyarzun B, Lee M. Learning and development roles and competency domains in higher education: a content analysis of job announcements. *J Comput High Educ.* 2022;34(2):297–320. doi:10.1007/s12528-021-09290-2.
11. Lee P, Bubeck S, Petro J. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. *N Engl J Med.* 2023;388(13):1233–9. doi:10.1056/nejmsr2214184.
12. Mäkelä E, Bone M, Sehrer M, Nanji F, Stephany F. Complement or substitute? How AI increases the demand for human skills. *arXiv:2412.19754.* 2024.
13. Linnyk O, Tetz I. Counteracting the global labor shortage risk through the human–AI collaboration in digital recruiting. *IEEE Technol Soc Mag.* 2023;42(2):42–7. doi:10.1109/mts.2023.3277108.
14. Squicciarini M, Nachtigall H. Demand for AI skills in jobs: evidence from online job postings. *OECD Sci Technol Ind Work Pap.* 2021;2021(3):1–74. doi:10.1787/3ed32d94-en.
15. Lukauskas M, Šarkauskaitė V, Pilinkienė V, Stundžienė A, Grybauskas A, Bruneckienė J. Enhancing skills demand understanding through job ad segmentation using NLP and clustering techniques. *Appl Sci.* 2023;13(10):6119. doi:10.3390/app13106119.
16. Brynjolfsson E, Li D, Raymond L. Generative AI at work. *Q J Econ.* 2025;140(2):889–942. doi:10.1093/qje/qjae044.
17. Noy S, Zhang W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science.* 2023;381(6654):187–92. doi:10.1126/science.adh2586.
18. Vaccaro M, Almaatouq A, Malone T. When combinations of humans and AI are useful: a systematic review and meta-analysis. *Nat Hum Behav.* 2024;8(12):2293–303. doi:10.1038/s41562-024-02024-1.
19. Marguerit D. Augmenting or automating labor? The effect of AI development on new work, employment, and wages. *arXiv:2503.19159.* 2025.
20. Frank MR, Ahn YY, Moro E. AI exposure predicts unemployment risk: a new approach to technology-driven job loss. *PNAS Nexus.* 2025;4(4):pgaf107. doi:10.1093/pnasnexus/pgaf107.

21. Acemoglu D, Autor D, Hazell J, Restrepo P. Artificial intelligence and jobs: evidence from online vacancies. *J Labor Econ.* 2022;40(S1):S293–340. doi:10.1086/718327.
22. Bone M, González Ehlinger E, Stephany F. Skills or degree? The rise of skill-based hiring for AI and green jobs. *Technol Forecast Soc Change.* 2025;214(S1):124042. doi:10.1016/j.techfore.2025.124042.
23. Ferraro C, Demsar V, Sands S, Restrepo M, Campbell C. The paradoxes of generative AI-enabled customer service: a guide for managers. *Bus Horiz.* 2024;67(5):549–59. doi:10.1016/j.bushor.2024.04.013.
24. Koehler M, Sauermann H. Algorithmic management in scientific research. *Res Policy.* 2024;53(4):104985. doi:10.1016/j.respol.2024.104985.
25. Grewal D, Satornino CB, Davenport T, Guha A. How generative AI is shaping the future of marketing. *J Acad Mark Sci.* 2025;53(3):702–22. doi:10.1007/s11747-024-01064-3.
26. Dorta-González P, López-Puig AJ, Dorta-González MI, González-Betancor SM. Generative artificial intelligence usage by researchers at work: effects of gender, career stage, type of workplace, and perceived barriers. *Telemat Inform.* 2024;94(2):102187. doi:10.1016/j.tele.2024.102187.
27. Zhou E, Lee D. Generative artificial intelligence, human creativity, and art. *PNAS Nexus.* 2024;3(3):pgae052. doi:10.1093/pnasnexus/pgae052.
28. Haltaufderheide J, Ranisch R. The ethics of ChatGPT in medicine and healthcare: a systematic review on Large Language Models (LLMs). *npj Digit Med.* 2024;7(1):183. doi:10.1038/s41746-024-01157-x.
29. Takita H, Kabata D, Walston SL, Tatekawa H, Saito K, Tsujimoto Y, et al. A systematic review and meta-analysis of diagnostic performance comparison between generative AI and physicians. *npj Digit Med.* 2025;8(1):175. doi:10.1038/s41746-025-01543-z.
30. Hou J, Wang L, Wang G, Wang HJ, Yang S. The double-edged roles of generative AI in the creative process: experiments on design work. *Inf Syst Res.* 2025;70(9):95. doi:10.1287/isre.2024.0937.
31. Wessel M, Adam M, Benlian A, Majchrzak A, Thies F. Generative AI and its transformative value for digital platforms. *J Manag Inf Syst.* 2025;42(2):346–69. doi:10.1080/07421222.2025.2487315.
32. Raza M, Jahangir Z, Riaz MB, Saeed MJ, Sattar MA. Industrial applications of large language models. *Sci Rep.* 2025;15(1):13755. doi:10.1038/s41598-025-98483-1.
33. Giuntella O, König J, Stella L. Artificial intelligence and the wellbeing of workers. *Sci Rep.* 2025;15(1):20087. doi:10.1038/s41598-025-98241-3.
34. Al Naqbi H, Bahroun Z, Ahmed V. Enhancing work productivity through generative artificial intelligence: a comprehensive literature review. *Sustainability.* 2024;16(3):1166. doi:10.3390/su16031166.
35. Wang X, Chen M, Chen N. How artificial intelligence affects the labour force employment structure from the perspective of industrial structure optimisation. *Heliyon.* 2024;10(5):e26686. doi:10.1016/j.heliyon.2024.e26686.