LLM Derived Evidence on Skill Portfolios, Complementarities, and Wage Growth in IT Jobs

Shanjukta Nath
University of Georgia

Jinsoo Yeo
University of Tennessee Chattanooga

Anandhi Bharadwaj
Emory University

Sundar Bharadwaj
University of Georgia*

February 3, 2026

Abstract

We quantify the impact of Generative AI on labor demand and wage returns across IT skill portfolios by applying transformer-based text analysis to job postings. Our novel multivariate skill measures outperform existing AI exposure metrics in predicting salaries. We find strong complementarities between AI and IT development skills. Workers in the 90th percentile of IT development skills gain a 18.5% wage premium from AI expertise, compared to just 5.1% for those with lower IT development skills (10th percentile). In contrast, AI substitutes for routine IT operations skills; however, operational expertise regains value when bundled with both AI and IT development skills. These complementarities are most pronounced among entry-level and mid-career workers, but diminish for experienced professionals. The results hold across work arrangements and locations, with the highest returns observed in tech hubs. Our findings highlight the importance of integrating AI and IT development skills to maximize wage gains in the evolving IT labor market.

Keywords: AI skill complementarities, IT labor market, Wage premiums, Skill portfolios, Transformer-based text analysis, Job posting data.

*Department of Agricultural and Applied Economics (shanjukta.nath@uga.edu), Gary W. Rollins College of Business (jinsoo-yeo@utc.edu), Goizueta Business School (a.bharadwaj@emory.edu), Terry College of Business (sundar@uga.edu). All authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

The rapid diffusion of generative artificial intelligence (AI), enabled by large language models, represents a new general-purpose technology with applicability across multiple applications and industries. This new technology would likely have significant implications for productivity changes (Achiam et al., 2023; Bresnahan and Trajtenberg, 1995). As with prior general-purpose technologies, its adoption is reshaping not only production processes and firm boundaries but also labor demand. Amazon CEO Andy Jassy's announcement of 30,000 corporate layoffs (Bensinger, 2025), coupled with a $31.4 billion investment in AI infrastructure (Amazon.com, Inc., 2025), characterizes the scale of this transformation. Large technology firms have simultaneously reduced employment in traditional roles while making substantial investments in AI infrastructure and capabilities, signaling a reconfiguration rather than a simple contraction of high-skill labor. These changes are especially salient in information technology (IT) occupations, which are intensive users of cognitive labor and among the earliest adopters of AI-enabled tools. Despite growing evidence of shifts in employment in these occupations, the specific mechanisms by which generative AI reshapes labor demand and wages remain an open empirical question.

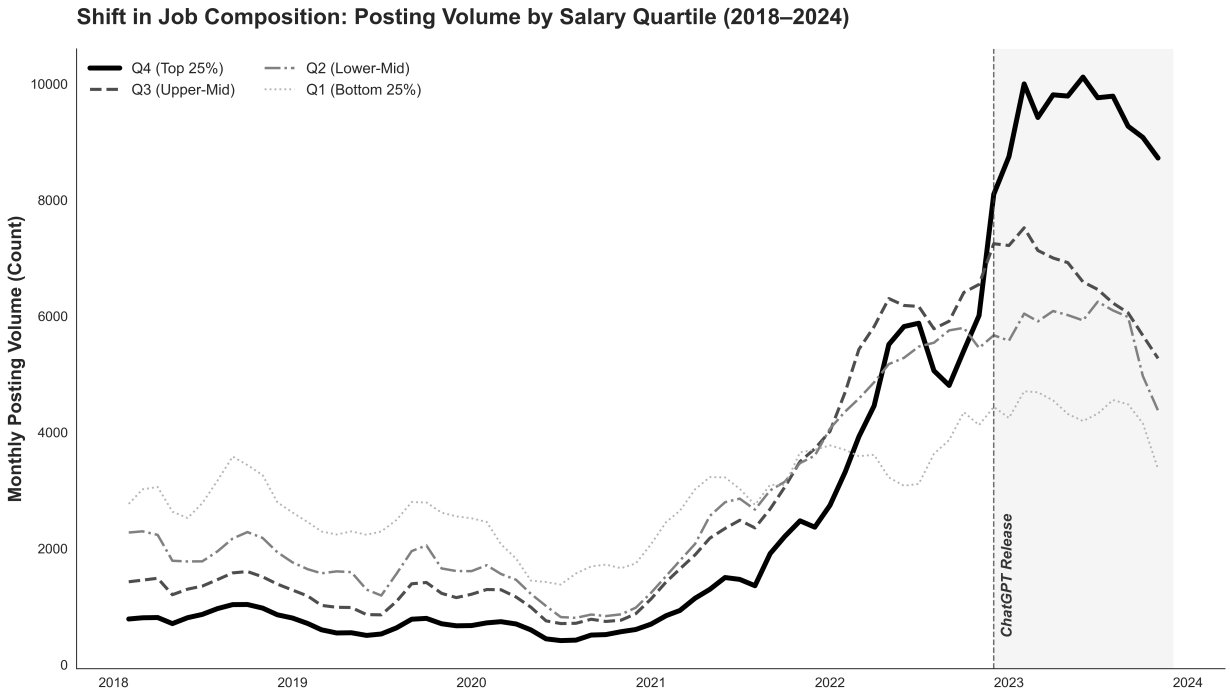
A task-based perspective provides a useful framework for analyzing these dynamics. Previous research emphasizes that automation technologies tend to substitute for routine cognitive tasks while complementing non-routine, judgment-intensive activities, thus altering the internal allocation of tasks within jobs rather than eliminating entire occupations (Acemoglu et al., 2023; Acemoglu and Restrepo, 2018; Dixon et al., 2021). Consistent with this view, the literature documents a nuanced pattern of labor market adjustment to AI and related technologies, shaped by task reallocation, skill complementarities, and frictions in human capital accumulation (Acemoglu and Restrepo, 2018; Agrawal et al., 2019; Autor and Salomons, 2018; Bessen, 2019; Eloundou et al., 2024; McAfee and Brynjolfsson, 2018). Importantly, the displacement effects associated with automation can be partially offset by the reinstatement effects that generate new tasks and new forms of labor demand (Acemoglu and Restrepo, 2019). Although this literature has made substantial progress in characterizing aggregate employment effects, it provides limited insight into how labor markets price specific skills as technologies evolve.

This limitation is particularly consequential in IT occupations, where generative AI can simultaneously depreciate some skills while increasing the value of others (Deming and Noray, 2018). Existing studies typically rely on the content of tasks at the occupation-level or coarse measures of technological exposure, which obscure the heterogeneity of the eco-

nomic returns to specific combinations of skills. As a result, we know relatively little about how wages reflect evolving skill portfolios within occupations. For example, within software development roles, proficiency in manual coding may become less valuable even as expertise in system architecture, AI integration, or model supervision becomes more valuable. Aggregated analyses risk averaging over these divergent trajectories, thereby masking the mechanisms through which skill-biased technological change operates.

Figure 1 shows a reallocation of IT job postings toward higher salary quartiles between 2018 and 2024. While this shift begins prior to the release of generative AI, it accelerates sharply after late 2022, raising questions about the mechanisms through which AI-enabled tools reshape labor demand and wages. Importantly, the figure serves as a visual summary of this phenomenon and illustrates that aggregate posting and wage trends would largely mirror overall wage growth, thereby masking divergent trajectories across salary segments.

Figure 1: Shift in Job Composition: IT Job Postings by Salary Quartiles (2018–2024)



In this paper, we address these gaps by deploying state-of-the-art large language models on the universe of online job postings to construct multivariate, data-driven measures of demand for AI and legacy IT skills. Unlike prior approaches that rely on univariate indicators or predefined taxonomies, our method captures the joint structure of traditional IT skills and emerging AI capabilities. We benchmark the predictive performance of these measures against leading approaches used in the literature and show substantial gains in explanatory

power (Babina et al., 2024). We then link skill composition to economic outcomes by estimating wage premia associated with specific skill clusters, using observed salary data as a direct measure of market valuation. This framework allows us to isolate the marginal return to AI-related skills conditional on existing IT expertise and to assess whether complementarities between traditional IT and AI skills drive wage differences.

Our analysis reveals that wage dispersion in IT occupations is driven not by AI skills in isolation but by their interaction with distinct components of the IT skill bundle. AI exhibits strong complementarities with IT development skills, whereas its interaction with IT operations skills is negative, indicating that operational roles are more exposed to substitution. These interaction effects are largely stable across geographic regions and work arrangements but vary systematically with worker experience. Wage premia associated with AI-enabled development are concentrated among junior- and mid-career workers, consistent with AI reshaping early-stage tasks toward design, integration, and system-level problem solving (Brynjolfsson et al., 2025; Dell'Acqua et al., 2023). In contrast, the marginal impact of AI diminishes in senior roles. Taken together, these findings suggest that the labor market consequences of generative AI are governed by the composition of skill portfolios rather than by occupational titles, with implications for workforce investment, training strategy, and career progression.

2 Related Work

Our paper leverages Large Language Models (LLM) to construct multivariate measures of AI skill demand, facilitating a granular analysis of labor market restructuring and the estimation of wage premia driven by the complementarities between emerging AI capabilities and established IT expertise. Traditional methodologies, such as those in Autor and Dorn (2013) and Acemoglu and Restrepo (2018), often rely on O*NET task descriptions or predefined keyword lists to proxy for automation exposure. These "top-down" approaches are inherently static and struggle to adapt to the rapid evolution of generative AI. By deploying LLMs on the raw universe of job postings, our model captures the joint structure of skills, identifying latent skill demand that traditional taxonomies miss, and emphasizing the nuanced ways in which AI capabilities are integrated into existing IT roles.

While Babina et al. (2024) demonstrates that AI-investing firms experience shifts in workforce composition, our multivariate measure allows us to isolate the marginal return to AI skills conditional on existing expertise. Whereas univariate measures treat AI as a

monolithic shock, our approach treats it as a component of a broader skill portfolio. This allows for the explicit estimation of wage premia driven by the complementarity between legacy IT skills and new AI competencies, rather than relying solely on aggregate firm-level effects.

Conceptually, our paper builds on the broader labor economics literature, which posits that technological change does not merely augment or replace workers based on their education, but rather reshapes the specific tasks required in production. Adopting the lens of routine-biased technological change (RBTC), Autor and Dorn (2013) show that although computers substitute for routine tasks, they complement the analytical and problem-solving tasks of high-skill workers, leading to gains in employment and wages. Deming (2017) expands the task-based framework to account for the growing importance of social skills. His work shows that the modern economy increasingly rewards flexibility and tacit knowledge, which remain significant barriers to automation. Building on the shift toward interpersonal and cognitive competencies, Deming and Kahn (2018) utilizes the dataset of online job postings to examine how these skill requirements vary within and across firms. They categorize thousands of job-ad keywords into ten general skills, finding that cognitive and social skills are not only the most frequently demanded but are also powerful predictors of firm productivity and wage premiums. Their research moves beyond the occupational level to show substantial heterogeneity even within narrowly defined roles; companies requiring greater cognitive and social abilities are typically bigger, more efficient, and offer higher pay. Significantly, they find strong complementarity between cognitive and social skills, suggesting that the most productive modern work environments leverage both analytical depth and collaborative efficiency.

Recent work documenting the dramatic increase in demand for AI skills finds that it is no longer confined to IT sectors but has diffused into management, finance, and professional services (Alekseeva et al., 2021). Crucially, the authors document a substantial wage premium for job postings that require AI skills, even after controlling for other high-level skills, such as software proficiency. This indicates that while AI may automate high-skill tasks, the initial stage of its diffusion is creating a new tier of elite workers capable of managing and implementing these systems. Although this study provides macro evidence that AI adoption drives wage growth, the authors treat AI as a monolithic skill and do not consider how AI skills integrate with or complement the existing portfolio of IT skills demanded by the labor markets.

In Acemoglu et al. (2022), the authors document the rapid takeoff in AI job vacancies

starting in 2010 and accelerating around 2015-16. AI exposure catalyzes a fundamental reconfiguration of task structures, driving a decline in legacy skills while generating new tasks and new skill demands. Although the authors document that AI-adopting establishments reduce their hiring of non-AI workers, suggesting a broad substitution effect for legacy skills, in the current study, we are able to further dissect this trend to show that the substitution is specifically concentrated in IT Operations roles, whereas IT Development roles capture significant wage premiums through complementarity with both Operations and AI.

Hampole et al. (2025) provides the most granular empirical assessment of AI's current impact by constructing firm-specific measures of task exposure. They distinguish between "mean exposure," which drives substitution and reduces labor demand, and "dispersion in exposure," which allows workers to reallocate effort toward non-displaced tasks. Their findings indicate that while AI substitution is strong at the task level, its aggregate employment effects are muted because AI adoption also generates significant firm-level productivity gains, thereby increasing overall labor demand. However, we show that these aggregate findings mask a critical divergence in skill valuation: AI actively substitutes for routine IT operations while simultaneously acting as a strong complement to IT development, generating significant wage premiums when bundled with high-level architectural skills.

Collectively, these papers illustrate a labor market in a state of continuous task-level restructuring. As technology moves from automating rules to augmenting judgment, the overarching takeaway from the extant literature on technology, jobs, and task restructuring is that while technology may substitute for specific tasks, the value of human labor increasingly resides in the ability to flexibly navigate a multidimensional landscape of coordination and specialized, non-routine tasks.

Against this background, our work, focusing on the changing nature of IT jobs across industry sectors, contributes in several ways. Conceptually, our work moves beyond viewing AI as a monolithic technological shock to identify how wage growth is driven by the specific complementarities between IT skills (IT Development and IT operations) and emerging AI capabilities. By isolating how specific skills within the same occupation can simultaneously appreciate (e.g., system architecture) or depreciate (e.g., manual coding), we provide a mechanism to explain skill-biased technological changes that aggregated occupational studies often mask.

From a methodological perspective, we deploy transformer-based models on raw job posting text to construct data-driven, multivariate measures that capture latent skill clusters ignored by static taxonomies, such as those from Lightcast and O*Net. We benchmark our

new measures against established approaches, such as Babina et al. (2024), and demonstrate superior predictive performance by directly linking granular skill bundles to observed salary data.

3 Empirical Setting

3.1 Data

This section describes the datasets used to construct the sample and the complete methodological pipeline for obtaining all variables required for the analysis. We focus on measuring latent multivariate skill attributes from raw job-posting text and combining them with firm-level controls and occupational variables. To process unstructured job-posting text, we use generative AI to empirically derive IT skill categories, thereby enabling a structured representation of the latent IT skill landscape.

Lightcast The online job postings dataset for this paper is from Lightcast (formerly known as Burning Glass Technologies), a leading company in labor market analytics. Lightcast is a reputable and frequently cited source in the literature, widely known as offering a "near universe" of online job vacancies (Deming and Kahn, 2018). As one of the first firms to pioneer large-scale online job posting data, Lightcast provides a robust framework of posting-level occupational variables derived through extensive data collection, cleaning, and deduplication. Relative to official occupational data sources such as O*NET, which is maintained by the U.S. Bureau of Labor Statistics (BLS) and relies on expert input and infrequent periodic updates, posting-level data provide more timely, granular, and market-responsive insights into evolving labor demand.

Lightcast data is particularly well-suited for our scope of analysis. First, Information Technology is one of the most robustly represented sectors in the dataset. While Lightcast is known to under-represent certain sectors such as agriculture and blue-collar labor jobs, IT roles are consistently well-captured (Hershbein and Kahn, 2018). Second, our analysis is restricted to postings with explicitly reported salary information, as wage information is central to the scope of our study. This distinction is important as other labor market data providers often rely heavily on imputed salary measures, which prior research has shown can introduce substantial measurement error (Dai et al., 2025).

One of Lightcast's primary competence lies in its comprehensive and dynamic skill taxonomy. While government frameworks like the BLS-managed O*NET categorize the labor

market through 35 broad, expert-defined skills at the aggregate occupation level, our Lightcast sample identifies more than 30,000 unique, granular skills. This enables identifying nuanced demand for both technical labor and specific software tools in real-time. Given the rapid technological advancements and shifting skill requirements, this granularity is essential.

Beyond its skill taxonomy, Lightcast provides a comprehensive set of cleaned and structured posting-level variables, including posting dates, standardized job titles, employer identities, and detailed geographic regions. While extant literature has primarily focused on these structured information to study AI adoption, geographic patterns, market value, and labor market dynamics (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024), our empirical strategy focuses on the raw job postings to measure multivariate latent IT skill attributes.

This shift in focus is necessitated by the need to quantify the marginal value of combining traditional IT skills with emerging AI-related capabilities, which requires moving beyond explicitly stated univariate skill measures. Substantial reconfiguration of IT employment increasingly depends on "algorithmic literacy" and "domain expertise" (Tambe, 2025) that are fundamental to task execution but are frequently left implicit in job descriptions. As employers often omit foundational skills they assume are implied by the job title or the specific technologies required, univariate skill measures systematically underrepresent the combined structure of workers' core IT capabilities.

For instance, our dataset includes a job posting for a "Senior React Native Developer." *React Native* is a *JavaScript*-based mobile development framework, implying that the mobile developer role fundamentally requires expert-level JavaScript proficiency. However, for this job posting, JavaScript is never explicitly mentioned in skill requirement[1]. As Lightcast primarily relies on traditional NLP methods to extract explicit tags, such implicit dependencies are not captured. To address this, we use a Generative AI approach to extract latent skill attributes from the raw text, capturing implicit technological requirements that traditional keyword-based extraction typically misses.

Supplementary data sources To control firm-level characteristics, which are essential for analyzing the relationship between wages and skill requirements, we use data from Compustat-Capital IQ, collected via the Wharton Research Data Services. These employer-year records include the official company name (CONM), which serves as the primary iden-

[1]To illustrate the extent of implicit dependencies, we conduct a gap analysis on our full sample. We find that 23.72% of postings that explicitly require React Native do not include a corresponding JavaScript skill tag.

tifier for matching the Lightcast dataset.

Sample construction Our initial dataset comprises more than 234 million job postings spanning 72 months, from January 2018 to December 2023. From this universe, we construct a focused sample of 870,810 IT job postings from 2,709 unique U.S. publicly traded firms that report salary information. This refinement allows us to examine the granular wage outcomes required to identify the marginal returns of multidimensional skill portfolios.

To link job-level skill requirements with firm-level financial characteristics, our sample construction process requires mapping individual job postings to firm-year level data in Compustat. We implement a hybrid matching approach that integrates semantic vector search, LLM review, and human validation to map Lightcast and Compustat data.

Specifically, we generate dense text embeddings for company names using a BERT-family sentence transformer model (all-MiniLM-L6-V2). These embeddings are indexed with Facebook AI Similarity Search (FAISS) to identify the k-nearest Compustat candidates based on cosine similarity. Compared to extant approaches based on fuzzy string matching and syntactic overlap, this embedding-based method enables semantic matching, allowing us to identify corporate aliases and parent-subsidiary relationships with limited lexical similarity.

We then use an LLM (GPT-4o) to review the candidate mappings, prompting it to assume the role of an occupational domain specialist and to evaluate firm mappings as of August 2024 an LLM (GPT-4o) to review the candidate mappings, prompting it to assume the role of an occupational domain specialist and to evaluate firm mappings as of August 2024[2]. Any mappings flagged by the LLM with a low confidence score undergo a final round of validation by a human expert.

This mapping procedure reflects the recent institutional shift towards hybrid AI-expert systems consistent with the National Center for O*NET Development (2025). By following this emerging standard for maintaining large-scale occupational databases, we are able to process millions of job postings efficiently while maintaining high mapping precision. Table 1 presents the mapping results and summarizes the job postings data sample.

3.2 Transformer-Based Multi-dimensional skills for IT Jobs

Given the rapid changes in the demand for skills in the current labor market, mapping the job postings' textual data to different skills presents a critical challenge for researchers. To

[2]Mappings are anchored to August 2024 to align with extraction date of Lightcast and Compustat data. This supports ensuring consistency across corporate structures and M&A or divesture activity.

Table 1: Lightcast Job Postings Data: Coverage and Linkage Summary

	Raw Lightcast Data		Compustat	IT Jobs	Salary
Year	Postings	Skills	Mapped	ONET15	Included
2018	33,125,641	417,279,417	9,459,777	1,115,564	87,263
2019	34,545,101	450,644,471	8,796,142	1,024,247	70,228
2020	35,072,775	435,835,184	8,528,923	852,387	53,469
2021	43,711,124	557,527,894	10,369,670	1,069,912	113,299
2022	48,727,337	652,293,853	10,986,588	1,323,175	235,335
2023	39,035,912	511,299,116	8,071,361	694,128	311,216
Total	234,217,890	3,291,410,613	56,212,461	6,079,413	870,810

Notes: This table presents the coverage of Lightcast job postings data from 2018–2023. Raw Lightcast Data includes all job postings and extracted skills. Compustat Mapped indicates successful linkage to public firm records. IT Jobs are defined by O*NET occupation code 15. Salary Included indicates verified salary information.

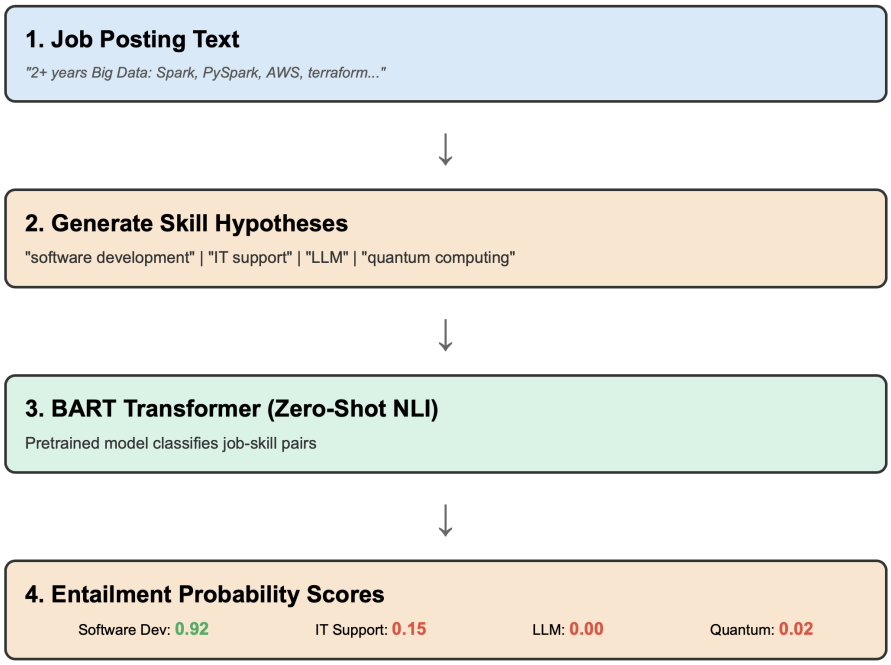
address this challenge, we introduce a new measure of AI and traditional skills in IT jobs using a zero-shot text classification (Yin et al., 2019) method. The classifier uses a transformer-based framework pretrained on natural language inference tasks to map unstructured job posting text to user-defined skill labels. The most significant advantage of this approach is that it does not require any task-specific training. In other words, the model does not need to be trained on the job postings data to develop the classifier. This, in turn, means that the researcher does not need to perform the costly task of relabeling and retraining the model with rapidly changing skills in the current labor market transformation following the introduction of Generative AI.

Precisely, we explain the steps for zero-shot classification using an example from our data-set. Consider the following abbreviated version of a job text "*BS in Computer Science or Engineering with at least 2 years of professional software work experience. 2+ years of experience in Big Data Platform technologies required Spark or PySpark or Scala. Experience with tools such as Databrick or DeltaLake preferred. Working knowledge of Amazon cloud technologies. Strong collaboration and teamwork skills..*" and to keep the explanation simple we work with four skill labels: Software development, IT support, LLMs, and Quantum computing for this section. [3]

[3]In our empirical application, we work with 30 skill labels encompassing AI and traditional skills in IT jobs for publicly listed US firms. This job text is an abbreviated version of a longer body of a job posting from the Lightcast dataset.

In the first step, the method creates four hypotheses using the text and the labels: i) This job requires Software development skills, ii) This job requires IT support skills, iii) This job requires LLM skills, and iv) This job requires Quantum computing skills. Next, it formats the text input (premise) and the hypothesis into a premise-hypothesis pair, which can be used as an input into the transformer model. In the third step, a pretrained transformer model such as Facebook BART/BERT will use the attention mechanism (Vaswani et al., 2017) to extract the underlying semantic connections between the text and the labels. If there is a strong semantic connection between a label and the text, it will be assigned a higher probability score. Figure 2 summarizes each step for the classification pipeline.

Figure 2: Zero-shot classification pipeline



Our classification method differs from related work that has used job postings and resumes to estimate skill demand in several ways. First, the attention mechanism in Vaswani et al. (2017) and the transformer framework Devlin (2018) do not use keyword-matching/dictionary methods that search for exact strings of skill labels in job-posting text. Related papers that use job-posting data for AI skill measurement are Babina et al. (2024) and Acemoglu et al. (2022). Babina et al. (2024) uses the skill categories provided by Lightcast for each job posting and constructs AI-relatedness at the skill level by the co-occurrence of these granular skills with four core skills: artificial intelligence, machine learning, computer vision, and natural language processing, and then aggregates the AI-relatedness over skills at the job

posting level. Similarly, Acemoglu et al. (2022) uses the Lightcast skills associated with each job posting to construct its AI measure based on a finite set of AI skill names, such as ML and CV, that explicitly show up in Lightcast's skill categories for each job posting.

Keyword matching techniques or dictionary methods completely ignore the context of the text, and hence, they cannot learn about the underlying semantic relationships between text and labels. The attention mechanism uses bidirectional context (left and right) to capture semantic connections in the text. Even when skill labels are not explicitly present in the job text, the attention mechanism can capture the text's underlying semantic connections to the labels. This context-awareness is essential for measuring the demand for modern heterogeneous IT labor, in which skill expression is often implicit, rapidly evolving, and loosely standardized.

As the attention mechanism leverages semantic depth rather than literal string matching, the selection of input labels constitutes a pivotal design choice in our empirical framework. These labels serve as the semantic anchors for our zero-shot classification algorithm, allowing us to delineate the boundaries of modern IT labor. Given the inherently fluid nature of these boundaries, which legacy taxonomies often fail to capture in detail, we adopt a systematic, data-driven approach to label construction. Specifically, we derive a set of 30 IT skill categories using principled, iterative prompting from a Large Language Model (LLM). Our objective is to induce an empirically robust taxonomy that outlines the core pillars of both AI and conventional IT work at the job-posting level.

Existing vendor-provided taxonomies highlight this challenge. For example, Lightcast currently characterizes the IT sector using 69 detailed categories, each intended to represent a distinct IT specialization. Yet even such fine-grained classifications are not stable. As an illustration, Computer Vision, a core AI domain and a central component of AI exposure in Babina et al. (2024), is classified under an "Image Analysis" subcategory, positioned far from AI and even from IT more broadly. Such inconsistencies underscore the limitations of static taxonomies for capturing contemporary IT skill demand.

To construct a structured representation of the latent multivariate IT skill domain, we utilize an occupational-specialist prompting framework with LLM. Specifically, the LLM is instructed to assume the role of an IT occupational expert and to propose distinct skill categories for traditional IT and emerging AI domains. For each category, the model must provide: (1) a formal domain name, (2) a brief rationale for its relevance in labor market analysis, and (3) a list of representative software tools or technical infrastructures associated with that domain. We choose to derive 25 AI-specific and 25 non-AI-specific IT skill categories

separately. This distinction is necessitated by the role of AI as a general-purpose technology (Goldfarb et al., 2023), where AI capabilities increasingly permeate traditional IT functions and the specific nuances of emergent AI specializations are subsumed by broader legacy IT categories. We further impose negative constraints, explicitly excluding non-technical "soft skills" and restricting attention to technical IT skill categories and frontier AI domains. To assess the robustness of these categories, we conducted sensitivity testing across multiple LLM architectures and prompt variations.

Our goal is to ensure a framework that approximates mutually exclusive and collectively exhaustive skill domains, without relying on predefined vendor taxonomies. This requirement is essential for robustly quantifying the heterogeneity of IT labor and capturing the latent complementarities between existing IT expertise and AI capabilities. Specifically, mutual exclusivity minimizes semantic overlap, ensuring that each skill intensity represents a distinct technical domain. This clarity is a prerequisite for econometric identification, as it mitigates multicollinearity and enables the isolation of wage premia associated with specific skill clusters. Simultaneously, collective exhaustiveness ensures full coverage of the skill distribution within our sample. By providing comprehensive label space, we account for the full technical composition of each vacancy and mitigate misalignment in IT labor demand.

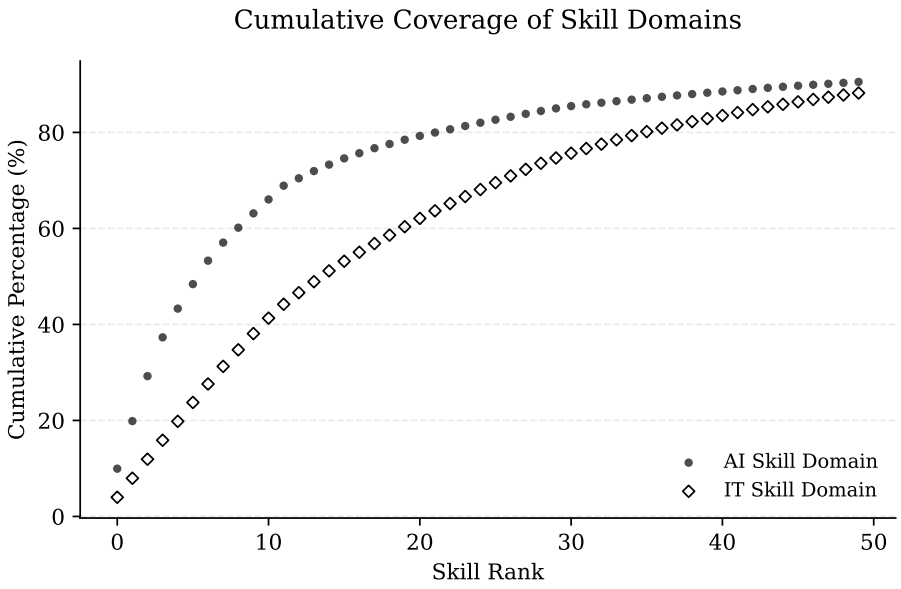
To support stability, we conduct large-scale iterative testing by running the induction prompt 1,000 times. While sampling randomness is theoretically suppressed under greedy decoding ($T = 0$), residual stochasticity may arise in LLM outputs, precluding strict deterministic results. We address this variability through frequency-based refinement and semantic consistency checks, ensuring that our final taxonomy represents a stable, LLM consensus-driven framework.

Each induced category is accompanied by a technical justification and subsequently subjected to a Human-in-the-Loop (HITL) cross-validation process. Using the resulting frequency distributions, we iteratively refine the taxonomy by removing infrequently appearing categories, consolidating semantically overlapping domains, and excluding overly granular labels that correspond to specific tools or individual skills rather than broader skill categories. Analogous to the elbow method in cluster analysis, we evaluate the trade-off between category granularity and semantic distinctness to arrive at a final set of 30 skill domains.[4]

[4]The 30 skill categories include: Artificial Intelligence, Big Data, Blockchain, Blockchain Technology, Cloud Computing, Computer Vision, Cybersecurity, Data Engineering, Data Science, Database Management, DevOps, Embedded Systems, Generative AI, Internet of Things (IoT), IT Support, Large Language Models, Machine Learning, Mobile App Development, Mobile Development, Natural Language Processing, Network Administration, Networking, Operating Systems, Quantum Computing, Reinforcement Learning, Robotics, Software Development, Systems Administration, Virtualization, and Web Development.

This frequency-guided pruning process yields a parsimonious taxonomy that balances semantic precision with empirical coverage, ensuring that our measures of IT skill demand are both empirically robust and stable across the near-universe of observed job postings (Figure 3).[5]

Figure 3: Cumulative Explanatory Power of Skill Domains



3.3 Summary Statistics

This section reports summary statistics for the key variables used in our analysis, drawn from the merged Lightcast-Compustat dataset. The statistics describe the distribution of job-posting-level outcomes, firm characteristics, and the zero-shot-generated skill measures that form the basis of our empirical exercises. All job postings are restricted to publicly traded IT firms with valid Compustat identifiers, ensuring consistent linkage between labor demand and firm-level information. Summary statistics are reported at the job-posting level unless otherwise noted. The longitudinal data on job postings reveal a stark decoupling of skill demands within the IT sector. As shown in Figure 4, aggregate AI, development, and operations skill intensities follow distinct trajectories throughout the sample period. While AI intensity rises steadily before reaching a high-level plateau, demand for traditional IT

[5]The current taxonomy represents the baseline induction from our prompting framework. While certain categories exhibit semantic proximity, they represent distinct nuances in the model's consensus induction. We remain open to consolidating these into broader clusters for further sensitivity testing in future iterations.

Table 2: Summary Statistics

	Mean	SD	Min	P25	P50	P75	Max
Panel A: Outcome Variable							
Salary ($'000$)	113.15	497.29	10.03	81.13	110.63	140.50	978.89
Log(Salary)	11.53	0.50	9.21	11.30	11.61	11.85	13.79
Panel B: Skill Scores							
IT Operations Skills	0.28	0.08	0.00	0.25	0.28	0.30	0.97
IT Development Skills	0.35	0.09	0.01	0.31	0.34	0.37	0.99
AI Skills	0.37	0.07	0.01	0.35	0.38	0.40	0.97
Panel C: Education							
Education: High School	0.06	0.24	0.00	0.00	0.00	0.00	1.00
Education: Associate	0.04	0.19	0.00	0.00	0.00	0.00	1.00
Education: Bachelor	0.58	0.49	0.00	0.00	1.00	1.00	1.00
Education: Master	0.20	0.40	0.00	0.00	0.00	0.00	1.00
Education: PhD	0.07	0.25	0.00	0.00	0.00	0.00	1.00
Panel D: Experience							
Minimum Years Experience	3.67	3.50	0.00	0.00	3.00	5.00	15.00
Panel E: Work Arrangement							
Remote: None	0.71	0.46	0.00	0.00	1.00	1.00	1.00
Remote: Remote	0.24	0.43	0.00	0.00	0.00	0.00	1.00
Remote: Stated None	0.03	0.16	0.00	0.00	0.00	0.00	1.00
Remote: Hybrid	0.03	0.16	0.00	0.00	0.00	0.00	1.00
Employment: Full-time	0.95	0.23	0.00	1.00	1.00	1.00	1.00
Employment: Part-time	0.03	0.18	0.00	0.00	0.00	0.00	1.00
Employment: Either	0.02	0.14	0.00	0.00	0.00	0.00	1.00
Internship Position	0.02	0.15	0.00	0.00	0.00	0.00	1.00

Number of Unique Firms: 2,709

Total Observations: 870,797

Notes: This table presents summary statistics for the main analysis sample. Panel A shows salary variables. Panel B reports skill scores aggregated into three categories: IT Operations (8 skills), IT Development (10 skills), and AI (12 skills). Panels C-E present covariates including education levels, work experience, and work arrangement characteristics. Additionally, the sample includes the NAICS code for the job postings obtained from Compustat data. The remote category none is for postings that do not specify a remote job type, and 'stated none' indicates postings that explicitly state that remote work is not allowed. The education categories do not add up to 1, as a job posting can demand one of multiple education levels. Regression also includes a missing indicator for the minimum years of experience. It also includes a measure of maximum years of experience (missing if not reported) for a very small fraction of postings that report this information.

development and IT operations declines. This divergence emerges prior to late 2022 and persists thereafter, underscoring substantial heterogeneity beneath aggregate trends. The portfolio-level panels further show that skill demand is dynamically recomposing within IT occupations, with sharp differences in timing and magnitude across specific domains rather than a uniform expansion of IT demand. Consistent with these patterns, Figure 5 documents a concurrent shift in the composition of IT job postings toward higher salary quartiles, suggesting a reallocation of labor demand toward higher-paying IT roles. Collectively, these

Table 3: Representativeness Check: Salary Disclosure vs. Full IT Universe

	Salary Sample		Full Universe		
	Mean	SD	Mean	SD	SMD
Panel A: Job Requirements					
Min Exp Required	5.19	3.08	5.26	3.01	-0.024
Education (Level 1-5)	1.82	0.68	1.86	0.63	-0.066
Skill Count	24.22	15.18	24.65	14.02	-0.031
Panel B: Firm Characteristics					
Log Total Assets ($\ln at$)	9.70	2.14	10.10	2.21	-0.178
Log Market Cap ($\ln mkvalt$)	9.71	2.01	10.07	2.04	-0.172
Panel C: Work Arrangement					
Remote: Not Stated	0.71	0.46	0.86	0.35	-0.150
Remote: On-site	0.03	0.16	0.01	0.10	0.015
Remote: Hybrid	0.03	0.16	0.01	0.10	0.014
Remote: Fully Remote	0.24	0.43	0.12	0.33	0.121
Panel D: Geography					
Transparency Law State	0.39	0.49	0.36	0.48	0.032
Total Observations: 870,810 (Salary) / 5,208,603 (Non-Salary)					

Notes: This table assesses the representativeness of the salary-disclosing sample by comparing it to the non-disclosing population of IT postings. Standardized Mean Difference (SMD) is calculated as $(\mu_{sal} - \mu_{full})/\sigma_{full}$. SMDs below 0.1 indicate statistical balance. Panel A covers human capital requirements, Panel B covers firm financial scale from Compustat, and Panels C-D cover work arrangements and regional characteristics.

figures highlight substantial heterogeneity in the evolution of skill demand across the full set of granular, zero-shot-generated skill measures. While this high-dimensional representation

captures rich variation in skill demand, it complicates interpretation and obscures recognition of broader economic patterns. To identify the latent structure of IT labor demand and map technical requirements to economic functions, we aggregate the granular skill measures into three functional categories, IT Operations, IT Development, and AI skills, corresponding to infrastructure and maintenance tasks, software engineering and system design, and advanced modeling and data-driven capabilities. These three functional categories constitute the core explanatory variables in our analysis, and salary information is the core dependent variable.

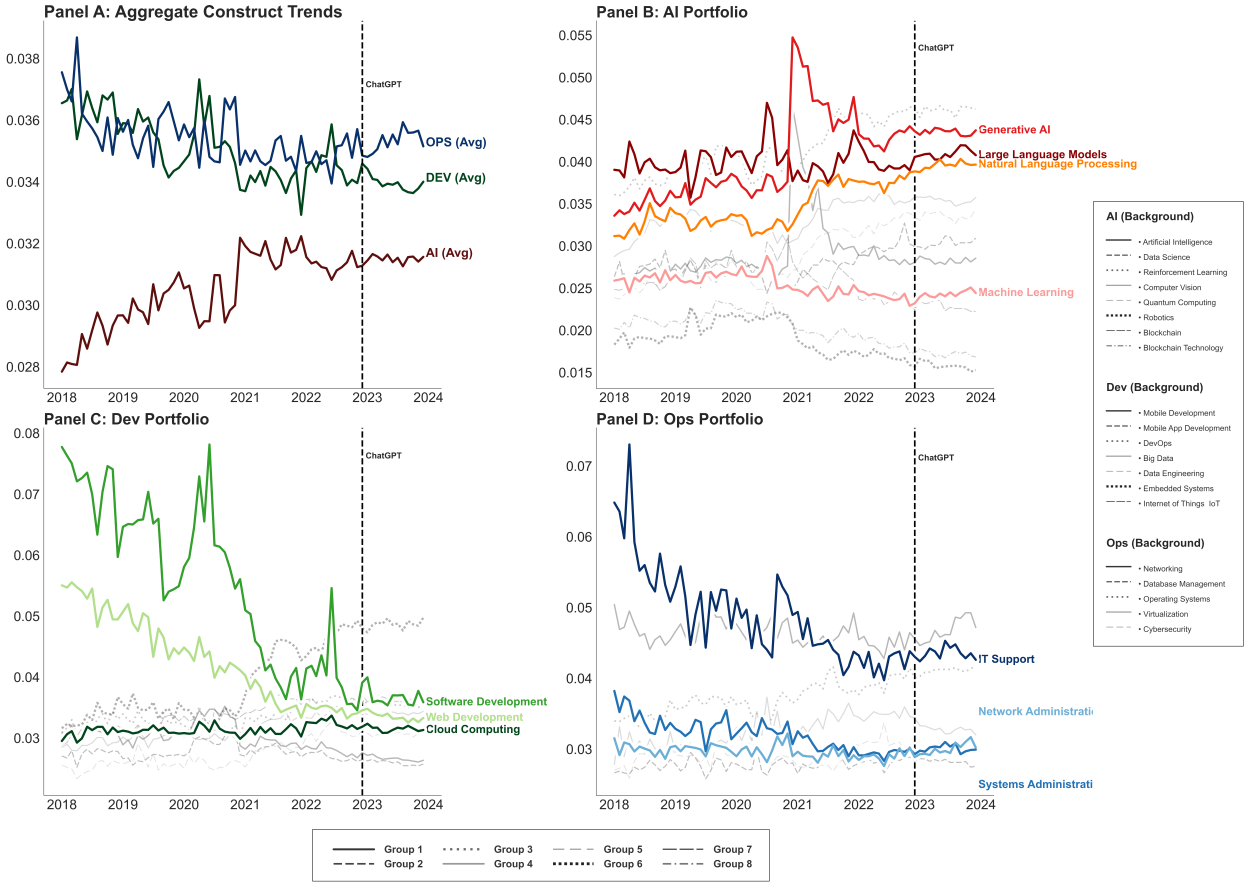
A potential concern using salary data is the non-random nature of salary disclosure. Our primary analysis sample is necessarily smaller than the full universe of postings, as it restricts attention to publicly traded firms, IT occupations (broadly defined by O*NET code first two-digit 15), and postings with usable salary information. To assess whether this selection induces systematic bias, we conduct a comprehensive representativeness audit comparing the analysis sample (n=870,000) with the full universe of IT job postings (n=5,208,603) and the full universe of N=6M. Following standard econometric diagnostics, we compute standardized mean differences (SMD) across key job-level and firm-level characteristics.

The balance results indicate that salary-disclosing postings are highly comparable to non-disclosing postings along core dimensions. Specifically, the SMDs for minimum experience requirements (0.02), educational thresholds (0.06), and skill Count (0.03) are all well below the conservative 0.1 threshold typically used to denote meaningful imbalance. While salary disclosure is more frequent among fully remote roles and within firms of modestly smaller scale-likely reflecting strategic transparency by agile firms competing for national talent-the underlying human-capital requirements remain invariant. To ensure these modest differences do not confound our results, our primary specifications in Section 5 include stringent controls for firm size (by incorporating firm fixed effects), remote-work status, and geographic-transparency mandates. Taken together, these results confirm that our analysis sample is highly representative of the broader IT labor market, alleviating concerns that the estimated wage premia are driven by selective disclosure or unobserved job-level heterogeneity.

Table 2 reports summary statistics for the main analysis sample. Panel A presents the distribution of posted salaries, which serve as the dependent variable in the wage regressions. The statistics indicate substantial variation in wages across job postings. Panel B summarizes the zero-shot-generated skill scores aggregated into three categories: IT Operations, IT Development, and AI skills, which constitute the primary explanatory variables of interest. The reported means and dispersion show meaningful variation in skill intensity across post-

ings. Panels C through E report summary statistics for education requirements, experience thresholds, and work arrangements, which are included as control variables in the baseline specifications.

Figure 4: Temporal Trends in IT Skill Portfolio (2018–2023)

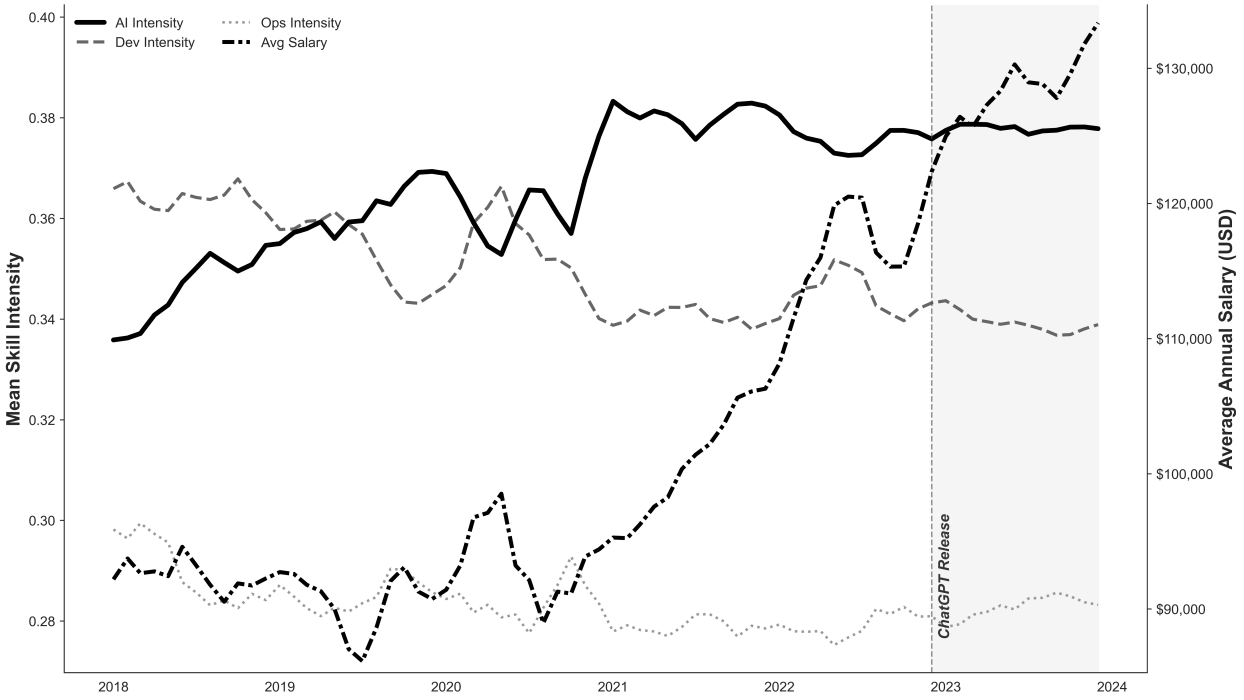


3.4 Validating the Granular LLM skills Measure

In this section, we validate our LLM-based granular measures of traditional IT and AI skills through a wage-prediction exercise. To provide an empirically robust and theoretically consistent benchmark for wage-based validation, we construct a job-posting-level AI-relatedness measure that replicates the co-occurrence approach used by Babina et al. (2024). Specifically, AI relatedness is defined based on the co-occurrence of skills with four foundational AI pillar domains: Artificial Intelligence, Machine Learning, Computer Vision, and Natural Language Processing.

Our construction follows the identical bottom-up aggregation logic as the original index:

Figure 5: Wage and Skill Trend (2018–2023)



AI Overlap Measure is first computed at the individual skill level and subsequently aggregated. However, whereas Babina et al. (2024) aggregate these indices to the firm level to capture AI human capital investment, we retain the aggregation at the job-posting level to ensure consistency with the unit of observation used in our zero-shot skill measures and validation framework. We verify that the resulting co-occurrence values closely match those reported in Babina et al. (2024). For example, the AI index score for the skill *TensorFlow* in our data is 0.918, compared to 0.9 in the original study, indicating near-perfect alignment.

We use this job-posting-level AI-overlap measure (Babina et al., 2024) in Table 4 as a univariate benchmark against our 30-dimensional Centered Log Ratio (CLR) skill portfolio. By applying the same train-test split to both the LASSO-based multidimensional model and the linear Babina-style baseline, we directly assess the incremental predictive value of our granular skill representation relative to a standard AI intensity measure.

The first row in Table 4 reports the results for the full sample from 2018 to 2023. We observe that the LLM-based measure reduces mean squared error by 10.4%, and the improvement in out-of-sample R^2 is 74.9% relative to that achieved with the AI overlap measure in Babina et al. (2024). Moreover, this analysis is repeated separately for each year, and we see MSE improvement between 7-10%. We also observe that the out-of-sample R^2 with the LLM-based measure is 1.5 to 2 times what can be achieved with Babina et al. (2024)'s AI

Table 4: Out-of-Sample Prediction Performance

Year	N	MSE (Out-of-Sample)			R-squared (Out-of-Sample)		
		LLM based Measure	AI Overlap Measure	% Reduction	LLM based Measure	AI Overlap Measure	% Increase
2018–2023	866370	0.195792	0.216268	10.37	0.1964	0.1123	74.88
2018	25683	0.229918	0.272711	15.69	0.2765	0.1418	94.99
2019	20608	0.225532	0.243067	7.21	0.2553	0.1974	29.33
2020	15509	0.217456	0.243115	10.55	0.3023	0.2200	37.41
2021	32736	0.194145	0.217703	10.82	0.2015	0.1046	92.64
2022	68808	0.145895	0.157792	7.54	0.1437	0.0739	94.45
2023	90726	0.152839	0.165852	7.85	0.1617	0.0903	79.07

Notes: Out-of-sample prediction results for log salary using 70-30 train-test split. LLM based Measure denotes LASSO regression applied to centered log-ratio-transformed 30 skill scores. The AI Overlap Measure refers to the OLS regression with the AI overlap measure, following the steps provided in Babina et al. (2024) to construct the AI index. % Reduction denotes the percentage reduction in MSE for the LLM-based measure relative to the AI Overlap Measure. % Increase shows the percentage increase in R^2 for the LLM-based measure relative to the AI Overlap Measure.

measure.

4 Identification and Empirical Strategy

We aggregate the 30 skills used as input to the transformer model for zero-shot learning into three categories. First *IT Operations (O)*, which includes skills related to maintenance and operations of systems. Precisely, this includes IT Support, Systems Administration, Network Administration, Networking, Database Management, Operating Systems, Virtualization, and Cybersecurity. These skills will enable the performance of routine IT tasks. Second is *IT Development (D)*, which involves skills in digital engineering and software development and requires more creative and complex problem-solving. This includes Software Development, Web Development, Mobile Development, Mobile App Development, Cloud Computing, DevOps, Big Data, Data Engineering, Embedded Systems, and Internet of Things (IoT). Finally, all the *AI skills (AI)* are grouped into AI, including Artificial Intelligence, Machine Learning, Reinforcement Learning, Computer Vision, Quantum Computing, Robotics, Blockchain, Blockchain Technology, Generative AI, Natural Language Processing, Large Language Models, and Data Science. Precisely, for each job posting i, we construct skill scores as: $s_{ik} = \sum_{j \in \mathcal{C}_k} p_{ij}$ where $k \in \{\text{O, D, A}\}$ and $\mathcal{C}_k$ denotes the set of skills in category k.

Each job posting is denoted by $i \in \{1, 2, \ldots, n\}$, in firm $j \in \{1, 2, \ldots, J\}$, state $s \in \{1, 2, \ldots, S\}$, in month-year $t \in \{1, 2, \ldots, T\}$. For each posting i, let $\mathbf{s}_i = (s_{iO}, s_{iD}, s_{iA})$ denote the skill composition, where s_{ik} represents the sum of probability scores over skills mentioned in category $k \in \{\text{O, D, A}\}$ corresponding to IT Operations, IT Development, and AI skills, respectively.

We cannot use the raw skill scores directly as covariates, since they are compositional data (Aitchison, 1982; Greenacre, 2021). An increase in one skill category mechanically implies changes in the relative importance of other categories. Consequently, regression coefficients cannot be interpreted in the standard way, as an increase in one skill does not hold other skills constant in relative terms. To address this, we apply the Centered Log-Ratio (CLR) transformation.

The CLR transformation treats each category symmetrically by centering on the geometric mean. This is preferable to additive log ratios (ALRs), which require choosing an arbitrary baseline category, and easier to calculate and interpret than isometric log ratios,

which involve a complex partitioning step. For each category k, we define:

$$\text{CLR}_{ik} = \log(s_{ik}) - \frac{1}{3}\sum_{k=1}^{3}\log(s_{ik}) = \log\left(\frac{s_{ik}}{g_i}\right) \tag{1}$$

where $g_i = \left(\prod_{k=1}^{3} s_{ik}\right)^{1/3}$ is the geometric mean of the three skill scores. By construction, CLR coordinates sum to zero: $\sum_{k=1}^{3}\text{CLR}_{ik} = 0$. This creates perfect multicollinearity, so regression removes one of the categories which henceforth we will call the baseline category. Our primary specification with main effects, two-way and three-way interactions is:

$$\begin{aligned}\log(w_{ijst}) = &\sum_{k=1}^{2}\beta_k \text{CLR}_{ik} + \sum_{k=1}^{2}\sum_{\ell>k}^{3}\gamma_{k\ell}(\text{CLR}_{ik}\times\text{CLR}_{i\ell}) \\ &+ \delta_{ODA}(\text{CLR}_{iO}\times\text{CLR}_{iD}\times\text{CLR}_{iA}) + \mathbf{X}'_{ijst}\boldsymbol{\theta} + \alpha_j + \mu_s + \lambda_t + \epsilon_{ijst}\end{aligned} \tag{2}$$

where w_{ijst} is the wage for job posting i at firm j in state s at time t. Fixed effects include α_j (firm), μ_s (state), and λ_t (year-month). The vector $\mathbf{X}_{ijst}$ includes education requirements, experience requirements, employment type, and remote work indicators. Standard errors are clustered at the firm level.

4.1 Identification

Our identifying assumption is that, conditional on firm, state, time fixed effects, and posting level covariates, variation in skill requirements across postings is not correlated with the error term:

$$\mathbb{E}[\epsilon_{ijst}|\text{CLR}_{iO}, \text{CLR}_{iD}, \text{CLR}_{iA}, \mathbf{X}_{ijt}, \alpha_j, \mu_s, \lambda_t] = 0 \tag{3}$$

The parameters of interest are identified from within-firm, within-state, within-time variation in skill composition across job postings. The firm fixed effects absorb: α_j captures time-invariant firm characteristics (productivity, wage policy structure, industry specifics); μ_s captures state-level wage differences (cost of living, local labor markets, minimum wage policies); and λ_t captures aggregate time shocks (business cycles, Generative AI introduction, inflation).

4.2 Interpretation of Parameters

In the model with two-way and three-way interactions, the marginal effect of skill category k on wages is:

$$\frac{\partial \log(w_{ijst})}{\partial \text{CLR}_{ik}} = \beta_k + \sum_{\ell\neq k}\gamma_{k\ell}\cdot\text{CLR}_{i\ell} + \delta_{ODA}\cdot\text{CLR}_{im}\cdot\text{CLR}_{in} \tag{4}$$

where $m, n \neq k$ denote the two other skill categories.

The coefficient β_k represents the marginal wage return to skill category k when all other CLR-transformed skills equal zero (i.e., at their geometric mean). Since this configuration may not represent typical job postings in our sample, we report marginal effects evaluated at meaningful skill compositions (e.g., sample mean, specific percentiles) in our results section.

In a fully interacted model, the two-way interaction coefficients $\gamma_{k\ell}$ cannot be interpreted in isolation. The complementarity or substitutability between skills k and ℓ depends on the level of the third skill. Specifically, the cross-partial derivative measuring the interaction between skills k and ℓ is:

$$\frac{\partial^2 \log(w_{ijst})}{\partial \text{CLR}_{ik} \partial \text{CLR}_{i\ell}} = \gamma_{k\ell} + \delta_{ODA} \cdot \text{CLR}_{im} \tag{5}$$

where m denotes the third skill category not indexed by k or ℓ.

Thus, $\gamma_{k\ell}$ represents the complementarity between skills k and ℓ only when the third skill is at its geometric mean ($\text{CLR}_{im} = \log(\frac{g_i}{g_i}) = 0$). For example, γ_{DA} captures AI-Development complementarity when Operations skills are at their geometric mean. The three-way interaction term δ_{ODA} captures moderation effects: how the complementarity between any two skills depends on the level of the third skill. A positive $\delta_{ODA} > 0$ indicates that skill synergies are amplified when all three skills are present. For instance, if $\delta_{ODA} > 0$, the AI-Development complementarity ($\gamma_{DA} + \delta_{ODA} \cdot \text{CLR}_O$) is stronger in jobs that also emphasize Operations skills. Given this interaction structure, we present our main results through marginal effect visualizations that show how returns to each skill vary across the full distribution of complementary skills, rather than relying solely on coefficient magnitudes, which cannot be interpreted in isolation in the presence of higher-order interactions.

5 Results

5.1 Main Specification

We start by reporting the results for the regression with no interactions, and then we introduce the higher-order interactions. The outcome variable wages is expressed in logs, and we transform the three aggregated skills, Operations, Development, and AI skill measure, into CLRs based on the discussion in the previous section. Table 5 provides the results with the main effects only without interactions in column (1), and then we gradually add two way interactions (Column (2)) and finally the three way interactions in Column (3). Each specification includes firm-level, state-level, and year-month fixed effects, and the standard

errors are clustered at the firm level. We also include a set of controls: dummy variables for the education degree required in the job ad, minimum and maximum experience, remote, in-person, or hybrid, an internship dummy, and employment type (full-time and part-time).

In Column (1) in Table 5 we observe that relative to the baseline of AI CLR there is a positive premium for both IT Operations and Development. However, the magnitude for IT Development is much higher and the 95% confidence interval does not include 0. On the contrary, the premium to IT Operations is not significant at the conventional level of significance (5%). Column (2) presents the results with two way interactions of AI with IT Operations and Development. We find that the main effects on Operations and Development relative to the baseline for AI are positive, and that AI skills are complementary to IT development. However, it suggests a substitution between Operations and AI. Finally, we present the results of the fully interacted 3 way interactions. We find that main effects and two way interactions are in the same direction as Columns (1) and (2). However, there appears to be a premium to having all 3 skills.

We calculate the marginal effects for Column (3) using the specification provided below.

$$\frac{\partial \log(w)}{\partial \text{CLR AI}} = -0.233 \cdot \text{CLR Ops} + 0.342 \cdot \text{CLR Dev} + 0.189 \cdot \text{CLR Ops} \cdot \text{CLR Dev} \qquad (6)$$

where the coefficients correspond to the two-way interactions (Operations $\times$ AI and Development $\times$ AI) and the three-way interaction (Operations $\times$ Development $\times$ AI) from Column (3) in Table 5. The marginal effect of AI is a function of the other two CLRs. Therefore, we vary the Operations and Development CLR between their values of 10th, 25th, 50th, 75th and 90th percentile and calculate the marginal effect of AI for each pair of these combinations (25 in total).

Figure 6 provides a heatmap to visualize these marginal effects. The heatmap reveals three key patterns. First, the negative coefficient on Operations $\times$ AI indicates substitution: moving from the 10th to the 90th percentile of Operations reduces the marginal effect, with the marginal effect declining from 0.05 to -0.07 in the bottom row (low Development). Second, the positive coefficient on Development $\times$ AI indicates complementarity: moving from the 10th to the 90th percentile of Development increases the marginal effect, with values rising from 0.05 to 0.17 in the left column (low Operations). Third, and most importantly, the positive three-way interaction (+0.189) manifests as non-parallel rate changes across the grid. Specifically, when moving from low to high Operations, the decline in marginal effect is smaller at high Development levels (-0.08, from 0.17 to 0.09 in the top row) than at low Development levels (-0.12, from 0.05 to -0.07 in the bottom row). This 0.04 difference in log points demonstrates that the substitution effect of Operations on AI value is attenuated

Table 5: Wage Premiums using CLR Transformations

	Dependent variable: Log(Salary)		
	Main Effects	2-Way Interactions	3-Way Interactions
	(1)	(2)	(3)
CLR: Operations	0.046	0.157***	0.137***
	(0.061)	(0.045)	(0.046)
CLR: Development	0.318***	0.543***	0.536***
	(0.097)	(0.090)	(0.085)
Operations × Development		0.149***	0.206***
		(0.017)	(0.033)
Operations × AI		−0.296***	−0.233***
		(0.034)	(0.032)
Development × AI		0.231***	0.342***
		(0.055)	(0.060)
Operations × Development × AI			0.189***
			(0.054)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	870,797	870,797	870,797
R^2	0.531	0.538	0.538
Adjusted R^2	0.530	0.537	0.537

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Remote work baseline: Not Remote. Employment type baseline: Either Full-time or Part-time. For education the baseline is dummy for no education listed. Controls include education level, years of experience, employment type, remote work status and internship indicator. Standard errors clustered at the firm level are reported in parentheses.

when Development skills are also present. The non-linear curvature in the color transitions across the heatmap visualizes this moderation effect, confirming that the return to AI skills depends not only on the presence of complementary Development skills but also on how these interact with Operations expertise. Moreover, the premium of Operations with Development on the marginal effect of AI is evident from the diagonal when both skills are moved from the 10th to the 90th percentile.

Figure 6: Marginal Effect of AI Skill on Wages

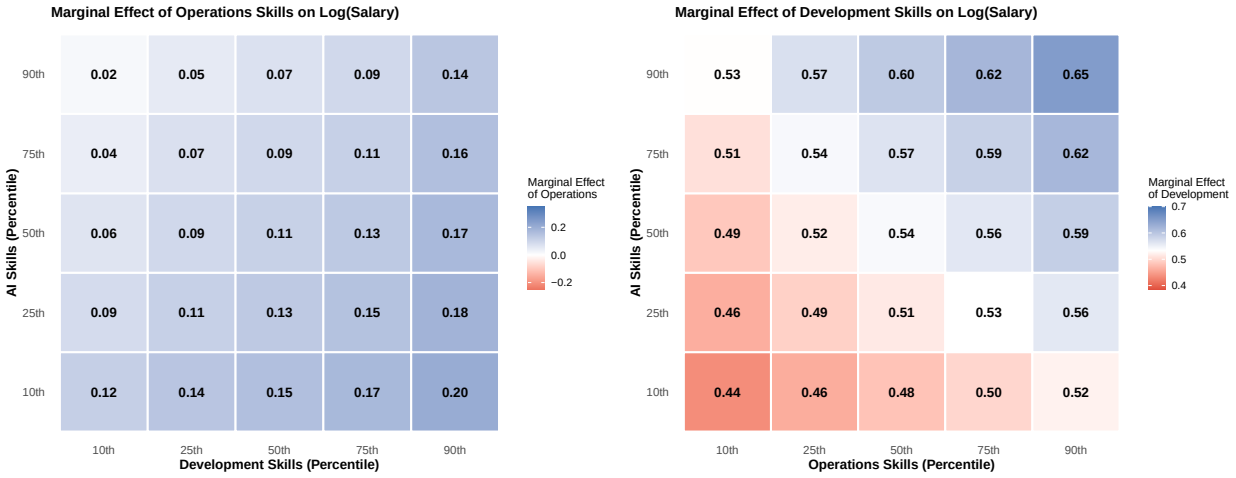


Next, we present the heatmaps for the marginal effects of IT operations and IT development. Panels (a) and (b) in Figure 7 provide these graphs. Panel (a) in Figure 7 displays the marginal effect of Operations skills on log salary across different combinations of Development and AI skill levels. The marginal effect is calculated as $\frac{\partial \log(w)}{\partial \text{CLR Ops}} = 0.137 + 0.206 \cdot \text{CLR Dev} - 0.233 \cdot \text{CLR AI} + 0.189 \cdot \text{CLR Dev} \cdot \text{CLR AI}$, with values ranging from 0.02 to 0.20 across the skill space. Operations skills exhibit universally positive re-

turns, though the magnitude varies systematically with complementary skills. The vertical gradient reveals the AI substitution effect: when AI skills are low (10th percentile, bottom row), Operations commands the highest marginal returns of 0.12 to 0.20, but as AI increases to the 90th percentile (top row), these returns diminish substantially to 0.02 to 0.14. This pattern confirms that AI and Operations are substitutes, as AI tools likely automate routine IT operations, thereby decreasing the incremental value of Operations expertise. The horizontal gradient demonstrates Development complementarity: Operations returns are lowest at 0.02 to 0.12 when Development is at the 10th percentile (left column) and increase to 0.14 to 0.20 when Development reaches the 90th percentile (right column), representing an approximately 0.08 log-point increase. The three-way interaction manifests as non-parallel changes in rate across the grid. Specifically, the wage penalty from high AI on Operations returns is attenuated when Development skills are also elevated. In the upper-right corner, where both Development and AI are at the 90th percentile, the Operations marginal effect is 0.14, substantially higher than would be predicted by the two-way interactions alone, demonstrating that Development expertise moderates the AI-Operations substitution.

Panel (b) in Figure 7 displays the marginal effect of Development skills on log salary across different combinations of Operations and AI skill levels. The marginal effect is calculated as $\frac{\partial \log(w)}{\partial \text{CLR Dev}} = 0.536 + 0.206 \cdot \text{CLR Ops} + 0.342 \cdot \text{CLR AI} + 0.189 \cdot \text{CLR Ops} \cdot \text{CLR AI}$, with values ranging from 0.44 to 0.65. Development skills command uniformly positive and substantial returns across all skill combinations, demonstrating that software development capabilities are universally valued in the IT labor market regardless of other skills possessed. The vertical gradient illustrates the strong AI complementarity: Development returns increase from 0.44 to 0.53 when AI is at the 10th percentile (bottom row) to 0.52 to 0.65 when AI reaches the 90th percentile (top row), representing wage premiums of approximately 0.09 to 0.13 log points. This suggests that AI and Development are complements, as building machine learning pipelines and AI-powered applications requires both AI expertise and software development capabilities. The horizontal gradient shows the Operations complementarity. Development returns rise by approximately 0.08 log points as Operations increases from the 10th to 90th percentile across all rows, reflecting the positive coefficient on the IT Operations $\times$ IT Development interaction term. The three-way interaction creates a multiplicative synergy most visible in the upper-right corner where the marginal effect reaches 0.65. The optimal IT skill portfolio thus integrates all three dimensions, with Development serving as the critical foundation that enables productive application of both AI and Operations knowledge.

Figure 7: Marginal Effect of Operations and Development



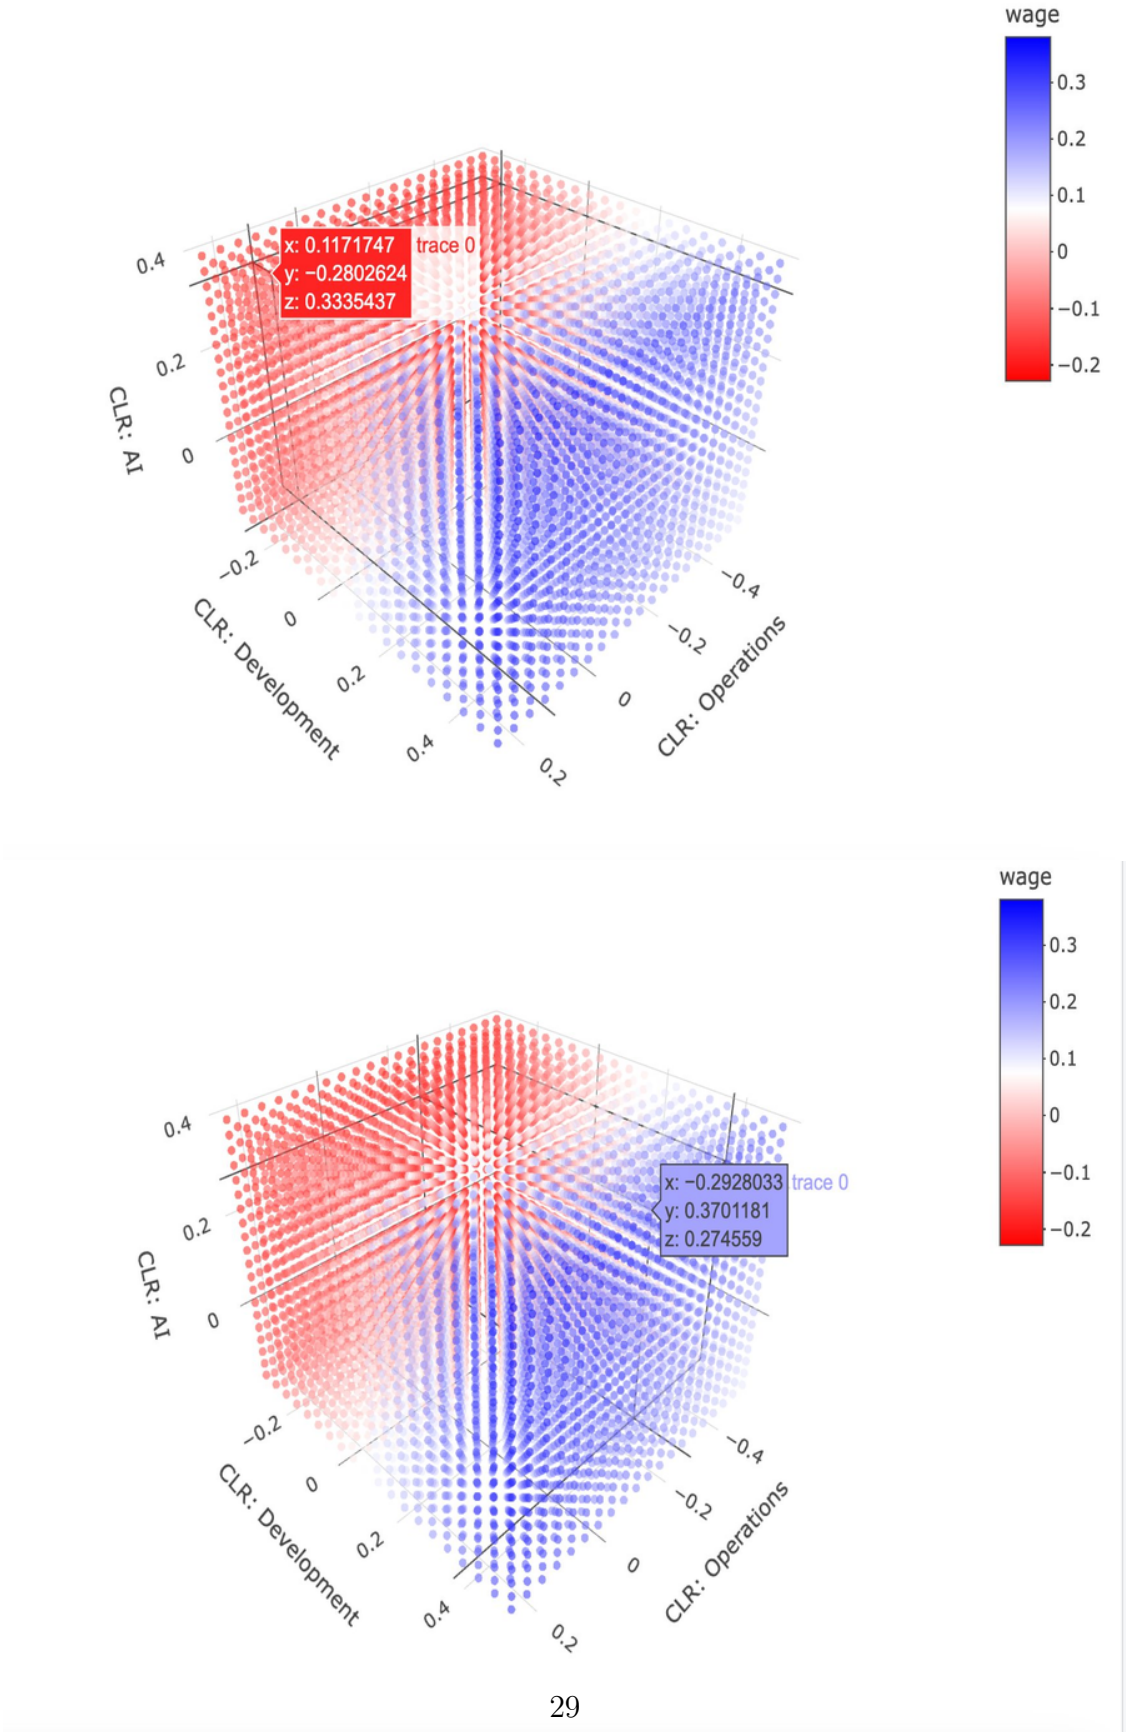
a. Operations b. Development

Figure 8 visualizes the estimated wage contribution from skill portfolios across all three dimensions of IT human capital. The estimated log wage contribution is calculated as:

$$\begin{aligned}\widehat{\log(w)} = {} & 0.137 \cdot \text{CLR Ops} + 0.536 \cdot \text{CLR Dev} \\ & + 0.206 \cdot \text{CLR Ops} \times \text{CLR Dev} \\ & - 0.233 \cdot \text{CLR Ops} \times \text{CLR AI} \\ & + 0.342 \cdot \text{CLR Dev} \times \text{CLR AI} \\ & + 0.189 \cdot \text{CLR Ops} \times \text{CLR Dev} \times \text{CLR AI}\end{aligned} \tag{7}$$

where each point in the three-dimensional space represents a unique skill portfolio combination, with position determined by CLR-transformed Operations (x-axis), Development (y-axis), and AI (z-axis) skill levels, and color intensity indicating the estimated log wage contribution from that skill bundle. The visualization reveals several key patterns in the wage structure. The highest predicted wages (blue region) concentrate in the upper-back-right octant where all three skills are elevated, particularly when Development and AI skills are both high. This clustering confirms the strong complementarity between Development and AI skills ($\beta_{\text{D,AI}} = 0.342$), as workers who combine software development capabilities with AI expertise command substantial wage premiums. The lowest estimated wages (red region) occur in the lower-front-left octant, where all skills are below their medians, indicating that the absence of a well-rounded IT skill portfolio results in significant wage penalties. Examining the vertical gradients along the AI dimension (z-axis), the rate at which wages

Figure 8: Three-Dimensional Skills and Contribution to Estimated Wages



increase with AI skills varies systematically with Development levels: when Development skills are high (positive y-values), increasing AI generates steeper wage gains, while when Development skills are low (negative y-values), AI additions yield smaller returns or even wage decreases due to the substitution effect with Operations skills ($\beta_{\text{O,AI}} = -0.233$). This three-dimensional wage landscape illustrates that IT workers face a complex optimization problem. While high levels of any single skill dimension generate positive returns, the highest wages accrue to workers who strategically combine complementary skills, with Development serving as the critical skill that unlocks the value of both AI and Operations.

5.2 Heterogeneous Impact of Skill Type on Wages

In this section, we examine heterogeneity in the wage premiums for these three skills by experience, work arrangement type, and company location within tech hubs.

Experience We begin by splitting the sample by the required experience level in the job posting. We construct three groups: junior hire with 0-2 years of experience, mid tier hire with 3-7 years of experience and finally senior tier with atleast 8 years of experience.[6] The wage premium to the three types of IT skills likely vary by experience as the senior roles often require non technical skills such as managerial and cognitive skills. Consequently, we hypothesize that the complementarity premium between the skills is highest for the junior roles and the premium declines with seniority.

Table 6 reports the results for the three groups separately in Columns (1)–(3). As for the coefficients on the interactions, there appears to be a positive premium for all three skills only among entry-level and mid-tier workers, and none among senior workers. The premium of combining development skills with AI appears to be driven primarily by entry- and mid-tier jobs, with the highest premium for entry-level jobs.

However, in a fully interacted model, one cannot interpret the coefficients on higher-order interactions in isolation. We use the same calculation as before, in which we estimate the marginal effect of AI on wages as a function of both operations and development skills, since they enter the marginal effect formula through higher-order interactions, as shown below.

$$\frac{\partial \log(w)}{\partial \text{CLR AI}} = \beta_{O,A} \cdot \text{CLR Ops} + \beta_{D,A} \cdot \text{CLR Dev} + \beta_{O,D,A} \cdot \text{CLR Ops} \cdot \text{CLR Dev} \quad (8)$$

[6] While constructing the covariates, about 29% of postings did not explicitly mention minimum years of experience required. We include a missing experience dummy for the main regressions. However, for the heterogeneity analysis, the sample without any minimum experience required is added to 0-2 years to keep the samples comparable to the main analysis.

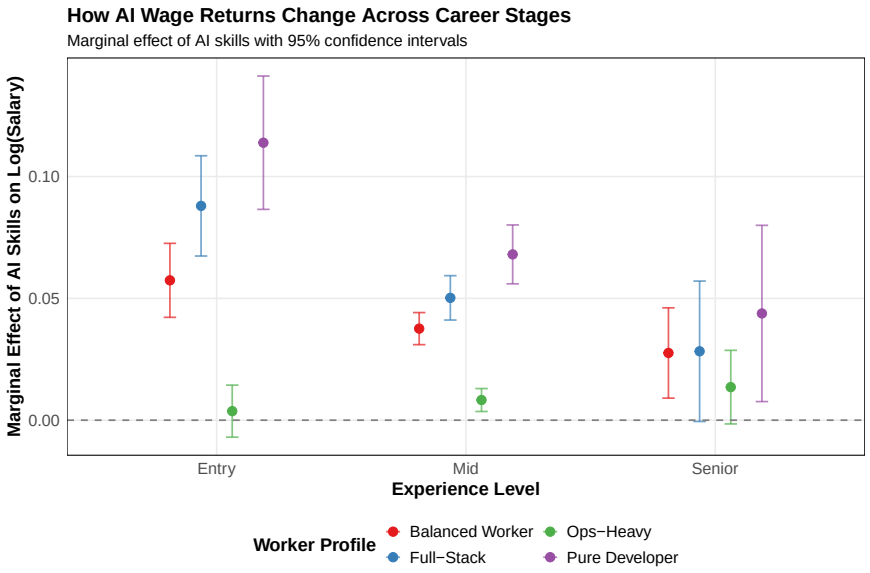
Table 6: Wage Premiums by Experience Level: Three-Way Skill Interactions

	Dependent variable:		
	Log(Salary)		
	Entry (0-2 yrs)	Mid (3-7 yrs)	Senior (8+ yrs)
	(1)	(2)	(3)
CLR: Operations	0.195***	0.035	0.057*
	(0.063)	(0.030)	(0.029)
CLR: Development	0.659***	0.323***	0.184***
	(0.107)	(0.041)	(0.036)
Operations × Development	0.201***	0.184***	0.085**
	(0.035)	(0.023)	(0.034)
Operations × AI	−0.281***	−0.195***	−0.155***
	(0.043)	(0.019)	(0.052)
Development × AI	0.384***	0.168***	0.003
	(0.066)	(0.030)	(0.106)
Operations × Development × AI	0.156**	0.119***	−0.030
	(0.072)	(0.045)	(0.101)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	367,750	369,784	133,263
R^2	0.551	0.452	0.451
Adjusted R^2	0.548	0.448	0.445

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Remote work baseline: Not Remote. Employment type baseline: Either Full-time or Part-time. Controls include education level, years of experience, employment type, remote work status, and an indicator for internship. Note that there are 29% job postings do not explicitly require a minimum number of years of experience. This is part of column (1) sample entry-level experience for the heterogeneity analysis. Standard errors clustered at the firm level are reported in parentheses.

We can now consider a grid of values for operations and development. We extract the percentiles for these two skills and calculate the marginal effect of AI on wages for every pair of these two skill percentiles. This helps us visualize changes in marginal effects as we vary one skill while holding the other constant, or as we increase both skills simultaneously. To implement this analysis, we consider four portfolios. We call these pure developer (High Dev (p75), Low Ops (p25)), balanced worker (Medium Dev (p50), Medium Ops (p50)), Ops heavy (Low Dev (p25), High Ops (p75)), and full stack (High Dev (p75), Medium Ops (p50)). Figure 9 shows the marginal effects for AI skills by experience and worker skill portfolios. We learn three key insights from this analysis. First, the marginal effect of AI skills is highest for entry-level jobs across all portfolios, except for Ops-heavy roles, relative to mid-tier and senior jobs. Second, this premium increases with higher development skills, peaking for a pure developer (entry and mid-tier). Third, this skill complementarity is absent among senior workers.

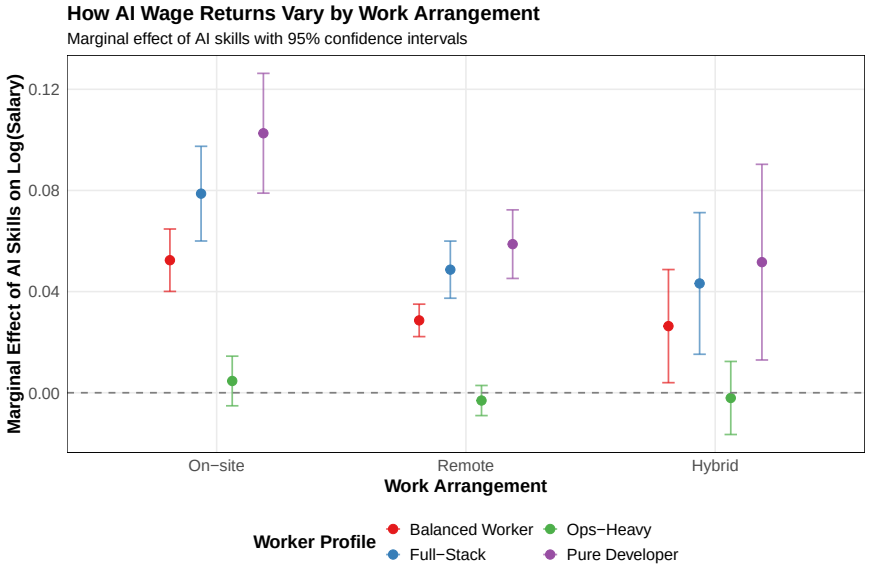
Figure 9: Marginal Effect of AI Skill on Wages by Experience



Work Arrangement Next, we explore heterogeneity by work arrangement. Note that we consider data from 2018 to 2023, including the onset of remote work during the 2020 pandemic. In our sample, job postings mention three types of worker arrangements. Such remote arrangements are much more likely in IT jobs than in other occupations. This includes the largest proportion of on-site jobs (73.1%), 23.8% remote jobs, and the remaining jobs

that allow hybrid in-person and remote work arrangements.[7] Table 7 presents the results of the three-way interaction model, disaggregated by work arrangement.

Figure 10: Marginal Effect of AI Skill on Wages by Worker Arrangement



As in the analysis above, we calculate the marginal effect of AI on log wages across four worker portfolios as used in the previous section on experience. Figure 10 shows the marginal effects separately for the skill portfolios and worker arrangement. This indicates a positive marginal effect of AI on wages across all work arrangements and worker types, except for Ops-heavy workers. However, these AI premiums are highest when combined with strong development skills across the full stack and balanced portfolios. Because the sample size for hybrid workers is small, we observe larger standard errors for this group than for the other two groups.

Tech Hubs. Finally, we examine the heterogeneity of skill demand for wage premia by whether the firm's headquarters is a tech hub. This is critical because labor market characteristics likely vary substantially across tech hubs relative to other regions. For this analysis, we use Core Based Statistical Area (CBSAs) mapped to each firm in the Compustat dataset. We identify five major tech hubs: San Jose, SF, Seattle, Austin, and Boston. Additionally, we classify four CBSAs as secondary tech hubs that include Denver, Atlanta, NYC, and

[7]The "on-site" sample for the heterogeneity analysis is constructed using the remote: none and remote: stated none categories, where the former captures job postings without any mention of remote work, and the latter are postings that explicitly mention remote work is not allowed.

Table 7: Wage Premiums by Work Arrangement: Three-Way Skill Interactions

	Dependent variable:		
	Log(Salary)		
	On-site	Remote	Hybrid
	(1)	(2)	(3)
CLR: Operations	0.176***	−0.030*	−0.032
	(0.045)	(0.018)	(0.069)
CLR: Development	0.589***	0.283***	0.273***
	(0.080)	(0.040)	(0.070)
Operations × Development	0.198***	0.219***	0.261***
	(0.032)	(0.026)	(0.049)
Operations × AI	−0.259***	−0.133***	−0.125*
	(0.035)	(0.018)	(0.064)
Development × AI	0.333***	0.284***	0.261***
	(0.068)	(0.045)	(0.080)
Operations × Development × AI	0.144**	0.283***	0.365***
	(0.057)	(0.041)	(0.102)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	636,928	211,622	22,247
R^2	0.533	0.604	0.619
Adjusted R^2	0.531	0.601	0.603

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Remote stated none has been combined with remote none for contructing the "on-site" sample. Controls include education level, years of experience, employment type, and an indicator for internship. Standard errors clustered at the firm level are reported in parentheses.

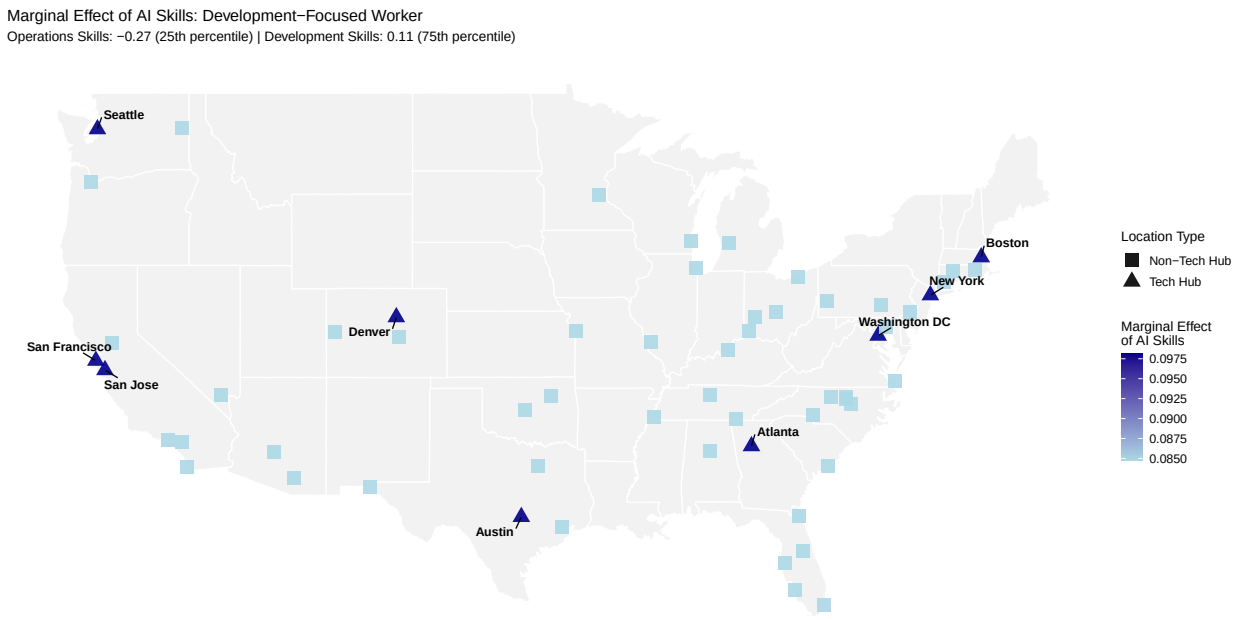
Washington DC. Together, these nine CBSAs will be called "tech hubs," and the remaining CBSAs will be called "other regions". In the sample used in this paper, 267,367 jobs are listed for firms headquartered in tech hubs, and the remainder are in other regions.

Table 8: Wage Premiums: Tech Hubs vs Non-Tech Hubs

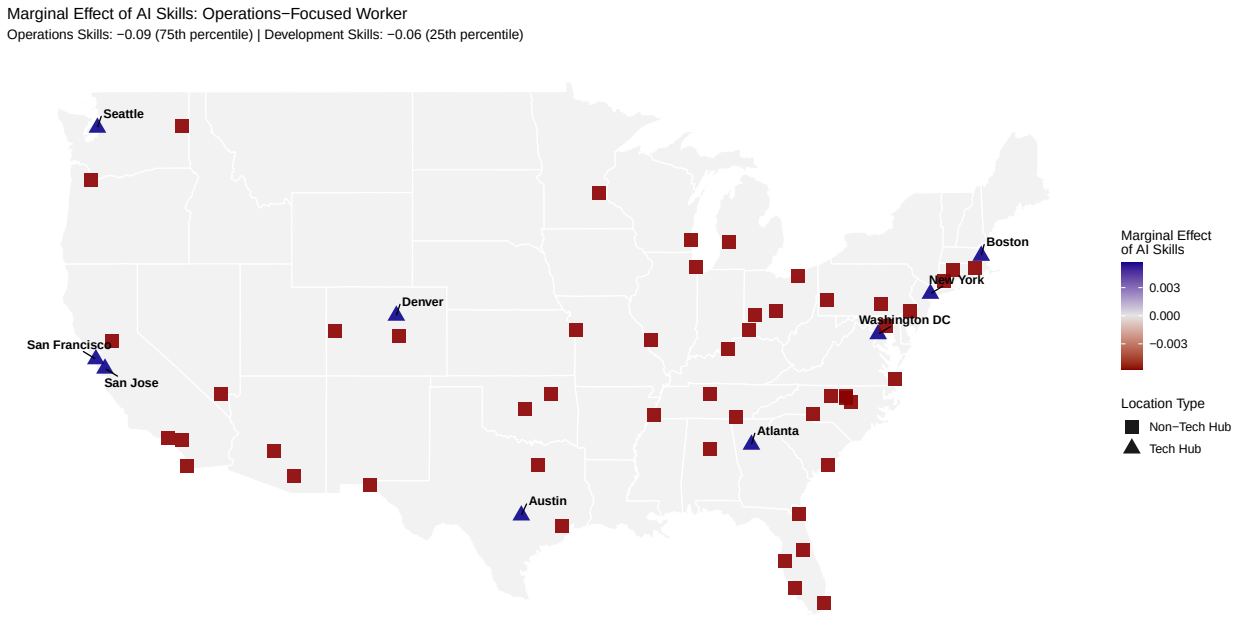
	Dependent variable:	
	Log(Salary)	
	Tech Hub	Non-Tech Hub
	(1)	(2)
CLR: Operations	0.109*	0.140***
	(0.056)	(0.037)
CLR: Development	0.492***	0.539***
	(0.104)	(0.076)
Operations × Development	0.213***	0.195***
	(0.044)	(0.036)
Operations × AI	−0.257***	−0.176***
	(0.047)	(0.031)
Development × AI	0.322***	0.387***
	(0.080)	(0.058)
Operations × Development × AI	0.240***	0.176***
	(0.059)	(0.059)
Firm FE	Yes	Yes
State FE	Yes	Yes
Year-Month FE	Yes	Yes
Controls	Yes	Yes
SE Clustered by Firm	Yes	Yes
Observations	267,367	543,089
R^2	0.478	0.560
Adjusted R^2	0.475	0.558

Table 8 shows the regression results separately done for the job postings from firms in tech hubs and others in Columns (1)-(2). We observe that the overall direction of the

Figure 11: Marginal Effect of AI and Tech Hubs



a. Development Focus



b. Operations Focus

Notes: This graph displays only 50 non-tech CBSAs for visualization purposes. The sample comprises 904 unique CBSAs.

coefficients remains consistent with the results from our full-sample specification. However, there appear to be differences in the magnitudes of the higher-order interactions between tech and non-tech hub jobs. To shed more light on the AI premium for wages, we create two worker profiles: i) Dev heavy profile (75th percentile), low operations (25th percentile), and ii) Ops heavy profile (25th percentile), low development (75th percentile). Using these, we calculate the marginal effect of AI and display it using a gradient scale in Panels (a) and (b) in Figure 11. Panel (a) shows that, for development-focused workers, the AI wage premium is higher for jobs listed by firms in tech hubs than in non-tech hubs. On the contrary, the premium for the Ops-heavy profile is negligible in both types of areas, as evidenced by very small marginal effects in panel (b) in terms of magnitude.

6 Discussion

In this paper, we leverage large language models to construct multivariate, data-driven measures of AI skill demand and to estimate wage premiums, driven by the complementarities between emerging AI skills and established IT expertise. First, we show that our multivariate, skill-based measures substantially outperform existing unidimensional AI exposure indices (Babina et al., 2024) in predicting wages. Next, our results demonstrate that wage premia are driven not by AI skills alone, but by complementarities between AI and traditional IT skills in development and operations. AI weakens the skill valuation of IT operational skills while amplifying returns to IT development capabilities. These effects vary systematically with worker experience. AI-related skill bundles command relatively larger wage premia among less-experienced workers (or postings that do not explicitly state any experience requirement) but generate more muted marginal returns at higher experience levels. The results are relatively stable across geography and work arrangements. These findings provide a granular mechanism through which AI-driven technological change affects IT labor markets, moving beyond occupation-level displacement narratives toward a portfolio-based view of skill repricing.

Implications Our study highlights the limitations of univariate AI exposure measures and static skill taxonomies commonly used in research on technology and labor (Acemoglu et al., 2022; Autor, 2015; Babina et al., 2024). By leveraging large language models to extract latent, multidimensional skill demand and embedding these measures in a compositional framework, we offer a scalable approach for studying technological change in settings char-

acterized by rapid innovation and implicit skill requirements. A central insight is that AI induces heterogeneous valuation of skills within narrowly defined IT jobs. Even within jobs classified under the same occupational classification, we observe that some skill domains receive wage premiums, whereas others pay lower wages. Skills associated with IT development exhibit positive wage associations, whereas more routine operations-oriented skills show weaker or even negative wage associations.

Our results extend the task-based literature by demonstrating that such reallocation operates at a more nuanced level than previously documented. Occupation-level or task-category analyses implicitly assume homogeneity within roles. By contrast, our compositional framework reveals that AI simultaneously raises and lowers the value of different skills within the same job family. This helps explain why aggregate employment and wage effects of AI often appear muted despite substantial restructuring at the task level (Acemoglu et al., 2022; Bessen, 2019). Aggregated measures average over offsetting forces that are clearly visible when skills are examined jointly.

Beyond documenting complementarities between AI and traditional IT skills, our results reveal systematic asymmetries in how AI interacts with different components of the IT skill bundle, as well as heterogeneity by worker experience and the more limited variations across organizational and geographic contexts. While AI exhibits strong complementarities with IT development skills, its interaction with IT Operations skills is negative, so that "pure" operations roles are at the highest risk of stagnation or displacement, as AI tools increasingly automate routine system monitoring and troubleshooting (Acemoglu and Restrepo, 2018; Autor and Dorn, 2013).

Importantly, this negative interaction does not imply a uniform devaluation of operations expertise. Rather, AI reshapes the relative contribution of operational skills, reducing their marginal wage association in AI-intensive environments. Consistent with this interpretation, we observe a positive three-way interaction among AI skills, IT development skills, and IT operations skills, suggesting that operations expertise regains economic value when embedded within broader, integrative skill bundles. Workers who can design systems while understanding their operational constraints benefit from AI even as routine execution tasks are automated. As a useful case in point, Anthropic AI's internal research (Huang et al., 2025) on software engineers revealed that AI enables workers to become more "full-stack," moving beyond their traditional expertise. However, it noted that while AI handles routine coding and debugging, engineers are increasingly focused on high-level design, planning, and oversight, echoing the reinstatement of judgment-intensive work. This finding is also

validated by PwC AI Jobs Barometer Report (2025).

A critical finding from our analysis is the heterogeneity of AI returns across career stages. The concentration of AI development wage premia among junior and mid-career workers suggests that AI primarily reshapes early task bundles within jobs, reallocating effort from execution toward design, integration, and system-level problem-solving, with diminishing marginal effects once workers transition into senior, coordination-intensive roles (Brynjolfsson et al., 2023).

The spatial heterogeneity observed in the results implies that AI-driven task reallocation is mediated by local labor market thickness and knowledge spillovers, with high-density tech regions better able to convert AI into reinstated, high-value IT Development-AI task bundles Andreadis et al. (2025). This suggests that the economic returns to AI skills are not purely at the individual level but depend on regional ecosystems that support rapid recombination, learning, and deployment of complementary IT capabilities. Our results are consistent with Andreadis et al. (2025), who have shown that the characteristics of the local labor market combined with other factors such as education will determine the overall gains of AI.

Limitations and Future Research Despite the robustness of our LLM-based metrics, this study is limited in several respects. First, our analysis is restricted to publicly traded firms and job postings with explicitly listed salaries. This may over-represent established firms and overlook more volatile wage dynamics in startup and private equity-backed sectors. Second, while our Generative AI approach successfully identifies latent skills (e.g., React Native implies JavaScript), it cannot account for skills that are omitted due to extreme commonality or obviousness in certain niches, potentially introducing residual measurement noise. Third, wage premia derived from job postings are a real-time indicator of demand but may not fully reflect long-term equilibrium wages as the supply of AI-trained workers increases. Our study does not include any supply-side data. Future research can leverage LinkedIn profiles to examine how workers are responding to labor market changes through reskilling. Finally, our model does not directly observe the actual use of AI within firms; it only observes the posted demand for AI skills, as reflected in job advertisements. Future work should link these skill-based measures to internal productivity data to determine whether these wage premia translate into realized firm-level productivity gains.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D., Autor, D., and Johnson, S. (2023). Can we have pro-worker ai? *CEPR Policy Insight*, 123.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30.

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

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

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44(2):139–160.

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

Amazon.com, Inc. (2025). Amazon.com announces second quarter results. Earnings release, Amazon Investor Relations.

Andreadis, L., Kalotychou, E., Chatzikonstantinou, M., Louca, C., and Makridis, C. A. (2025). Local heterogeneity in artificial intelligence jobs over time and space. In *AEA Papers and Proceedings*, volume 115, pages 29–34. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Autor, D. and Salomons, A. (2018). Is automation labor-displacing? productivity growth, employment, and the labor share. Technical report, National Bureau of Economic Research.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, 29(3):3–30.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review*, 103(5):1553–1597.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth,

and product innovation. *Journal of Financial Economics*, 151:103745.

Bensinger, G. (2025). Amazon targets as many as 30,000 corporate job cuts, sources say. *Reuters*. Accessed: January 22, 2026.

Bessen, J. (2019). Automation and jobs: When technology boosts employment. *Economic Policy*, 34(100):589–626.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. Technical report, National Bureau of Economic Research.

Dai, R., Liu, T., and Wang, C. R. (2025). Employee turnover in the spotlight: Evidence from mergers and acquisitions. Wharton research paper, Wharton Research Data Services (WRDS). Posted June 3, 2025; revised August 20, 2025.

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

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

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

Deming, D. J. and Noray, K. L. (2018). Stem careers and the changing skill requirements of work. Technical report, National Bureau of Economic Research.

Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dixon, J., Hong, B., and Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Goldfarb, A., Taska, B., and Teodoridis, F. (2023). Could machine learning be a general purpose technology? evidence from online job postings. *Research Policy*, 52(1):104653.

Greenacre, M. (2021). Compositional data analysis. *Annual Review of Statistics and its Application*, 8(1):271–299.

Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.

Huang, S., Seethor, B., Durmus, E., Handa, K., McCain, M., Stern, M., and Ganguli, D. (2025). How ai is transforming work at anthropic.

McAfee, A. and Brynjolfsson, E. (2018). *Machine, Platform, Crowd: Wie wir das Beste aus unserer digitalen Zukunft machen*. Plassen Verlag.

ONET Development (2025). Using a hybrid artificial intelligence-expert method to develop work style ratings for the O*NET database. Technical report, U.S. Department of Labor, Employment and Training Administration.

PwC AI Jobs Barometer Report (2025). Ai jobs barometer: Industry. https://pwc.turtl.co/story/ai-jobs-barometer-industry/page/6/1.

Tambe, P. B. (2025). Reskilling the workforce for ai: Domain expertise and algorithmic literacy. *Management Science*. Forthcoming.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Yin, W., Hay, J., and Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. *arXiv preprint arXiv:1909.00161*.