Penetration of Large Language Models across Demographic and Occupational Subgroups

Riccardo Omenti[1,*], Nicola Barban[1], Saverio Minardi[1], and Maria Ludovica Isani[1]

[1]*Department of Statistical Sciences, University of Bologna*

[*]*Corresponding email: riccardo.omenti2@unibo.it*

October 2025

Abstract

This paper aims to investigate the penetration of large language models across demographic, geographic, and occupational groups in the United States through the development of new ad-hoc indicators. By relying on evaluations from LLM, we quantify on a 0–100 scale the extent to which specific tasks can be complemented by LLMs, aggregating these task-level scores to the occupation level with task-relevance weights. Merging these occupational scores with micro-data from the American Community Survey, we construct two indicators: the LLM Penetration Index and the LLM High-Exposure Index. Preliminary findings show distinct gender patterns: exposure peaks for women with some college education and median earnings, while among men, it is highest for those with a doctoral degree and top wages. These results suggest that LLM diffusion may exacerbate existing labor market inequalities, amplifying advantages for high-skilled male workers and providing limited gains to women in middle-skilled occupations.

1 Introduction

The proliferation of Generative Artificial Intelligence (GenAI) systems, especially large language models (LLMs), has radically changed the labor market landscape. Unlike earlier waves of automation, such as the spread of industrial robots that primarily displaced manual and low-skilled workers (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020, 2022), these novel technological systems have affected occupations once thought to be exempted from automation (Felten et al., 2021). Their capacity to generate, summarize, and analyze text, code and data implies that many high-skilled and knowledge-intensive occupations are now directly exposed to the technological transformation determined by the diffusion of GenAI. Beyond this, these technological systems have the potential to reconfigure cognitive workflows, change the organization of professional work, and challenge traditional boundaries between human expertise and machine capability.

Although GenAI systems have become more and more accessible in the last couple of years (Stanford Institute for Human-Centered Artificial Intelligence, 2025), their diffusion has been uneven. From a global perspective, adoption has been concentrated in high-income countries, while low- and middle-income economies continue to lag behind. This uneven diffusion underlines the possibility that GenAI may not only transform the nature of work, but also reinforce existing inequalities in access to AI technologies and its associated benefits (Demombynes et al., 2025). At the same time, recent research on the United States (Schendstok and Wertz, 2024) and Europe (Guarascio et al., 2025) highlights emerging within-country disparities in occupational exposure to GenAI. From a demographic point of view, women have been shown to be relatively more exposed to GenAI than men. Likewise, individuals in highly-educated, highly-paid, white-collar occupations have been shown to be disproportionately exposed to GenAI (Felten et al., 2023; Schendstok and Wertz, 2024). Regarding the role of age, (Schendstok and Wertz, 2024) have documented that

older workers (55+) are relatively more exposed to GenAI than prime-age workers (15-

54). Nonetheless, (Humlum and Vestergaard, 2025), which relied on individual-level data on GenAI usage based a representative sample of 18,000 Danish workers, found that younger individuals and men are more likely to use GenAI tools in their jobs.

One of the major challenges in measuring occupational and socio-demographic exposure to generative AI (GenAI) is the absence of large-scale population surveys capturing its use. Consequently, recent research has relied on indirect estimation methods to infer the penetration of AI across occupations. Broadly speaking, two main strategies have emerged: the skill-based and the task-based approaches.

The skill-based approach assesses the penetration of AI within an occupation by examining how closely abilities required for that occupation relates to specific AI applications, such as reading comprehension or language modeling. A pioneering effort in this direction was made by Felten et al. (2018, 2021), who linked AI applications to workplace abilities using the US Department of Labor's Occupational Information Network (O*NET) database. For each ability, they assigned a measure of its relatedness to AI applications, based on expert evaluations collected through the crowdsourcing platform MTurk. These ability-level measures were then aggregated into the AI Occupational Exposure (AIOE) index, constructed as a weighted average of the relevance of occupational abilities to AI applications.

By contrast, the task-based approach takes a different angle by assessing the extent to which the concrete tasks carried out in an occupation can be performed by AI systems. This frame- work, originally developed by (Acemoglu and Autor, 2011) and Acemoglu and Restrepo (2022), evaluates the susceptibility of occupations to automation based on their task composition rather than broader skill categories. More recently, Eloundou et al. (2024) extended this perspective to generative AI, combining large language models with human annotation to estimate the share of occupational tasks that could be directly affected by LLM capabilities. In a similar vein, Benítez and Parrado (2024) employed LLMs to provide expert-like assessments of the degree to which specific tasks could be substituted by AI systems, thereby demonstrating the potential of these models themselves as analytical tools for evaluating technological exposure.

Another relevant challenge refers to the distinction between complementarity and the re- placement of tasks by GenAI. As pointed by Pizzinelli et al. (2023), GenAI systems are rarely "pure substitutes" or "pure complements". Instead, they typically combine both roles in ways that blur the boundary between complementarity and substitution. While a wide range of tasks can be, in principle, performed by GenAI Felten et al. (2021, 2023), others re- main shielded by social, ethical, physical, or contextual constraints. This implies that some occupations, such as judges, may experience high exposure to GenAI but remain protected by legal frameworks, professional norms, and societal expectations, making complementarity significantly more likely than replacement. By contrary, occupations such as clerical workers are characterized by high exposure to GenAI and lower institutional shielding, leaving them more vulnerable to job displacement.

While numerous studies have proposed measures of occupational exposure to AI, considerably less attention has been devoted to understanding how GenAI diffuses across different demographic groups. To address this gap, this paper aims to construct novel indicators of GenAI exposure that capture the extent of GenAI penetration across occupational and socio-demographic subgroups in the US. Our approach combines occupation-specific exposure scores with micro-level population survey data, enabling a detailed assessment of how GenAI adoption varies by gender, age, education, and geography.

The U.S. provides an ideal context for this analysis for two main reasons. First, the availability of large-scale, publicly accessible population surveys containing detailed information on occupation, geographic location, and socio-demographic characteristics allows for a granular examination of population-level heterogeneity in GenAI exposure. Second, the integration of our occupation-level indicators with individual-level data is facilitated by standardized Standard Occupational Classification (SOC) codes, which provide a consistent framework for linking GenAI exposure to respondents' occupational profiles. Overall, the remainder of the paper is structured as follows. We begin by describing the data sources used in the analysis. We then outline the methodology employed to estimate the indicators of GenAI penetration across occupations and subgroups. Finally, we present the preliminary results and conclude by outlining directions for future research.

2 Data

This section presents the main data sources for the current study. First, we describe the main features of the Occupational Information Network (O*NET) database, which is used to retrieve the tasks associated to each occupation. Second, we display the main characteristics of the American Community Survey that is harnessed to obtain demographic information associated to each occupation.

To evaluate exposure to LLMs across occupations, we rely on the Occupational Information Network (O*NET) 30.0 database developed by the US Department of Labor. The O*NET data define and describe all the main professions within the United States and have been frequently employed to measure occupational work (e.g., see Felten et al. (2021) and Felten et al. (2023)). For the purpose of our analysis, we focus on over 17,000 distinct tasks that are associated to 711 occupations. O*NET provides both the importance and prevalence of each task within each occupation.

In order to evaluate the differences in terms of exposure to AI across different demographic groups, we rely on data from the American Community Survey (ACS). The ACS was launched in 2005 by the US Census Bureau to collect detailed information about the US population on a continuous basis, providing more frequent updates than the decennial census while not replacing it. Each year, it gathers extensive social, economic, housing, and demographic data from households across all 50 states, the District of Columbia, and Puerto Rico. Occupations in the ACS are coded according to the Standard Occupational Classification (SOC) system, which facilitates linkage with external datasets such as O*NET. For this analysis, we use data from the 2023 American

Community Survey (ACS), the most recent year available.

3 Methodology

3.1 Task-specific evaluation

Our evaluation of AI exposure is based on the *Meta Llama 3.3* multilingual large language model, an instruction generative model with 70 billion tuned parameters. Additionally, *Meta Llama 3.3* was pretrained on approximately 15 trillion tokens of data from publicly avail- able sources.[1] Following the procedure outlined by Schendstok and Wertz (2024), we asked Llama 3.3 to evaluate, on a 0–100 scale, the extent to which large language models (LLMs) can complement or augment each task, providing a brief justification for its rating. Collecting these textual explanations is a crucial part of the procedure, as they enhance the transparency of the evaluation process and enable the identification of potential inconsistencies, biases, or misunderstandings in the model's reasoning.

Specifically, we used the following prompt:

"You are an automation-risk rater. For each task, estimate the probability (on the scale 0–100) that the task could be complemented/augmented by Large Language Models (LLMs). Provide no more than two sentences explaining your reasoning for the score."

After evaluating the task-level complementarity, we compute the overall complementarity probability for each occupation (expressed by its SOC code) as a weighted average of the task scores, where weights indicate the relevance of each task within a given occupation.

[1]For information on the model, visit the web page https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct.

In contrast to earlier work by Felten et al. (2021) and Felten et al. (2023), which evaluated exposure to AI according to the relevance of skills within a specific occupation, we adopt a task-based evaluation framework. This approach relies on the premise that advances in generative AI are more likely to affect specific tasks within jobs rather than entire jobs. Similar task-based methodologies have been widely adopted both in the study of the automation of robotics (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2022; Autor and Dorn, 2013; Autor, 2013) and, more recently, generative AI (Benítez and Parrado, 2024; Schendstok and Wertz, 2024).

Table 1 reports the ten occupations with the highest and lowest complementarity probabilities as evaluated by the LLM *Meta Llama 3.3*. At the lower end, manual jobs, such as janitors, dishwashers and engine assembles, experience extremely low degrees of LLM complementarity. These jobs are characterized by physically embodied and context-dependent tasks that cannot be meaningfully augmented by LLMs and generally lie outside the realm of digital automation. In contrast, the high-ranking occupations include disc jockeys, clerical jobs and financial analysts, all of which record a complementary probability above 80%. These occupations rely on tasks, such as pattern recognition, data analysis, and information retrieval, which can be easily complemented by LLMs. In general, manual and routine service jobs appear to remain untouched by Generative AI augmentation, while data-intensive and low-skilled cognitive jobs exhibit high complementarity with LLMs.

Table 1: Ten Occupations with Lowest and Highest Complement Probability

Lowest Occupation	Complement Probability	Highest Occupation	Complement Probability
Janitors and Building Cleaners	4.909	Disc Jockeys, Except Radio	95.000
Maids and Housekeeping Cleaners	6.696	Billing and Posting Clerks	89.683
Dishwashers	9.654	Accounting and Auditing Clerks	89.118
Aircraft Mechanics and Service Technicians	11.481	Financial and Investment Analysts	86.993
Engine and Other Machine Assemblers	11.583	Insurance Sales Agents	86.480
Helpers, Construction Trades	12.381	Data Entry Keyers	85.820
Plasterers and Stucco Masons	14.550	Gambling Services Workers	85.159
Automotive Body and Related Repairers	15.363	Computer Operators and Programmers	83.361
Fence Erectors	15.553	Statistical Assistants	83.264
Small Engine Mechanics	15.757	Medical Transcriptionists	83.026

3.2 Computation of Indices of AI exposure across socio-demographic subgroups

To compute subgroup-specific indicators of exposure to LLMs, we matched our estimated occupation-level scores of LLM complementarity to micro-data from the 2023 ACS, using SOC codes. Military occupations were excluded as they are not covered in O*NET.

We restricted our analysis to employed individuals between ages 15 and 69. To capture het- erogeneity in exposure across socio-demographic, geographical and occupational groups[2], we constructed two indicators: the *LLM Penetration Index* and the *LLM High-exposure Index*.

The *LLM Penetration Index* for a subgroup g , denoted by I_g^{pen}, is the weighted average of the occupation-specific LLM scores across individuals in a given subgroup.

$$I_g^{pen} = \frac{\sum_{i=1}^{n_g} w_i \cdot \hat{P}_{o(i)}}{\sum_{i=1}^{n_g} w_i}$$

where $\hat{P}_{o(i)}$ is the LLM complementarity score of occupation o of individual i , w_i is the ACS survey weight for individual i , and n_g is the total number of respondents in subgroup g. This indicator can be interpreted as the probability of LLM complementarity for an

employed individual within the subgroup g.

The *LLM High-exposure Index* for a subgroup g , denoted by I_g^{high} measures the share of individuals in that subgroup g employed in occupations that are highly augmented by LLMs. We classify an occupation as high exposure if its complementarity score lies above the 80-th percentile of the overall distribution of scores, denoted by $\hat{\tau}_{0.8}$.

$$I_g^{high} = \frac{\sum_{i=1}^{n_g} w_i \cdot I(\hat{P}_{o(i)} > \hat{\tau}_{0.8})}{\sum_{i=1}^{n_g} w_i}$$

[2]Subgroups refer to categories of socio-demographic and geographical variables, including sex, age, education, wage, area of residence, and occupation.

Where I() is an indicator function that is equal to 1 if the occupation of individual i falls above the previous threshold and 0 otherwise. This indicator therefore represents the percentage of employed individuals within subgroup g working in occupations highly exposed to LLMs.

Figure 1 illustrates a summary of the proposed methodological approach. We rely on LLMs to evaluate the extent to which LLMs can complement tasks by requiring the model to pro- vide with a probability of complementarity for each prompted task. Afterwards, we aggregate the task-specific probabilities by occupation using a weighted average with weights mirroring the relevance of each task. We then match the aggregated occupation-specific scores to records from the 2023 ACS.

After this match, we can exploit the occupation-specific scores based on the LLM evaluations and the socio-demographic information from the ACS to compute our indicators of LLM penetration across population subgroups in the US.

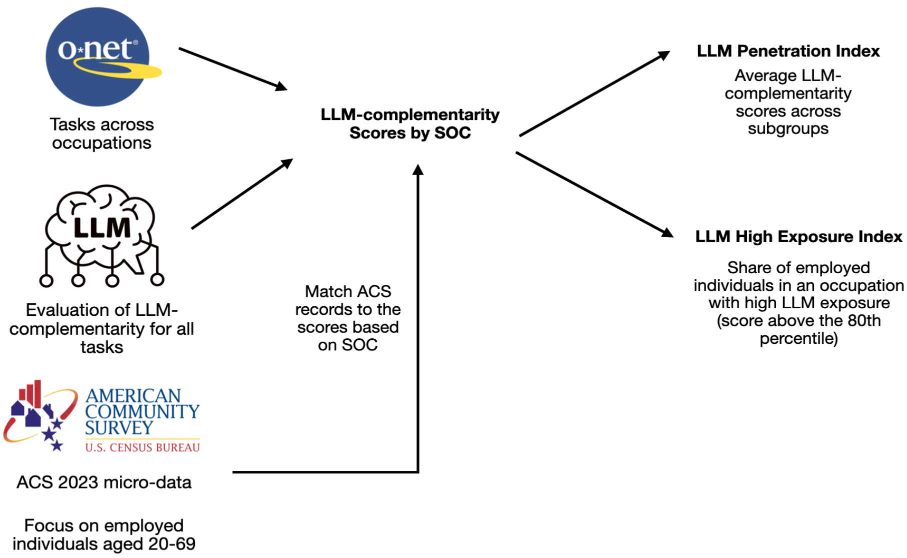

Figure 1: Illustration of the methodological framework for the computation of the LLM Penetration Index and the LLM High Exposure indexes.

4 Results

By computing the two indicators across various subgroups, we are able to investigate the geographical, occupational and demographic heterogeneities in LLM penetration. Table 2 reports the estimated values of both indicators for male and female workers by subgroup. Clear gender disparities emerge by education, wage, and occupation. Overall, women tend to work in jobs more likely to be complemented by LLMs. Among men, expo- sure to LLMs increases steadily with both education and wages, peaking among those with advanced degrees and higher earnings. For women, instead, exposure reaches its maximum among those with some college education and median incomes. Racial patterns also differ: Asian men show the highest exposure to LLMs, while among women, White workers are the most exposed. Nonetheless, across genders, Black workers and those belonging to other racial groups tend to experience significantly lower occupational exposure to LLMs, reflecting persistent racial disparities in the labor market and access to roles with high exposure to LLMs.

Figure 2 and Figure 3 illustrate the differences between men and women in terms of education and wages among workers in occupations highly exposed to AI, both for LLM penetration index and high exposure penetration index. The main differences between the two sexes are observed for the high exposure penetration index. The share of women in highly exposed occupations is consistently higher than that of men, except among those holding a master's degree. For individuals with doctoral degree, the proportion of men and women employed in highly exposed occupations is roughly similar. Regarding wages[3], women also account for a higher proportion of employment in highly exposed occupations compared to men across all wage levels. This difference can be explained by the higher tendency of women to be employed in office and administrative support occupations, which display a high degree of augmentation to LLMs but also relative lower wages in comparison to jobs in computer, engineering, and science occupations. Second, we do not observe striking differences by age in terms of probability of having a job complemented by LLMs.

[3]Wage refers to respondent's total pre-tax wage and salary income, that is money received as an employee, for the year before the survey. In this paper, we consider wage levels that are defined by percentiles of the distribution of salary income reported by employed respondent.

However, if we focus on the proportion of workers in jobs with high complementarity, we see that these proportions increase across ages in women, while among men it peaks in the age group 30-34 to decrease afterwards.

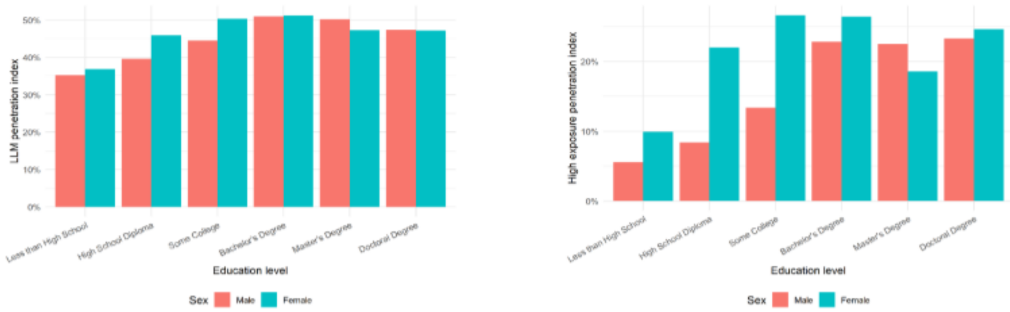

(a) LLM penetration index by education level and sex

(b) High exposure penetration index by education level and sex

Figure 2: AI exposure by education level and sex

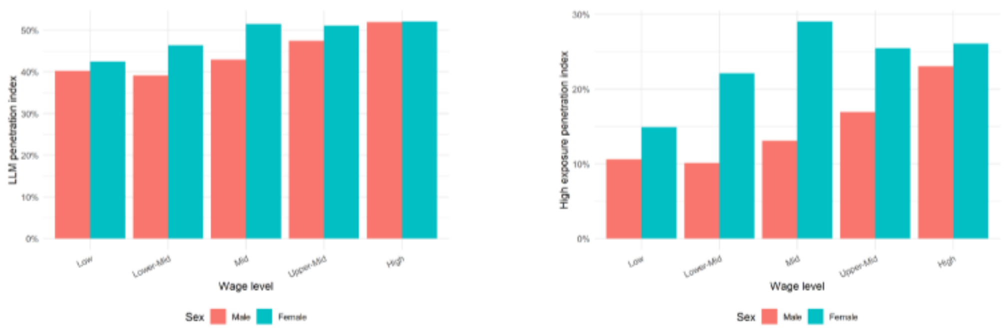

(a) LLM Penetration Index by wage level and sex

(b) LLM High-Exposure Penetration Index by wage level and sex

Figure 3: AI exposure by wage level and sex

Table 2: AI Indices across various subgroups by Sex

	Female		Male	
Wage Class	I_g^{pen}	I_g^{high}	I_g^{pen}	I_g^{high}
Low	43.26 (0.059)	16.44 (0.117)	40.76 (0.063)	11.04 (0.113)
Lower-Mid	47.98 (0.062)	24.84 (0.138)	39.88 (0.061)	10.78 (0.109)
Mid	51.70 (0.064)	28.55 (0.162)	43.98 (0.061)	13.84 (0.122)
Upper-Mid	51.15 (0.062)	25.30 (0.163)	48.14 (0.056)	17.63 (0.130)
High	52.09 (0.062)	26.09 (0.178)	52.07 (0.045)	23.24 (0.125)
Education				
Less than High School	36.22 (0.138)	10.64 (0.214)	35.54 (0.084)	5.57 (0.123)
High School Diploma	46.29 (0.072)	22.68 (0.150)	39.81 (0.050)	8.48 (0.085)
Some College	50.74 (0.068)	27.36 (0.162)	44.69 (0.059)	13.54 (0.121)
Bachelor's Degree	51.23 (0.043)	26.49 (0.112)	51.04 (0.043)	22.87 (0.110)
Master's Degree	47.33 (0.059)	18.57 (0.149)	50.19 (0.066)	22.52 (0.177)
Doctoral Degree	47.26 (0.167)	24.66 (0.451)	47.39 (0.165)	23.34 (0.420)
Age				
20-24	45.74 (0.091)	18.47 (0.209)	41.27 (0.091)	12.56 (0.176)
25-29	48.91 (0.087)	23.40 (0.210)	45.35 (0.084)	16.91 (0.179)
30-34	48.96 (0.083)	23.04 (0.197)	45.99 (0.079)	17.09 (0.168)
35-39	48.85 (0.085)	23.21 (0.199)	46.21 (0.077)	16.68 (0.166)
40-44	48.57 (0.086)	22.82 (0.197)	45.85 (0.077)	15.36 (0.160)
45-49	48.35 (0.091)	23.39 (0.206)	45.63 (0.079)	15.29 (0.166)
50-54	49.09 (0.088)	25.40 (0.203)	45.73 (0.077)	15.12 (0.159)
55-59	49.30 (0.090)	26.23 (0.203)	45.27 (0.079)	15.03 (0.161)
60-64	49.68 (0.096)	27.41 (0.217)	45.34 (0.084)	15.78 (0.171)
65-69	49.49 (0.133)	27.77 (0.297)	45.50 (0.110)	16.23 (0.221)
Race				
White	49.81 (0.034)	24.90 (0.083)	46.56 (0.032)	16.98 (0.069)
Asian	48.87 (0.100)	24.49 (0.241)	48.98 (0.090)	19.77 (0.218)
Black	47.55 (0.093)	22.11 (0.218)	43.36 (0.098)	13.01 (0.194)
Other	45.58 (0.073)	21.00 (0.150)	41.33 (0.060)	11.84 (0.110)
Occupation				
STEM Occupations	59.97 (0.071)	39.30 (0.352)	57.43 (0.043)	25.46 (0.192)
Construction	33.64 (0.311)	0.00 (0.000)	34.57 (0.057)	0.00 (0.000)
Education, Legal, Arts, Media	40.92 (0.038)	4.31 (0.076)	43.04 (0.059)	5.54 (0.123)
Farming, Fishing, Forestry	36.00 (0.328)	0.00 (0.000)	30.95 (0.102)	0.00 (0.000)
Healthcare	39.92 (0.052)	6.33 (0.113)	36.51 (0.108)	6.79 (0.214)
Installation	37.42 (0.393)	0.00 (0.000)	33.28 (0.075)	0.00 (0.000)
Management	58.48 (0.043)	32.37 (0.167)	57.11 (0.044)	32.35 (0.160)
Office Support	70.56 (0.045)	82.07 (0.153)	62.75 (0.101)	64.90 (0.327)
Production	35.72 (0.083)	0.46 (0.059)	35.83 (0.063)	1.43 (0.067)
Sales	54.02 (0.080)	16.16 (0.203)	57.37 (0.078)	26.78 (0.230)
Service	32.21 (0.056)	1.46 (0.047)	32.05 (0.079)	4.39 (0.092)
Transportation	33.59 (0.084)	0.00 (0.000)	37.49 (0.055)	0.00 (0.000)

Notes: Indicators are expressed in percentages; standard errors in parentheses. Subgroup indices are computed as weighted averages using ACS person weights, as defined in the Methods (I^{pen} and I^{high}).

Figure 4 reports the estimates values for the LLM Penetration Index and the LLM High-Exposure Penetration Index across (Public Use Microdata Areas) PUMAS within the continental US. In general, we observe a high heterogeneity both within and across states. The District of Columbia stands out as the area with the highest LLM penetration among employed individuals, with all its PUMAs showing over 25% of workers in highly exposed occupations. In contrast, New York State exhibits marked internal disparities: PUMAs in Manhattan record the highest concentration of workers in LLM-exposed occupations for the whole continental US, whereas in lower-income areas, such as the Bronx and Northwestern counties, the share drops below 15%. These spatial patterns underscore the uneven diffusion of LLM exposure, shaped by local labor market composition and regional economic inequalities.

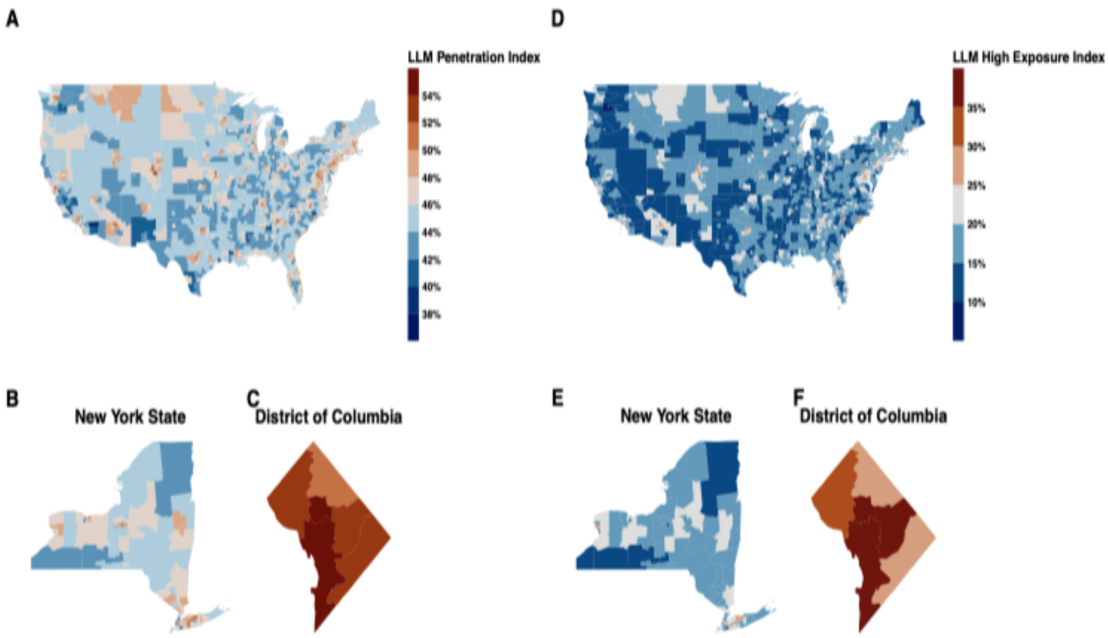

Figure 4: Geographical distribution of LLM penetration and high-exposure indices across Public Use Microdata Areas (PUMAs).

Note: Panel A illustrates the spatial distribution of the LLM penetration index across the continental United States. Panels B and C provide detailed views for the State of New York and the District of Columbia, respectively. Panel D presents the corresponding LLM high-exposure index at the national level, while Panels E and F show the respective distributions for State of New York and the District of Columbia.

5 Preliminary Conclusions

This paper introduces a simple methodological framework to indirectly estimate AI penetration across demographic subgroups by combining data generated from an LLM with large-scale population surveys. The preliminary analysis reveals notable gender disparities. While employed women are, on average, more exposed to LLMs in their occupations, they are also more vulnerable to potential job displacement. This pattern reflects the concentration of highly exposed women in clerical and mid-income roles. In contrast, highly exposed men are typically found among a smaller group of highly educated, high-earning workers who are less susceptible to job substitution. Furthermore, the results highlight persistent racial disparities: White and Asian workers are disproportionately employed in high-skill occupations, making them more likely to benefit from LLM complementarity, while other racial groups remain underrepresented in such roles. Broadly speaking, these patterns suggest that the diffusion of LLMs may reinforce existing labor market inequalities, amplifying returns for already advantaged groups while offering only limited benefits to workers from marginalized groups.

6 Next Steps

This paper is a work in progress and not yet ready for publication. By the time of the conference, we aim to advance it in several key directions.

First, we plan to refine our methodology by experimenting with alternative prompts when asking the LLM to assess task-level complementarity with LLMs. This extension will explicitly distinguish complementarity (augmentation) from replacement (automation). To date, our analysis has focused on complementarity; separating the two will allow us to identify which occupations and population subgroups are most likely to benefit from LLM augmentation and which face elevated displacement risk.

Second, the US context has been widely explored in the literature thank to the public availability of rich micro-data. For this reason, we aim to conduct a similar analysis within the context of Europe. To achieve this, we have submitted a request for access to the micro-level data from the European Labor Force Survey and are currently awaiting authorization to use these data.

Third, we will develop a Bayesian hierarchical framework to estimate LLM penetration across occupations and socio-demographic subgroups while explicitly accounting for measurement and sampling error. In this setup, occupation-level complementarity is treated as la tent, with LLM evaluations serving as informative (but noisy) priors that allow partial pooling across related occupations.

Fourth, we will validate our measures by comparing the LLM-based task evaluations with expert human assessments and by benchmarking the resulting complementarity scores across multiple LLMs.

Finally, we will estimate the number of workers exposed to LLMs in Europe and the United States, disaggregated by socio-demographic subgroup. We will also conduct cross-country comparisons to assess how LLM diffusion is reshaping the composition of the employed population across different institutional settings.

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016.

Autor, D. H. (2013). The “task approach” to labor markets: an overview. *Journal for Labour Market Research*, 46(3):185–199.

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.

Benítez, M. and Parrado, E. (2024). Mirror, mirror on the wall: Which jobs will ai replace after all?: A new index of occupational exposure.

Demombynes, G., Langbein, J., and Weber, M. (2025). *The Exposure of Workers to Artificial Intelligence in Low-and Middle-Income Countries*. World Bank.

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

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

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

Felten, E. W., Raj, M., and Seamans, R. (2023). Occupational Heterogeneity in Exposure to Generative AI. *SSRN Electronic Journal*.

Guarascio, D., Reljic, J., and Stöllinger, R. (2025). Diverging paths: Ai exposure and employment across European regions. *Structural Change and Economic Dynamics*, 73:11–24.

Humlum, A. and Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1):e2414972121.

Pizzinelli, C., Panton, A. J., Tavares, M. M. M., Cazzaniga, M., and Li, L. (2023). *Labor market exposure to AI: Cross-country differences and distributional implications*. International Monetary Fund.

Schendstok, M. and Wertz, S. S. (2024). Occupational exposure to artificial intelligence by geography and education. *Office of Economic Policy Working Paper*, 2.

Stanford Institute for Human-Centered Artificial Intelligence (2025). AI Index Report 2025. Technical report, Stanford University. Accessed: September 10, 2025.