femington

Int. J. Artif. Intell. Agent Syst. (2026) 1(1):1–15
Article ID: IJAIAS-2026-800

A Large-Scale Audit of Worker Preferences for AI Agent Automation and Augmentation

Eeshwar Pasula
eshwarpasula@gmail.com
Independent Researcher

Abstract

As AI agents are increasingly being integrated into work, little is known about the systematic knowledge among workers. We conduct a large audit of 1,500 workers across 104 US occupations, examining the preferred AI involvement in 844 occupational tasks. Our main contribution is the Human Agency Scale (HAS), which comprises 5 levels of human involvement between fully automated (H1) and essential human agency (H5). HAS seeks to circumvent the automation debate. The findings show that there is a 46.1% positive tendency to auto-automate tasks. According to research, this was largely to free up time for higher-value work (69.4%). Moreover, resistance has come from the creative sector as well. A significant mismatch is indicated by the comparison between what workers want and what 52 AI specialists think. We illustrate these mismatches as a desire-capability landscape consisting of four zones. The zones are Automation "Green Light", Automation "Red Light", R&D Opportunity, and Low Priority. Currently, investment is mostly concentrated in low-priority zones with unmet labour needs. H3 (Equal partnership) occupies 45.2% of the jobs, indicating an upcoming phase of collaboration. Interpersonal skills are becoming increasingly important. The implementation of AI technology and focused research should be guided by this worker-centred framework.

Keywords

• AI Agents • Human Agency Scale (HAS) • Human–AI Collaboration • Task Automation and Augmentation • Future of Work • Human-Centered AI

1. Introduction

The past few years have witnessed remarkable progress in foundation models, especially large language models. These advances have sparked growing interest in AI agents systems that can pursue goals independently, tap into various tools, and execute multiple steps in sequence. Unlike conventional models that simply respond to prompts, these agents can handle intricate workflows. Many researchers and practitioners now see them taking on meaningful roles across a wide range of professions. Early evidence already shows these technologies reshaping labor markets in tangible ways. For instance, Eloundou and colleagues [1] estimated that around four-fifths of U.S. workers could see language models affecting at least one tenth of their tasks, with nearly one fifth potentially experiencing impacts on over half of their work. Usage statistics from major AI platforms analyzed by Handa et al. [2] indicate that by early 2025, workers in more than a third of occupations were already using AI for at least a quarter of their tasks.

Despite these rapid developments, our understanding of how workers actually feel about AI agents remains surprisingly shallow. Most existing studies approach the question from a purely technical angle. They ask what AI can do, not what workers want AI to do [3]. This gap matters because the successful integration of AI into workplaces depends on more than raw capability. It depends on acceptance, trust, and alignment with human values [4]. Workers who feel threatened or ignored may resist even technically superior solutions. Conversely, systems that genuinely address worker needs could unlock substantial productivity gains while improving job satisfaction. Autor [5] argued that automation historically has not eliminated work but rather changed its nature, yet the current wave of AI agents may differ qualitatively.

The current literature suffers from several blind spots. Many assessments rely on high-level occupation analysis, treating each job as a monolithic entity [1]. This approach misses the considerable variation across tasks within the same occupation. A software engineer might welcome help with debugging but resist automation of architectural design. An accountant might want AI to handle data entry but not strategic planning. Occupation-level averages obscure these important distinctions. Additionally, most existing work emphasizes capital interests by focusing on tasks that generate the most profit, such as coding or financial analysis. Worker values rarely enter the equation [6, 7]. Some researchers have analyzed usage logs from chatbots, but this approach captures only what people already do, not what they might want or what could be possible [2].

To address these gaps, we developed a principled auditing framework centered on worker perspectives. Our approach operates at the task level rather than the occupation level, allowing us to capture the nuanced, contextual nature of real work [8]. We draw task definitions from the U.S. Department of Labor's O*NET database, which provides standardized descriptions of work activities across hundreds of occupations. This gives us a comprehensive view of the entire workforce that could potentially be affected by digital AI agents. Unlike prior work that treats automation as a binary outcome, we recognize a spectrum of possibilities between full automation and human augmentation, a distinction highlighted by Brynjolfsson in his discussion of the Turing trap [3].

A key contribution of our work is the Human Agency Scale, or HAS. This five-level framework quantifies how much human involvement different tasks require or deserve. The levels range from H1, where an AI agent handles everything independently, to H5, where continuous human participation remains essential. This scale complements existing automation taxonomies by centering human agency rather than machine capability [9, 10]. It provides a shared language for discussing when automation makes sense and when augmentation is more appropriate. For instance, tasks at H1 or H2 favor fully autonomous approaches, while tasks at H3 through H5 benefit from collaborative designs where humans and agents work together [11].

We collected data from two complementary sources. First, we surveyed 1,500 domain workers actively performing tasks across 104 occupations. Using an audio-enhanced survey interface, we captured their automation desires and preferred HAS levels. Second, we gathered assessments from 52 AI experts with experience in agent research and development. These experts rated current technical feasibility and realistic HAS levels for the same tasks. Combining these perspectives gives us a unique view of both social demand and technical reality. We call the resulting resource the Worker Outlook & Readiness Knowledge Bank, or WORKBank. To our knowledge, this represents the first large-scale audit of AI agent capabilities alongside worker preferences.

Our analysis yields several important findings. Domain workers express positive attitudes toward automation for 46.1% of tasks, even after considering potential job loss and reduced enjoyment. Their primary motivation is freeing up time for higher-value work. However, preferences vary dramatically across sectors, with creative fields showing the strongest resistance. When we map worker desires against expert capability assessments, we identify four distinct zones: Automation "Green Light" (high desire, high capability), Automation "Red Light" (high capability but low desire), R&D Opportunity (high desire but low capability), and Low Priority (low desire and low capability). Current investment patterns show concerning misalignment, with 41% of startup activities concentrated in low-priority or red-light zones.

The Human Agency Scale reveals diverse patterns across occupations. Nearly half of all occupations show H3 as the dominant worker-desired level, indicating strong potential for equal partnership between humans and agents. Workers generally prefer higher levels of human agency than experts deem technically necessary, suggesting potential friction as capabilities advance. Finally, our skill analysis indicates a possible shift in valued competencies. Traditional high-wage skills like data analysis and information processing appear less emphasized in high-agency tasks, while interpersonal and organizational skills gain importance. These early signals suggest how AI integration might reshape core human roles in the workplace.

2. Related Work

The aspiration to build autonomous digital agents traces back to the earliest days of artificial intelligence research. Pioneers like McCarthy and later Genesereth and Nilsson envisioned systems that could dynamically direct their own processes toward complex goals. Recent advances in foundation models, particularly large language models, have brought this vision closer to reality. Modern AI agents can plan sequences of actions, interface with external tools, and adapt to new information. Researchers have demonstrated their effectiveness across diverse domains, including software engineering, analytical writing, and customer support [12].

However, the technical pursuit of autonomous agents has often proceeded without adequate attention to the human context in which these systems will operate. Research in computer-supported cooperative work (CSCW) has long emphasized that the success of collaborative technologies depends not only on technical capability but also on alignment with work practices, social norms, and user values [13]. This perspective is particularly relevant for AI agents, which operate in contexts where human judgment, expertise, and agency remain central. The CSCW literature has documented how technologies that disregard worker perspectives often fail to achieve adoption, regardless of their technical sophistication [14].

Human-automation interaction (HAI) research offers established frameworks for understanding how people interact with automated systems. Parasuraman and Wickens [9] developed a comprehensive model of human-automation trust, identifying factors including reliability, predictability, and transparency that influence whether people accept or reject automation. Their work demonstrates that trust is not simply a function of technical performance but is shaped by expectations, experience, and the nature of the task [4]. Lee and See [4] extended this work by distinguishing between dispositional, situational, and learned trust, providing a nuanced framework that helps explain the variation we observe in worker automation desires across different occupational contexts.

A growing body of work recognizes the potential benefits of human-agent collaboration over full automation. Studies have shown that for certain tasks, joint human-agent performance can surpass what either party achieves alone, even when the agent could technically complete the task independently [10, 15]. This finding underscores the potential of AI to augment rather than simply replace human labor. Dellermann et al. [11] studied AI-assisted decision-making in professional service firms, finding that the effectiveness of human-AI collaboration depends on task characteristics, user expertise, and the design of the interaction interface. These findings motivate our Human Agency Scale, which distinguishes among different patterns of human-AI interaction rather than treating automation as a binary outcome.

The literature on algorithmic management has examined how AI-driven systems reshape authority, control, and worker autonomy in organizational settings. Kellogg et al. [6] documented how algorithmic management practices can fragment work, intensify monitoring, and reduce worker discretion, raising concerns about worker welfare and job quality. Wood et al. [7] showed how algorithmic management in gig work platforms can create experiences of 'algorithmic bullying' and reduce worker autonomy. These studies provide important context for understanding worker resistance to automation, as our findings show that concerns about job replacement and loss of autonomy are significant factors shaping automation preferences.

Labor economists have analyzed the employment effects of automation across multiple technological waves. Acemoglu and Restrepo [8] developed a task-based framework for understanding how automation displaces workers from some tasks while creating new tasks that complement human labor. Their framework emphasizes that the net employment effect depends on the balance between displacement and reinstatement effects. Autor [5] argued that automation historically has not eliminated work but rather changed its nature, though he cautioned that current AI capabilities may differ qualitatively from prior technological shifts. These economic perspectives inform our task-level approach, recognizing that automation affects specific activities rather than whole occupations.

The economic impacts of generative AI have attracted substantial research attention. Early studies examined machine learning and computer vision systems, documenting their effects on productivity and employment. More recent work has focused on large language models and their potential to reshape labor markets. Following the public release of ChatGPT, Eloundou et al. [1] estimated that approximately 80% of the U.S. workforce has at least some tasks exposed to language model capabilities. These estimates, while influential, did not incorporate worker desires or preferences, focusing instead on technical feasibility alone.

Field studies have begun to document actual AI usage in workplace settings. Research in customer support organizations has shown that AI-assisted chatbots can improve worker productivity, particularly for novice employees [12]. More recently, Handa et al. [2] analyzed real user data from commercial chatbots to identify which economic tasks people actually perform with AI. These usage-based approaches offer valuable ground truth about current adoption patterns. However, they cannot tell us about tasks workers might want to automate but cannot yet, nor can they reveal resistance that leads workers to avoid AI even when it is technically capable.

Brynjolfsson [3] introduced the concept of the "Turing trap," warning that evaluating AI systems solely by their ability to mimic human behavior may lead to the development of systems that replace rather than augment human capabilities. This perspective directly motivates our Human Agency Scale, which centers human agency rather than machine capability. By asking not just what AI can do but what workers want AI to do, we provide an alternative framing that prioritizes human values and preferences in the design and deployment of AI agents.

Together, these diverse threads of research inform our understanding of AI agents in the workplace. However, none provide the systematic, worker-centered audit that we undertake. Our contribution lies in integrating technical capability assessment with explicit measurement of worker desires and preferences, using a standardized framework that spans occupations and tasks.

3. Methodology

3.1 Audit Granularity and Scope

We designed our audit to focus on complex, multi-step tasks associated with specific occupations. The O*NET database provides standardized task statements that reflect actual job responsibilities. For example, "Credit Analysts: Analyze credit data and financial statements to determine the degree of risk involved in extending credit or lending money" represents a realistic work activity involving data analysis, risk evaluation, and decision-making. These tasks differ fundamentally from isolated low-level activities like "track goods or materials" or "translate information." The latter might appear in laboratory studies but do not capture the richness of actual work.

Task-level analysis offers substantial advantages over occupation-level analysis, aligning with task-based models of labor markets [8]. Within a single profession, different tasks can vary enormously in their suitability for automation. A marketing manager might spend some time compiling lists of product offerings, which could be easily automated, and other time developing creative campaign strategies, which probably should not be. Occupation-level averages would obscure this variation. By working at the task level, we can identify precisely which activities workers want help with and which they want to keep for themselves.

We limited our scope to computer-compatible tasks. This decision reflects the nature of digital AI agents, which operate through software rather than physical hardware. Drawing on established definitions from the agent literature, we defined AI agents as systems capable of autonomously performing tasks on behalf of a user by designing their own workflows and utilizing available software tools. Physical actions, such as assembling products or performing surgery, fall outside this definition. The O*NET database contains many tasks that require physical manipulation, and we excluded those from consideration.

3.2 The Human Agency Scale

A central innovation of our work is the Human Agency Scale, which provides a shared language for discussing automation versus augmentation. The scale has five levels, each describing a different pattern of human-agent interaction. At H1, the AI agent handles the task entirely on its own, with no human involvement required. At H2, the agent needs minimal human input to perform optimally, perhaps for initial setup or occasional oversight. H3 describes an equal partnership where human and agent work together and their combined performance exceeds what either could achieve alone [11]. At H4, the agent requires substantial human input to complete the task successfully. Finally, at H5, the agent cannot function without continuous human involvement throughout the process.

This scale differs from existing automation taxonomies in important ways. Traditional scales, such as those developed for automated vehicles, adopt an "AI-first" perspective [9]. They describe what the machine does and how much human monitoring is required. The Human Agency Scale flips this framing. It asks how much human involvement is appropriate or desirable, centering worker agency rather than machine capability. Higher HAS levels are not inherently better than lower levels. Different levels suit different tasks and different contexts. Fully autonomous agents are appropriate for H1 scenarios, while agents designed for H3 scenarios must support meaningful two-way communication and coordination.

The distinction between automation and augmentation maps naturally onto the HAS. Tasks at H1 and H2 favor automation approaches, where the goal is to remove human effort entirely. Tasks at H3 through H5 favor augmentation approaches, where the goal is to enhance human capabilities while preserving meaningful agency [3]. Understanding where a task falls on this spectrum helps both workers adapting their skills and developers building context-appropriate systems.

3.3 Worker-Centric Data Collection

Our auditing framework centers on the needs and perspectives of workers. To support domain workers in providing well-calibrated feedback, we developed an audio-enhanced survey system. This approach allows participants to share their experiences verbally rather than typing everything out. Speaking tends to be more natural and less burdensome than writing, especially for workers who spend their days doing things other than composing text. The audio interface also encourages more thoughtful reflection before participants move to quantitative ratings.

The survey begins with a mini-interview section containing five open-ended questions. Participants describe what they do for work, what tasks occupy their time, what tools and software they use, how they complete their most time-consuming tasks, and how they envision using AI in their daily work. These audio responses serve multiple purposes. They help workers contextualize subsequent ratings within their actual experience. They also provide rich qualitative data that illuminates the reasons behind quantitative patterns.

For each task associated with a participant's occupation, we collect several types of ratings. First, we measure automation desire using a five-point Likert scale. The question asks: "If an AI system can do this task for you completely, how much do you want it to do it for you?" Responses range from "Not at all" to "Entirely." Second, we measure desired Human Agency Scale level using an analogous five-point scale anchored to the HAS definitions. Third, we collect ratings on task characteristics including physical action requirements, uncertainty and decision stakes, domain expertise needs, and interpersonal communication demands.

Quality control happens through several mechanisms. Participants must confirm they have performed each task before rating it. We also ask about their experience level and how much time they typically spend on the task. These filters ensure that ratings come from genuine experience rather than speculation. Before rating automation desire, participants consider job security concerns and task enjoyment, factors known to influence automation attitudes [4, 9]. Before rating HAS level, participants consider task characteristics drawn from the literature on human-automation interaction [10, 15].

3.4 Expert Assessments

Worker insights, while vital, offer a partial lens on integration possibilities. While they bring deep knowledge of task execution, their familiarity with the capabilities and limitations of present-day AI systems may be incomplete. For example, workers might seek the automation of a task the current technology is not yet equipped to perform consistently, or oppose the automation of a task that is easily automatable due to their engagement or concern about future roles. To align worker wishes with the possibilities afforded by current technologies, we further augment the former with perspectives from our AI experts.

Our subject pool of 52 AI experts comprises PhD researchers and experienced industry practitioners knowledgeable in agent research and deployment. Participants met one or more of the following conditions: (1) current PhD students focusing on language models, large language models, or AI agents; (2) individuals with computer science doctorates with relevant skills; or (3) industry professionals with real-world experience constructing agent systems. They hail from institutions such as Google, MIT, xAI, Stanford, and others.

We assigned each of our tasks to two or more experts for independent evaluation. They rated each task in terms of its automation ability on the identical 5-point scale as the workers and on the feasible level of Human Agency Scales possible today with currently available technologies (the 'feasible HAS' level). We established protocols calling for extra annotations whenever two annotators differed on a rating by more than 1 point. The Krippendorff's alpha values for the two ratings - 0.539 for automation capability and 0.511 for HAS - were characteristic of moderate inter-annotator reliability for the inherently complex nature of this task, consistent with similar studies in subjective annotation tasks.

To understand the sources of disagreement, we conducted a qualitative analysis of tasks where expert ratings diverged by two or more points on the five-point scale (constituting 12.4% of all task-expert pairs). The primary source of disagreement (52% of divergent cases) concerned the handling of uncertainty and edge cases. Some experts focused on the current capability of state-of-the-art systems for typical scenarios, while others emphasized performance in atypical or high-stakes situations where system failures could have serious consequences. This divergence was most pronounced for tasks involving financial decision-making, health-related assessments, and legal judgments.

A second source of disagreement (31% of divergent cases) related to the assessment of integration complexity specifically, whether a task required interacting with legacy systems, proprietary software, or physical infrastructure that would complicate agent deployment. Experts with industry experience tended to rate these tasks as less feasible than academic researchers, reflecting different exposure to real-world technical constraints.

A third source of disagreement (17% of divergent cases) centered on the appropriate HAS level for tasks that could technically be automated but where experts disagreed about whether automation would be desirable given the social or organizational context. This pattern was most common in tasks involving supervision, mentoring, or client relationship management, where technical feasibility and normative appropriateness diverged.

These sources of disagreement highlight important nuances in expert assessments. The moderate agreement scores reflect genuine uncertainty in the expert community about both technical capabilities and appropriate levels of human oversight, particularly for tasks at the boundaries of current AI capabilities. We addressed this uncertainty through our requirement that tasks with high inter-expert variance receive additional reviews, ensuring that our final capability estimates reflect a balanced assessment across different expert perspectives.

3.5 WORKBank Construction

We sourced tasks from the O*NET database version 29.2. The complete database contains 18,796 task statements spanning 923 occupations. Using language model-based filtering, we identified tasks that are computer-compatible and performed at least monthly. We used GPT-4 (version GPT-4-0125-preview) with a temperature setting of 0.0 to ensure deterministic outputs. The filtering process involved two stages. First, we asked the model to identify whether the occupation primarily involves working with computers by presenting the prompt: "Given the following occupation title and description, does this occupation primarily involve working with computers and digital tools, or does it primarily involve physical manipulation, in-person interaction, or outdoor work? Respond with 'computer-based' or 'non-computer-based' and provide a brief justification." Second, for occupations classified as computer-based, we filtered individual tasks using the prompt: "For the following task statement, can this task be completed entirely on a computer or digital device without requiring physical manipulation of objects or in-person interactions? Respond with 'yes' or 'no' and provide a brief justification." Tasks were retained only when both the occupation was classified as computer-based and the task was classified as completable on a computer.

This filtering process retained 2,131 tasks across 287 occupations. To audit the filtering process and assess potential biases, we conducted manual review of a stratified random sample of 500 tasks (approximately 25% of the excluded set and 25% of the retained set). Two researchers independently reviewed each sampled task, applying the same inclusion criteria. Inter-rater agreement for the manual audit was 94.2%, indicating high consistency with the LLM-based filtering decisions. Disagreements were resolved through discussion. The manual review identified 47 tasks that were incorrectly excluded and 23 tasks that were incorrectly retained. All tasks with manual classification disagreements were reviewed and corrected based on consensus judgment. This audit procedure ensures that the filtering process did not introduce systematic biases that would substantially affect the composition of the final dataset.

Data was collected between January 2025 and May 2025. We recruited for our survey through Prolific, Upwork, and targeted LinkedIn efforts. Our Prolific survey featured prescreening options based on employment, country of nationality, and fluency in a language (including computer).

Upwork sourcing focused on workers located in the United States.

LinkedIn efforts consisted of direct contact and flyers. Collectively we recruited a total of 1,678 participants resulting in 7,016 task ratings. After further limiting this sample to include only workers whose job title had at least ten individual participants (N = 1,500 worker/job titles across 104 distinct jobs). For each task we determined the average score that individual workers rated an occupation's desire for automation and its desired HAS score. We assessed representativeness based on U.S. Workforce statistics for the different sectors derived from the Bureau of Labor Statistics, compared to our database's sector distribution. These distributions show a robust comparison, but reflect certain sectoral over- and underrepresentation.

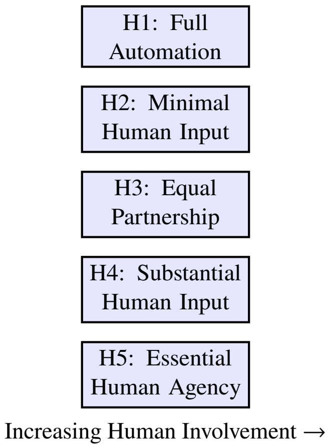

Figure 1: The Human Agency Scale (HAS) provides five levels describing different patterns of human-agent interaction. Higher levels indicate greater required human involvement, though no level is inherently superior.

4. Results

4.1 Worker Automation Desires

We first examined where workers want AI agent automation and where they resist it. Figure 1 (the Human Agency Scale) provides the framework for understanding the different levels of human involvement that workers may desire, while our analysis below describes the distribution of automation desire scores across tasks. Among the 844 tasks analyzed, 46.1% received positive automation desire ratings, with substantial variation across occupational sectors. This finding holds even though our survey explicitly prompted participants to consider job loss concerns and potential reductions in work enjoyment. The distribution shows mixed attitudes overall, with 7.11% of tasks receiving strong positive desire (score of 4 or higher) and 6.16% receiving strong negative desire (score of 2 or lower).

Table 1: Top 10 tasks with the highest worker preference for automation

Task	Automation Desire
Tax Preparers: Schedule client appointments	5.00
Public Safety Telecommunications: Maintain records of emergency calls	4.67
Payroll and Timekeeping Clerks: Process pay adjustments	4.60
Desktop Publishers: Convert files for print or web use	4.50
Online Merchants: Maintain customer account databases	4.50
Quality Control Systems Managers: Track defects and test results	4.50
Statisticians: Present analysis results using graphs and charts	4.50
Computer User Support Specialists: Maintain records of transactions and issues	4.50
Online Merchants: Calculate revenue, sales, and expenses	4.40
Data Entry Keyers: File completed documents	4.33

To investigate drivers of the desire for automation, we examined the expressed reasons for positive sentiments. By far the most frequently selected reason was that automation of a task would result in "saving my time for higher-value work" (selected in 69.38% of all expressions of positive automation desire), indicating that many view AI as a delegate rather than a replacement. Task repetitiveness (46.6% of expressions), opportunity for quality improvement (46.6% of expressions) and stressfulness (25.5% of expressions) followed, which indicates workers generally seek assistance for cumbersome tasks that are less a product of unique human skills and judgment [4, 9].

When we analyzed these in terms of broad occupational classes, many of our highest desire for automation occupations represented only a tiny fraction of the users of a large commercial chatbot.

For example, our ten highest desire-occupation categories represent only 1.26% of all usage instances for one major chatbot [2]. The pattern of who desires automation is dramatically different from current AI use patterns. This demonstrates the value of asking worker expectations rather than relying on observed behavior. Automation avoidance We addressed avoidance of automation using topic models applied to our open-ended audio responses about AI in our survey respondents (28% had expressed some fears or concerns about using AI for their day-to-day tasks) The most prevalent concerns were in relation to a lack of trust in the AI system, whether they can rely on its results and whether it performs competently (45.0% of worried workers), consistent with prior work on human-automation trust [4].

The second highest concern related to job displacement (23.0% of worried workers), followed by a concern that the AI system lacks any human capabilities (16.3% of worried workers), reflecting themes from the algorithmic management literature [6, 7].

In these responses about how a lack of human capabilities manifests itself, workers spoke about the desire not to lose the "human touch," the erosion of creative autonomy, and the maintenance of meaningful agency when AI is integrated into the work. It is in creative sectors where this becomes most apparent. In the Arts, Design and Media domain, for example, only 17.1% of its tasks expressed high automation desire.

The audio responses in the sector provide nuanced explanations: for example, one art director commented that "I want it to be used for seamlessly maximizing workflow.

4.2 Desire-Capability Landscape

Though workers' preferences are extremely informative about deployment that leads to social good, we can only achieve social good by bringing those preferences in sync with technical capabilities. We paired workers' self-reported preferences for task automation with experts' ratings of how readily available automation would be for the task. When we graph desire on the y-axis and capability on the x-axis, the desired automation–technological capability plot has four zones based on being either high or low on each dimension: the "Green Light" Zone (high on both desirability and capability); the "Red Light" Zone (low on desirability, high on capability); the R&D Opportunity Zone (high on desire, low on capability); and the "Low Priority" Zone (low on both).

The tasks in WORKBank are pretty evenly distributed across these four zones. Desire and capability are not correlated, with a Spearman rank of 0.17, indicating that workers are often expressing a desire for what's technologically unattainable currently [1]. Overall, people had lower expressed desires for tasks they liked or feared being displaced from, which is consistent with past research [4, 9].

We then mapped Y Combinator companies to tasks based on technology available, or predicted by available text. The matches were pretty evenly dispersed among the four groups and clustered in the software and business intelligence/analytics domain.

Table 2: Distribution of Y Combinator companies across desire-capability zones

Zone	Percentage of Mapped Companies
Automation Green Light	29.5%
Automation Red Light	18.3%
R&D Opportunity	29.5%
Low Priority	22.7%

Alarmingly, 41.0% of Y Combinator companies Map into the "Low Priority Zone" or "Automation Red Light Zone" it means that serious money is being spent for work where people don't need or desire automation, or it would solve trivial problems. Instead, some exciting opportunities for automation within the "Green Light" and "Opportunity Zone" are significantly underfunded. The pattern implies a significant mismatch between investment trends and worker demands.

To complement this study, I performed a similar study for AI Agent research Papers On The Internet I conducted a similar study of AI agent research papers on arXiv.

The research patterns here looked much better, with a higher cluster in the R&D "Opportunity Zone." Nonetheless the field was mostly dominated by computer and engineering fields. Three most-popular task category topics were all in the computer and information related category: (a) applying of new technological knowledge, (b) analysis to determine problems solved by computers and developing software programs and, (c) analysis to determine problems solved by computers and revising programs. Therefore, work need to begin by conducting research in more occupational field to aid in task within the R&D Opportunity zone and making certain to address the highest demand work.

4.3 Human Agency Scale Spectrum

While AI agents could also help in the workplace beyond automation, our goal is to identify which task augmentations are most likely. To analyze opportunities for such augmentations, we visualized worker-desired and expert-assessed feasible HAS levels over tasks across different occupations. Out of the 844 tasks, we found that for 26.9% of them, AI experts agreed with the worker on the optimal HAS level.

We observe workers preferring a greater role of human agents than the one identified as technically feasible by experts, where 47.5% of the tasks land in the lower triangle region of the comparison plot.

We measured this difference using the Jensen-Shannon Distance between the worker-desired and expert-assessed HAS level distributions at the occupational level. We notice a great disagreement between them at the lower end of the HAS axis, where the task can be fully automated but humans expect higher engagement in such activities. This implies a point of possible conflicts in the workplace as AI is adopted further.

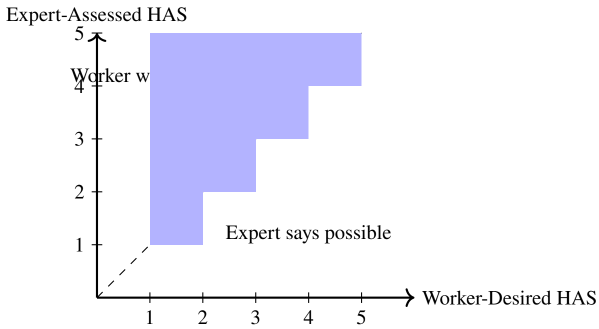

Figure 2: Comparison of worker-desired and expert-assessed HAS levels across 844 tasks. The x-axis shows expert-assessed HAS levels (H1-H5), and the y-axis shows worker-desired HAS levels (H1-H5). Points above the diagonal indicate tasks where workers desire more human agency than experts deem technically necessary, while points below the diagonal indicate tasks where experts assess higher capability than workers desire. Point size reflects the number of tasks at each combination. The figure reveals systematic divergence between worker preferences and expert technical assessments, particularly at lower HAS levels where workers often prefer more human involvement than experts consider necessary.

Exploring which levels of HA are dominant across different occupations produces striking observations. Several occupation's are characterized by an inverted U-shape for both worker-desired and expert-determined levels of HAS. Indeed, H3 is the most dominant worker-desired level for almost half of the analyzed occupations (47 out of 104).

This points to immense possibility for human-AI collaboration in many areas [11].

We do not want AI to run everything, nor do we want AI to be irrelevant; we want to work with it. Some occupations lie in extremes regarding the scale. For 16 out of 104, a significant amount of jobs is primarily HA 1 in our experts' ratings (based on our estimates of what level of HA currently is or soon will be available). Computer Programmers, Proofreaders and copy-makers and Travel agents are examples.

Some occupations may lie in completely opposite regions within same industry.

For instance while computer programmer is considered mostly HA 1, Information Technology Project Manager falls into HA 4 among our experts. Lastly, very few are characterized by dominantly H5: In the worker-dominated settings there are only 1 job that wants us primarily H5, the Editors category. Regarding the expert ratings, only two professions make it to H5: Mathematicians and Aerospace Engineers.

Worker evaluations emphasize face to face communication as its defining characteristic, experts point out both face to face communication and deep expertise on some topics. It may be that it requires of us interaction and specialized skills that are hard to simulate.

5. Discussion

5.1 Shifts in Core Human Skills

We find that HAS levels vary quite significantly across occupations. Some types of work performed by humans are likely much more vulnerable to automation, whereas others might see more augmentation. These findings can inform the kinds of training and education that might be necessary in the future. We aligned the various skills to particular types of work using generalized work activities (GWAs), the human skills for every task. For example, Financial Managers' task: approve, reject or coordinate the approval or rejection of credit lines. Correlates with making decisions and solving problems; managing and motivating subordinates.

Next, we used estimates of experts' average human agency level for the GWA to quantify the level of human involvement required in each. We calculated wages for each of the GWAs as well using the US Bureau of Labour Statistics as a proxy for the current economic value of each skill set. Comparing rankings by wage versus by required human agency reveals three patterns of particular importance. First, Information processing skills are declining. Skills related to learning and being updated with information; analysing data; all appear quite frequently among tasks within those jobs characterized by current higher wages. However, they appear far less often among high-agency skills. As the sophistication of AI agents increases, we might see the economic penalty on information-processing decline, consistent with predictions from task-based models [5, 8]. Those jobs that are currently high wage for their ability to process information may have to add complementary skills to maintain wages. Second, an increased demand for interpersonal and organisational skills. Those skills relating to: management/ supervision; organising, planning, or strategizing; negotiating; as well as coordinating with and guiding other people, and managing and utilizing resources more and more frequently accompany the highest level of human agency; and importantly, none of these are among those for which currently wages is higher based. This discrepancy between wages and future demand presents both challenges and opportunities. Organizations, and the workers who can anticipate these changes early, may benefit by emphasizing these skills earlier than the economywide consensus catches up [6]. Third, skills requiring higher human agency are diverse. Those skills that currently are ranked highest in required human agency consist of both those related to human and organizational functions as well as skills like analysing data and judgment. No type of skill alone is a"safe haven" of necessity to human skills-it comes from the range of uncertainty, decision-making, coordination, and adaptation to change.

We expect that continuing work identifying these evolving human skill needs may highlight how new types of agents interact with existing job roles and also shed light on the types of new tasks that will emerge that require humans.

These initial signs point to several changes in the core competencies that individuals should develop to remain valuable in the workplace and offer guidance to organization leaders considering workforce development.

5.2 Implications for Responsible AI Deployment

Our findings have several implications for organizations deploying AI agents. First, worker desires matter and vary systematically across tasks. Organizations should not assume that technical feasibility alone determines appropriate automation targets. A task that could be automated might still be one that workers want to keep, either because they enjoy it or because it provides a sense of purpose. Imposing automation on such tasks risks resistance, reduced morale, and possibly even sabotage [6, 7].

Second, the Human Agency Scale provides a practical tool for deployment planning. Organizations can map their workflows onto the HAS framework, identifying which tasks fall into which zones. Tasks in the Green Light zone offer immediate opportunities for automation with likely worker support. Tasks in the Red Light zone warrant careful change management or reconsideration of automation plans. Tasks in the R&D Opportunity zone might be worth tracking as technology improves, but current attempts at automation may disappoint.

Third, the strong preference for H3 partnership in many occupations suggests that organizations should invest in collaborative systems rather than pursuing full autonomy. Systems designed to work alongside humans, supporting their judgment rather than replacing it, align better with worker desires in most domains [10, 11]. This finding contradicts the common narrative that AI development inevitably trends toward full automation. In many workplace contexts, the optimal system may be one that empowers workers rather than eliminating them [3].

For AI developers, our findings highlight under-addressed opportunities. The R&D Opportunity zone contains tasks workers desperately want help with but current technology cannot handle. Many of these tasks fall outside traditional computer science domains. Expanding research efforts into areas like creative work, complex coordination, and nuanced communication could yield both commercial value and social benefit. The concentration of current investment in low-priority and red-light zones suggests that market signals alone may not guide development toward socially beneficial directions.

5.3 Limitations and Future Directions

Worker Attitudes and Technological Readiness: Limitations First, our quantitative work takes its cues from current, defined occupations within the ONET taxonomy. These ONET tasks do not encompass the range of new tasks that could be created as we move to widespread agent integration with existing tasks [8]. Future work can probe worker transcripts in natural language to find other new job tasks not in ONET.

Second, despite our efforts to direct participants to consider task desirability and the possibility of losing their job to a machine, it is not possible for these workers to have a complete understanding of the capabilities and limitations of agent technology.

Thus, responses might shift once the workers gain greater experience with these capabilities. We partially addressed this issue through ensuring each occupation received a minimum of 10 participant responses, and through inclusion of expert ratings. However, a gap between perceived and real capabilities always exists for any survey methodology. Third, when workers feel that they might be displaced or surveilled by AI agents, they may not always be forthcoming with critical feedback, for any number of good reasons [7].

Indeed, this very fear or concern of worker resistance must be seen as a critical input in a human-centered approach.

It guides us to design for more worker participation. Our approach gives the workers the opportunity to shape rather than react to the workplace technologies being developed. Finally, WORKBANK currently has coverage for just 104 computer-based jobs out of 287 computer occupations identified in ONET, with the last of data collection ending in May, 2025 in order to ensure the temporal integrity of the dataset.

Our subset of 104 occupations was the one for which we obtained adequate coverage, having a minimum of 10 responses per occupation. This database represents substantial and representative coverage, though not all workplaces may be covered. The data in WORKBANK is reflective of generic generative AI and agentic tools in use in early 2025.

Future advances in technology may make a broader array of jobs feasible for agents to assist or undertake, and future editions of this audit would be well positioned to monitor such changes.

6. Conclusion

AI agent development represents progress in possibilities for job roles that can radically alter the workplace. In this paper, we reported the first-ever audit of both worker interest and technological feasibility of AI agents for tasks, focusing on automation and augmentation in the work. From surveys completed by 1,500 domain workers and interviews conducted with 52 AI domain experts between January and May 2025, we built the WORKBank database of worker preferences and expert knowledge of these capabilities.

Overall, the domain workers favoured automation of AI agents for tasks where those agents handled repetitive, low-value tasks and could free up workers' time for more meaningful work [5].

We integrated worker and expert perspectives to create the automation desire-capability map, a guide toward prioritising AI agent research, investment, and deployment; based on worker preferences and existing tech capabilities, there are three zones that highlight opportunities to maximize positive impacts of AI agent technologies, including two highly prioritized ones (green and light green zones). The existing development, however, showed high investment in other zones, including the two on the "red light" section of the map that are currently least desirable to workers for automation, and not feasible with existing capabilities, such that significant work is needed for automation opportunities. Using the Human Agency Scale, the development of these AI agents was found to not always follow a path of increased automation toward eventual job elimination [3]. In addition to exploring whether jobs are fully automated or just aided, we discovered that within job role areas there are trends toward different types of technology-human collaboration through AI agents [11].

Many jobs expressed an inverted U-shape toward high levels of interaction, where there was an optimal amount of AI engagement, where neither maximum nor minimum assistance was desired.

Our work focused on using these results for the betterment of jobs: asking what we should do as the workplace is evolving, rather than what we can do?

These results suggest that worker agency may be preserved. In the longer term, in workforce development and education, these results have indications for how jobs might shift toward more human, relational and organizational skills from the more information processing jobs, and workers in these roles are most prepared for these shifts [8]. Using our worker-centred approach to analyse workers' desires to engage with AI and the current technological capabilities of AI to do this work is one such approach that allows for a continued ongoing evaluation of technology use.

References

[1] T. Eloundou, S. Manning, P. Mishkin, and D. Rock. GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint*, 2023.

[2] K. Handa et al. Which economic tasks are performed with AI? evidence from millions of Claude conversations. *arXiv preprint*, 2025.

[3] E. Brynjolfsson. The Turing trap: The promise & peril of human-like artificial intelligence. *Daedalus*, 151(2):272–287, 2022.

[4] J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1):50–80, 2004.

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

[6] K. C. Kellogg, M. A. Valentine, and A. Christin. Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1):366–410, 2020.

[7] A. J. Wood, M. Graham, V. Lehdonvirta, and I. Hjorth. Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1):56–75, 2019.

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

[9] R. Parasuraman and C. D. Wickens. Humans: Still vital after all these years of automation. *Human Factors*, 50(3):511–520, 2008.

[10] M. L. Cummings. Automation and accountability in decision support system interface design. *Journal of Technology Studies*, 32(1):23–31, 2006.

[11] D. Dellermann, P. Ebel, M. Söllner, and J. M. Leimeister. Hybrid intelligence. *Business & Information Systems Engineering*, 61(5):637–643, 2019.

[12] E. Brynjolfsson, D. Li, and L. R. Raymond. Generative AI at work. Working paper, National Bureau of Economic Research, 2023.

[13] M. S. Ackerman. The intellectual challenge of CSCW: The gap between social requirements and technical feasibility. *Human-Computer Interaction*, 15(2-3):179–203, 2000.

[14] G. M. Olson and J. S. Olson. Distance matters. *Human-Computer Interaction*, 15(2-3):139–178, 2000.

[15] M. T. Dzindolet, S. A. Peterson, R. A. Pomranky, L. G. Pierce, and H. P. Beck. The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6):697–718, 2003.