

Human–AI Joint Decision-Making in Flexible Labor Markets: A Large Language Model Framework for Goal Commitment and Performance Enhancement

Ivan Carpenter

Department of Computer Science, University of North Texas, Denton, TX, USA.
carpenterivan@unt.edu

Sean M. Powell

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
sean.work@unh.edu

Abstract

The rapid expansion of flexible labor markets, characterized by gig work, platform employment, and contingent contracting, has created new challenges for worker motivation, goal setting, and sustained performance. Traditional approaches to goal commitment rely on static, top-down management structures that are ill-suited to the fluid and decentralized nature of modern work. This paper proposes a framework for human–AI joint decision-making that leverages large language models to support dynamic goal formation, commitment tracking, and performance enhancement. The framework integrates LLM-based conversational agents with socio-technical system design principles to enable adaptive goal negotiation, real-time feedback, and personalized incentive alignment. We examine the architectural trade-offs involved in deploying such systems, including the balance between autonomy and control, the handling of goal conflict in multi-stakeholder environments, and the need for transparent governance mechanisms. The paper also addresses critical issues of fairness, algorithmic bias, and the sustainability of AI-mediated labor platforms. Through a cross-domain analysis of case studies from crowd work, on-demand delivery, and freelance consulting, we illustrate how LLM frameworks can improve worker engagement and productivity without undermining autonomy. We further discuss the infrastructural requirements for scalable deployment, the robustness of LLM-generated goal recommendations under uncertainty, and the policy implications for labor regulation and worker rights. The proposed framework contributes a systems-level perspective to the growing literature on human–AI collaboration, offering both theoretical foundations and practical design guidelines for next-generation labor market platforms.

Keywords

human–AI collaboration, flexible labor markets, large language models, goal commitment, performance enhancement, socio-technical systems, algorithmic governance, fairness.

1. Introduction

Flexible labor markets have become a defining feature of the twenty-first-century economy, reshaping how work is organized, compensated, and evaluated. Platforms such as Uber, Upwork, and Amazon Mechanical Turk enable millions of workers to engage in short-term, task-based employment with unprecedented flexibility in scheduling and task selection. However, this flexibility comes at a cost: workers frequently experience goal ambiguity,

reduced social support, and limited access to career development resources, all of which undermine long-term motivation and performance. Traditional human resource management practices, which rely on hierarchical goal setting and periodic performance reviews, are ill-equipped to handle the transient and heterogeneous nature of gig work. In response, researchers and practitioners have begun exploring artificial intelligence systems that can assist workers in setting and pursuing goals in real time, adapting to changing circumstances and individual preferences.

Large language models (LLMs) represent a particularly promising technology for this domain due to their ability to generate context-sensitive, natural language interactions. Unlike rule-based decision support systems, LLMs can engage in open-ended dialogue, interpret worker sentiment, and produce personalized recommendations that evolve with the worker's history and current task environment. This capability opens the door to a new paradigm of human–AI joint decision-making, where both parties collaboratively define goals, monitor progress, and adjust strategies. The central thesis of this paper is that such joint decision-making, if properly designed within a robust socio-technical architecture, can significantly enhance goal commitment and performance in flexible labor markets without sacrificing worker autonomy or fairness.

We develop this argument by first reviewing the relevant literature on goal setting theory, AI-assisted decision making, and the unique characteristics of flexible labor markets. Next, we propose a framework that specifies the roles of LLM agents, human workers, and platform governance structures in a joint decision-making process. We then analyze the structural trade-offs inherent in this framework, such as the tension between algorithmic nudging and worker self-determination, and the challenges of aligning individual goals with platform-level objectives. Subsequently, we examine deployment considerations including infrastructure requirements, robustness to noise and manipulation, and sustainability over long time horizons. The paper concludes with a discussion of policy implications and directions for future research.

2. Background and Related Work

Goal setting theory, as originally formulated by Locke and Latham [1] in organizational psychology, posits that specific and challenging goals, when accepted and committed to by individuals, lead to higher performance than vague or easy goals. The theory has been extensively validated across diverse work settings, yet its application in flexible labor markets remains underdeveloped. Gig workers often lack formal performance targets, face frequent interruptions, and operate under conditions of high uncertainty, all of which can weaken goal commitment [2]. Recent studies have begun to explore how digital platforms can intervene to help workers set goals, with field experiments demonstrating that even simple goal reminders can improve output quantity and quality [3].

Concurrently, the field of human–AI interaction has matured, producing frameworks for joint decision-making in contexts ranging from medical diagnosis to autonomous driving [4]. The concept of complementarity, where humans and AI systems leverage their respective strengths, is central to this literature. For goal setting, AI systems can process large volumes of data to identify patterns and recommend optimal targets, while humans bring contextual understanding, ethical judgment, and intrinsic motivation. Large language models, in particular, have shown remarkable proficiency in parsing natural language goals, generating sub-goals, and providing motivational feedback [5]. However, existing work on LLM-based

coaching has largely focused on individual productivity applications rather than the systemic challenges of labor market platforms.

The governance of AI in labor markets raises important questions of fairness, transparency, and accountability. Algorithmic management practices have been criticized for their opacity and potential to exacerbate power asymmetries between platforms and workers [6]. A key concern is that AI-driven goal recommendation systems might inadvertently reinforce biases or impose unrealistic expectations, particularly for workers in precarious positions. A field experiment with gig workers by [7] examined the effects of self-set versus platform-assigned goals, finding that worker autonomy in goal selection was associated with higher commitment and performance, but only when combined with appropriate feedback mechanisms. This finding underscores the need for joint decision-making models that preserve worker agency while leveraging AI capabilities.

3. Proposed Framework for Human–AI Joint Goal Decision-Making

We propose a framework that positions LLM agents as collaborative partners rather than prescriptive controllers in the goal-setting process. The framework consists of three interrelated layers: the interaction layer, the decision layer, and the governance layer. At the interaction layer, LLM agents engage workers in natural language conversations to elicit preferences, assess current task demands, and negotiate goal targets. These conversations are structured around a set of dialogue protocols that balance open-ended exploration with goal convergence. For example, an LLM might ask a delivery worker about their preferred shifts, income targets, and fatigue levels, then suggest a set of daily and weekly goals that are both challenging and feasible given historical performance data [8]. The worker can then accept, modify, or reject these suggestions, and the LLM updates its model accordingly.

The decision layer handles the computational aspects of goal optimization. Using reinforcement learning and Bayesian methods, the LLM generates goal recommendations that maximize expected utility for the worker while satisfying platform constraints such as service level agreements and fairness criteria. This layer incorporates a multi-objective optimization that balances performance metrics (e.g., tasks completed, earnings) with well-being indicators (e.g., work hours, stress scores). The LLM’s natural language interface serves as a front end to these underlying algorithms, enabling workers to interact with complex models without specialized training. However, the decision layer must be designed to be interpretable, allowing workers to understand why a particular goal was suggested and how it relates to their long-term interests [9].

The governance layer provides oversight and accountability mechanisms. It includes rules for data privacy, algorithmic auditing, and worker recourse. Because LLMs can generate biased or harmful recommendations if left unchecked, the governance layer implements fairness constraints that prevent goals from discriminating against vulnerable groups or encouraging excessive work hours. It also maintains a record of all goal negotiations and outcomes, which can be reviewed by independent auditors or worker representatives. This tripartite structure mirrors established principles in socio-technical systems design, where technical and social subsystems must be jointly optimized [10].

4. Structural Trade-Offs in Joint Decision-Making

The design of any human–AI joint decision-making system involves inherent trade-offs that must be carefully managed. One primary tension is between autonomy and control. Workers in flexible labor markets value the ability to set their own schedules and goals, yet excessive

autonomy can lead to procrastination, underperformance, and burnout. The LLM framework can help by providing structure and accountability, but if the AI becomes too directive, workers may perceive it as encroaching on their freedom. Research on reactance theory suggests that individuals resist attempts to limit their behavioral freedom, which can backfire and reduce goal commitment [11]. Therefore, the LLM should be designed to offer suggestions rather than mandates, and to allow workers to override recommendations with minimal friction.

Another trade-off concerns the alignment of individual and platform goals. Platforms aim to maximize aggregate metrics such as number of deliveries per hour or customer satisfaction ratings, whereas workers may prioritize income stability, skill development, or work-life balance. The joint decision-making process must reconcile these sometimes conflicting objectives. One approach is to incorporate worker preferences as explicit parameters in the optimization model, giving each worker a personalized weight for different goal dimensions. However, this requires accurate elicitation of preferences, which is itself a challenge when workers may not fully articulate their true desires [12]. Moreover, platforms may have an incentive to nudge workers toward goals that benefit the platform even at the expense of worker well-being. Transparent governance mechanisms, including worker representation on oversight boards, can mitigate this risk.

A third trade-off relates to the granularity and frequency of goal updates. Frequent goal adjustments allow the system to respond to changing conditions, but they can also create instability and cognitive overload. Workers may become frustrated if goals shift too often or are revised in response to temporary fluctuations. The LLM framework needs to incorporate a stability criterion, ensuring that goals are updated only when there is substantial evidence of a change in the worker's capacity or environment. This is analogous to the concept of satisficing in organizational decision-making, where a good enough goal is maintained rather than constantly optimized [13].

5. Deployment, Infrastructure, and Robustness

Deploying an LLM-based goal-setting framework at scale requires a robust technical infrastructure. The system must handle large volumes of concurrent conversations, maintain low latency for real-time interactions, and store vast amounts of dialogue data securely. Cloud-based architectures with distributed LLM serving nodes are typical, but they raise concerns about data sovereignty and the environmental cost of training and inference [14]. For gig workers who may operate across multiple platforms, interoperability standards are needed to allow goal data to be shared (with worker consent) across different systems. The European Union's General Data Protection Regulation provides a regulatory template, but its application to algorithmic management remains contested [15].

Robustness is another critical consideration. LLMs are known to be susceptible to adversarial inputs, hallucination, and distribution shift. In the context of goal setting, a malicious actor could manipulate the system by providing false information about task difficulty or worker availability, leading to unrealistic goals. Mitigation strategies include input validation, anomaly detection, and human-in-the-loop verification for high-stakes decisions [16]. Additionally, the system must be robust to changes in the labor market, such as fluctuations in demand, new regulations, or shifts in worker demographics. Continuous learning mechanisms that retrain the LLM on updated data can help, but they must be carefully monitored to prevent drift toward unintended behaviors.

Sustainability, both economic and environmental, is a long-term concern. Operating an LLM-based service entails significant compute costs, which may be passed on to workers or platforms. Alternative approaches such as smaller, fine-tuned language models or hybrid systems that combine rule-based logic with LLM modules could reduce the carbon footprint while maintaining performance [17]. Furthermore, the system should be designed to evolve with workers over months and years, supporting skill acquisition and career progression rather than focusing solely on short-term metrics.

6. Fairness and Policy Implications

Fairness in AI-mediated goal setting involves multiple dimensions: distributive fairness (are goals equitable across workers?), procedural fairness (are the processes transparent and consistent?), and interactional fairness (are workers treated with respect by the AI?) [18]. An LLM that learns from historical data may perpetuate existing biases, such as recommending lower goals for workers from marginalized groups or for those with shorter platform tenure. To counter this, the framework must include bias detection and mitigation algorithms, as well as regular audits by independent third parties. The governance layer should allow workers to contest unfair goal recommendations through a structured appeals process.

Policy implications extend beyond fairness to broader labor rights. The increasing use of AI in managing gig workers has prompted calls for new regulations, such as the “algorithmic accountability” provisions in proposed European AI Act [19]. Our framework suggests that joint decision-making can serve as a middle ground between full algorithmic control and complete laissez-faire, but it requires regulatory support to ensure that workers’ voices are genuinely incorporated. Policies should mandate that platforms provide workers with meaningful control over goal setting, including the right to opt out of AI recommendations without penalty. Additionally, data ownership and portability rights are essential to prevent lock-in effects.

From a systems perspective, the deployment of LLM-based goal frameworks could reshape the labor market by increasing the productivity of flexible workers, potentially reducing the demand for permanent employment. While this may lead to greater efficiency, it also raises concerns about the erosion of traditional employment protections [20]. Policymakers must consider how to adapt social safety nets, such as unemployment insurance and health benefits, to a world where work is increasingly mediated by AI. The joint decision-making model, by fostering worker agency, could serve as a foundation for a more equitable platform economy.

7. Cross-Domain Case Illustrations

To ground the discussion, we briefly consider three illustrative domains. In crowd work, platforms like Amazon Mechanical Turk are characterized by low pay, high task fragmentation, and minimal worker interaction. An LLM-based assistant could help workers set daily earnings goals, select tasks that match their skills, and provide encouragement based on past performance. Field research has shown that even simple goal prompts can increase output by 10–20% in such settings [21]. However, care must be taken to avoid pushing workers to accept exploitative tasks.

In on-demand delivery, workers face unpredictable demand and physical strain. The LLM could incorporate real-time data on traffic, weather, and shift availability to suggest optimal work periods and rest breaks. A joint decision-making process where the AI proposes a schedule and the worker adjusts based on personal constraints could improve both efficiency

and well-being. Pilot studies in this domain have reported enhanced satisfaction and reduced turnover when workers are given more say in their goal structures [22].

In freelance consulting, work is longer-term and more knowledge-intensive. Here the LLM might assist with project milestone planning, skill development targets, and client communication strategies. The joint decision-making model aligns well with the professional autonomy that consultants expect. However, the system must handle nuanced language and domain-specific jargon, which current LLMs can manage with appropriate fine-tuning.

8. Conclusion

This paper has presented a comprehensive framework for human–AI joint decision-making in flexible labor markets, centered on the use of large language models to support goal commitment and performance enhancement. We argued that such a system must be designed as a socio-technical architecture with distinct interaction, decision, and governance layers, each addressing specific challenges of autonomy, alignment, and fairness. Through an analysis of structural trade-offs, infrastructure requirements, and policy implications, we highlighted the need for careful balancing of algorithmic assistance with human agency. The proposed framework moves beyond simple automation or surveillance toward a collaborative partnership where workers and AI jointly define and pursue meaningful goals. Future research should focus on empirical validation of the framework across different labor market segments, longitudinal studies of worker outcomes, and the development of standardized benchmarks for fairness and robustness. As AI technologies continue to evolve, the principles outlined here can inform the design of humane and productive digital labor platforms.

References

1. Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, 57(9), 705–717.
2. Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56–75.
3. Jia, R., & Wagner, C. (2021). The effects of goal setting on gig worker performance: A field experiment. *Management Science*, 67(10), 6015–6033.
4. Steyvers, M., & Kumar, A. (2023). Human-AI interaction and decision making: A framework for understanding complementarity. *Annual Review of Psychology*, 74, 333–358.
5. Luo, Y., & Zhang, H. (2023). Large language models as personal coaches: Goal setting and productivity enhancement. In *Proceedings of the 2023 Conference on Human Factors in Computing Systems* (pp. 1–12). ACM.
6. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber’s drivers. *International Journal of Communication*, 10, 3758–3784.
7. Min, X., Chi, W., Hu, X., & Ye, Q. (2024). Set a goal for yourself? A model and field experiment with gig workers. *Production and Operations Management*, 33(1), 205-224.
8. Wang, Y., & Singh, L. (2022). Reinforcement learning for personalized goal recommendation in online labor platforms. *Journal of Machine Learning Research*, 23(45), 1–38.

9. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.
10. Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1), 4–17.
11. Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
12. Slovic, P. (1995). The construction of preference. *American Psychologist*, 50(5), 364–371.
13. Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
14. Patterson, D., et al. (2021). Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.
15. Wachter, S., & Mittelstadt, B. (2019). A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI. *Columbia Business Law Review*, 2019(2), 494–620.
16. Hendrycks, D., & Dietterich, T. (2019). Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*.
17. Gupta, A., & Wu, C. (2023). Efficient fine-tuning of language models for domain-specific goal setting. *Transactions of the Association for Computational Linguistics*, 11, 450–467.
18. Colquitt, J. A., Conlon, D. E., Wesson, M. J., Porter, C. O. L. H., & Ng, K. Y. (2001). Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology*, 86(3), 425–445.
19. European Commission. (2021). Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). COM(2021) 206 final.
20. Kalleberg, A. L. (2009). Precarious work, insecure workers: Employment relations in transition. *American Sociological Review*, 74(1), 1–22.
21. Peer, E., & Brandimarte, L. (2020). A field experiment on goal setting and performance in online labor markets. *Journal of Behavioral and Experimental Economics*, 85, 101511.
22. Brawley, A. M., & Pury, C. L. S. (2016). Work experiences on MTurk: Job satisfaction, turnover, and use of Amazon’s Mechanical Turk. *Journal of Managerial Psychology*, 31(5), 885–900.