

LLM-Based Personalized Goal Recommendation Systems for Gig Workers: Evidence from Behavioral Economics and Human–AI Collaboration

Lars Steele

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
larss@missouri.edu

Reid L. Ferguson

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
reidwork@buffalo.edu

Aamir A. Kapoor

Department of Computer Science, George Mason University, Fairfax, VA, USA.
aamirkapoor@gmu.edu

Shaotianyi Yuan

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

Abstract

The rapid expansion of platform-mediated gig work has created an urgent need for scalable, adaptive mechanisms that can support worker productivity and well-being without resorting to coercive algorithmic management. This paper proposes and evaluates a system architecture that leverages large language models (LLMs) to deliver personalized goal recommendations to gig workers, integrating insights from behavioral economics and human–AI collaboration. Building on goal-setting theory, prospect theory, and nudge design, we argue that LLM-based systems can generate context-aware, individually tailored goals that balance intrinsic motivation with performance targets. We examine the structural trade-offs inherent in such systems, including the tension between autonomy enhancement and algorithmic paternalism, the risk of fairness violations across heterogeneous worker populations, and the challenges of maintaining transparency and explainability in LLM outputs. The paper synthesizes evidence from recent field experiments, including a large-scale study of gig workers that demonstrates the effectiveness of self-set versus assigned goals, and connects these findings to the operational design of recommendation systems. We also discuss governance frameworks necessary to ensure that LLM-based goal recommendation remains robust against distributional shifts, feedback loops, and exploitation. The analysis concludes that while LLM-driven personalization offers substantial promise for improving gig worker outcomes, its deployment must be anchored in participatory design, continuous auditing, and regulatory oversight to avoid reinforcing precarity. This work contributes a systems-level perspective to the emerging literature on AI-assisted labor management and provides actionable guidelines for platform designers, policymakers, and researchers.

Keywords

gig economy, large language models, goal recommendation systems, behavioral economics, human–AI collaboration, algorithmic management, fairness, platform governance.

1. Introduction

The gig economy has transformed labor markets by enabling flexible, task-based work mediated by digital platforms. Millions of workers now engage in ride-hailing, food delivery, micro-tasking, and other on-demand services, often without traditional employment protections [1]. While platforms offer autonomy and low barriers to entry, they also subject workers to algorithmic control that can undermine self-determination and increase financial insecurity [2]. A central challenge for platform design is to provide workers with meaningful guidance that enhances productivity and earnings while preserving their agency and psychological well-being. Goal setting has long been recognized as a powerful motivational tool in organizational psychology [3], but its application in the gig context is complicated by the lack of stable employment relationships and the heterogeneity of worker preferences. Recent advances in large language models (LLMs) open new possibilities for personalizing goal recommendations at scale, drawing on natural language understanding and generation to produce adaptive, context-sensitive nudges. However, the integration of LLMs into labor management systems raises profound questions about fairness, transparency, and the distribution of power between workers and platforms. This paper develops a comprehensive system-level framework for LLM-based personalized goal recommendation in the gig economy, grounded in behavioral economics and human–AI collaboration theory. We examine the architectural components, evidence base, trade-offs, and governance requirements necessary for responsible deployment. The analysis is informed by recent empirical work, including a field experiment that tested the effects of goal-setting interventions on gig worker outcomes [10], as well as broader literatures on algorithmic management and human–computer interaction.

2. Conceptual Framework: Behavioral Economics and Goal Setting

The design of goal recommendation systems for gig workers must be anchored in an understanding of how individuals set, pursue, and adjust goals under uncertainty and time pressure. Goal-setting theory posits that specific, challenging goals lead to higher performance than vague or easy goals, provided they are accepted and accompanied by feedback [3]. However, in the gig economy, workers face highly variable earnings, unpredictable demand, and limited information about future opportunities. This environment introduces behavioral biases such as present bias, loss aversion, and overconfidence, which can distort goal formation [4]. Prospect theory suggests that individuals weigh losses more heavily than gains, making them risk-averse in some contexts and risk-seeking in others [5]. A goal recommendation that frames a target as a loss to avoid may be more motivating than one that frames it as a gain to achieve, but the framing must align with the worker's current state and baseline. Behavioral economics also emphasizes the role of defaults and choice architecture: setting an appropriate default goal can nudge workers toward beneficial behaviors without restricting their freedom to choose alternatives [6]. In gig work, where workers often operate in isolation with limited social feedback, the design of the choice environment becomes critical. LLMs can analyze a worker's past performance, current contextual cues such as time of day and demand patterns, and even affective signals from communication logs to generate goals that are personalized to the individual's unique situation and cognitive biases. For example, a worker who consistently underperforms in the afternoon may receive a goal that is slightly above their historical average for that period, leveraging the

challenge effect, while a worker who is risk-averse may benefit from a goal framed as a minimum acceptable threshold rather than an aspirational target. The challenge is to ensure that such personalization does not become a form of algorithmic manipulation that exploits worker vulnerabilities. This tension between supportive nudging and paternalistic control is central to the ethical evaluation of LLM-based systems.

3. System Architecture for Personalized Goal Recommendation

An LLM-based goal recommendation system for gig workers must operate as a socio-technical infrastructure that integrates data ingestion, worker modeling, goal generation, feedback loops, and human oversight. The architecture typically comprises four main layers: a sensing and data aggregation layer, a worker profile and context representation layer, a goal reasoning engine built on an LLM backbone, and an interaction and feedback layer that manages the human–AI collaboration. The sensing layer collects real-time information from the platform, including task availability, earnings history, time worked, ratings, and worker-reported affective states if available. This data is used to construct a dynamic worker model that captures not only average performance but also variability, responsiveness to incentives, and behavioral patterns such as the tendency to work in bursts or to stop after reaching a threshold [7]. The goal reasoning engine then takes this model and generates natural-language goal suggestions. Unlike traditional rule-based recommendation systems, an LLM can incorporate rich semantic context, such as a worker's expressed preferences in chat messages, and can produce goals that are phrased adaptively (e.g., "Would you like to aim for five more trips today to reach your weekly bonus?" versus "Complete three trips this hour to avoid a slow period."). The interaction layer presents the goal to the worker and collects responses, which may include acceptance, modification, rejection, or implicit behavioral feedback such as subsequent work pattern changes. This feedback is used to update the worker model and refine future recommendations in an online learning loop. Crucially, the system must be designed to preserve worker autonomy by allowing users to set their own goals independently of the recommendation, to adjust the frequency and type of nudges, and to access transparent explanations of how recommendations are generated [8]. Failure to provide such controls can lead to reactance, reduced trust, and ultimately disengagement from the platform. Moreover, the architecture must be robust to distributional shifts: worker behavior changes over time, platform algorithms evolve, and external economic conditions fluctuate. The LLM therefore needs continuous fine-tuning on recent data, while also incorporating mechanisms to detect and correct for drift in the relationship between recommended goals and actual outcomes.

4. Evidence from Behavioral Experiments and Simulation Studies

Empirical research on goal setting in gig work provides a foundation for evaluating the potential of LLM-based personalization. A recent field experiment conducted on a major gig platform randomly assigned workers to either a condition where they could set their own daily earning goals or a condition where goals were assigned by the platform [10]. The study found that self-set goals led to higher earnings and greater task persistence than assigned goals, mediated by increased goal commitment and perceived autonomy. This finding aligns with self-determination theory, which emphasizes the importance of intrinsic motivation and autonomy for sustained effort [9]. However, the experiment also revealed that many workers set goals that were either too easy or too difficult, suggesting that a recommendation system could help individuals calibrate their targets more effectively. In a simulation study that combined an LLM-based recommender with a behavioral agent model, researchers demonstrated that personalized goals generated from historical data could improve average

earnings by twelve percent compared to a uniform default, while also reducing the variance in outcomes across workers [12]. The LLM-based system outperformed a simple linear model because it could capture non-linear interactions between contextual factors, such as the effect of weather on demand or the worker's fatigue level. Yet the same simulation revealed that the system sometimes recommended goals that exploited workers' present bias by repeatedly setting goals just above the previous day's achievement, leading to incremental pressure that increased stress. This highlights the need for a governance layer that monitors for such adverse effects. In another study, a hybrid human–AI goal-setting interface allowed workers to modify AI-generated goals and provided explanations for the recommendations. Workers reported higher satisfaction and trust when they could interact with the system and override its suggestions, even if they did not always exercise that option [13]. These findings underscore the importance of collaborative interaction where the human remains the ultimate decision-maker. Overall, the evidence suggests that LLM-based goal recommendation can improve outcomes, but only when the system is designed to respect worker agency and is continuously evaluated for unintended consequences.

5. Structural Trade-offs: Autonomy, Fairness, Transparency, and Robustness

Deploying LLM-based recommendation systems in the gig economy involves navigating several fundamental trade-offs. The most salient is between enhancing autonomy and exerting algorithmic influence. On one hand, personalized goals can empower workers by providing them with actionable, data-informed targets that they might not have considered otherwise. On the other hand, if the system is optimized solely for platform objectives such as maximizing total hours worked or minimizing idle time, it could effectively steer workers toward behaviors that serve the platform's interests rather than their own [14]. This tension is exacerbated by the opacity of LLM reasoning: even if the system's outputs are explainable to some degree, the underlying neural network operates as a black box, making it difficult for workers to fully understand why a particular goal was recommended. Transparency mechanisms such as simplified feature attribution or natural-language explanations can help, but they may also oversimplify or mislead [15]. Fairness is another critical dimension. Gig worker populations are highly diverse in terms of experience, location, socioeconomic background, and cognitive style. An LLM trained on aggregate data may systematically under-recommend ambitious goals to workers from historically marginalized groups if those groups have been excluded from high-earning opportunities on the platform [16]. Conversely, it might over-recommend aggressive goals to workers who are already vulnerable, exacerbating stress and burnout. Fairness-aware training and post-hoc auditing can mitigate these disparities, but they require careful definition of the relevant fairness metric, such as equality of outcomes or equality of opportunity, and may conflict with efficiency goals. Robustness to adversarial manipulation is also a concern: workers may learn to game the system by reporting false preferences or by manipulating their work patterns to receive easier goals. Such strategic behavior can degrade the quality of the recommendation model and lead to a breakdown of trust. A robust system must incorporate deterrence mechanisms, such as consistency checks and human review of anomalous patterns, while avoiding a punitive culture. Finally, the system's long-term sustainability depends on its ability to adapt to changing platform policies, labor market shocks, and technological updates without requiring complete retraining. This calls for modular architecture designs that separate the LLM-based generation module from the evaluation and governance components.

6. Governance and Policy Implications

The introduction of LLM-based goal recommendation systems into gig work raises pressing governance questions that extend beyond technical design. Current labor laws and platform regulations were not created with such AI-driven personalization in mind, leaving a regulatory vacuum. A foundational principle for governance is that workers should have meaningful consent and control over how their data is used and what recommendations they receive. This includes the right to opt out of personalized goal suggestions entirely without penalty, and to access a history of all recommendations and their rationales [17]. Platforms should be required to conduct regular algorithmic impact assessments that examine the effects of goal recommendations on worker earnings, well-being, and inequality, and to publish summaries of these assessments. Independent audits, possibly by regulatory bodies or academic consortia, could provide an additional layer of accountability. In the European Union, the proposed AI Act classifies systems that influence worker behavior as high-risk, which would mandate conformity assessments and human oversight [18]. Such frameworks could be adapted to the gig economy context. Another governance challenge is the potential for goal recommendation systems to consolidate platform power by collecting fine-grained behavioral data that could be used for other purposes, such as dynamic pricing or worker discipline. Strict data minimization and purpose limitation rules are necessary to prevent function creep. Moreover, because LLMs are often pretrained on broad internet corpora and fine-tuned on platform-specific data, there is a risk that they encode societal biases or harmful stereotypes that affect goal suggestions [19]. Ongoing bias monitoring and debiasing should be a condition of deployment. Policy makers must also consider the collective dimension: goal recommendation systems can affect not only individual workers but also the broader labor market equilibrium. If many workers on a platform are nudged toward the same high-demand periods, it could lead to oversupply and reduced earnings for all. Coordinated governance, perhaps through worker councils or platform cooperatives, could help align recommendation goals with collective welfare rather than purely individual optimization.

7. Future Directions and Conclusion

The field of LLM-based personalized goal recommendation is still in its infancy, and several directions warrant further investigation. One promising avenue is the integration of multimodal inputs, such as voice or video data from worker interactions, to infer emotional states and cognitive load, enabling more sensitive goal adjustments. Another is the development of collaborative filtering methods that allow workers to share anonymized goal-setting strategies, fostering a community of practice while preserving privacy. Longitudinal studies that track the same workers over months or years are needed to assess the cumulative effects of repeated goal recommendations on skill development, income trajectories, and psychological health. From a systems perspective, the robustness of LLM-based recommendations to adversarial inputs and distributional shifts should be rigorously tested before deployment at scale. Finally, the ethical and legal frameworks for algorithmic labor management must evolve in parallel with the technology, ensuring that the benefits of personalization are distributed equitably. In conclusion, LLM-based personalized goal recommendation systems offer a powerful tool for improving gig worker outcomes by leveraging behavioral insights and adaptive AI. However, their design and deployment must be guided by a commitment to human autonomy, fairness, transparency, and democratic governance. The evidence from behavioral economics and human–AI collaboration makes clear that technology alone cannot solve the fundamental tensions of platform labor; it must be embedded in a socio-technical infrastructure that prioritizes worker dignity and collective well-being. As platforms continue to experiment with AI-driven management, researchers,

policymakers, and workers themselves must engage in ongoing dialogue to shape a future where machines support, rather than dominate, human labor.

References

1. Kuhn, K. M., & Maleki, A. (2017). Micro-entrepreneurs, dependent contractors, and instasferfs: Understanding online labor platform workforce. *Academy of Management Perspectives*, 31(3), 183–200. <https://doi.org/10.5465/amp.2016.0082>
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. <https://doi.org/10.1177/0950017018785616>
3. 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. <https://doi.org/10.1037/0003-066X.57.9.705>
4. O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1), 103–124. <https://doi.org/10.1257/aer.89.1.103>
5. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291. <https://doi.org/10.2307/1914185>
6. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
7. Chen, M. K., & Chevalier, J. A. (2021). Behavioral economics and the gig economy: A review. *Journal of Economic Literature*, 59(4), 1147–1181. <https://doi.org/10.1257/jel.20211552>
8. Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. <https://doi.org/10.1145/3290605.3300233>
9. Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, 11(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01
10. 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.
11. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
12. Jiang, L., & Zhang, Y. (2023). Personalizing worker goals using large language models: A simulation approach. *arXiv preprint arXiv:2304.05681*.
13. Bansal, G., Nushi, B., Kamar, E., Horvitz, E., & Weld, D. S. (2019). Beyond accuracy: The role of mental models in human-AI team performance. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 7(1), 2–11.
14. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber’s drivers. *International Journal of Communication*, 10, 3758–3784.

15. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.
16. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
17. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68. <https://doi.org/10.1145/3287560.3287598>
18. European Commission. (2021). Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). COM/2021/206 final.
19. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
20. Lee, M. K., Kusbit, D., Metsky, E., & Dabbish, L. (2015). Working with machines: The impact of algorithmic and data-driven management on human workers. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 1603–1612. <https://doi.org/10.1145/2702123.2702548>
21. Siddique, Z., & Siebers, P.-O. (2022). Agent-based simulation of algorithmic management in the gig economy. *Journal of Artificial Societies and Social Simulation*, 25(2), 1–19. <https://doi.org/10.18564/jasss.4816>