

# Digital Goal Architectures: An LLM-Driven Framework for Predicting Persistence and Earnings in Gig Economy Workforces

Rremdan Gurns

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

Xavier Horton

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.  
xhorton@uc.edu

Varun Murehy

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.  
varun.murthy602@oregonstate.edu

Milos L. Horton

Department of Computer Science, Binghamton University, Binghamton, NY, USA.  
milos.work@binghamton.edu

## Abstract

The rapid expansion of gig economy platforms has created urgent demands for systems that can sustain worker engagement and predict income trajectories under highly variable conditions. Traditional econometric models and rule-based management tools often fail to capture the dynamic interplay between worker autonomy, platform incentives, and evolving task environments. This paper introduces a conceptual framework termed Digital Goal Architectures, which integrates large language models (LLMs) into a predictive infrastructure designed to forecast persistence and earnings among gig workers. The framework positions LLMs as central reasoning engines that process heterogeneous data streams including worker goal statements, platform interaction logs, temporal effort patterns, and exogenous market signals. By leveraging the contextual understanding and generative capabilities of LLMs, the architecture moves beyond static goal-setting theory toward adaptive, personalized prediction that respects the situated nature of gig work. We examine the structural trade-offs inherent in deploying such systems, focusing on the tension between predictive accuracy and interpretability, the risks of algorithmic feedback loops that may exacerbate earnings inequality, and the governance challenges associated with real-time behavioral inference. The paper also addresses infrastructure requirements for scalable deployment, including latency constraints, data privacy safeguards, and the need for continuous model updating in non-stationary labor markets. Policy implications are discussed with an emphasis on transparency standards, worker consent mechanisms, and the design of fallback procedures that prevent harmful decisions. By situating LLM-driven predictive architectures within the broader socio-technical landscape of platform labor, this work provides a foundational perspective for researchers and practitioners aiming to build fair, robust, and sustainable systems for the future of work.

## **Keywords**

gig economy, large language models, predictive modeling, goal architectures, algorithmic management, worker persistence, earnings prediction, socio-technical systems, fairness, governance.

## **1. Introduction**

The gig economy has emerged as a defining feature of contemporary labor markets, characterized by short-term, task-based engagements mediated by digital platforms. Workers in these environments operate with high degrees of temporal flexibility yet face substantial income volatility, limited social protections, and asymmetric information about task availability and compensation. Platforms such as ride-hailing services, food delivery applications, and micro-task marketplaces must simultaneously attract a sufficient labor supply while ensuring that workers remain active and productive over time. Understanding the determinants of worker persistence and earnings is therefore a central challenge for both platform operators and policymakers. Traditional approaches have drawn on labor economics, organizational behavior, and operations research to model worker decisions as responses to wages, surge pricing, and schedule flexibility [1], [2]. However, these models often assume rational expectations and stable preferences, assumptions that conflict with the highly contextual and emotionally nuanced reality of gig work.

Recent advances in artificial intelligence, particularly the development of large language models, have opened new possibilities for capturing the rich semantic signals embedded in worker communications, self-reported goals, and platform feedback. LLMs possess the ability to reason over unstructured text, incorporate long-range dependencies, and generate context-aware predictions that can be personalized at scale. This paper proposes a framework called Digital Goal Architectures, which uses LLMs as the core predictive engine for anticipating how individual workers' persistence and earnings evolve over time. The framework is grounded in goal-setting theory, which has long established that specific, challenging goals enhance task performance and motivation [3]. Yet in the gig economy, workers set their own goals, often implicitly, and these goals shift in response to platform dynamics, personal circumstances, and social comparisons. A static goal model fails to capture this fluidity. By contrast, an LLM-driven architecture can continuously ingest narrative data such as a worker's expressed intention to earn a certain amount in a week, their frustration about low demand, or their plans to take a break, and integrate these signals with behavioral logs to produce real-time predictions of persistence and earnings.

The purpose of this paper is not to present a fully implemented system but to articulate the conceptual design, examine the structural trade-offs and governance challenges that accompany such an architecture, and discuss the broader implications for fairness, robustness, and policy. We begin by situating the framework within existing literature on goal-setting, algorithmic management, and predictive labor analytics. We then detail the components of the proposed architecture, describing how LLMs can be orchestrated to process multimodal inputs and generate forecasts. Subsequent sections analyze the trade-offs between accuracy and interpretability, the infrastructure demands for deployment at platform scale, and the ethical and regulatory considerations that must guide design choices. The paper concludes with a synthesis of forward-looking perspectives on the role of LLM-driven systems in shaping the future of gig work.

## **2. Theoretical Foundations of Goal Architectures in Gig Work**

Goal-setting theory has been one of the most robust frameworks in organizational psychology, demonstrating across hundreds of studies that specific and challenging goals lead to higher performance compared to vague or easy goals [3]. In traditional work settings, goals are often assigned by managers and are embedded within performance appraisal systems. The gig economy inverts this arrangement: workers are largely self-directed, setting their own targets for hours worked, income, and task completion. This autonomy introduces both opportunities for intrinsic motivation and risks of underperformance due to lack of structure. Empirical research has begun to examine how goal-setting interventions can be implemented in gig platforms. For instance, a field experiment with delivery workers revealed that prompting workers to set personal daily earnings goals increased their subsequent effort and earnings, though the effect was moderated by workers' baseline motivation and the salience of the goal prompt [8]. Such findings underscore the potential of goal-based interventions but also highlight the need for adaptive systems that can recognize when and how to intervene.

Existing platform management tools rely predominantly on algorithmic control, where workers' behavior is influenced through dynamic pricing, task allocation, and performance ratings [2], [20]. These mechanisms are often opaque and non-negotiable, leading to worker dissatisfaction and turnover. An LLM-driven goal architecture offers a complementary approach that privileges worker agency by interpreting their own stated objectives and aligning platform recommendations accordingly. Rather than imposing external constraints, the system can suggest personalized strategies, such as optimal times to work or task types that align with a worker's income goals, thereby fostering a sense of partnership. This shift from algorithmic control to algorithmic support requires a deep understanding of how goals are formed, maintained, and revised. Cognitive science suggests that goals are not static but dynamically updated based on feedback and environmental cues [7]. LLMs, with their capacity for sequential reasoning and context maintenance, are particularly suited to model such dynamic processes.

The theoretical foundation of the proposed architecture also draws from research on predictive analytics in labor markets. Traditional regression-based models of worker earnings and tenure have used variables such as experience, location, and platform rating [10], but these models lack the granularity to capture short-term fluctuations driven by motivational states. Machine learning approaches, including gradient boosting and neural networks, have improved predictive accuracy but often sacrifice interpretability and require fixed feature engineering [5]. LLMs introduce a paradigm shift because they can directly process natural language as input, eliminating the need for handcrafted features and enabling the incorporation of rare or idiosyncratic signals, such as a worker's mention of a family event or a change in preferred delivery area. This capability significantly expands the dimensionality of the predictive space while maintaining the ability to explain predictions through attention mechanisms and generated rationales.

### **3. An LLM-Driven Framework for Persistence and Earnings Prediction**

The Digital Goal Architecture proposed here consists of several interconnected modules that together form a pipeline from raw data collection to actionable predictions. The first module is the data ingestion layer, which aggregates three primary streams: worker-generated text (including goal statements, status updates, support tickets, and social network posts), platform telemetry (login times, task acceptance rates, trip durations, earnings history), and external context (weather, local demand, traffic patterns, public holidays). Each stream is timestamped and stored in an event log that serves as the input to the LLM. Because gig workers often

interact with platforms via mobile devices, text inputs are typically short and informal; LLMs fine-tuned on conversational data are well-suited to handle such brevity and nuance.

The second module is the goal extraction and representation component. Here, the LLM is prompted to identify and formalize a worker's current goals from their textual statements. For example, a worker might type "I need to make \$200 today to cover rent." The system extracts the goal metric (earnings), the target value (\$200), and the temporal horizon (today). Multiple goals may coexist, such as a weekly earnings target combined with a desire to minimize idle time. The LLM can represent these goals as structured objects while retaining the original language embeddings for downstream reasoning. This dual representation allows the system to reason about goal hierarchy, potential conflicts, and priority shifts.

The core predictive module uses the extracted goals and historical telemetry to forecast two target variables: short-term persistence, defined as the probability that the worker continues to accept tasks over the next several hours, and earnings over a specified horizon, typically the remainder of the day or week. The prediction is generated through a chain-of-thought reasoning process [4] where the LLM iteratively considers the worker's stated goal, current progress (e.g., \$50 earned so far), platform conditions (surge multipliers, task density), and behavioral patterns (e.g., typical quitting time, response to low demand). By generating intermediate reasoning steps, the model provides not only a probability or point estimate but also a textual explanation that can be surfaced to platform managers or to the worker themselves. This interpretability is crucial for building trust and enabling oversight.

The third module is the feedback and adaptation loop. After a prediction is made, the system can optionally recommend actions to the worker, such as extending shift duration or moving to a higher-demand area. The worker's response (whether they follow the recommendation) is logged, and the LLM updates its internal state accordingly. Over time, the system learns the effectiveness of different recommendation types for each worker, personalizing the goal architecture. Importantly, the model must be continuously retrained or updated to adapt to non-stationary environments, such as seasonal demand changes or shifts in platform policies. Because fine-tuning large LLMs is computationally expensive, a more practical approach is to use a frozen base model with a lightweight adapter that is updated online, or to employ retrieval-augmented generation that pulls from recent worker-specific history.

#### **4. Systemic Trade-offs: Accuracy, Interpretability, and Fairness**

No predictive architecture is without trade-offs, and the integration of LLMs into gig economy management introduces several critical tensions. The first trade-off is between accuracy and interpretability. LLMs, especially large-scale models, are often described as black boxes because their internal computations are difficult to trace. While chain-of-thought prompting can improve explainability by generating explicit reasoning steps, these explanations are not guaranteed to be faithful to the model's actual reasoning process [5]. In a high-stakes context where predictions may influence a worker's access to bonuses, schedule recommendations, or even platform standing, the lack of verifiability is concerning. One response is to design hybrid systems that combine an LLM's output with a simpler, transparent model that acts as a validator. For example, a logistic regression trained on a small set of interpretable features could be used to flag predictions that deviate significantly from historical norms, prompting human review. Such redundancy reduces the risk of catastrophic errors but also increases system complexity and cost.

The second trade-off involves fairness and the potential for algorithmic bias. LLMs are trained on internet-scale data that reflects societal stereotypes and inequalities. When applied to gig workers, these biases could manifest in systematically lower predictions of persistence or earnings for demographic groups that are historically underrepresented or stereotyped as less reliable [6], [12]. Furthermore, the very act of making a prediction can create a self-fulfilling prophecy: if the system predicts low future earnings for a particular worker and recommends fewer shifts or lower pay, that worker may indeed reduce their effort, confirming the initial prediction. This feedback loop is a form of algorithmic harm that is especially pernicious because it appears data-driven and neutral. To mitigate this, the architecture must incorporate fairness constraints at the design stage, such as requiring that the distribution of predicted outcomes across protected groups be calibrated to observed outcomes, and that recommendations do not differ based on race, gender, or other sensitive attributes unless justified by objective task performance. Regular audits of the system's outputs, both offline and in production, are essential.

A third trade-off concerns the balance between personalization and privacy. To generate accurate predictions, the system must ingest large amounts of worker-specific text and behavior. This creates a detailed digital profile that could be used to infer health status, emotional state, or political views if not carefully anonymized. Workers may feel surveilled, which undermines trust and can lead to strategic behavior, such as withholding genuine goal statements [14], [15]. The architecture should therefore be designed with privacy-preserving techniques, including on-device processing where possible, differential privacy mechanisms for aggregated statistics, and clear policies regarding data retention and third-party sharing. Moreover, workers should have the right to access, correct, and delete their goal data, and to opt out of predictive recommendations without penalty. These requirements are not merely technical; they are ethical and legal imperatives that shape the legitimacy of the entire system.

## **5. Infrastructure and Deployment Challenges**

Deploying an LLM-driven goal architecture at the scale of a major gig platform introduces significant infrastructure challenges. Latency is a primary concern: predictions and recommendations must be generated in near real-time, often within seconds, to be useful during a worker's decision moments, such as when deciding whether to log in or accept a task. Large language models, especially those with hundreds of billions of parameters, have substantial inference costs and may require specialized hardware such as high-end GPUs or custom accelerators. A viable approach is to use a distilled or quantized version of the model for real-time serving, reserving the full model for offline batch predictions or periodic retraining. Alternatively, the system can employ a tiered architecture where a lightweight rule-based trigger determines when to invoke the LLM, thereby reducing the average compute load.

Model updating is another significant challenge. Gig labor markets are non-stationary: worker populations shift, platform policies change, and economic conditions fluctuate. A model trained on data from one quarter may become stale within weeks. Continuous learning pipelines must be established that monitor prediction drift and trigger retraining when performance degrades below a threshold. However, retraining large language models from scratch is prohibitively expensive. Parameter-efficient fine-tuning methods, such as low-rank adaptation (LoRA) or adapter layers, offer a practical alternative by updating only a small fraction of the model weights. These methods allow the core LLM to retain its general language understanding while adapting to specific platform dynamics. Furthermore, the

system should incorporate a rollback mechanism so that if a new model version introduces unexpected biases or errors, the previous version can be reinstated quickly.

Robustness to adversarial inputs is equally important. Workers may attempt to game the system by fabricating goals or manipulating their expressed sentiment in order to receive more favorable recommendations. The architecture must include anomaly detection modules that flag inconsistent statements, such as a worker claiming an intention to work ten hours while simultaneously logging out of the platform. Additionally, the system should be designed to be tolerant of missing data; not all workers will provide text inputs regularly. In such cases, the LLM can fall back on historical telemetry and a default goal model derived from population averages. These fallback procedures must be documented and transparent to regulators.

## **6. Governance, Sustainability, and Policy Implications**

The deployment of an LLM-driven goal architecture raises profound questions about governance and accountability. Who is responsible when a prediction leads to harm? In many jurisdictions, platform companies are considered intermediaries and are insulated from liability for algorithmic decisions. However, as predictive systems become more autonomous and personalized, the line between platform and employer blurs. The European Union’s proposed AI Act and various data protection regulations increasingly require that high-risk AI systems undergo conformity assessments, maintain human oversight, and provide explanations for decisions that significantly affect individuals [17]. The Digital Goal Architecture, by virtue of its influence on workers’ earnings and persistence, would likely fall under such regulations. Platforms must therefore prepare for compliance by implementing audit trails, logging all predictions and recommendations, and ensuring that human operators can override automated decisions in cases of apparent error or unfairness.

Sustainability in the context of gig work often refers to the long-term viability of the labor model, but it also applies to the technical system itself. The energy consumption of large language models is a growing concern. Frequent retraining and high-volume inference contribute to carbon emissions that conflict with corporate sustainability goals. Researchers have proposed various mitigation strategies, including model pruning, knowledge distillation, and the use of renewable energy for data centers. Moreover, the architecture should be designed to minimize unnecessary invocations. For instance, if a worker has a stable pattern of behavior and has not expressed any change in goals, the system could rely on a simple autoregressive model rather than calling the LLM for every prediction. Such economies reduce both operational costs and environmental impact.

Policy implications extend beyond the platform to the broader regulatory environment. One key area is data portability. Workers should be able to extract their own goal histories and prediction logs and transfer them to competing platforms, fostering labor mobility and reducing lock-in effects. Another area is the right to explanation. While chain-of-thought reasoning provides a narrative explanation, it may not meet the legal standard of “meaningful information” required under the General Data Protection Regulation. Platforms should invest in research to improve the faithfulness and comprehensibility of LLM-generated explanations. Finally, there is a need for independent oversight bodies that can audit algorithmic management systems in the gig economy. These bodies would have access to model specifications, training data, and deployment logs, and would be empowered to issue corrective orders when systemic bias or unfair practices are identified.

## 7. Conclusion

This paper has presented Digital Goal Architectures as a conceptual framework for integrating large language models into the prediction of persistence and earnings in gig economy workforces. The architecture leverages the unique capabilities of LLMs to process natural language goal statements, contextualize worker behaviors, and generate personalized forecasts with explanatory reasoning. We have examined the structural trade-offs between accuracy, interpretability, and fairness, and discussed the infrastructure challenges of deploying such systems at scale, including latency, model updating, and robustness. Governance and policy considerations were addressed, emphasizing transparency, worker autonomy, and regulatory compliance. While the framework is not yet implemented, it provides a roadmap for researchers and engineers who seek to build systems that enhance rather than undermine the agency and wellbeing of gig workers. Future work should focus on empirical validation through pilot studies on real platforms, the development of fairness-aware training objectives, and the design of participatory governance models that give workers a voice in how their data and predictions are used. The intersection of large language models and labor market dynamics is a rapidly evolving domain; Digital Goal Architectures offer a principled foundation for navigating this complex socio-technical frontier.

## References

1. Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30.
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. Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation. *American Psychologist*, 57(9), 705-717.
4. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Le, Q., Zhou, D., & Chi, E. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35, 24824-24837.
5. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.
6. Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudík, M., & Wallach, H. (2019). Improving fairness in machine learning systems: What do industry practitioners need? *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-16.
7. Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
8. 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.
9. Chen, M. K., & Sheldon, M. (2016). Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform. Working Paper.
10. Hall, J. V., & Krueger, A. B. (2018). An analysis of the labor market for Uber's driver-partners in the United States. *ILR Review*, 71(3), 705-732.

11. Kuhn, P. J., & Malchow-Møller, N. (2022). Online labor markets: A review. *Journal of Economic Literature*, 60(3), 895-945.
12. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning: Limitations and opportunities*. MIT Press.
13. 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.
14. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
15. Bucher, T. (2018). *If... then: Algorithmic power and politics*. Oxford University Press.
16. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30.
17. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
18. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., & Zhao, W. X. (2023). A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*.
19. Möhlmann, M., & Zalmanson, L. (2017). Hands on the wheel: Navigating algorithmic management and Uber drivers. *Proceedings of the International Conference on Information Systems*.
20. Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410.