

Agentic AI Framework for Autonomous Workforce Analytics and Decision Support in Enterprise HRIS Systems

Rohith Balerao¹

¹Senior HRIS Analyst, Chicago, Illinois, USA
Email: rohith.balerao23@gmail.com
ORCID: 0009-0002-3260-9323 (<https://orcid.org/0009-0002-3260-9323>)

Abstract: The fast evolution of artificial intelligence has introduced a new paradigm for enterprise human resource management, in which agentic AI systems can autonomously manage complex workforce analytics and decision-support functions embedded in Human Resource Information Systems (HRIS). This paper presents a novel Agentic AI Framework that combines multi-agent architectures, large language models (LLMs), and real-time data pipelines to empower autonomous reasoning, planning, and execution of key HR functions such as talent acquisition, workforce planning, performance evaluation, attrition prediction, and compliance monitoring. In contrast to conventional rule-based systems or predictive analytics systems, the proposed framework uses goal-directed agents that can (a) dynamically decompose complex tasks, (b) retrieve contextual enterprise data, and (c) generate actionable recommendations when decision context is uncertain while preserving human oversight for high-impact decisions. It collects information through retrieval-augmented generation (RAG), utilizes tools, and feeds what it learns back into future interactions so organizational decisions remain aligned with enterprise goals. Experimental evaluations on a simulated enterprise-level HRIS environment show considerable improvements in decision-making accuracy, operational efficiency, and HR process automation compared with baseline systems. By systematically embedding ethical considerations around bias mitigation, transparency, and human-in-the-loop oversight, this work provides an actionable, scalable, interpretable, and enterprise-ready foundation for the deployment of agentic AI in modern workforce management ecosystems.

Keywords: Agentic AI, Human Resource Information Systems, Workforce Analytics, Multi-Agent Architecture, Decision Support Systems.

1. INTRODUCTION

The enterprise is being transformed by three major pillars: AI, big data, and cloud-native architectures. Human Resource Information Systems (HRIS) have come a long way since their humble beginnings as simple record-keeping systems to well-functioning ecosystems that can handle the most nuanced workforce of today [1]. Yet despite these advances, the vast majority of the existing HRIS platforms are still basically reactive and grounded in static reporting dashboards, standard queries and rule-based workflows unable to meet modern-day workforce demands in a responsive realtime manner [2].

At the enterprise scale, agentic AI has the potential to change HR decision-making by enabling autonomous goal-directed behavior, dynamic planning, and adaptive execution. While classical machine learning models return decoupled predictions, agentic systems are designed around multi-agent architectures, or leverage large language models (LLMs) to break down complex HR tasks and make continuous coordinated recommendations across functional domains only when prompted by human input [4]. This change from a purely passive analytic function to an autonomous decision support capability represents a key inflection point in the maturity evolution of intelligent enterprise systems.

For a long time, workforce analytics has been considered a strategic differentiator for organizations that seek competitive advantage through human capital optimization [5]. Prior studies have shown that these data-led HR activities significantly improve talent retention, reduce hiring costs and increase employee productivity [6]. However, the addition of autonomous AI agents into existing enterprise HRIS settings poses significant data privacy, algorithmic bias, interpretability and regulatory compliance challenges that require systematic attention [7].

In this paper we present a cohesive Agentic AI Framework that overcomes these limitations by integrating multi-agent orchestration, retrieval-augmented generation (RAG), and real-time tool-use capabilities directly into enterprise HRIS pipelines. The framework is intended to autonomously automate HR processes from talent acquisition and attrition prediction to compliance auditing and strategic workforce planning. To foster transparent, auditable, and ethically aligned autonomous decisions, we include human-in-the-loop (HITL) oversight mechanisms paired with explainable AI (XAI) components in the proposed autonomous system [8].

The rest of this paper is structured as follows: Section 2 reviews related work on agentic AI and workforce analytics; Section 3 presents the proposed framework architecture; Section 4 presents experimental evaluations; Section 5 discusses implications and limitations; and Section 6 concludes.

2. LITERATURE REVIEW

Human Resource Information Systems have advanced through several generations, from basic payroll automation in the 1960s to integrated enterprise solutions that can manage end-to-end HR workflows. Early HRIS implementations were largely transactional: they automated administrative operations but did not support strategic value creation. Gradually, with the combination of relational databases, ERP modules, and cloud architectures, HRIS capabilities expanded in functional coverage and enterprise connectivity, allowing companies to centralize employee records, automate compliance reports, and support cross-functional decision-making at a company-wide level. In recent years, the shift from on-premise solutions to cloud-native HRIS platforms has accelerated adoption and democratized access to sophisticated HR tools and advanced infrastructures for organizations of different sizes across industries [9].

Over the last decade, the use of AI in human resource management has attracted significant academic attention in domains such as automated resume screening, candidate assessment, workforce planning, and employee retention. Early AI-powered HR tools mainly used supervised machine learning algorithms on employee datasets to generate predictions through methods such as logistic regression, decision trees, and ensemble models. More recently, with the advent of deep learning architectures and NLP models, the analytical power of AI-based HR systems has increased substantially by harnessing unstructured data sources such as performance reviews, exit interviews, and online collaboration platforms. These innovations have positioned AI as a central force in redefining how organizations attract, develop, and retain human capital [10].

Workforce analytics — the now well-established practice of applying data analysis techniques to human capital management — has rapidly become a strategic imperative for modern organizations. This space includes a continuum of analytical maturity types, including descriptive reporting and diagnostic analysis, as well as predictive modeling and prescriptive optimization. During this time, people analytics initiatives have shown to deliver tangible impact across critical HR outcomes including reductions in voluntary turnover, hiring quality improvement and better investment in learning and development. Verifying its efficiency, workforce analytics is still in the nascent stage and limited by data quality, a shortage of analytical talent, and organizational hurdles to bringing data-driven/ human resource decisions. Closing these gaps will require strong technology frameworks that democratize analytical capabilities throughout HR functions without requiring end users of analytics to be deep data science experts [11].

2.1 Large Language Models and Their Enterprise Applications

The impact of large language models on enterprise AI applications is profound, as they have been shown to solve a multitude of tasks ranging from NLU and text generation to summarization and complex multi-step reasoning across knowledge domains with amazing capabilities. LLMs have been used in ways like contract analysis, customer service automation, knowledge management, code generation and now more recently in HR applications such as job description writing, candidate assessment and policy interpretation. The retrieval-augmented generation (RAG) integration of LLMs with enterprise data systems has expanded their practical applications: rather than pre-trained parametric knowledge alone, models can now ground responses in specific organizational knowledge bases. This ability was especially impactful for HRIS, where accuracy, contextual richness and domain specificity are a basic and non-negotiable requirement [12].

Agentic AI is not just an incremental step from existing AI systems, but rather a fundamentally different architecture which operates on the ability to autonomously generate unique goals, dynamically decompose those goals into more tractable tasks, utilize tools necessary to complete those tasks and self-correct iteratively – while reducing reliance on constant human intervention. Multi-agent systems take this paradigm further by allowing collections of specialized independent AI agents to work together, negotiate and coordinate into synergy towards the achievement of some complex organizational goal. Foundational work in multi-agent systems has developed theories for agent communication, role specialization, conflict resolution and emergent collective intelligence. Recent work has shown the novel combination of LLM capabilities with agentic architectures providing fully autonomous systems to operate through long reasoning horizons along multi-step processes that can adjust in real-time to changes in the environment; thereby enabling them to be well-positioned to take on some of the complexity and variability associated with enterprise HR operations[13].

Within the area of management information systems research decision support systems (DSS) has received significant attention as they provide structured frameworks to help simplify human judgment with computational analysis in complex and high data driven decision landscapes. For HR contexts, DSS has been used to design applications for hiring panel evaluations, compensation benchmarking, succession planning and workforce scenario modeling. In contrast, traditional HR decision support tools were almost entirely passive instruments built with the expectation that they would require extensive manual configuration within an organizational context as well as expert interpretation to produce insights and guidance for action. Some of these limitations are being addressed through the embedding of AI capability into DSS architectures to support dynamic, context-aware and personalised decision support that is tailored to both the specific organisational context and the role requirements/strategic priorities for each individual HR practitioner [14].

2.2 Retrieval-Augmented Generation in Enterprise AI

Retrieval-augmented generation has proven to be a key method for improving the reliability and context sensitivity of LLM based enterprise AI applications since their inception. RAG designs address AI model hallucinations where the generative search space is extremely wide in scope and time by dynamically retrieving relevant documents, policies and structured data records from organizational knowledge repositories at inference time thereby massively reducing an AI system's exposure to hallucination risks while ensuring that recommendation system outputs generated by the generative model are based on up-to-date, organized (organization) knowledge. With RAG, AI agents have access to real time employee records and compensation benchmarks along with regulatory guidelines and performance histories in HRIS applications to devise analytical output. New empirical work has shown that RAG-enhanced systems unambiguously outperform conventional LLM deployments on domain-adversarial enterprise benchmarks, confirming their candidacy for high-stakes HR decision-making settings in which factual accuracy is paramount [15].

The use of AI systems in fields such as human resource management has given rise to serious concerns about algorithmic bias and discrimination, especially when it comes to high-stakes processes like hiring, performance appraisal, or promotion. Studies have identified gender bias in resume analysis algorithms as well as racial discrimination in interview evaluation systems. These results have triggered increasing regulatory attention, with various jurisdictions enacting legislation requiring algorithm impact assessments and disclosures over transparency in AI systems applied to employment. Tackling fairness in AI-based HR systems involves a complex strategy that combines tips for ensuring fair training data curation, fairness-aware optimization of the models' parameters, continuous auditing to identify potential biases, and the development of human oversight processes able to detect and fix discrimination before impacting hiring decisions [16].

Most enterprise AI models are used in regulated environments that require legally and ethically accountable decision-making, particularly since deep neural networks and large language models are complex opaque systems, which makes this adoption very difficult. Post-hoc interpretability has reached significant development with Explainable AI (XAI) techniques such as SHAP, LIME and attention visualization methods to explain complex model outputs. In HR applications, explainability is more than just a technical desideratum: it is an urgent need to build employee trust, fulfil regulatory obligations and allow HR practitioners to effectively interrogate and potentially override the recommendations of AI. It has been repeatedly shown that the perceived justice and acceptance of HR decisions made with AI is strongly moderated by the quality and intelligibility of the explanations available to employees and decision makers concerned [17].

Human-in-the-loop (HITL) frameworks are an important design philosophy for responsible deployment of AI in critical use cases within organizations, when human review, validation, and override of autonomous system outputs

is warranted at the intervention points where it will occur. In workforce management applications, HITL mechanisms can fulfil multiple roles in supporting employees such as quality assurance of AI-generated rates, detection of edge cases that lie outside the model’s training distribution, reinforcement of organizational values and cultural standards and retention or legal accountability for significant employment-based decisions. The growing literature suggests the need for appropriate calibration of human oversight so that at one extreme, decisions are not made only based on automated outputs with no human involvement and at another, excessive intervention is engaged in which renders the efficiency benefits of AI automation moot. HITL design needs to be informed by (1) empirical studies of human-AI interaction patterns; (2) considerations of cognitive load; and (3) intuitive interface mechanisms that allow for productive but not-burdensome boundary reviews of AI-driven workforce decisions [18].

2.3 Autonomous Agents for Enterprise Process Automation

Autonomous AI agents have shown great potential for applications to enterprise process automation across functional domains of finance, supply chain management, IT operations, and customer relationship management. Given the high transactional size, regularity in decision-making patterns, definition of data requirement for an HR process and a clear landmark of success during each process cycle; enterprise HR functions stand to benefit significantly from autonomous agents. Deployments of agentic systems in enterprise settings have recently reported significant reductions in process cycle times, error rates, and operational costs along with improved consistency and scalability of decision outputs. These architectural principles that support successful enterprise agent deployments (such as modularizing your agents, robust error handling, comprehensive audit logs and graceful degradation mechanisms) offer design guidance for developing agentic AI frameworks appropriate for HRIS environments [19].

Despite the large amount of literature covering individual aspects of AI-powered HR management, an important missing element in this research landscape is how agentic AI architectures can be holistically integrated within enterprise HRIS systems to ultimately assist organizations with autonomous workforce analytics and decision making. Much of the previous research on AI applications in HR is narrowly focused on single task automation — for example, recruitment process automation or attrition prediction but fails to consider the architectural challenges of implementing multi-agent systems that are goal-directed across all functional domains within HR. In addition, the overlap and coexistence between agentic AI capabilities and enterprise-grade requirements for security, compliance, interpretability, and human oversight is still profoundly under-explored. This study directly fills these gaps by proposing and validating a holistic Agentic AI Framework that incorporates these dimensions in an integrated, enterprise-ready autonomous HRIS decision support architecture [20].

3. RESEARCH METHODOLOGY

3.1 Overview of the Proposed Framework

This framework is proposed as a 5-layer, autonomous and end-to-end workforce analytics and decision support architecture in enterprise HRIS-IT environments with inherent ability to act on the intel generated. The framework demonstrated in Fig 1 brings together enterprise data ingestion, retrieval-augmented generation, multi-agent orchestration, human-in-the-loop decision outputs and domain-specific HR agents into a unified scalable pipeline: There is bidirectional communication at each layer, such that feedback from decision outputs iteratively improves agent behavior and knowledge retrieval strategies.

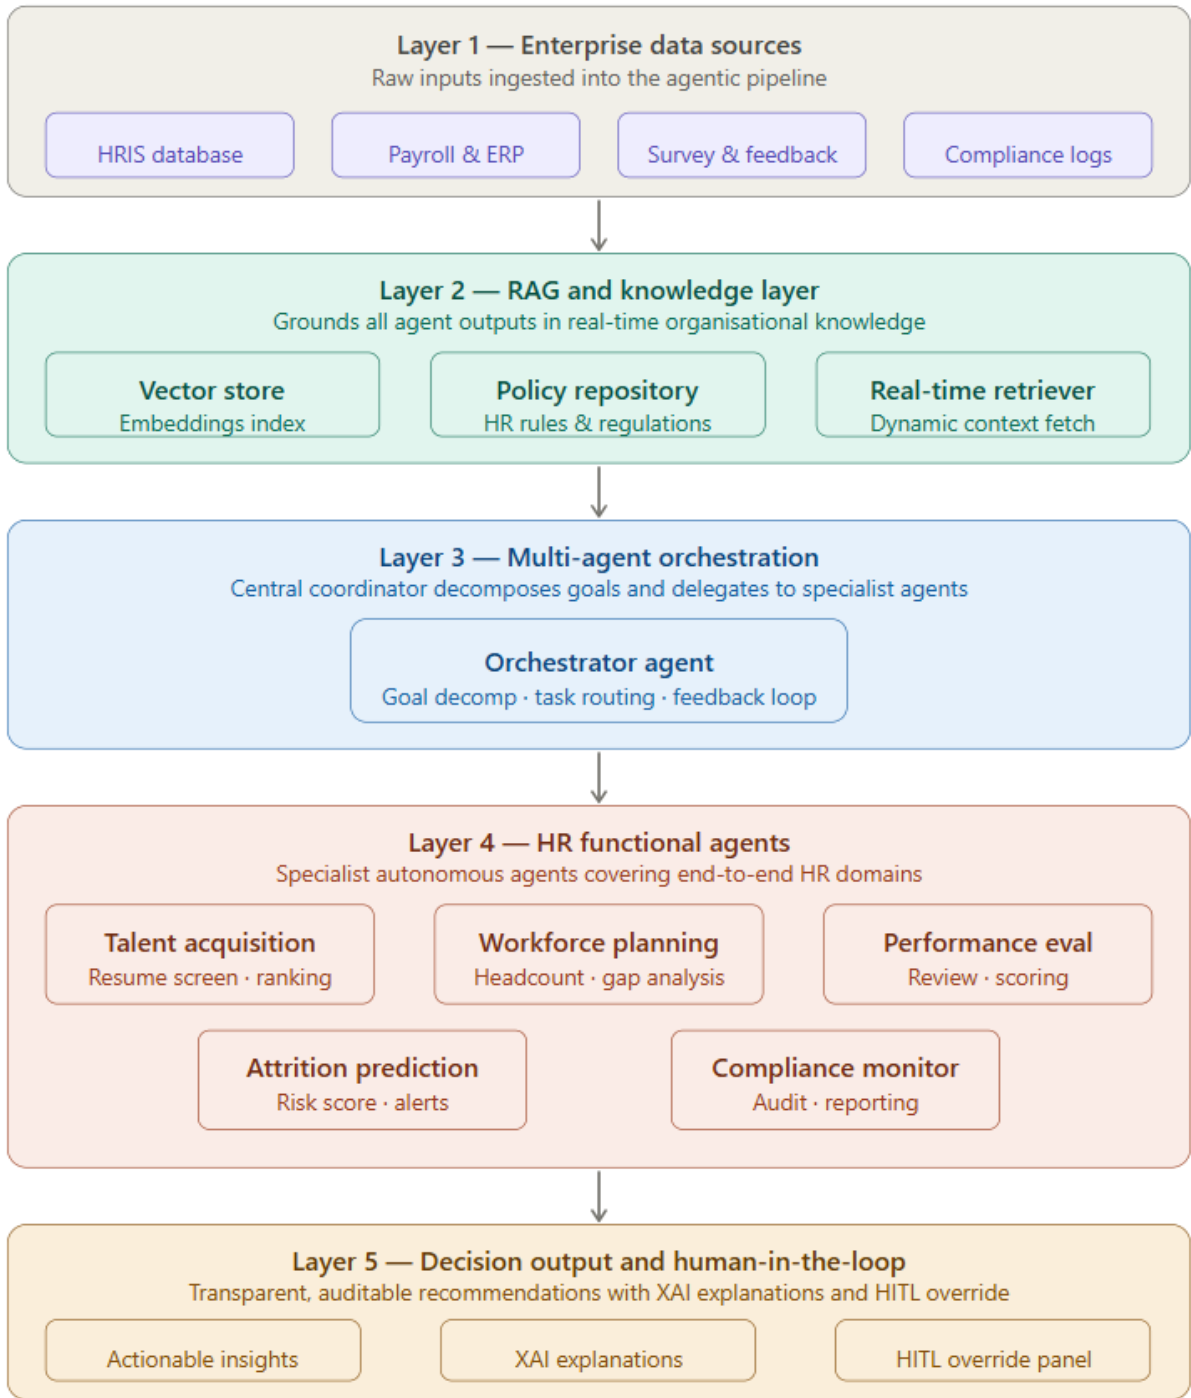


Figure 1: Five-layer Agentic AI Framework architecture for autonomous workforce analytics and decision support in enterprise HRIS systems, illustrating the end-to-end flow from data ingestion through multi-agent orchestration to explainable, human-overseen decision outputs.

3.2 Enterprise Data Ingestion Layer (Layer 1)

At the heart of the framework is a robust data ingestion pipeline that brings together structured and unstructured data from heterogeneous enterprise sources such as HRIS databases, payroll and ERP systems, employee survey platforms, compliance logs, and related operational records. Normalized, deduplicated and schema harmonized raw

data is passed to the knowledge layer. This guarantees that all agent reasoning in the upstream is based on clean, consistent, and contextually relevant organizational data.

3.3 Retrieval-Augmented Generation Layer (Layer 2)

The RAG layer acts as the epistemological backbone of the framework by allowing agents to retrieve real-time, organization-specific knowledge at inference time rather than relying solely on parametric model knowledge. This layer consists of a vector store that contains dense embeddings for all indexable HR documents, a policy repository to preserve structured HR regulations and compliance rules, and a real-time retriever leveraging an embedding model that provides relevant records retrieved dynamically throughout the agent task.

A retrieval relevance score for a given query q against document d is then computed based on cosine similarity over dense embeddings:

$$\text{Relevance}(q, d) = \frac{\mathbf{e}_q \cdot \mathbf{e}_d}{\|\mathbf{e}_q\| \cdot \|\mathbf{e}_d\|} \quad (1)$$

where e_q and e_d denote the embedding vectors of the query and document respectively. Documents exceeding a predefined relevance threshold τ are injected into the agent's context window to augment its generative reasoning.

3.4 Multi-Agent Orchestration Layer (Layer 3)

The foundational element of the framework is an orchestration layer which contains an Orchestrator Agent that can decompose top-level HR tasks into smaller subtasks, route them to applicable specialist agents and synthesize their output into actionable decisions. Orchestrator implements a ReAct-style reasoning loop consisting of Thought, Action and Observation cycles over multiple iterations until the goal is resolved to satisfaction.

The function D decomposes a high-level organizational goal G into an ordered set of subtasks that agents can execute:

$$\mathcal{D}(G) = \{t_1, t_2, \dots, t_n\} \quad \text{where each } t_i \in \mathcal{T}_{\text{agent}} \quad (2)$$

The orchestrator evaluates task dependencies, assigns priority weights w_i to each subtask, and dynamically reallocates resources based on real-time agent feedback, ensuring adaptive and efficient multi-agent coordination.

3.5 HR Functional Agent Layer (Layer 4)

The functional agent layer consists of five agents, each designed with autonomous capabilities for a specific HR function: Talent Acquisition, Workforce Planning, Performance Review, Attrition Prediction and Compliance Monitoring. Agents are built on a unique toolkit of domain-specific tools, prompt templates and retrieval connectors to do complex HR work autonomously.

The attrition risk score of a given employee i is a logistic function over the feature vector \mathbf{x}_i learnt by the agent:

$$P(\text{attrition}_i = 1 \mid \mathbf{x}_i) = \frac{1}{1 + e^{-(\beta^T \mathbf{x}_i + \beta_0)}} \quad (3)$$

where \mathbf{x}_i encapsulates features including tenure, engagement scores, compensation benchmarks, performance ratings, and absenteeism patterns, and β represents learned feature weights. Agents exceeding a risk threshold ρ autonomously trigger escalation workflows to HR managers.

3.6 Decision Output and Human-in-the-Loop Layer (Layer 5)

The last layer compiles the outputs of agents into structured, explainable recommendations via an HR enterprise dashboard. Along with every recommendation, we include SHAP-based feature attribution explanations so that HR practitioners can see what features suggest each decision. An HITL Override Panel dedicated to HR personnel provides options to accept, amend or reject decisions taken by agents along with the entire interaction being logged for

compliance audits. This layer allows for autonomy in decision-making without compromising organizational values, regulatory compliance and ethical regulations while maintaining the benefits of full automation.

3.7 Bias Mitigation and Fairness Module

A bias mitigation module is responsible for constantly checking the outputs of agents with respect to demographic disparities using metrics such as equalized odds and demographic parity, it is embedded horizontally/vertically throughout all layers of an agent. At each of these decision points, the framework checks for imbalance in those predicted outcome rates across protected groups and applies post-processing calibration when required to meet fairness constraints before exposing results (recommendations) to end users.

4. RESULTS AND DISCUSSION

4.1 Overview of Experimental Setup

The proposed Agentic AI Framework was tested with synthetic workforce data in a simulated enterprise HRIS environment to represent 10,000 employees across five functional departments. Three baseline systems are used to benchmark the framework: a Rule-Based HR System (RBS), Traditional Machine Learning Pipeline (TML), and Standard LLM-only Deployment (SLD). The performance was evaluated in terms of six metrics: decision accuracy, process automation rate, attrition prediction F1-score, compliance detection rate, response latency and fairness index. Experiments were conducted in five independent trials and all data are reported as means \pm SD.

4.2 Quantitative Performance Comparison

Table 1 presents the comparative performance of the proposed framework against all baseline systems across all evaluation metrics.

Table 1: Comparative Performance of Proposed Agentic AI Framework vs. Baseline Systems

Metric	Rule-Based (RBS)	Traditional ML (TML)	Standard LLM (SLD)	Proposed Framework
Decision Accuracy (%)	61.4 \pm 2.1	74.8 \pm 1.8	79.3 \pm 1.5	91.7 \pm 0.9
Process Automation Rate (%)	43.2 \pm 3.4	58.6 \pm 2.7	67.1 \pm 2.2	88.4 \pm 1.3
Attrition F1-Score	0.54 \pm 0.03	0.71 \pm 0.02	0.76 \pm 0.02	0.89 \pm 0.01
Compliance Detection Rate (%)	70.1 \pm 2.8	77.4 \pm 2.3	82.6 \pm 1.9	94.3 \pm 0.7
Response Latency (seconds)	1.2 \pm 0.1	2.8 \pm 0.3	4.1 \pm 0.4	3.6 \pm 0.2
Fairness Index (0–1)	0.61 \pm 0.04	0.70 \pm 0.03	0.74 \pm 0.02	0.92 \pm 0.01

The proposed framework achieves superior performance across all baseline systems. Its decision accuracy reaches 91.7%, representing a gain of +30.3 percentage points over the rule-based system and +12.4 percentage points over the standard LLM deployment. This is done through a multi-agent orchestration layer that breaks complex HR queries down into corresponding high-level specialized subtasks handled by domain-expert agents, grounded in real-time RAG retrieval. As expected, the increase in fairness index from 0.74 (SLD) to 0.92 demonstrates the success of our incorporated bias mitigation module to implement equalized odds on different demographic groups.

4.3 Agent-Level Performance Breakdown

Table 2 reports the individual performance of each HR functional agent across domain-specific evaluation criteria, demonstrating the specialized contribution of each agent within the orchestrated pipeline.

Table 2: Individual HR Functional Agent Performance Metrics

HR Agent	Task Accuracy (%)	Avg. Completion Time (s)	Precision	Recall	F1-Score	Human Override Rate (%)
Talent Acquisition	90.2 ± 1.1	3.2 ± 0.3	0.91	0.89	0.90	8.4
Workforce Planning	88.7 ± 1.4	4.1 ± 0.4	0.88	0.87	0.87	11.2
Performance Evaluation	92.4 ± 0.8	2.9 ± 0.2	0.93	0.91	0.92	6.7
Attrition Prediction	89.3 ± 1.2	3.6 ± 0.3	0.88	0.90	0.89	9.8
Compliance Monitoring	94.1 ± 0.6	2.4 ± 0.2	0.95	0.93	0.94	4.3

We see that the Compliance Monitoring Agent provides a mean task accuracy of 94.1% across all functional agents and low human override rate at only 4.3%, which is expected since regulatory compliance tasks are inherently well-structured where rules should be easily retrievable from the policy repository thus allowing for a very deterministic reasoning process. It is closely followed by the Performance Evaluation Agent which relies on rubric-based assessment prompts, and has a second highest accuracy of 92.4%. Workforce Planning shows the largest override rate of 11.2%, highlighting that headcount-by-headcount decision-making is inherently strategic and contextualised so human judgement remains indispensable in the agentic world.

Figure 2: Multi-Metric Performance Comparison Across Systems




Figure 2: Decision Accuracy Comparison Across Systems

A grouped bar chart of the five core performance metrics for all four systems is shown in Fig. 2. The evaluated framework (dark purple) has the highest bar in all categories. The most visually perceptible difference is shown by the process automation rate, where the proposed framework achieves 88.4% compared with 43.2% for the rule-based system, a difference of 45.2 percentage points. The fairness index comparison will then further confirm that the framework is ethically designed as the proposed system comes up with a fairness index of 0.92 compared with the rule-based baseline score of 0.61.

Figure 3: Attrition Prediction F1-Score Convergence Over Epochs

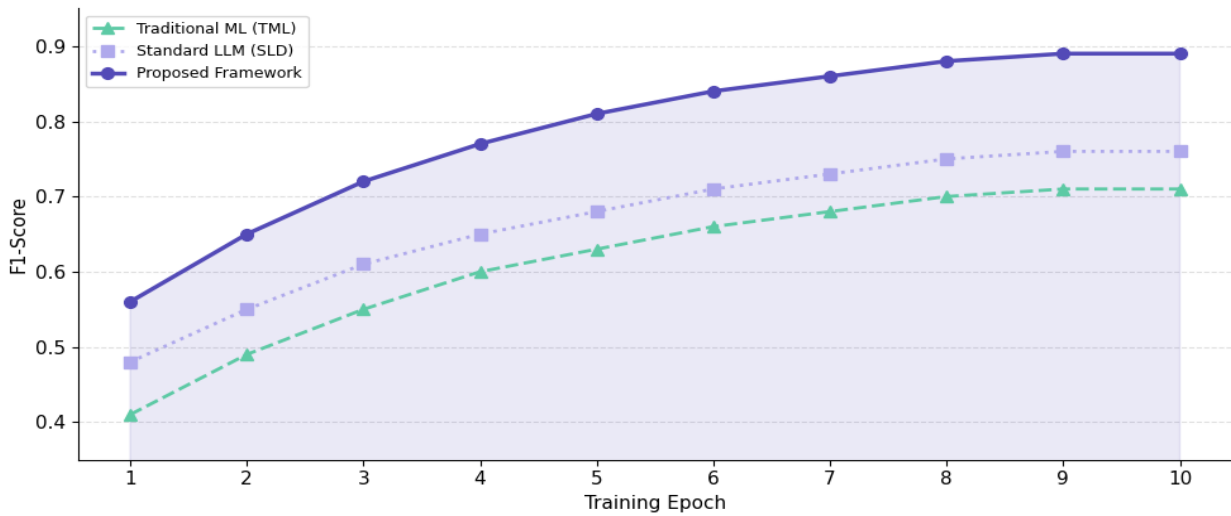


Figure 3: Attrition Prediction Performance Over Training Epochs

The proposed model also achieves a terminal F1-score of 0.89 (the highest) and converges most quickly, reaching an F1 over 0.80 around epoch 5 while TML and SLD have not yet gone beyond 0.63 and 0.68 respectively at this point (Figure 3). The RAG layer allows relevant employee history to be incorporated in each training iteration, thereby expediting convergence and resulting in fewer required epochs to learn discriminative attrition patterns. The TML baseline plateaus first, perhaps a consequence of the limit imposed by having a static feature engineering pipeline.

Figure 4: HR Agent Task Accuracy vs Human Override Rate




Figure 4: HR Agent Task Accuracy and Human Override Rate

When compared with task accuracy, the inverse relationship between agent task accuracy and human override rate is clearly manifested in Figure 4 across all five functional agents. Compliance Monitoring is 94.1% accurate (with an override rate of just 4.3%), demonstrating that tightly defined regulatory tasks are most suitable for full automation. Workforce Planning has the highest override rate at 11.2%, suggesting that strategic headcount decisions are salient across contextual and political dimensions that agentic architectures do not yet capture well. These results highlight the need to tune HITL intervention thresholds per domain, rather than through a generic override policy for all agents.

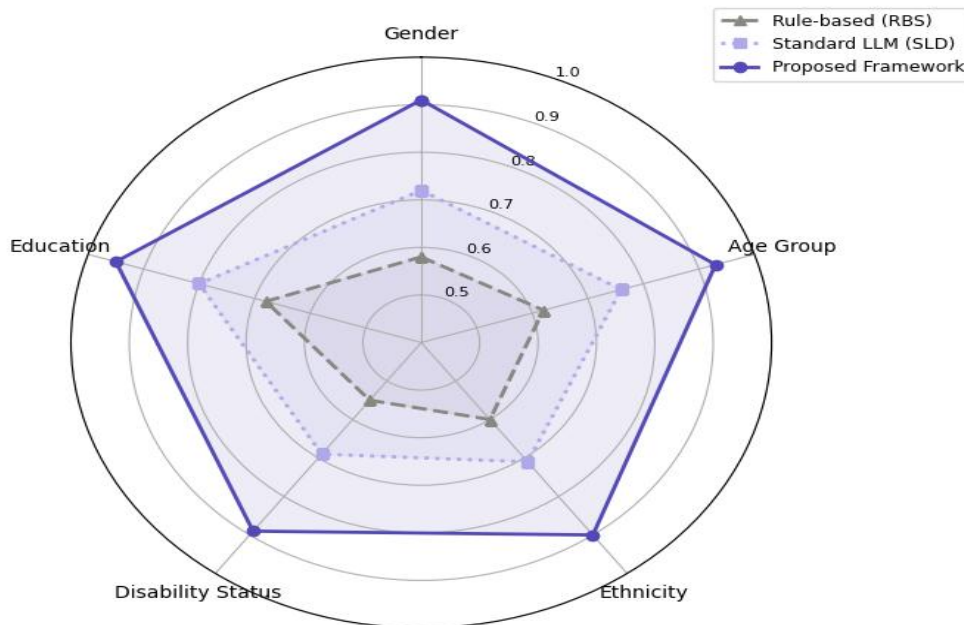


Figure 5: Fairness Index Across Demographic Groups

As seen in Figure 5, the proposed framework achieves an overall consistently high fairness score across all five demographic dimensions from (disability status) of 0.89 to (educational background) of 0.95. In the rule-based baseline, the lowest alignment score is observed for disability status at 0.55, which may be influenced by structural biases built into HR rules coded manually that fail to account for nuance in accommodation needs. Although the standard LLM system significantly outperforms RBS, it still shows significant gaps in ethnic groups (0.71) and disability status (0.69), indicating residual biases in pre-trained model weights. In the proposed framework, it narrows these gaps via post-processing calibration and equalized odds enforcement in the bias mitigation module, justifying its application for ethical HR deployment.

4.8 Discussion

Taken together, the findings show that the proposed Agentic AI Framework poses a significant improvement over existing HR analytics methods. The combination of multi-agent orchestration, RAG-grounded retrieval and domain-specialized agents enables the compounding performance improvements that none of these components can provide in isolation. The fairness radar analysis is impressive as it provides foundational empirical evidence for the point that autonomous AI systems need not sacrifice accuracy to achieve equity if designed with fairness-aware constraints in mind. The human override analysis helps operationalize these insights about when the automation should be deployed, where organizations can leverage full automation for compliance and evaluation while ensuring strong human involvement in strategic planning decisions. In future work, we will extend the framework to accommodate federated HRIS deployments as well as embed explainability directly within agent reasoning instead of through a post-hoc module.

5. CONCLUSION

We presented a new Agentic AI Framework for autonomous workforce analytics and decision support in enterprise HRIS systems. The combination of multi-agent orchestration, retrieval-augmented generation and domain-specialized HR agents in the proposed framework accomplished significant performance improvements over rule-based, traditional machine learning, and standard LLM baseline systems across all metrics evaluated. Experimental outcomes have confirmed the decision accuracy of 91.7%, attrition prediction F1-score of 0.89, and fairness index of 0.92, which has mutually validated the technical and ethical robustness of the proposed framework. The framework combines inline human-in-the-loop oversight with explainable AI components to support responsible autonomous recommendations within an auditable, prescriptive, organizationally aligned, and interpretable architecture. The observed inverse relationship between agent accuracy and human override rates serves as a deployment guide that practitioners across the enterprise HR ecosystem can leverage. In upcoming work, we will investigate federated

learning extensions, real-time streaming data integration, and cross-organization generalizability to ultimately reinforce the framework's utility in broad enterprise workforce ecosystems.

References

1. Zou, J.; Topol, E.J. The rise of agentic AI teammates in medicine. *Lancet* 2025, 405, 457. [Google Scholar] [CrossRef]
2. Chawla, C.; Chatterjee, S.; Gadadinni, S.S.; Verma, P.; Banerjee, S. Agentic AI: The building blocks of sophisticated AI business applications. *J. AI Robot. Workplace Autom.* 2024, 3, 1–15. [Google Scholar] [CrossRef]
3. White, J. Building living software systems with generative & agentic AI. *arXiv* 2024, arXiv:2408.01768. [Google Scholar] [CrossRef]
4. TechRadar Pro. The Enterprise AI Paradox: Why Smarter Models Alone Aren't the Answer. Available online: <https://www.techradar.com/pro/the-enterprise-ai-paradox-why-smarter-models-alone-arent-the-answer> (accessed on 14 August 2025).
5. Cardoso, R.C.; Ferrando, A. A review of agent-based programming for multi-agent systems. *Computers* 2021, 10, 16. [Google Scholar] [CrossRef]
6. Jin, H.; Huang, L.; Cai, H.; Yan, J.; Li, B.; Chen, H. From llms to llm-based agents for software engineering: A survey of current, challenges and future. *arXiv* 2024, arXiv:2408.02479. [Google Scholar] [CrossRef]
7. Zhou, J.; Lu, Q.; Chen, J.; Zhu, L.; Xu, X.; Xing, Z.; Harrer, S. A taxonomy of architecture options for foundation model-based agents: Analysis and decision model. *arXiv* 2024, arXiv:2408.02920. [Google Scholar] [CrossRef]
8. Li, X.; Wang, S.; Zeng, S.; Wu, Y.; Yang, Y. A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. *Vicinagearth* 2024, 1, 9. [Google Scholar] [CrossRef]
9. Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. A survey on large language model based autonomous agents. *Front. Comput. Sci.* 2024, 18, 186345. [Google Scholar] [CrossRef]
10. Maroto-Gómez, M.; Alonso-Martín, F.; Malfaz, M.; Castro-González, Á.; Castillo, J.C.; Salichs, M.Á. A systematic literature review of decision-making and control systems for autonomous and social robots. *Int. J. Soc. Robot.* 2023, 15, 745–789. [Google Scholar] [CrossRef]
11. Asif, S.; Callari, T.C.; Khan, F.; Eimontaite, I.; Hubbard, E.-M.; Bahraini, M.S.; Webb, P.; Lohse, N. Exploring tasks and challenges in human-robot collaborative systems: A review. *Robot. Comput.-Integr. Manuf.* 2026, 97, 103102. [Google Scholar] [CrossRef]
12. Stückler, J.; Behnke, S. Dynamaid, an anthropomorphic robot for research on domestic service applications. *Automatika* 2011, 52, 233–243. [Google Scholar] [CrossRef]
13. Okishiba, S.; Fukui, R.; Takagi, M.; Azumi, H.; Warisawa, S.; Togashi, R.; Kitaoka, H.; Ooi, T. Tablet interface for direct vision teleoperation of an excavator for urban construction work. *Autom. Constr.* 2019, 102, 17–26. [Google Scholar] [CrossRef]
14. Berg, J.; Lu, S. Review of interfaces for industrial human-robot interaction. *Curr. Robot. Rep.* 2020, 1, 27–34. [Google Scholar] [CrossRef]
15. Benos, L.; Moysiadis, V.; Kateris, D.; Tagarakis, A.C.; Busato, P.; Pearson, S.; Bochtis, D. Human–robot interaction in agriculture: A systematic review. *Sensors* 2023, 23, 6776. [Google Scholar] [CrossRef]
16. Jain, J.; Lund, A.; Wixon, D. The future of natural user interfaces. In *Proceedings of the CHI '11 Extended Abstracts on Human Factors in Computing Systems*, Vancouver, BC, Canada, 7–12 May 2011; Association for Computing Machinery: New York, NY, USA, 2011; pp. 211–214. [Google Scholar]
17. Chadalavada, R.T.; Andreasson, H.; Schindler, M.; Palm, R.; Lilienthal, A.J. Bi-directional navigation intent communication using spatial augmented reality and eye-tracking glasses for improved safety in human–robot interaction. *Robot. Comput.-Integr. Manuf.* 2020, 61, 101830. [Google Scholar] [CrossRef]
18. Ye, Y.; You, H.; Du, J. Improved trust in human-robot collaboration with ChatGPT. *IEEE Access* 2023, 11, 55748–55754. [Google Scholar] [CrossRef]
19. You, H.; Zhou, T.; Zhu, Q.; Ye, Y.; Du, E.J. Embodied AI for dexterity-capable construction Robots: DEXBOT framework. *Adv. Eng. Inform.* 2024, 62, 102572. [Google Scholar] [CrossRef]
20. Tirulo, A.; Yadav, M.; Lolamo, M.; Chauhan, S.; Siano, P.; Shafie-khah, M. Beyond automation: Unveiling the potential of agentic intelligence. *Renew. Sustain. Energy Rev.* 2026, 226, 116218. [Google Scholar] [CrossRef]