The American Journal of Management and Economics Innovations
ISSN 2693-0811 | Open Access

TYPE Original Research
PAGE NO. 38-45
DOI 10.37547/tajmei/Volume07Issue08-04

Check for updates

AI in Turnover Risk Assessment: Early Warning Algorithms and Employee Retention Strategies

Zvezdilin Anatoly

PhD in Economics, Lomonosov Moscow State University (Russia). Member of the Association for Talent Development (USA)
San Diego, California, USA

OPEN ACCESS
SUBMITTED 21 July 2025
ACCEPTED 25 July 2025
PUBLISHED 12 August 2025
VOLUME Vol.07 Issue 08 2025

CITATION
Zvezdilin Anatoly. (2025). AI in Turnover Risk Assessment: Early Warning Algorithms and Employee Retention Strategies. The American Journal of Management and Economics Innovations, 7(8), 38–45. https://doi.org/10.37547/tajmei/Volume07Issue08-04

COPYRIGHT
© 2025 Original content from this work may be used under the terms of the creative commons attributes 4.0 License.

Abstract: This paper reviews artificial intelligence approaches to predicting the risks of employee turnover and managing strategies designed to retain them. The purpose of the current study is to carry out a systematic review and practical assessment of existing algorithms used as early warnings for personnel turnover in corporate environments and to recommend ways through which the derived models could be incorporated into HR management processes. The relevance of this work is determined by organizations' enormous costs associated with replacing specialists, the rapid growth of the HR analytics market, and the need to shift from a reactive turnover management model to a proactive talent-retention system. The novelty of the research lies in the comprehensive comparison of classical statistical methods (logistic regression, CoxRF) and modern machine learning algorithms (XGBoost, LSTM-RNN, Bidirectional-TCN, graph neural networks) on both proprietary and open datasets, as well as in the incorporation of interpretability criteria (SHAP, LIME), organizational and ethical barriers, MLOps requirements, and EU regulatory standards into the architecture of predictive HR systems. The findings demonstrate that basic statistical models provide a rapid start and clear interpretability on small samples; however, as data volumes grow, gradient boosting emerges as the "gold standard," and recurrent and convolutional networks become preferable for analyzing temporal communications. Graph neural networks improve flight-risk detection quality by accounting for social connections, while interpretability tools enable the translation of a score into a concrete retention plan. The key takeaway is the need for an

integrated approach: starting from detailed data prep and cleanup, building a cross-functional team, setting up an MLOps loop, designing solutions ethically, training end-users, and monitoring success metrics regularly. This paper will be helpful to HR directors, people analytics specialists, AI-in-HR project managers, as well as academic researchers in the field of human capital management.

Keywords: artificial intelligence, turnover prediction, early warning algorithms, employee retention strategies, HR analytics, machine learning, model interpretability, MLOps

Introduction

Unpredictable employee departures remain one of the costliest managerial risks: conservative SHRM estimates indicate that replacing a single specialist cost between half and twice their annual salary due to recruitment, training, and lost productivity, not counting hidden costs such as morale decline and loss of expertise (SHRM Labs, 2023). Globally, such losses amount to billions of dollars: according to PR Newswire, the HR analytics segment alone grew from $3.7 billion in 2024 to nearly $10 billion in 2025, with a substantial share of demand driven by turnover-prediction solutions (PR Newswire, 2025).

The modern labor market intensifies the need for a proactive model: hybrid work formats have expanded behavioral patterns, and the shortage of digital skills has increased talent competition, making every hiring mistake critical. At the same time, organizations have accumulated vast arrays of structured and unstructured signals—from personnel transaction histories to corporate messenger activity—that are suitable for machine learning. “Early warning” algorithms enable managers to intervene precisely months before a potential resignation. Yet, the shift from pilot experiments to full-scale implementation remains hindered: a recent European study by Vlerick Business School found that nearly six out of ten HR directors report minimal AI usage and cite the absence of a clear plan and competencies as the main barriers (Buyens & Quataert, 2025). Thus, the cost of inaction is high, and the technological and market prerequisites are ripe: now is the time for integrated AI analytics and retention strategies to transform turnover management from a reactive function into a source of sustainable competitive advantage.

Materials and Methodology

The study is based on the analysis of 14 key sources, including academic articles on predictive turnover modeling, industry reports on the HR analytics market, implementation case studies, and regulatory reviews. The theoretical framework relies on works about early warning algorithms such as logistic regression and CoxRF–based survival analysis (Ma et al., 2024; Zhu et al., 2019) and the XGBoost gradient booster (Leidner, 2024), LSTM-RNN recurrent networks (Ganapathisamy, 2023), Bidirectional-TCN (Shiri et al., 2025), graph neural networks for social connection modeling (Shiri et al., 2025), and interpretability methods SHAP and LIME (Varkiani et al., 2025).

Industry context is enriched with SHRM Labs reports on replacement costs and turnover risks (SHRM Labs, 2023), PR Newswire data on HR analytics market growth from $3.7 billion in 2024 to nearly $10 billion in 2025 (PR Newswire, 2025), Global Growth Insights research on AI trends in recruitment and training (Global Growth Insights, 2025a; 2025b), and HR Future (2024) and Visier (2024) data on AI tool integration levels among HR directors.

Methodologically, the work comprises three interrelated stages. First, a systematic literature review covering the period 2019–2025 was conducted, selecting English-language publications on turnover prediction and retention model architectures (Ma et al., 2024; Zhu et al., 2019; Leidner, 2024). Second, a comparative analysis of experimental results was performed: contrasting F1-score, AUC, and accuracy of logistic regression, CoxRF, XGBoost, LSTM-RNN, and Bidirectional-TCN on open IBM HR datasets and in reported corporate cases (Ganapathisamy, 2023; Shiri et al., 2025). To assess model robustness to class imbalance, practices such as SMOTE and ROSE were considered.

The third stage involves qualitative content analysis: survey results from Vlerick Business School HR directors (Buyens & Quataert, 2025) helped identify organizational and competency barriers, and a structured case analysis of Hitachi (Kapadia, 2025) and an Italian bank (Varkiani et al., 2025) elucidated MLOps loop implementations, ethical design, and algorithmic transparency practices. Regulatory and ethical requirements were analyzed based on international responsible-AI principles and anticipated EU regulations.

Results and Discussion

Classical statistical methods remain the starting point of any early-warning project for employee turnover. Logistic regression on small samples yields a rapidly interpretable result: in a recent experiment with fine-tuned GPT models, its baseline version achieved an F1-score of 0.78 on the open IBM HR dataset, ranking as the closest contender to more complex ensembles, as shown in Fig. 1 (Ma et al., 2024).

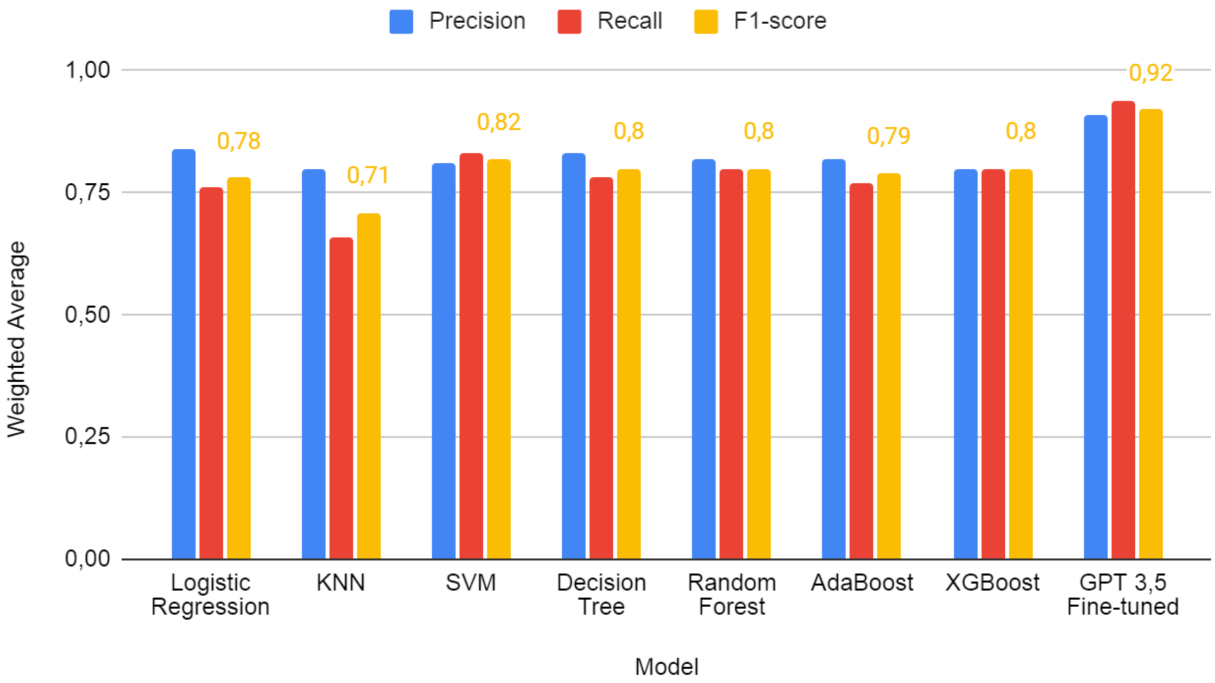

Fig. 1. Comparison of Machine Learning Models (Ma et al., 2024)

When the task shifts from the binary "will leave/will stay" to the question "when exactly might a departure occur," survival analysis comes into play. The CoxRF modification combines Cox proportional hazards with random forests and, operating on longitudinal data from a professional network, predicts not only the fact of departure but also the hazard-function curve shape, thus enabling HR managers to plan the timing of intervention (Zhu et al., 2019).

On real corporate datasets comprising tens of thousands of records, the "gold standard" has become the gradient-boosting family, notably XGBoost. The classic study by Punnoose & Ajit on data from a global retailer showed that XGBoost achieves AUC = 0.86, outperforming SVM and bagging algorithms while training in mere minutes (Leidner, 2024). Its high sensitivity is coupled with built-in feature-selection mechanisms and robustness to class imbalance, especially when paired with SMOTE or ROSE.

Communication and sequencing data demand networks that are aware of temporal sequencing. Thus, a Long Short-Term Memory Recurrent Neural Network optimized by Butterfly achieved 96.7 % accuracy upon the same IBM data, further proving how well recurrent architectures can pick up "micro-signals" from dynamic features like frequencies over time or messaging rhythms. Temporal convolution extended to Bidirectional TCN delivered 89% accuracy with substantially less overfitting in recent work, hence found to be suitable network architectures for streaming HR dashboards, as depicted in Fig.2.

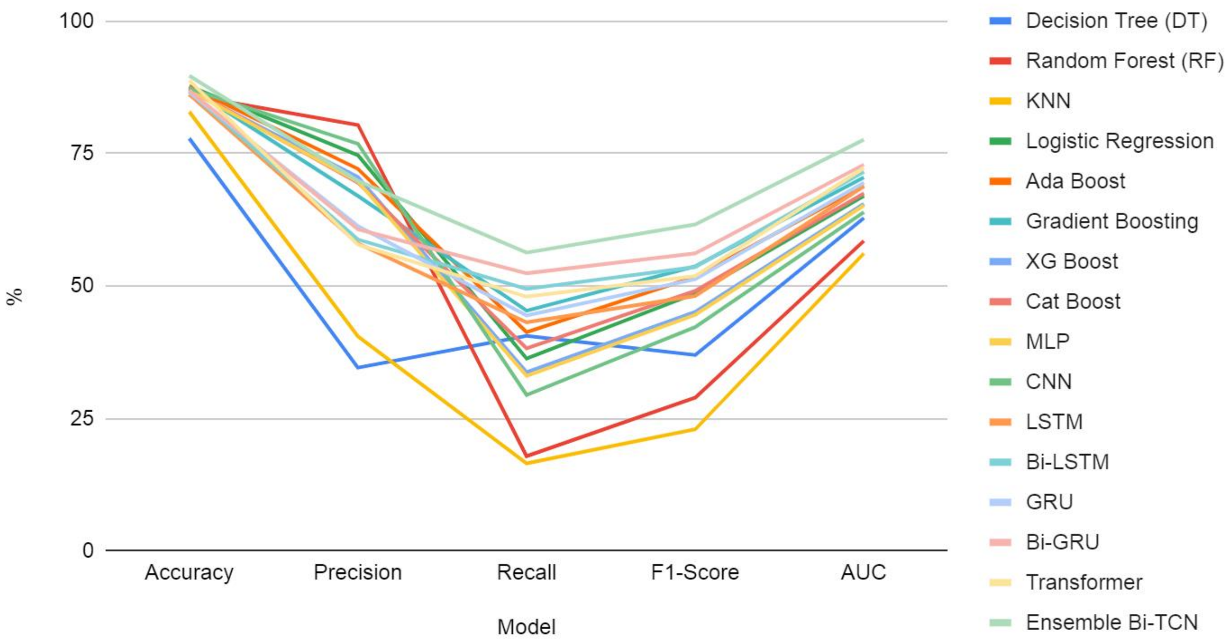

Fig. 2. Performance of models compared to various models on the IBM dataset (Shiri et al., 2025)

Where not only sequences but also interpersonal connections matter, graph neural networks come to the fore. A GCN model that maps tabular HR attributes into a graph of colleague interactions demonstrated that accounting for an employee's structural role—for example, peripheral position in a project network—increases recall for "flight risk" by double-digit percentages at a comparable computational cost (Shiri et al., 2025). This approach is instrumental in hybrid teams, where formal and informal ties diverge.

Regardless of the chosen algorithm, the industry demands transparency. SHAP maps have become the de facto standard: in the Italian-bank case, they were applied not only to rank factors but also to determine their direction of influence—for example, it was found that overtime elevates turnover risk only after the fourth year of tenure, rather than linearly (Varkiani et al., 2025). Combined with local LIME methods, this enables managers to convert an abstract score into a concrete action plan without forfeiting employee trust or ethical compliance.

Artificial intelligence is gradually constructing a continuous "attract → onboard → develop → retain" chain. The forefront of this chain is recruitment: the proportion of companies using AI tools in Talent Acquisition rose from 26% to 53% over one year, indicating a shift from pilots to widespread adoption (HR Future, 2024).

In recruitment, not only speed but also value-alignment quality matters. That is why 66% of recruiting teams already employ AI evaluations that assess skills and "culture fit," increasing hiring accuracy to 48% and reducing early-departure risk (Global Growth Insights, 2025a). Meanwhile, the Global AI Recruitment Market size was USD 0.69 billion in 2024 and is projected to reach USD 0.73 billion in 2025 and further grow to USD 1.07 billion by 2033, registering a compound annual growth rate of 4.89 % during the forecast period from 2025 to 2033, as shown in Fig. 3.

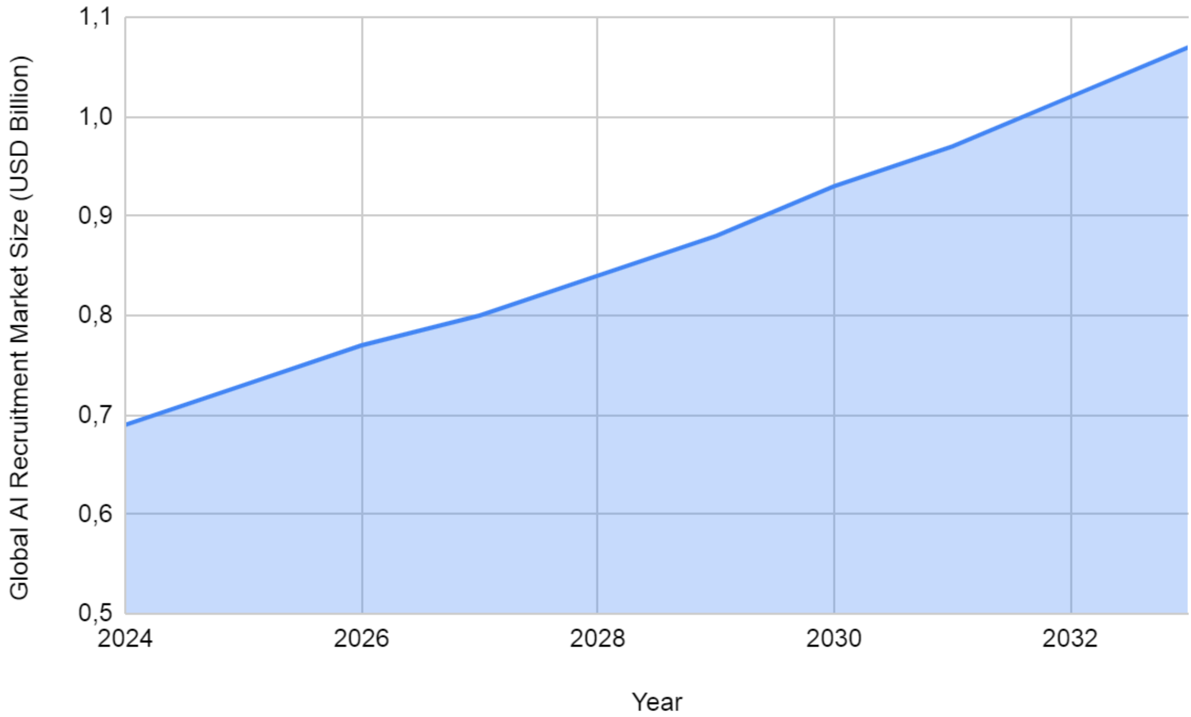

Fig. 3. The Global AI Recruitment Market Size (Global Growth Insights, 2025a)

Classical statistical methods remain the starting point of any early-warning system for employee turnover. Algorithms rank candidates by their probability of long-term success, then automatically generate a short list, leaving space for HR specialists to conduct expert interviews, thus allowing data and human judgment to complement one another.

The next critical moment in the employee lifecycle is onboarding. Real-world cases illustrate the economics of automation: at Hitachi, shifting the process to a chatbot powered by a large language model reduced time-to-productivity by four days and cut HR effort from 20 to 12 hours per new hire, while providing employees with a 24/7 channel for routine questions (Kapadia, 2025). When administrative barriers are removed by automation, the likelihood of premature departure during the first months decreases significantly.

Next comes training. Adaptive Learning Experience Platforms are replacing closed LMS platforms: by 2024, over 60% of Fortune 500 companies had adopted them, and market analysis shows that personalized LXP scenarios increase employee retention by 46% and remain the preferred format for 67% of corporate learners (Global Growth Insights, 2025b). Recommendation algorithms select content aligned with career trajectories and instantly update skill profiles, thereby narrowing the motivational gap between expectations and actual development.

The cycle is closed by contextual people analytics: 87 % of line managers already openly request AI tools for rapid team-related decisions, and 96 % acknowledge that access to up-to-date metrics would boost their confidence (Visier, 2024). Consequently, turnover-risk data become part of the manager's daily interface, closing the loop from early signal to precise action.

Early-warning models are valuable only insofar as they transform a list of high-risk names into a comprehensible map of underlying causes. Cognitive segmentation of motives reveals that identical indicators can mask different intentions: some employees seek new challenges, others external validation, and still others flexibility in scheduling. When each employee's motivational profile is specified, risk data function as an interface for personalized interventions rather than as an alarm background.

From this logically follows a package of measures tied to key turnover drivers. For some teams, financial recognition remains the primary incentive, and therefore, compensation discussions should occur before the external market generates offers. Others prioritize freedom of time and place; with them, discussions of hybrid work or temporary project rotations are more effective. A third group seeks rapidly evolving challenges, making accelerated development programs, micro-courses for in-demand skills, and clear career roadmaps paramount. The algorithm advises the

HR partner on the combination of measures likely to yield the most significant effect, and the partner then tailors actions to the team's cultural context.

To ensure these actions are not one-offs, companies shift the dialogue's entry point: instead of traditional exit interviews, they implement stay interviews conducted before performance reviews or bonus seasons. This format helps managers focus on individual needs using fresh model data rather than post-departure retrospectives. Crucially, the conversation is framed not as a control tool but as a collaborative search for mutually beneficial working conditions.

The cycle culminates in transforming the line manager's role. The algorithm can only highlight red flags; it is the manager-as-coach who translates them into individual development plans, recognizes achievements promptly, and reduces team stress. To support this transformation, companies embed simple coaching prompts within work platforms: feedback scripts, sample stay-interview questions, and micro-recognition suggestions. Thus, the loop between data and behavior is closed: a model signal leads to conversation, the conversation to concrete action, the action to a new engagement metric, which in turn refines the model. The shorter this cycle, the more resilient the organization is to unplanned talent loss.

The main obstacle to implementing predictive turnover-management systems emerges at the conception stage: many organizations perceive the solution as an off-the-shelf product, forgetting that without a long-term plan, even a brilliant model remains isolated. It is necessary to define which specific business questions the algorithm must answer, which metrics will constitute success, how findings will be embedded in the "recruit → develop → retain" cycle, and who will be responsible for action post-signal. Without this target loop, data scientists remain mere suppliers of numbers, and line managers—spectators with no guidance on how to act on those numbers.

Low readiness of the HR function for an analytics culture exacerbates the problem. Interpreting model outputs requires fluency in probabilistic reasoning, and integrating results into processes demands understanding HR-system architectures. These competencies are often dispersed across departments, necessitating a cross-functional team in which data specialists, change-management experts, and legal advisors operate as a unified project group. Such teams must conduct joint workshops, shared training, and regularly crystallize knowledge into methodological guides; otherwise, the project quickly devolves into a collection of isolated solutions incapable of scaling.

Even with strategy and skills in place, organizations face regulatory and ethical constraints. International responsible-AI principles and impending EU regulation classify employee-assessment models as high-risk. This automatically triggers requirements for transparency, human-in-the-loop oversight, bias audits, and avenues for contesting outcomes. Formally, these manifest as documentation of processing purposes, bias-testing protocols, and intervention procedures. Still, in practice, they become a matter of trust: employees must understand which data is used, why, and how this affects their careers. Consequently, companies increasingly establish internal ethics committees, publish public algorithm-governance principles, and appoint an "AI ombudsman" to handle employee inquiries.

Data quality is directly linked to transparency. Information sources about employees have traditionally been kept in disparate systems, resulting in incomplete attributes with significant duplication and outdated information. If the date of last promotion or exact job title is not harmonized between the core HR system, the project-tracking system, and the learning platform, then the model starts building its logic based on an inconsistent view. In the best case, this reduces forecast accuracy; in the worst, it generates false alarms and undermines managers' trust in analytics. An effective response includes a centralized attribute glossary, automatic anomaly checks, regular "data–prediction–outcome" validation cycles, and detailed lineage tracking so that each metric can be traced back to its source.

High-quality data alone does not guarantee prediction stability over time. Business needs evolve, new forms of hybrid employment emerge, teams undergo mergers, and all these factors induce model drift: its rules become stretched between the old and the new reality. Managing drift risk requires a complete MLOps loop: automatic metrics monitoring, retraining triggers, test environments for version comparison, and a formalized reaction to deviations. Although such practices are widespread in financial and technology domains, HR seldom engages with them; without a similar approach, solution longevity remains in question.

Finally, any attempt to algorithmically model people touches on psychological aspects: fear of surveillance and concerns about discrimination. The level of resistance depends directly on how early and transparently the company initiates dialogue. Practice shows that including employee representatives, unions, and community-group members in the project team reduces tension and helps adapt communication language, emphasizing development opportunities rather than "monitoring." When employees see that data is used to offer suitable courses, flexible schedules, or career steps, trust in analytics grows and defensive reactions disappear.

Thus, implementation barriers lie not in the technologies themselves but in organizational readiness, regulatory literacy, and ethical maturity. If strategy is integrated, competencies distributed, rules transparent, and data clean and verifiable, risks shift from blocking to manageable. This opens a way for a roadmap in which the AI system would be a source of reasoned decisions, not an experimental showcase. Breaking these barriers should start with a clear roadmap that kicks off with a full data inventory. The HR leader should piece together a unified registry of all HR, learning, and operational sources; judge their currency, completeness, and legal status; then pick a minimal set of features sufficient to build a pilot model. The key principle at this stage is "small but reliable": it is better to restrict oneself to a dozen well-cleaned attributes and obtain a rapid prototype than to try to cover the entire heterogeneous system landscape and drown in cleaning. Deploying the first model version in a test environment with real users is essential not so much to verify mathematical accuracy as to assess the practical applicability of outputs in live processes.

The next step is the ethical design of the solution, which a single function cannot realize. In addition to HR analysts, the cross-functional team should include data specialists, information-protection lawyers, operations-change experts, and employee representatives. Together, they define processing objectives, fairness criteria, algorithm-review procedures, and feedback channels. Such a composition not only distributes responsibility but also ensures diverse perspectives, so that the model from inception accounts for both business value and the social dimension.

After this preparatory work, the continuous "data → insight → action → outcome" loop is launched. For it, one must formalize rules by which new records enter the reservoir, how often and by whom retraining occurs, which risk thresholds trigger which type of intervention, and how impact is measured. The loop only becomes operational when each transition between stages is automated or prescribed: analytics becomes an alert, an alert becomes a task, a task becomes a completed intervention, the outcome becomes updated data, and the cycle repeats.

Finally, the initiative's sustainability depends on end-user proficiency. Managers and HR partners must understand the model's logic, interpret key factors, and know which tools to employ in response to a signal. To this end, a program of brief modules with case studies, interactive simulations, and access to an "FAQ" reference guide is developed. Its goal is not merely to transfer knowledge but to change habits: the model must become a support for daily decisions, not an exotic add-on. Concurrently, a set of success metrics is introduced, capturing not only turnover reduction but also reaction speed, the proportion of managers using recommendations, and employee satisfaction with intervention quality. With regular monitoring of these indicators, the implementation roadmap transforms from a one-off project into a self-renewing system that adapts to new data and the organization's strategic goals.

Conclusion

It shows that the newest methods of machine learning and artificial intelligence analytics can greatly improve the accuracy of predictions regarding turnover risk, thereby repositioning the HR function from being reactive to an agent in a proactive retention system. The classical statistical algorithms-logistic regression and CoxRF-based survival analysis-have served as reliable starting solutions by providing interpretable forecasts at the pilot project's early stages. As data volume and complexity increase, gradient-boosting models (XGBoost) emerge as the new gold standard, while communication sequences are captured using recurrent and convolutional neural networks (LSTM-RNN, Bidirectional-TCN). Graph neural networks further enhance prediction quality by introducing social connections within teams, while interpretability tools (SHAP, LIME) ensure transparency and trust in algorithmic mechanics.

The most important conclusion is the necessity of a comprehensive approach to implementing predictive HR systems. Beyond selecting and tuning models, critical success factors include data preparation and cleaning, creation of a cross-functional team (HR analysts, data scientists, lawyers, and employee representatives), and development of MLOps processes for drift monitoring and timely retraining. Ethical and regulatory requirements—including transparency, bias management, and contestability—must be embedded in the solution's architecture from the outset, necessitating ethics committees and the appointment of an "AI ombudsman."

Finally, initiative sustainability largely depends on end-user readiness: only when users understand model logic, can interpret key factors, and have articulated response procedures do risk data become practical management tools. Developing training modules, simulations, and success metrics (reaction speed, recommendation usage rate, employee satisfaction) allows analytics to integrate into daily processes. It turns the AI-implementation roadmap into a self-renewing system that adapts to new data and the organization's strategic objectives.

References

1. Buyens, D., & Quataert, S. (2025, April 2). *Lack of expertise, time, and budget are the biggest stumbling blocks to implementing AI in HR.* Vlerick. https://www.vlerick.com/en/insights/lack-of-expertise-time-and-budget-are-the-biggest-stumbling-blocks-to-implementing-ai-in-hr/
2. Ganapathisamy, S. (2023). A Long Short-Term Memory with Recurrent Neural Network and Brownian Motion Butterfly Optimization for Employee Attrition Prediction. *International Journal of Intelligent Engineering and Systems*, *17*(1), 183–192. https://doi.org/10.22266/ijies2024.0229.18
3. Global Growth Insights. (2025a). *AI Recruitment Market Trends*. Global Growth Insights. https://www.globalgrowthinsights.com/market-reports/ai-recruitment-market-116973
4. Global Growth Insights. (2025b). *Learning Experience Platform Market Size*. Global Growth Insights. https://www.globalgrowthinsights.com/market-reports/learning-experience-platform-lxp-market-102110
5. HR Future. (2024). *HR.com's Future of AI and Recruitment Technologies 2024-25: Carefully leverage AI to reap greater success in recruitment*. HR Future. https://eightfold.ai/wp-content/uploads/hr_future_of_ai_and_recruitment_technologies_report-1.pdf
6. Kapadia, S. (2025, March 12). *AI is transforming employee onboarding, saving HR days*. Business Insider. https://www.businessinsider.com/generative-ai-employee-onboarding-human-resources-2025-3
7. Leidner, J. L. (2024, May 13). *Challenges and Opportunities of NLP for HR Applications: A Discussion Paper*. Arxiv. https://arxiv.org/pdf/2405.07766
8. Ma, X., Liu, W., & Zhao, C. (2024). Can Large Language Models Predict Employee Attrition? *ArXiv*. https://doi.org/10.48550/arxiv.2411.01353
9. PR Newswire. (2025, June). *HR Analytics Market Size to Surpass USD 9.89 Billion CAGR of 14.9% during 2025-2031*. PR Newswire. https://www.prnewswire.com/news-releases/hr-analytics-market-size-to-surpass-usd-9-89-billion-cagr-of-14-9-during-20252031--302492353.html
10. Shiri, F. M., Yamaguchi, S., & Ahmadon, M. A. B. (2025). A Deep Learning Model Based on Bidirectional Temporal Convolutional Network (Bi-TCN) for Predicting Employee Attrition. *Applied Sciences*, *15*(6), 2984. https://doi.org/10.3390/app15062984
11. SHRM Labs. (2023, September). *Employee Retention Technologies*. SHRM Labs. https://www.shrm.org/content/dam/en/shrm/shrm-labs/documents/222395-SHRMLabs-_-TechStars-Report-TechnologiesRetention_R3-2.pdf
12. Varkiani, S. M., Pattarin, F., Fabbri, T., & Fantoni, G. (2025). Predicting employee attrition and explaining its determinants. *Expert Systems with Applications*, *272*(126575), 126575–126575. https://doi.org/10.1016/j.eswa.2025.126575
13. Visier. (2024, June 4). *Visier's Research Finds 87% of People Managers Asking for AI Tools to Make Their Jobs Easier*. Visier. https://www.visier.com/company/news/hr-generative-ai-for-people-managers-na/
14. Zhu, Q., Shang, J., Cai, X., Jiang, L., Liu, F., & Qiang, B. (2019, August 1). *CoxRF: Employee Turnover Prediction Based on Survival Analysis*. IEEE Xplore. https://doi.org/10.1109/SmartWorld-UIC-ATC-SCALCOM-IOP-SCI.2019.00212