

AI-Driven People Analytics and Employee Retention: Assessing The Role of Data-Driven HR Decision-Making

Dr. Sameer Pawar

Assistant Professor, SVPM's, Institute of Management, Malegaon (Bk.), Baramati

Abstract

This paper examines how artificial intelligence (AI) can enhance employee retention amidst significant turnover issues in modern organizations. It conducts a systematic review of 58 sources, including peer-reviewed articles and practitioner reports, predominantly from 2015 to 2026, and proposes an integrative conceptual framework. The findings indicate a discrepancy between improved predictive accuracy for voluntary turnover and inconsistent retention gains, particularly in organizations with advanced analytics capabilities. Resistance from employees often stems from concerns about fairness, transparency, and surveillance. The study introduces a dual-pathway model: the enabling pathway, where analytics lead to interventions that enhance job embeddedness and retention, and the evaluative pathway, where perceptions of justice and privacy can undermine these benefits. The paper advances nine propositions detailing these relationships and uniquely integrates four existing research areas—HR analytics, algorithmic management, organizational justice, and turnover theory—reconceptualizing retention analytics as a matter of decision quality and legitimacy rather than mere predictive accuracy. It lays out a research agenda for future empirical inquiries and offers practical recommendations for ethical AI deployment in HR decision-making.

Keywords: people analytics; artificial intelligence; employee retention; HR analytics; data-driven decision-making; algorithmic management; turnover prediction; organisational justice.

1. Introduction

Employee turnover remains a persistent challenge for organisations despite decades of research identifying job satisfaction, organisational commitment, perceived alternatives, and quitting intentions as important predictors of employee exits (Griffeth et al., 2000; Mobley, 1977). High turnover can negatively affect organisational performance, productivity, and financial outcomes, while effective human resource practices can help reduce avoidable employee exits (Hancock et al., 2013; Huselid, 1995). However, traditional turnover prediction approaches have limitations because employee attitudes are dynamic and conventional annual surveys often provide only a temporary assessment of employee experiences (Kammeyer-Mueller et al., 2024). Furthermore, employees may leave following sudden or critical events rather than through a gradual decline in job satisfaction (Lee & Mitchell, 1994).

The digitalization of Human Resource Management (HRM) has introduced people analytics as an evidence-based approach to workforce decision-making. People analytics uses descriptive, visual, and statistical analysis of HR

and workforce data to support managerial decisions and demonstrate business value (Marler & Boudreau, 2017). Its effectiveness, however, depends on data quality, analytical maturity, strategic alignment, and appropriate organisational decision-making processes (Nocker & Sena, 2019; McCartney & Fu, 2022).

The emergence of Artificial Intelligence (AI) has further transformed people analytics through machine learning, predictive analytics, natural language processing, and generative AI applications in recruitment, performance management, and employee retention (Tambe et al., 2019; Vrontis et al., 2022). Research demonstrates that machine-learning-based people analytics can support employee attrition prediction and identify factors associated with employee turnover (Yahia et al., 2021). Recent research continues to develop machine-learning and explainable-AI approaches for improving employee attrition prediction.

Nevertheless, AI-driven HRM also creates concerns relating to privacy, algorithmic bias, transparency, employee autonomy, fairness, and trust (Tursunbayeva et al., 2022; Edwards et al.,

2022). Employees may be more receptive to algorithmic evaluation when organisations provide explanations, appeal mechanisms, and meaningful human involvement (Park et al., 2021; Park et al., 2022). Therefore, effective AI-driven employee retention requires not only predictive capability but also ethical governance, transparency, data quality, employee trust, and human oversight.

2. Literature Review

2.1 People Analytics and the Digitalisation of Human Resource Management

People analytics, used interchangeably in much of the literature with HR analytics and talent analytics, denotes the systematic application of data and statistical reasoning to workforce questions. The most widely adopted definition describes it as an HR practice enabled by information technology that uses descriptive, visual and statistical analyses of data related to HR processes, human capital, organisational performance and external benchmarks in order to establish business impact and support data-driven decision-making (Marler & Boudreau, 2017). The definitional emphasis on decision-making rather than reporting is significant: analytics is distinguished from conventional HR metrics by its orientation towards inference, prediction and action. Strohmeier (2020) locates this development within a broader conceptualisation of digital human resource management, in which digital technologies are understood as both instruments and constitutive elements of HRM, and in which the boundaries between the digital and the human elements of HR practice are increasingly difficult to draw. This framing is important for retention research because it implies that the introduction of analytical technology is not neutral with respect to the employment relationship it purports to manage. Adoption research consistently identifies a maturity gradient. Organisations progress from descriptive reporting through diagnostic and predictive analytics to prescriptive optimisation, and many stall at the earliest stages because of fragmented data architecture, weak analytical skills and limited managerial demand for evidence (Fernandez & Gallardo-Gallardo, 2021). More recent work has begun to model the capability configurations that support advancement, proposing that analytics capability be understood as an architecture of data infrastructure, analytical talent, governance and business partnering rather than as a single technology investment (Shet,

2025). Studies of national contexts reach congruent conclusions: applications cluster around recruitment, workforce planning and attrition, enablers include executive sponsorship and data quality, and barriers include legal uncertainty, cost and a shortage of analysts with domain knowledge (Di Lauro et al., 2025). Reviews that engage stakeholders directly argue that many adoption obstacles are relational rather than technical, and that overcoming them requires sustained dialogue between HR, information systems, legal and line-management constituencies (Alam et al., 2026).

The relationship between analytics and organisational performance has been specified more precisely in recent work. McCartney and Fu (2022) show that HR analytics affects performance through improved decision quality rather than through the mere production of metrics, and that the effect is conditioned by analytical maturity and by the organisational routines that convert insight into action. Consistent with this, information systems scholarship conceptualises the business value of people analytics as emerging from an interaction between technological capability, organisational process and managerial interpretation, and warns that value is easily dissipated when analytics artefacts are produced but not used (Bottesch et al., 2025). Bibliometric mapping of the field confirms rapid growth in publication volume while noting persistent fragmentation across sub-themes and a relative scarcity of studies that trace analytics through to employee outcomes (Rana & Sachdeva, 2025). Nocker and Sena (2019) provide an early and still relevant explanation for this pattern: talent analytics is adopted partly for legitimacy reasons, and legitimacy is fragile when analytical claims cannot be substantiated.

Finally, acceptance research explains why deployment does not guarantee use. Studies applying the unified theory of acceptance and use of technology to HR information systems show that performance expectancy, effort expectancy, social influence and facilitating conditions predict managerial and employee adoption, and that these relationships are robust across contexts (Venkatesh et al., 2003; Menant et al., 2021). A meta-analytic validation of the framework in the digital HRM domain reports substantial explanatory power, but also notes that HR-specific acceptance is shaped by trust in the system and by perceived fairness, variables that general technology acceptance

models treat peripherally (Theres & Strohmeier, 2023). These findings anticipate the argument developed later in this paper, namely that the fate of retention analytics depends as much on how the technology is appraised as on what it can compute.

2.2 Artificial Intelligence and Algorithmic Human Resource Management

AI extends analytics in three directions simultaneously: it increases predictive power, it enables automation of judgement, and it generates new forms of interaction between employees and HR systems. Tambe et al. (2019) provided an influential early assessment, arguing that the principal obstacles to AI in HRM are not algorithmic but infrastructural and institutional, including the fragmentation of employee data, the difficulty of constructing valid outcome measures, the ambiguity of constructs such as performance and potential, and the legal and ethical exposure created by automated decisions about people. Their analysis anticipated the current research agenda by identifying fairness, privacy and accountability as core design constraints rather than peripheral concerns.

Meijerink et al. (2021) subsequently formalised algorithmic HRM as a distinct object of study, defining it as the configuration of data, algorithms and human actors that jointly shape HR practices, and arguing that research must attend to the distribution of decision rights, the transparency of algorithmic logic, and the consequences for employee voice. This synthesis has been extended by multidisciplinary work that maps the evolution of algorithmic HRM across business, law, computer science and ethics, and that identifies a persistent asymmetry between technical capability and institutional capacity to govern it (Cameron et al., 2024). Reviews of AI across the HR function report diffusion in recruitment, learning, performance management and workforce planning, alongside recurring risks of bias propagation, depersonalisation of the employment relationship and reduced employee discretion (Vrontis et al., 2022; Madanchian et al., 2023).

Strategic perspectives connect these developments to the resource-based logic of the firm. Kim et al. (2025) argue that algorithmic technologies reconfigure strategic HRM by shifting the locus of competitive advantage towards data assets and analytical capability, while creating new dependencies that are not easily imitated but also

not easily controlled. Dima et al. (2024) show that AI does not replace the HR triad of employees, line managers and HR professionals but rather redistributes tasks and judgement among them, with consequences for role identity and for the perceived legitimacy of HR decisions. Taken together, this literature establishes that AI in HR is best understood as a socio-technical configuration whose outcomes depend on governance, interpretation and employee appraisal, and not as a deterministic productivity technology. That insight provides the starting point for the framework developed in Section 3.

2.3 Theoretical Foundations of Employee Retention

Retention theory has developed through several generations. Early models treated turnover as the terminal step in a sequential process in which dissatisfaction generates thoughts of quitting, evaluation of alternatives and intention, with intention as the strongest proximal predictor of behaviour (Mobley, 1977). Meta-analytic synthesis broadly supports this ordering while demonstrating that the magnitude of attitudinal effects is moderate, and that moderators such as labour market conditions, occupation and alternatives substantially alter observed relationships (Griffeth et al., 2000). The implication is that attitudinal prediction is possible but bounded, which is precisely the gap that richer data environments promise to fill.

The unfolding model challenged the assumption that turnover is always the product of accumulated dissatisfaction. Lee and Mitchell (1994) showed that many exits follow shocks, that is, jarring events that prompt a decision to leave, and that decision paths vary from scripted responses to deliberate comparison of alternatives. This insight has direct implications for analytics: shocks such as an unexpected appraisal outcome, a change in supervisor or an unanticipated algorithmic decision are neither evenly distributed nor easily captured by annual attitude surveys, yet they may be decisive. The job embeddedness construct extended the same tradition by asking why people stay rather than why they leave, identifying fit, links and sacrifice as mechanisms that bind employees to organisations and communities and that predict both intent to leave and voluntary turnover beyond satisfaction and commitment (Mitchell et al., 2001). Embeddedness is

analytically attractive in the present context because several of its components are, in principle, observable in workforce data, and because it suggests that retention interventions should build attachment rather than merely suppress dissatisfaction.

Two complementary theories explain the relational and energetic mechanics of these attachments. Social exchange theory holds that employees reciprocate favourable treatment with discretionary effort and loyalty, so that perceived support from the organisation produces felt obligation and reduces withdrawal (Harden et al., 2018). Organisational justice research specifies the evaluative criteria employees use to judge managerial action, distinguishing distributive, procedural, interpersonal and informational justice, each with distinct relationships to trust, satisfaction and turnover-related outcomes (Colquitt et al., 2001). The relevance to algorithmic decision-making is direct: informational justice concerns the adequacy of explanations provided for decisions, and procedural justice concerns the consistency and correctability of decision processes, both of which are precisely the dimensions on which algorithmic systems are most often criticised. The job demands-resources model, in turn, frames strain and engagement as outcomes of the balance between demands and available resources, so that interventions which increase job resources such as autonomy, feedback and support reduce exhaustion and increase engagement (Bakker & Demerouti, 2017). Conservation of resources theory generalises the same logic by treating employees as motivated to protect and accumulate valued resources, with resource loss spirals producing withdrawal and resource gain spirals producing investment in the employment relationship (Hobfoll et al., 2018).

Bringing these theories together suggests that retention analytics can only work through a small number of levers: it can raise the accuracy with which resource deficits are identified, it can increase the perceived fairness and supportiveness of managerial action, and it can strengthen the fit, links and sacrifice that embed employees in the organisation. Whether it does so depends on how the technology is experienced, which is the subject of the next subsection.

2.4 Predictive Modelling of Turnover: Evidence and Limitations

Empirical work on AI-driven turnover prediction has grown quickly and is summarised in Table 6. Applying deep learning to employee records, Yahia et al. (2021) report improved classification of leavers relative to baseline models and argue that combining structured HR data with richer behavioural indicators improves early detection of attrition risk. Avrahami et al. (2022) use a large internal dataset to show that machine learning can surface turnover determinants that are not prominent in the practitioner literature, and that model outputs can be translated into retention priorities, while cautioning that associations identified in this way are not automatically causal. Srivastava and Eachempati (2021) combine ensemble learning with multi-criteria decision-making to rank employees by intervention priority, an approach that shifts the analytical objective from classification to allocation of scarce retention effort. Yuan et al. (2024) address a methodological weakness that runs through much of this work: workforce data are hierarchically structured, and models that ignore grouping structure overstate generalisability. Their case study demonstrates how grouped-data modelling improves prediction of turnover intention and, importantly, how prediction can be used to test theoretical propositions rather than merely to rank employees. Transparency has been treated as a design variable rather than an afterthought. Chowdhury et al. (2023) embed explainability in turnover models and show that interpretable output changes managerial reasoning and increases the likelihood that algorithmic recommendations are acted upon, while also mitigating trust deficits. Text-based approaches expand the range of signals available for diagnosis. Chittiprolu et al. (2021) analyse thousands of online employee reviews to identify determinants of satisfaction in the hospitality sector, and Ding et al. (2025) combine structural topic modelling with supervised learning to classify employee experience narratives, demonstrating that unstructured commentary can be converted into actionable categories.

Three limitations recur. First, outcome validity: many studies predict turnover intention rather than actual exit, and the intention-behaviour gap is well documented in the turnover literature (Griffeth et al., 2000), so reported performance may overstate practical value. Second, temporal validity: models trained on historical data degrade as labour markets, organisational structures and job designs

change, and few studies report out-of-time validation. Third, decision validity: predictive accuracy is not the same as decision quality, because an accurate risk score can still generate discriminatory, demoralising or self-fulfilling interventions if it is deployed without attention to how employees will experience being classified as a flight risk. This third limitation is the analytical hinge of the present paper.

2.5 Synthesis and Research Gaps

Five gaps emerge from this review. First, evidence on retention outcomes is thin: studies demonstrate predictive performance on intention or on historical exit data, but few link analytics deployment to reductions in actual voluntary turnover in field settings. Second, the enabling and evaluative literatures are disconnected, so that the conditions under which analytical gains survive employee appraisal remain unspecified. Third, capability and maturity conditions are well theorised in the analytics literature but rarely integrated into retention models, producing recommendations that ignore organisational readiness. Fourth, causal identification is largely absent: the dominant designs are predictive and correlational, leaving open the possibility that observed relationships reflect selection, reverse causality or common causes. Fifth, the employee perspective has been studied intensively in recruitment and performance evaluation but rarely in the retention domain, even though retention interventions involve the most intrusive forms of employee data use, including flight-risk scoring, network analysis and sentiment monitoring. The framework developed next responds to these gaps.

3. Conceptual Framework and Propositions

The framework developed here rests on a dual-pathway logic. AI-driven people analytics influences retention through an enabling pathway, in which analytical capability produces more precise and timely diagnosis of retention risk, which in turn supports targeted interventions that restore or augment valued resources and thereby strengthen embeddedness and reduce voluntary exit. Simultaneously, the same technology

activates an evaluative pathway, in which employees form judgements about the fairness, transparency, privacy-respectfulness and trustworthiness of algorithmic decision-making. The evaluative pathway does not simply add to the enabling pathway; it moderates it, because interventions delivered through a system that employees regard as illegitimate are less likely to be accepted, reciprocated or sustained. The framework therefore predicts that identical analytical output can produce opposite retention outcomes depending on the social and institutional conditions of its use.

The central constructs are defined in Table 1, and the theoretical lenses that inform them are summarised in Table 2. AI-driven people analytics capability refers to the configuration of data infrastructure, analytical talent, governance and managerial data literacy that enables an organisation to generate and use workforce insight. Diagnostic precision denotes the accuracy, timeliness and actionability with which retention risks and their antecedents are identified. Decision quality refers to the extent to which an HR decision is informed by relevant evidence, consistent with organisational strategy, justifiable to affected employees and correctable in the light of new information. Data-driven HR decision-making is the routine use of analytical evidence in managerial and HR judgements, as distinct from the production of analytical artefacts. Targeted retention intervention describes a differentiated action, such as workload redesign, development opportunity, compensation adjustment or career conversation, directed at identified employees or groups. Perceived algorithmic fairness captures employees' evaluations of the distributive, procedural, interpersonal and informational justice of algorithmic decisions. Privacy legitimacy refers to the perceived appropriateness of data collection and use relative to employment purposes. Trust in HR analytics is the willingness of employees and managers to rely on analytical output under conditions of uncertainty. Job embeddedness comprises fit, links and sacrifice. Finally, enabling and evaluative pathways denote the two mechanisms described above.

Table 1. Core constructs and working definitions

Construct	Working definition
AI-driven people analytics capability	The configuration of data infrastructure, analytical talent, governance mechanisms and managerial data literacy that allows an organisation to generate and apply workforce insight

	(Shet, 2025; McCartney & Fu, 2022)
Diagnostic precision	The accuracy, timeliness and actionability with which retention risks and their antecedents are identified (Yahia et al., 2021; Avrahami et al., 2022)
Decision quality	The extent to which an HR decision is evidence-informed, strategically aligned, justifiable to affected employees and correctable in light of new information (McCartney & Fu, 2022)
Data-driven HR decision-making	The routine incorporation of analytical evidence into managerial and HR judgements, as distinct from the production of analytical artefacts (Marler & Boudreau, 2017; Bottesch et al., 2025)
Targeted retention intervention	A differentiated action such as workload redesign, development opportunity, pay adjustment or career conversation directed at identified employees or groups (Srivastava & Eachempati, 2021)
Perceived algorithmic fairness	Employees' distributive, procedural, interpersonal and informational justice evaluations of algorithmic decisions (Colquitt et al., 2001; Lavanchy et al., 2023)
Privacy legitimacy	The perceived appropriateness of employee data collection and use relative to legitimate employment purposes (Ebert et al., 2021)
Trust in HR analytics	Willingness of employees and managers to rely on analytical output under uncertainty (Xu et al., 2024; Theres & Strohmeier, 2023)
Job embeddedness	The fit, links and sacrifice that bind an employee to the organisation and community (Mitchell et al., 2001)
Enabling pathway	The sequence from analytical capability through diagnostic precision and targeted intervention to resource gain and retention
Evaluative pathway	The sequence from algorithmic deployment through justice, privacy and trust appraisals to acceptance or resistance

Table 2. Theoretical lenses informing the framework

Lens	Core logic	Application in this framework
Job embeddedness theory	Employees stay when fit, links and sacrifice accumulate (Mitchell et al., 2001)	Specifies the proximal retention outcome of targeted interventions
Unfolding model	Turnover often follows shocks and scripted decision paths (Lee & Mitchell, 1994)	Explains why event-based signals, including algorithmic decisions, matter
Social exchange and justice	Employees reciprocate treatment and judge it against justice criteria (Harden et al., 2018; Colquitt et al., 2001)	Explains how algorithmic decisions translate into reciprocity or withdrawal
Job demands-resources	Strain and engagement depend on the balance of demands and resources (Bakker & Demerouti, 2017)	Identifies intervention levers such as autonomy, feedback and support
Conservation of resources	Employees protect and accumulate valued resources (Hobfoll et al., 2018)	Explains resource gain and loss spirals triggered by retention interventions
Technology acceptance and use	Adoption depends on expectancy, influence, facilitation and trust (Venkatesh et al., 2003; Theres & Strohmeier, 2023)	Explains managerial and employee use of analytics tools
Algorithmic and strategic HRM	Data, algorithms and actors jointly configure HR practice (Meijerink et al., 2021; Kim et al., 2025)	Locates analytics capability within the firm's resource base and governance

The framework yields nine propositions. The first two specify the enabling pathway. Capability determines the quality of diagnosis: organisations with integrated data, analytical skills and governance can identify retention risk earlier and more accurately than organisations without them, and this advantage is amplified by managerial data literacy, because insight that cannot be interpreted cannot be acted upon.

Proposition 1. AI-driven people analytics capability is positively associated with the diagnostic precision and timeliness of retention risk identification.

Proposition 2. Managerial data literacy positively moderates the relationship between diagnostic precision and the adoption of targeted retention interventions.

Propositions 3 and 4 connect intervention to retention through established psychological mechanisms. When interventions are differentiated and address identified resource deficits, they increase perceived support and justice, which in turn raise fit, links and sacrifice, producing lower voluntary turnover. The argument here is deliberately consistent with the job demands-resources and conservation of resources traditions:

what retains employees is not the analytics itself but the resource gain that analytics makes possible. Proposition 3. Targeted retention interventions are positively associated with perceived organisational support and justice, which in turn increase job embeddedness.

Proposition 4. Job embeddedness mediates the negative relationship between targeted retention intervention and voluntary turnover.

Propositions 5 through 7 specify the evaluative pathway and its moderating role. Employees do not respond to analytics as a neutral instrument; they respond to being measured, scored and forecast. Where transparency, explanation, appeal and privacy protection are weak, the evaluative pathway competes with and can negate the enabling pathway, so that the relationship between intervention and retention weakens or reverses.

Proposition 5. Perceived algorithmic fairness and privacy legitimacy positively moderate the relationship between targeted retention interventions and perceived organisational support and justice.

Proposition 6. Perceived opacity and surveillance intensity weaken the relationship between

diagnostic precision and retention outcomes by reducing trust in HR analytics.

Proposition 7. Trust in HR analytics mediates the relationship between perceived algorithmic fairness and employees' acceptance of retention-related decisions.

The final two propositions concern institutional and contextual conditions. Governance quality determines whether the evaluative pathway is managed or neglected, and the relationship between analytics capability and retention outcomes is contingent on organisational context, including size, sector, unionisation and labour market conditions (Table 3).

Proposition 8. Human-in-the-loop decision design, explainability and privacy due diligence strengthen the enabling pathway and weaken the negative effects of the evaluative pathway.

Proposition 9. The relationship between AI-driven people analytics capability and retention outcomes is contingent on contextual factors including organisational size, sector, union presence and labour market tightness.

Table 3. Summary of propositions

Proposit ion	Statement
P1	AI-driven people analytics capability is positively associated with the diagnostic precision and timeliness of retention risk identification.
P2	Managerial data literacy positively moderates the relationship between diagnostic precision and targeted retention interventions.
P3	Targeted retention interventions are positively associated with perceived organisational support and justice, which increase job embeddedness.
P4	Job embeddedness mediates the negative relationship between targeted retention intervention and voluntary turnover.
P5	Perceived algorithmic fairness and privacy legitimacy positively moderate the intervention to support and justice relationship.
P6	Perceived opacity and surveillance intensity weaken the precision to retention relationship by reducing trust.
P7	Trust in HR analytics mediates the relationship between perceived fairness and decision acceptance.
P8	Human-in-the-loop design, explainability and privacy due diligence strengthen the enabling pathway and weaken negative evaluative effects.
P9	The capability to retention relationship is contingent on size, sector, union presence and labour market conditions.

4. Methodology

4.1 Review Design

The study adopts a systematic literature review design informed by the PRISMA reporting convention, combined with conceptual development. A systematic review was selected in preference to an empirical survey for three reasons. First, the research questions concern the mechanisms and boundary conditions of a relationship that has been studied in dispersed

literatures, and synthesis is therefore the appropriate method of inquiry. Second, the field lacks consensus measures of analytics capability and of retention outcomes, so pooling effect sizes would require untenable assumptions of construct equivalence. Third, a systematic review permits the explicit specification of conditions and mechanisms, which is what the research questions demand. The review is complemented by the construction of an integrative framework and the

derivation of testable propositions, a combination appropriate to a field in which theoretical fragmentation rather than evidence scarcity is the principal obstacle to progress.

The review proceeded through five stages: protocol specification; systematic searching; screening against explicit criteria; data extraction and quality appraisal; and thematic synthesis. The protocol,

including the search strings, criteria and extraction template, is summarised in Table 4. The approach follows established practice for reviews in the HR analytics and algorithmic management domains (Tursunbayeva et al., 2022; Alam et al., 2026), and the reporting structure reflects standard expectations of methodological transparency.

Table 4. Review protocol and search strategy

Element	Specification
Objective	To synthesise evidence on the mechanisms, conditions and outcomes of AI-driven people analytics in relation to employee retention
Design	Systematic literature review with conceptual framework development
Databases	Scopus, Web of Science, ScienceDirect, Emerald Insight, Springer Link, Taylor and Francis, IEEE Xplore, Google Scholar for citation chasing
Search strings	("people analytics" OR "HR analytics" OR "human resource analytics" OR "talent analytics" OR "workforce analytics") AND ("retention" OR "turnover" OR "attrition" OR "churn"); ("artificial intelligence" OR "machine learning" OR "algorithmic management" OR "algorithmic HRM") AND ("human resource management" OR "employee" OR "workforce") AND ("fairness" OR "ethics" OR "trust" OR "privacy"); plus seminal turnover theory searches
Timespan	Primary window 2015 to 2026, with seminal theoretical works included without date restriction
Language	English
Screening	Title and abstract screening followed by full-text assessment against inclusion and exclusion criteria
Extraction	Standardised template capturing context, method, constructs, findings, mechanisms and boundary conditions
Appraisal	Assessment of theoretical grounding, methodological transparency and relevance to retention outcomes
Synthesis	Thematic analysis with narrative synthesis and tabulation of evidence

4.2 Search Strategy and Selection Criteria

Searches were executed in Scopus, Web of Science, ScienceDirect, Emerald Insight, Springer Link, Taylor and Francis, IEEE Xplore and Google Scholar, supplemented by citation chasing of reference lists and by targeted searches of practitioner and regulatory sources for adoption and cost data. Search strings combined three concept blocks, namely analytics terminology, AI and algorithmic terminology, and employee outcome terminology, with a fourth block introduced to capture evaluative concepts such as fairness, trust, privacy and justice. The combination of blocks ensured that studies appearing in either the enabling or the evaluative literature would be retrieved, which was essential given the central argument of the paper.

Inclusion required that a source be published in English, be peer-reviewed or, in the case of a small number of practitioner and regulatory reports, be

issued by a recognised institution, and address at least one of the following: the adoption, capability or value of people analytics; the application of AI or machine learning to workforce decisions, especially turnover prediction; employee or applicant reactions to algorithmic decision-making; or the theoretical explanation of employee retention. Exclusion criteria removed purely technical papers on algorithm design without employee or organisational relevance, papers addressing customer rather than employee analytics, opinion pieces without evidentiary basis, and studies whose context or constructs could not be mapped onto the framework. The screening process yielded fifty-eight sources for full synthesis, comprising fifty-two peer-reviewed journal and proceedings articles, two peer-reviewed conference papers and four institutional or practitioner reports. The profile of the corpus is summarised in Table 5.

Table 5. Profile of the reviewed corpus

Dimension	Category	Sources (n)
Source type	Peer-reviewed journal and proceedings articles	52
	Peer-reviewed conference papers	2
	Institutional and practitioner reports	4

Publication period	2021 to 2026	38
	2016 to 2020	10
	2015 or earlier (seminal theory and practice reports)	10
Thematic domain	People and HR analytics adoption, capability and value creation	17
	AI and machine learning for turnover and attrition prediction	7
	Algorithmic management, fairness, trust and ethics	20
	Turnover and retention theory	14

4.3 Data Extraction, Appraisal and Synthesis

Data were extracted into a standardised template covering bibliographic details, study context and sample, research design, constructs and measures, analytic approach, principal findings, mechanisms proposed, and boundary conditions or limitations reported by the authors. Quality appraisal addressed theoretical grounding, transparency of method, construct validity in relation to retention outcomes, and explicit discussion of limitations. Appraisal was used to weight findings in the narrative synthesis rather than to exclude sources, since the objective was theory development rather than effect estimation. Two categories of evidence were distinguished throughout: sources that report primary empirical findings, and sources that report practitioner survey data. Where adoption statistics are reported in the findings, they are attributed to their originating institution and are complemented by the observation that such figures are self-reported and not independently audited.

Thematic synthesis followed the procedures associated with qualitative evidence synthesis. Extracted findings were coded inductively, and codes were clustered into themes through iterative comparison. Six themes emerged, concerning analytics capability and value capture, predictive performance and its limits, retention determinants surfaced by analytics, employee appraisal and resistance, governance and ethics, and organisational contingencies. These themes were then mapped onto the two pathways of the conceptual framework, with provision for findings that contradicted the framework, which were retained rather than suppressed. All propositions were formulated after synthesis and are grounded in the reviewed evidence.

5. Findings

5.1 Analytics Capability Determines Whether Value Is Realised

The first theme concerns the conditions under which people analytics produces organisational value. Across the corpus, the strongest and most

consistent finding is that value depends on capability configurations rather than on the presence of analytical tools. Fernandez and Gallardo-Gallardo (2021) find that organisations face a common set of barriers, including fragmented data architecture, limited analytical skills and insufficient managerial demand, and that progress stalls where these are not addressed simultaneously. Shet (2025) extends this by proposing that analytics capability is best conceived as an architecture in which data infrastructure, analytical talent, governance and business partnering must be developed in combination, and shows that organisations advancing along the maturity continuum reconfigure rather than simply add these elements. Di Lauro et al. (2025), studying Italian organisations, report that applications cluster in recruitment, workforce planning and attrition analysis; that enablers include executive sponsorship, data quality and analytical skills; and that barriers include legal uncertainty, cost and the scarcity of analysts with domain knowledge. Alam et al. (2026) add that many obstacles are relational, arising from misalignment between HR, information systems, legal and line-management stakeholders, and that stakeholder engagement is itself a capability.

Crucially, the relationship between analytics and performance is mediated rather than direct. McCartney and Fu (2022) show that HR analytics affects organisational performance through improved decision quality, and that this effect is conditioned by analytical maturity and by the routines that convert analysis into action. Bottesch et al. (2025) reach a convergent conclusion from an information systems perspective, arguing that business value emerges from the interaction of technological capability, organisational process and managerial interpretation, and that value dissipates when analytical artefacts are produced but not used. This theme is reinforced by the practitioner evidence, which documents a pronounced gap between deployment and realised

benefit. Whereas a substantial minority of organisations have implemented AI in HR functions and a further group plan to do so (SHRM, 2026), and whereas the proportion of HR leaders piloting or planning generative AI deployments rose sharply within a single reporting cycle (Gartner, 2024), the large majority of HR leaders report that their organisations have not realised significant business value from AI tools (Gartner, 2025). Read together, these findings indicate that the binding constraint on retention analytics is not analytical power but the organisational capacity to interpret, legitimise and act.

5.2 Predictive Performance Is Real but Conditionally Useful

The second theme concerns the analytical achievement itself. Table 6 summarises representative studies. Machine learning methods applied to employee records and behavioural data outperform conventional baselines in discriminating leavers from stayers, and richer feature sets derived from multiple systems improve early detection (Yahia et al., 2021). Analytics-driven modelling can also surface determinants that practitioners underweight, and can convert model output into retention priorities, although the authors caution that identified associations are not automatically causal (Avrahami et al., 2022). Ensemble methods combined with multi-criteria decision-making shift the objective from

classification to allocation, ranking employees by the expected benefit of intervention (Srivastava & Eachempati, 2021). Unstructured data extend the accessible signal space: analysis of online employee reviews identifies sector-specific satisfaction determinants (Chittiprolu et al., 2021), and structural topic modelling combined with supervised learning classifies employee experience narratives into actionable categories (Ding et al., 2025).

Three limits qualify these achievements. The first is outcome validity. A considerable share of studies predict turnover intention rather than actual exit, yet the turnover literature establishes that intention is an imperfect proxy and that the intention-behaviour relationship is moderated by alternatives and constraints (Griffeth et al., 2000; Mobley, 1977). Reported performance therefore likely overstates practical value. The second is temporal and structural validity. Models trained on historical data degrade as labour markets and job designs change, and workforce data are hierarchically structured, so models that ignore grouping structure overstate the generalisability of their predictions (Yuan et al., 2024). The third is decision validity, which is discussed in the next subsection. Explainability is presented in the reviewed literature as a response to these limits rather than as an optional refinement: interpretable models change managerial reasoning, increase the probability that recommendations are enacted, and mitigate trust deficits (Chowdhury et al., 2023).

Table 6. Representative studies of AI-based turnover prediction and analytics

Study	Context and data	Method	Principal contribution	Principal limitation
Yahia et al. (2021)	Employee records with behavioural indicators	Deep learning and comparative baselines	Improved classification of leavers and earlier detection of attrition risk	Limited external validity and reliance on historical labels
Avrahami et al. (2022)	Large organisational workforce dataset	Machine learning with feature interpretation	Surfaces turnover determinants overlooked in practitioner literature	Associational rather than causal inference
Srivastava and Eachempati (2021)	Employee records dataset	Ensemble learning with multi-criteria decision-making	Ranks employees by intervention priority rather than risk alone	Decision weights depend on managerial judgement
Yuan et al. (2024)	Grouped workforce data	Grouped-data prediction modelling	Demonstrates the necessity of respecting organisational nesting in turnover models	Single-organisation case limits generalisation

Chowdhury et al. (2023)	Turnover prediction modelling	Explainable machine learning	Embeds transparency and improves managerial use of algorithmic output	Interpretability gains depend on managerial analytical literacy
Ding et al. (2025)	Online employee reviews	Structural topic modelling with support vector machines	Converts unstructured experience narratives into actionable themes	Self-selection and platform bias in review data
Chittiprolu et al. (2021)	Employee reviews in hospitality	Text mining and determinant analysis	Identifies sector-specific satisfaction determinants at scale	Sampling and authenticity concerns in review platforms

5.3 Effects Are Contingent on Organisational and Institutional Context

The sixth theme concerns contingency. The reviewed evidence documents substantial variation across organisational size, sector, national regulatory context and labour market conditions. Smaller organisations report technological and skills barriers that larger organisations partially overcome, while larger organisations report coordination and governance challenges (Fernandez & Gallardo-Gallardo, 2021; Di Lauro et al., 2025). Sectoral differences appear in the determinants that analytics surfaces, with hospitality, technology and public sector contexts producing distinct retention profiles (Chittiprolu et al., 2021; Harden et al., 2018). Regulatory context shapes what data may be collected and how it may

be used (Ebert et al., 2021; Tursunbayeva et al., 2022), and institutional pressures influence whether analytics is adopted for substantive or symbolic reasons, with consequences for whether it is actually used (Nocker & Sena, 2019). Labour market tightness affects both the baseline turnover rate and the credibility of retention interventions, since employees with abundant alternatives respond differently to organisational inducements than those without (Griffeth et al., 2000). Collectively, these findings indicate that the relationship between analytics capability and retention outcomes is moderated rather than universal. Table 7 summarises the thematic synthesis, and Table 8 summarises the conditions associated with successful retention analytics.

Table 7. Thematic synthesis of the reviewed evidence

Theme	Core finding	Representative sources	Framework mapping
Analytics capability and value capture	Value depends on capability configurations and on the routines that convert insight into action, not on tool availability	Fernandez and Gallardo-Gallardo (2021); McCartney and Fu (2022); Shet (2025); Bottesch et al. (2025)	Enabling pathway, antecedent conditions
Predictive performance and its limits	Machine learning improves discrimination of leavers from stayers, but outcome, temporal and decision validity constrain practical value	Yahia et al. (2021); Avrahami et al. (2022); Yuan et al. (2024); Chowdhury et al. (2023)	Enabling pathway, diagnostic precision
Retention determinants surfaced	Analytics makes workload, development, supervision and fairness deficits visible, but intervention incidence is rarely documented	Bakker and Demerouti (2017); Hobfoll et al. (2018); Mitchell et al. (2001); (Gallup, n.d.)	Enabling pathway, intervention logic
Employee appraisal and resistance	Fairness, transparency, privacy and trust judgements determine whether analytics retains or repels employees	Lee (2018); Park et al. (2021); Lavanchy et al. (2023); Ochmann et al. (2024); Giermindl et al. (2022)	Evaluative pathway, moderation
Governance quality	Technical fairness criteria diverge from employee experience, and governance remains aspirational in practice	Köchling and Wehner (2020); Fabris et al. (2025); Tursunbayeva et al. (2022); Edwards et al. (2022)	Evaluative pathway, institutional conditions
Contextual contingency	Effects vary with size, sector, regulation and labour market conditions	Di Lauro et al. (2025); Ebert et al. (2021); Nocker and Sena (2019); Griffeth et al. (2000)	Moderating conditions

6. Discussion

6.1 Theoretical Implications

The synthesis supports a reframing of what AI-driven people analytics contributes to retention. The dominant framing in the literature treats analytical capability as an input and turnover as an output, with predictive accuracy as the intervening achievement. The evidence assembled here suggests instead that the decisive variable is decision quality, understood as the extent to which a decision is evidence-informed, strategically aligned, justifiable to affected employees and correctable in the light of new information. Predictive accuracy contributes to decision quality but does not constitute it, because an accurate risk score can still generate inequitable, demoralising or self-fulfilling interventions when deployed without interpretive discipline and procedural safeguards. This reframing aligns with evidence that HR analytics affects performance through decision quality rather than metric production (McCartney & Fu, 2022), and it explains why value capture lags deployment so markedly in survey evidence (Gartner, 2025; SHRM, 2026).

The dual-pathway model offers a second theoretical contribution by specifying how the enabling and evaluative mechanisms interact. In the enabling pathway, capability produces diagnostic precision, which enables targeted intervention, which restores or augments valued resources and thereby strengthens fit, links and sacrifice. This sequence is theoretically continuous with job demands-resources theory, conservation of resources theory, social exchange theory and job embeddedness research, each of which supplies a distinct link in the chain (Bakker & Demerouti, 2017; Hobfoll et al., 2018; Harden et al., 2018; Mitchell et al., 2001). The evaluative pathway operates in parallel, converting the act of algorithmic measurement into justice, privacy and trust appraisals. Its distinctive feature in the present framework is not that it produces negative outcomes directly, but that it moderates the enabling pathway. This specification responds to a recurring weakness in the algorithmic management literature, which documents resistance and distrust but rarely links them to retention outcomes, and to a recurring weakness in the retention analytics literature, which models turnover without modelling the experience of being predicted.

A third contribution concerns the conditions under which analytics produces retention gains. The

evidence indicates that capability and governance are configuration variables rather than additive investments: barriers cluster and must be addressed jointly, and organisations typically stall when one element lags (Fernandez & Gallardo-Gallardo, 2021; Shet, 2025). This finding extends resource-based accounts of algorithmic HRM by suggesting that the inimitable resource is not the algorithm but the governance and interpretive apparatus surrounding it (Kim et al., 2025; Cameron et al., 2024). It also qualifies technology acceptance accounts, which explain use through expectancy and facilitation but treat fairness and trust as peripheral, whereas HR-specific evidence suggests these are decisive (Theres & Strohmeier, 2023; Xu et al., 2024).

A fourth contribution concerns temporality. The unfolding model establishes that exits often follow shocks, and that some decision paths are scripted rather than deliberative (Lee & Mitchell, 1994). Attitudinal measurement is inherently weak at capturing shocks, and attitude-behaviour relationships depend heavily on measurement timing and frequency (Kammeyer-Mueller et al., 2024). Event-based signals, including changes in supervisor, unexpected evaluation outcomes and discontinuities in workload, are therefore theoretically privileged for retention analytics. Yet the reviewed prediction studies rely predominantly on static or slowly updating features, which suggests that the predictive advantage of AI is currently being realised in the wrong signal space. This is a testable proposition with direct design implications.

Finally, the framework reframes the ethics of retention analytics as a performance issue rather than a compliance issue. If perceived unfairness, opacity and surveillance reduce the effectiveness of retention interventions, then procedural safeguards are not constraints on analytics value but conditions for it. This argument connects governance scholarship (Tursunbayeva et al., 2022; Edwards et al., 2022), privacy analysis (Ebert et al., 2021) and evidence on resistance (Park et al., 2021; Park et al., 2022) to the retention outcome of interest, converting a normative claim into an empirical one.

7. Implications

7.1 Practical Implications

Five practical implications follow. First, organisations should treat retention analytics as a

decision system rather than a prediction system. The evidence indicates that analytical output converts into retention outcomes only through decisions that are taken, resourced and followed up, and that value dissipates when analysis is produced but unused (McCartney & Fu, 2022; Bottesch et al., 2025). Practical design should therefore specify, for each analytical insight, the owner, the decision it informs and the window in which action must occur.

Second, governance should be designed at the point of deployment rather than added after adverse events. Purpose limitation, data minimisation, documented provenance, human review of consequential decisions and accessible appeal mechanisms are not merely legal safeguards; they are the conditions under which employees accept algorithmic decisions, and acceptance is what makes retention interventions effective (Lavanchy et al., 2023; Park et al., 2022; Ebert et al., 2021).

Third, explanation must be treated as an operational requirement. Interpretable models increase the likelihood that managers act on recommendations and mitigate trust deficits, and transparency raises perceived fairness (Chowdhury et al., 2023; Ochmann et al., 2024). Explanation should be differentiated by audience: managers require decision-relevant reasoning, while employees require an account of what data were used, how they affected the decision and how the decision may be contested.

Fourth, analytics capability should be built as a configuration. Because barriers cluster, investments in algorithms without corresponding investment in data quality, analytical talent and managerial data literacy tend to produce analytical artefacts rather than outcomes (Fernandez & Gallardo-Gallardo, 2021; Shet, 2025). Managerial

data literacy deserves specific attention, since it is the mechanism through which precision becomes intervention.

Fifth, retention interventions should be designed to build embeddedness rather than only to remove dissatisfaction. The theoretical evidence indicates that fit, links and sacrifice are the proximate binding mechanisms (Mitchell et al., 2001), and that resource gain rather than deficit management drives engagement (Bakker & Demerouti, 2017; Hobfoll et al., 2018). Analytics that identifies who is at risk without strengthening attachment addresses the symptom rather than the mechanism.

7.2 Policy and Organisational Governance Implications

At the institutional level, the evidence supports three directions. Regulators and standard setters could require that consequential employment decisions supported by automated systems be accompanied by documentation of data provenance, validation of outcome measures and accessible mechanisms of contestation, since compliance-oriented approaches systematically under-protect employees (Ebert et al., 2021). Professional bodies could promote employee consultation and ethics review as standard components of analytics programmes, given that these mechanisms appear in ethical recommendations but rarely in practice (Tursunbayeva et al., 2022; Edwards et al., 2022). Finally, since technical fairness criteria and lay fairness intuitions diverge (Fabris et al., 2025; Köchling & Wehner, 2020), organisations should validate fairness claims against employee experience rather than relying exclusively on statistical tests. Table 11 maps these implications onto the framework's pathways.

Table 11. Implications mapped to framework pathways

Implication	Pathway addressed	Primary supporting evidence
Specify decision ownership and action windows for each insight	Enabling pathway, decision quality	McCartney and Fu (2022); Bottesch et al. (2025)
Design purpose limitation, human review and appeal at deployment	Evaluative pathway, governance	Ebert et al. (2021); Park et al. (2022)
Provide audience-specific explanations of algorithmic decisions	Evaluative pathway, fairness	Chowdhury et al. (2023); Lavanchy et al. (2023)
Build capability as a configuration including data literacy	Enabling pathway, capability	Shet (2025); Fernandez and Gallardo-Gallardo (2021)
Prioritise interventions that build fit, links and sacrifice	Enabling pathway, embeddedness	Mitchell et al. (2001); Hobfoll et al. (2018)
Validate fairness against employee experience, not only statistical tests	Evaluative pathway, legitimacy	Fabris et al. (2025); Köchling and Wehner (2020)

8. Limitations and Future Research

Four limitations impact the findings of this review. First, it is interpretive rather than meta-analytic due to a lack of consensus on analytics capability measures and varying outcome definitions, making narrative synthesis suitable but not providing aggregate effect estimates. Second, publication and selection bias may inflate the perceived benefits of retention analytics, as successful studies are more likely to be published. Third, the review is primarily based on English-language sources from North America, Europe, and partial data from Asia, limiting the applicability of conclusions to diverse institutional settings with different labor market frameworks. Fourth, the framework's propositions are conceptual and remain untested as a system, necessitating validated measurement tools for constructs like decision quality and privacy legitimacy.

To address these limitations, future research should focus on four priorities: causal identification, outcome validity, direct evaluative pathway measurement, and the evaluation of analytics-triggered interventions. Additionally, three questions emerge: the moderating effect of perceived fairness on the intervention-retention relationship; the diminishing returns of analytics intensity on retention; and the potential for standardized governance that is credible without becoming merely ritualistic. Addressing these questions will require research designs that integrate organizational data with employee perceptions using existing analytical tools.

9. Conclusion

This paper assessed the role of AI-driven people analytics in employee retention and developed an integrative framework to explain when data-driven HR decision-making produces its intended effects. The evidence supports a qualified conclusion. Analytics has materially improved the diagnosis of retention risk, offering earlier, richer and more granular identification of the antecedents that drive voluntary exit. Yet the translation of diagnostic capability into retention outcomes is conditional, uneven and, in a substantial proportion of documented instances, unrealised. Survey evidence showing that most HR leaders have not realised significant business value from AI tools is consistent with the theoretical argument advanced

here, namely that value depends on decision routines, capability configurations and governance rather than on analytical power alone.

The framework's central claim is that two pathways operate simultaneously. The enabling pathway converts capability into precision, precision into targeted intervention, and intervention into resource gain that strengthens fit, links and sacrifice and thereby reduces voluntary exit. The evaluative pathway converts the act of algorithmic measurement into fairness, privacy and trust judgements that moderate the first pathway. Where employees judge analytics to be opaque, unfair or intrusive, retention interventions delivered through it become less effective, and the analytical advantage is partially or wholly offset. This formulation explains why identical technologies produce divergent outcomes across organisations, and why governance and legitimacy belong in performance models rather than in a separate ethical domain.

Three contributions follow. Theoretically, the paper integrates four literatures that have developed in parallel into a single explanatory structure with nine testable propositions. Methodologically, it specifies the designs needed to move from prediction to identification and from intention to verified exit. Practically, it identifies decision quality, explanation and governance as the levers through which analytical capability becomes retention performance. The wider implication is that the value of data-driven HR decision-making is realised not at the moment of prediction but at the moment of enactment, and that the organisations most likely to benefit are those that treat analytical capability and procedural legitimacy as a single, indivisible resource.

Bibliography

1. Alam, S., Dong, Z., Kularatne, I., & Rashid, M. S. (2026). Exploring approaches to overcome challenges in adopting human resource analytics through stakeholder engagement. *Management Review Quarterly*. <https://doi.org/10.1007/s11301-025-00491-y>
2. Avrahami, D., Pessach, D., Singer, G., & Chalutz Ben-Gal, H. (2022). A human resources analytics and machine-learning examination of turnover: Implications for theory and practice. *International Journal of Manpower*, 43(6), 1405–1424. <https://doi.org/10.1108/IJM-12-2020-0548>

3. Bakker, A. B., & Demerouti, E. (2017). Job demands–resources theory: Taking stock and looking forward. *Journal of Occupational Health Psychology*, 22(3), 273–285. <https://doi.org/10.1037/ocp0000056>
4. Bottesch, S., Schwenke, C., Förster, M., & Klier, M. (2025). Driving business value through people analytics: Literature review and research agenda from an information systems perspective. *Electronic Markets*, 35(1). <https://doi.org/10.1007/s12525-025-00842-3>
5. Cameron, R., Herrmann, H., & Nankervis, A. (2024). Mapping the evolution of algorithmic HRM (AHRM): A multidisciplinary synthesis. *Humanities and Social Sciences Communications*, 11. <https://doi.org/10.1057/s41599-024-02786-z>
6. Chittiprolu, V., Singh, S., Bellamkonda, R. S., & Vanka, S. (2021). A text mining analysis of online reviews of Indian hotel employees. *Anatolia*, 32(2), 232–245. <https://doi.org/10.1080/13032917.2020.1856157>
7. Choung, H., Seberger, J. S., & David, P. (2024). When AI is perceived to be fairer than a human: Understanding perceptions of algorithmic decisions in a job application context. *International Journal of Human–Computer Interaction*. <https://doi.org/10.1080/10447318.2023.2266244>
8. Chowdhury, S., Joel-Edgar, S., Dey, P. K., & Bhattacharya, S. (2023). Embedding transparency in artificial intelligence machine learning models: Managerial implications on predicting and explaining employee turnover. *The International Journal of Human Resource Management*, 34(14), 2732–2764. <https://doi.org/10.1080/09585192.2022.2066981>
9. Colquitt, J. A., Conlon, D. E., Wesson, M. J., Porter, C. O. L. H., & Ng, K. Y. (2001). Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology*, 86(3), 425–445. <https://doi.org/10.1037/0021-9010.86.3.425>
10. Di Lauro, S., Tursunbayeva, A., Antonelli, G., & Moschera, L. (2025). Disrupting human resource management with people analytics: A study of applications, value, enablers and barriers in Italy. *Personnel Review*, 54(2), 697–721. <https://doi.org/10.1108/PR-11-2023-0927>
11. Dima, J., Gilbert, M.-H., & Dextras-Gauthier, J. (2024). The effects of artificial intelligence on human resource activities and the roles of the human resource triad: Opportunities and challenges. *Frontiers in Psychology*, 15, 1360401. <https://doi.org/10.3389/fpsyg.2024.1360401>
12. Ding, K., Li, R., Li, Z., & Hu, S. (2025). Uncovering employee insights: Integrative analysis using structural topic modeling and support vector machines. *Journal of Big Data*, 12(1). <https://doi.org/10.1186/s40537-025-01100-1>
13. Ebert, I., Wildhaber, I., & Adams-Prassl, J. (2021). Big data in the workplace: Privacy due diligence as a human rights-based approach to employee privacy protection. *Big Data & Society*, 8(1). <https://doi.org/10.1177/20539517211013051>
14. Edwards, M. R., Charlwood, A., Guenole, N., & Marler, J. (2022). HR analytics: An emerging field finding its place in the world alongside simmering ethical challenges. *Human Resource Management Journal*, 34(2), 326–336. <https://doi.org/10.1111/1748-8583.12435>
15. Fabris, A., Baranowska, N., Dennis, M. J., & Graus, D. (2025). Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Transactions on Intelligent Systems and Technology*, 16(3), 1–54. <https://doi.org/10.1145/3696457>
16. Fernandez, V., & Gallardo-Gallardo, E. (2021). Tackling the HR digitalization challenge: Key factors and barriers to HR analytics adoption. *Competitiveness Review*, 31(1), 162–187. <https://doi.org/10.1108/CR-12-2019-0163>
17. Gallup. (n.d.). 42% of employee turnover is preventable but often ignored. <https://www.gallup.com/workplace/646538/employee-turnover-preventable-often-ignored.aspx>
18. Gartner. (2024). Gartner survey finds 38% of HR leaders are piloting, planning implementation, or have already implemented generative AI. <https://www.gartner.com/en/newsroom/press-releases/2024-02-27-gartner-finds-38-percent-hr-leaders-piloting-generative-ai>
19. Gartner. (2025). Gartner survey shows 88% of HR leaders say their organizations have not realized significant business value from AI tools. <https://www.gartner.com/en/newsroom/press-releases/2025-10-28-gartner-survey-shows-88-percent-of-hr-leaders-say-their-organizations-have-not-realized-significant-business-value-from-ai-tools>
20. Giermindl, L. M., Strich, F., Christ, O., & Leicht-Deobald, U. (2022). The dark sides of people analytics: Reviewing the perils for organisations and employees. *European Journal of Information Systems*, 31(3), 410–435. <https://doi.org/10.1080/0960085X.2021.1927213>
21. Griffeth, R. W., Hom, P. W., & Gaertner, S. (2000). A meta-analysis of antecedents and correlates of employee turnover: Update, moderator tests, and research implications for the next millennium. *Journal of Management*, 26(3), 463–488. <https://doi.org/10.1177/014920630002600305>
22. Hancock, J. I., Allen, D. G., Bosco, F. A., McDaniel, K. R., & Pierce, C. A. (2013). Meta-analytic review of employee turnover as a predictor of firm performance. *Journal of Management*, 39(3),

- 573–603.
<https://doi.org/10.1177/0149206311424943>
23. Harden, G., Boakye, K. G., & Ryan, S. (2018). Turnover intention of technology professionals: A social exchange theory perspective. *Journal of Computer Information Systems*, 58(4), 291–300.
<https://doi.org/10.1080/08874417.2016.1236356>
24. Hobfoll, S. E., Halbesleben, J., Neveu, J.-P., & Westman, M. (2018). Conservation of resources in the organizational context: The reality of resources and their consequences. *Annual Review of Organizational Psychology and Organizational Behavior*, 5(1), 103–128.
<https://doi.org/10.1146/annurev-orgpsych-032117-104640>
25. Huselid, M. A. (1995). The impact of human resource management practices on turnover, productivity, and corporate financial performance. *Academy of Management Journal*, 38(3), 635–672.
<https://doi.org/10.2307/256741>
26. Kammeyer-Mueller, J. D., & Rubenstein, A. L. (2024). The role of attitudes in work behavior. *Annual Review of Organizational Psychology and Organizational Behavior*.
<https://doi.org/10.1146/annurev-orgpsych-101022-101333>
27. Kim, S., Khoreva, V., & Vaiman, V. (2025). Strategic human resource management in the era of algorithmic technologies: Key insights and future research agenda. *Human Resource Management*, 64(2), 447–464. <https://doi.org/10.1002/hrm.22268>
28. Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13(3), 795–848. <https://doi.org/10.1007/s40685-020-00134-w>
29. Lavanchy, M., Reichert, P., Narayanan, J., & Savani, K. (2023). Applicants' fairness perceptions of algorithm-driven hiring procedures. *Journal of Business Ethics*, 188(1), 125–150.
<https://doi.org/10.1007/s10551-022-05320-w>
30. Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5(1).
<https://doi.org/10.1177/2053951718756684>
31. Lee, T. W., & Mitchell, T. R. (1994). An alternative approach: The unfolding model of voluntary employee turnover. *Academy of Management Review*, 19(1), 51–89.
<https://doi.org/10.2307/258835>
32. Madanchian, M., Taherdoost, H., & Mohamed, N. (2023). AI-based human resource management tools and techniques: A systematic literature review. *Procedia Computer Science*, 229, 367–377.
<https://doi.org/10.1016/j.procs.2023.12.039>
33. Marler, J. H., & Boudreau, J. W. (2017). An evidence-based review of HR analytics. *The International Journal of Human Resource Management*, 28(1), 3–26.
<https://doi.org/10.1080/09585192.2016.1244699>
34. Mayer, R. C., & Davis, J. H. (1999). The effect of the performance appraisal system on trust for management: A field quasi-experiment. *Journal of Applied Psychology*, 84(1), 123–136.
<https://doi.org/10.1037/0021-9010.84.1.123>