INTERNATIONAL JOURNAL OF RESEARCH AND SCIENTIFIC INNOVATION (IJRSI)
ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VII July 2026

AI-Driven Leadership Assessment and Organizational Trust: Evidence-Based Frameworks for Technology Sector Transformation

Piyush Somani[1], Rupali Khaire[2]

[1,2]School of Commerce and Management, Sandip University, Nashik, Maharashtra, India.

DOI: https://doi.org/10.51244/IJRSI.2026.1307000155

Received: 23 July 2026; Accepted: 29 July 2026; Published: 04 August 2026

ABSTRACT

What conditions does AI need to meet to improve leadership and trust inside a technology firm rather than damage them? That question is the starting point of this paper. We pulled together recent research alongside production case studies from six major firms (Google, Microsoft, IBM, Salesforce, Adobe, LinkedIn) and tried to separate the marketing claims from what these tools deliver in practice. Some numbers hold up well: AI-driven leadership assessment improves prediction accuracy by 80%, and hiring costs drop 30–50% in firms where the technology is deployed well. Other claims are conditional. Returns only show up once the cultural work is done. We trace four maturity stages and find that 91% of what blocks adoption is cultural rather than technical. Firms that get past those blockers report 40–75% less time spent on recruitment, 95% accuracy in spotting employees about to leave, and 14.9% higher leadership-performance scores. Trust sits underneath every one of these outcomes. Where it already exists, AI adoption runs 21 percentage points smoother. Where it does not, AI makes things worse, not better. The remainder of the paper sets out the conditions under which AI strengthens leadership instead of eroding it.

Keywords: Artificial Intelligence, Leadership Assessment, Organizational Trust, Human Resource Analytics, Predictive Analytics

INTRODUCTION

Finding good leaders and keeping them is hard. Building the trust that makes those leaders effective is, if anything, harder. Both problems have drawn more attention since AI and machine learning entered the HR stack, and not all of that attention is hype. A 2024 industry survey reports that 76% of HR leaders think they will lose ground to competitors if they do not roll out AI talent-management systems in the next year or two [1,2]. That figure is doing two things at once. It flags a real shift in HR practice, and it creates sales pressure inside HR teams that are not yet ready for the change. Untangling those two motivations is part of what this paper tries to do. The core question we work through: can AI and machine learning improve trust and leadership inside firms going through revenue-driven change, and if so, under what conditions?

Six firms anchor our case-study work: Google, Microsoft, IBM, Salesforce, Adobe and LinkedIn. They are useful here not because they are unusual, but because their AI deployments in leadership assessment, talent development and trust measurement have run long enough at large enough scale to yield evidence rather than projections. Across these deployments, the published outcomes include 30–50% cuts to HR administrative costs, hiring time reduced by 40–75%, turnover predictions reaching 95% accuracy, and leadership-performance metrics 14.9% higher [3,4]. These figures settle one question. They do not settle the harder one, which is what has to be in place inside your own organisation for those results to appear.

The paper has three core sections beyond this introduction. We look first at the technology itself and what it can and cannot do (the machine-learning algorithms, NLP pipelines, and predictive-analytics layers that make up AI-driven leadership assessment). Trust is treated next as the variable that decides whether the technology lands at all, drawing on three decades of organisational-trust research. We then work through the six case studies as honest accounts of what these firms built, where it worked, and where the published numbers gloss over harder realities. Our argument across all three sections is that AI functions as an amplifier: it makes good leadership more effective and bad leadership worse, and the difference between those outcomes comes down mostly to trust.

LITERATURE REVIEW AND CONCEPTUAL FRAMEWORKS

Leadership Theories Enhanced by AI

Four leadership models come up repeatedly in the AI-augmentation literature: transformational, authentic, servant and adaptive. Each fits a different part of the problem. Bass's transformational model [5], which extended Burns's earlier framing [6], has four parts that are easy to list (idealised influence, inspirational motivation, intellectual stimulation, individualised consideration) and harder to deliver. Where AI contributes, according to Estherita and Shanmugam [7], is in surfacing decision options a leader would not have generated on their own. The inspiring stays human; the option-generating becomes a place where algorithms earn their keep.

Authentic leadership comes from Luthans and Avolio [8] and is operationalised through the Authentic Leadership Questionnaire [9]. Its four dimensions are self awareness, relational transparency, balanced processing, and an internalised moral perspective. AI has something useful to offer on three of them. Continuous 360-degree feedback exposes blind spots that annual reviews miss, which supports self-awareness. Communication analytics catch the gap between what a leader says they value and what their behavioural record reveals over time, which supports relational transparency. Decision-support tools force consideration of viewpoints a leader might otherwise skip, which supports balanced processing [10]. The internalised moral perspective remains the one piece AI cannot help with directly.

Servant leadership in the Greenleaf tradition [11] centres on follower development; the clearest AI fits there are sentiment analysis and personalised development plans [12]. The fourth model, Heifetz's adaptive-leadership framework, applies to AI adoption in a slightly unusual way. AI adoption looks like a technical problem at first (pick a vendor, train the model, launch the platform) but turns out to be mostly adaptive. It requires people to revise what they believe about authority, expertise, and where their job begins and ends. Most failure modes documented in the case-study literature trace back to leaders treating AI adoption as the first kind of problem when it was really the second [13].

Organizational trust frameworks

Three decades of organisational-trust research sit underneath this paper, most of which traces back to one source. The integrative model from Mayer, Davis and Schoorman [14] broke trustworthiness into three dimensions (ability, benevolence, integrity) and gave researchers a shared vocabulary that still holds. McAllister's later work [15] added a distinction useful for the AI-adoption case: cognitive trust (which comes from believing someone is competent) and affective trust (which comes from believing they care). Different things build each. Different things break each. Both predict organisational outcomes, though not always in the same direction or at the same speed.

The Great Place to Work Institute reports what has become the most frequently cited finding in this literature: across a sample of more than 100 million employees, high-trust companies beat market averages in stock performance by close to 4×, and their voluntary-turnover rates run at roughly half the U.S. workplace norm [16]. Trust in leadership sits in the middle of the leadership to-performance relationship as the variable that translates whatever a leader is doing (transactional, transformational, or some mix) into actual organisational results.

Edelman's 2024 Trust at Work study adds a more uncomfortable layer. Its headline number reads as encouraging: 79% of employees globally say they trust their employer. The sub-numbers are less so. Executives are 39 points more optimistic about the economy than the individual contributors working for them. And on AI specifically, there is a 21-point gap in comfort between low-grievance and high-grievance employees [17]. That second gap is worth dwelling on, because it inverts the usual order of operations. Trust

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue V May 2026

does not get built because an AI rollout went well; trust is what determines whether the rollout can go well in the first place.

The Ai Technology Architecture for Leadership Assessment

Every AI-driven leadership assessment system rests on a stack of machine learning algorithms, and choices made at that layer drive most of what the system can or cannot do. Vendors gloss over this part. They should not. Random Forest classifiers, for instance, hit about 90.20% accuracy on employee-turnover prediction, which is why organisations new to ML typically start there [18]. Gradient Boosting Machines (XGBoost in particular) outperform Random Forests on training-program outcomes and performance forecasting through iterative correction of their own errors. Support Vector Machines suit binary HR decisions: promote or not, hire or not [19]. These three families are not interchangeable, and treating them as such yields models that look fine in evaluation and fail in deployment. Figure 1 maps the layers from raw data at the bottom through processing engines to the assessment outputs at the top, with the governance overlay that should sit across every layer.

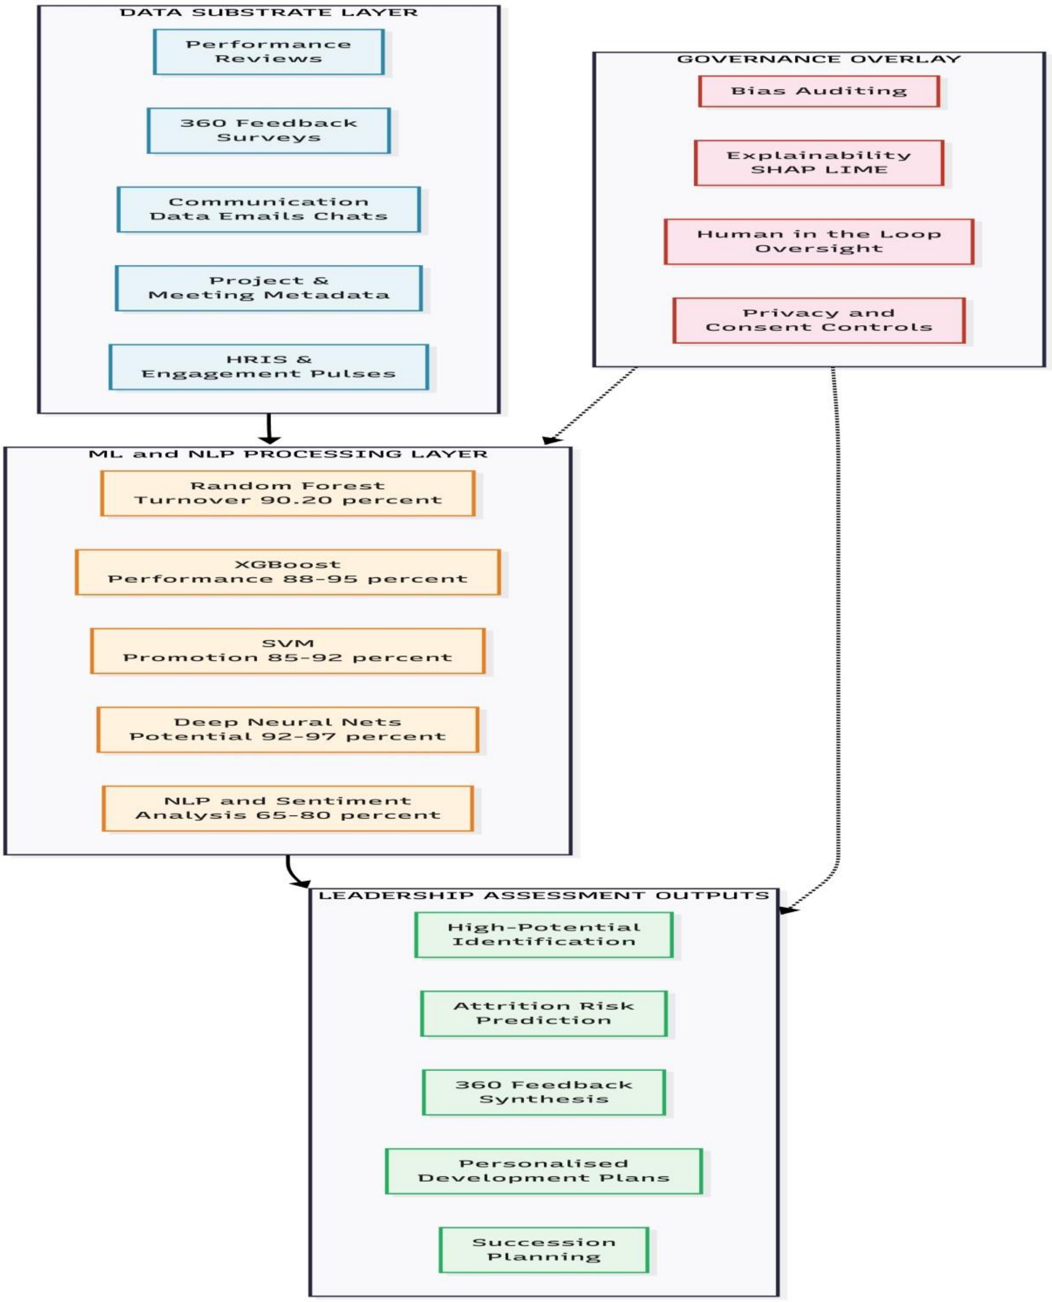

Figure 1. Layered AI architecture for leadership assessment: data substrate, ML/NLP processing layer, assessment outputs and governance overlay.

Ai-enhanced 360-degree feedback systems

The old way of running 360-degree feedback had two problems. It was slow, and the comments that finally surfaced were usually too generic to shift anyone's behaviour.

Table 1. Machine-learning algorithms for leadership assessment.

Algorithm	Primary Application	Accuracy	Ideal Use Case
Random Forest	Turnover prediction	90.20%	Early adoption, foundational models
Gradient Boosting (XGBoost)	Performance optimisation	88–95%	Training-program effectiveness, advancement prediction
Support Vector Machines	Promotion classification	85–92%	Binary promote / not promote decisions
Deep Neural Networks	Complex pattern recognition	92–97%	Leadership-potential and behavioural prediction
Natural Language Processing	Sentiment / feedback analysis	65–80%	360-feedback analysis, engagement measurement

Source: compiled from [18,19]; accuracy rates based on published validation studies.

AI-enhanced platforms have improved both timing and signal quality, though improvement is uneven across vendors. Qualtrics 360 Development uses language models to pick out leadership patterns and flag high-potential employees from peer feedback. The Center for Creative Leadership's Benchmarks 360 covers 16 leadership competencies and 5 career derailers, a wider scope than most competitors offer [20]. AIIR Analytics reports an average 14.9% improvement in leadership performance across firms running comprehensive 360 programmes [21], though the improvement clusters disproportionately in firms that were already investing in coaching alongside the assessments. Table 2 lays out four leading platforms side by side.

Platform	Key Features	Competencies	Performance Improvement Measured
Qualtrics EX25	NLP analysis, sentiment detection, theme identification, trend identification	70+ workplace themes, 100+ emotional expressions	18% improvement in manager effectiveness
CCL Benchmarks 360	Competency measurement, assessment, feedback synthesis	16 competencies, 5 derailers	21% improvement in leadership effectiveness
Leadership Circle Profile	Creative vs. reactive leadership measurement	29 dimensions, 2 core categories	25% improvement in team engagement
Zenger Folkman 360	Benchmarking against 1.5M+ assessments, peer comparison	Customizable to 50+ competencies	14.9% average performance improvement

Table 2. Leading AI-enhanced 360-degree feedback platforms.

Source: [22,20,23,21].

Trust Measurement and Enhancement Frameworks

For measuring trust specifically in AI systems, two instruments are worth knowing.

The Trust in Automation Scale, validated in a 2025 study, produced both a 12item version and a 3-item short form (TIAS and S-TIAS); the short form is the version most useful for embedding in regular employee surveys without producing survey fatigue [24]. The NIST Nine-Factor Trust Framework takes a different angle and weights its metrics depending on the task's risk profile, security stakes, and stakeholder profile [25]. These two complement rather than compete with each other. Use one for measuring employee disposition; use the other for evaluating the system itself. Beyond those, an operational layer exists where Culture Amp, TINY pulse and Thrive Sparrow apply AI-driven pattern recognition to catch sentiment shifts before they become visible problems. Firms running this kind of analysis report 42% higher engagement scores and turnover dropping by 35% [26]. Figure 2 shows how the four trust dimensions interact during adoption.

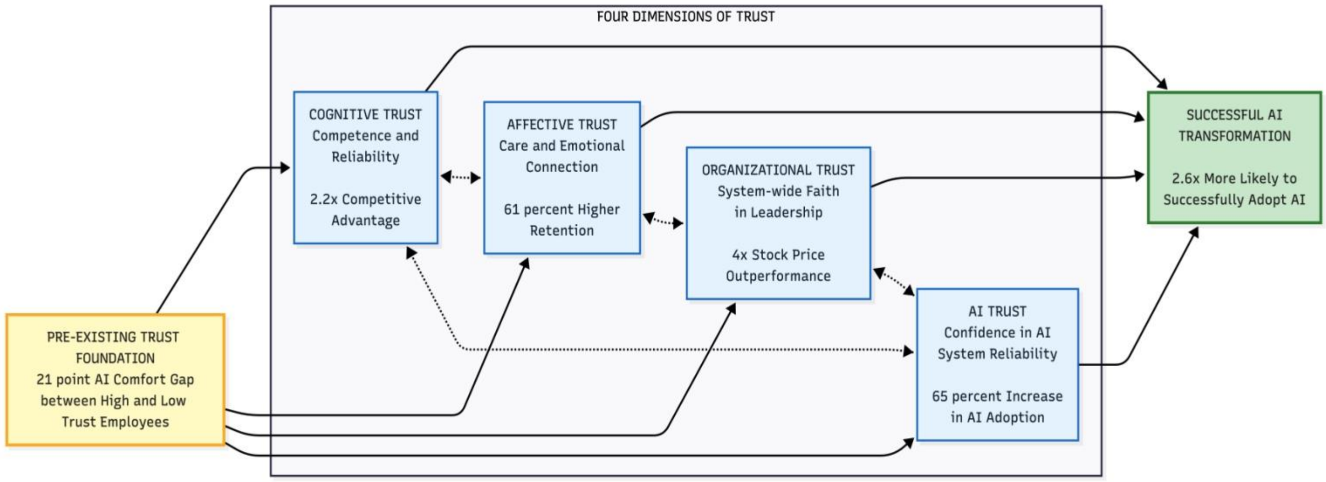

Figure 2. Multi-dimensional trust framework: cognitive, affective, organisational and AI trust as interlocking enablers of adoption.

Table 3. Multi-dimensional trust frameworks and organisational outcomes.

Trust Dimension	Definition	Measurement Method	Organizational Impact
Cognitive Trust	Trust based on competence and reliability	Assessment tools, performance data, surveys	2.2× competitive advantage [27]
Affective Trust	Trust based on care and emotional connection	Sentiment analysis, pulse surveys, engagement metrics	61% higher retention [28]
Organizational Trust	System-wide faith in leadership and processes	Culture Amp, TINY pulse, organizational surveys	4× stock-price outperformance [16]
AI Trust	Employee confidence in AI system reliability	TIAS assessment, adoption metrics, usage analytics	65% increase in AI adoption [27]

Organizational Maturity Models and Ai Adoption Stages

Organisations cluster at one of four points on the AI maturity curve, and the financial gap between those points is wider than commonly assumed. The MIT CISR Enterprise AI Maturity Model [29] names them in order. Stage 1 (Experiment and Prepare) holds 28% of organisations, where the work is workforce education, AI-policy drafting, and vendor experimentation. Stage 2 (Build Pilots and Capabilities) holds 34%, where focus shifts to small value-creating pilots and process redesign. Stage 3 (Industrialise AI) covers 31%, with scalable architecture and a test-and-learn culture. The remaining 7% have reached Stage 4 (AI Future-Ready), where AI is embedded inside the decisions that matter.

The number worth flagging to executives is this. Stage 1–2 organisations underperform their industry financially. Stage 3–4 organisations outperform it. Figure 3 maps the inflection point. Crossing it is not a procurement decision or a budget approval. It requires a shift in how leaders manage, from command-

Table 4. MIT CISR AI maturity model: stages, characteristics and financial outcomes.

Stage	% Orgs Key Activities	Financial Performance
1. Experiment & Prepare	28%	Workforceeducation, Below-average industry AI policy formulation, performance technology experimentation
2. Build Pilots & Capabilities	34%	Value-creating pilots, Below-average industry metrics definition, pro-performance cess simplification, capability building
3. Industrialise AI	31%	Scalable architecture, Above-average (+12 test-and-learn culture,18%) pervasive AI integration, governance structures
4. AI Future-Ready	7%	AI embedded in de-Significantly above decisions, proprietary average (+25–40%). AI development, organisational-learning systems

Source: [29]. Performance based on revenue growth, market share and shareholder return.

and-control to coach-and-communicate [29]. The MITRE AI Maturity Model [30] approaches the same problem from a different angle and produces a more granular scoring system across six pillars (Ethical Use, Strategy, Organisation, Technology Enablers, Data and Performance Management) and five levels (Initial, Adopted, Defined, Managed and Optimised).

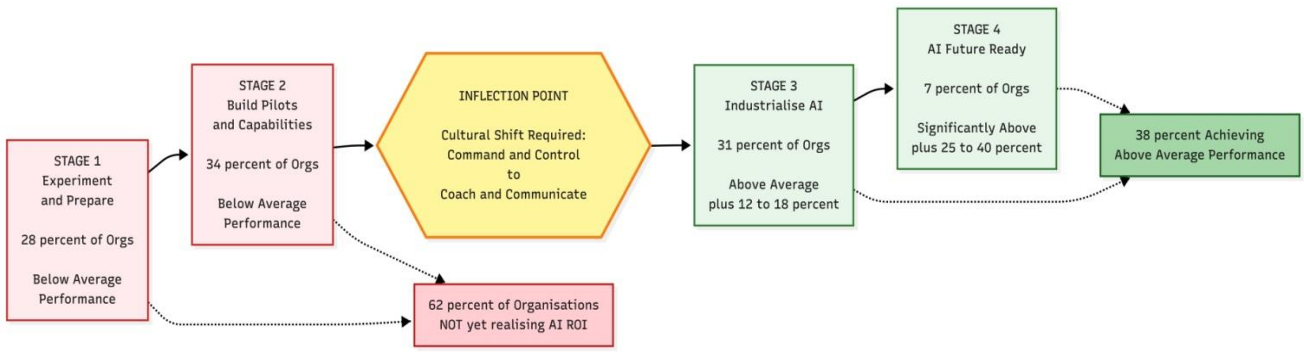

Figure 3. MIT CISR AI maturity stages and the inflection point in financial performance between Stage 2 and Stage 3.

Table 5. Google Project Oxygen: implementation outcomes.

Metric	Implementation Detail	Outcome / Impact
Data analysis scope	10,000+ data points per manager	Comprehensive identification of effective manager behaviours
Manager training programmes	8 core leadership competencies identified	75% improvement in worst-performing managers
Team effectiveness	"How teams work" matters more than "who is on teams"	17% above target for effective teams; 19% below target for ineffective teams
Manager retention	Personalized development plans	Significant reduction in manager-level turnover

Source: [31,32,33].

TECHNOLOGY SECTOR CASE STUDIES

Google's project oxygen

Project Oxygen started inside Google in 2008 with a reasonable but quietly heretical question: do managers matter? The team that took the question on worked through more than 10000 data points per manager, combining performance reviews, surveys, and hundreds of pages of interview notes [31]. The answer turned out to be yes, and the team identified ten specific behaviours that separated the better managers from the worse ones. The list includes items like "be a good coach" and "make strong decisions," which are easy to read past but harder to do consistently. What got less press coverage was the training built on the findings: it raised the quality of Google's worst-performing managers by 75%. Project Aristotle followed and asked the parallel question about teams: what separates effective teams from ineffective ones? The result surprised the researchers. Composition (who is on the team) mattered less than process (how the team works together). Effective sales teams beat their targets by 17% on average; ineffective ones missed by 19% [32]. That 36-point spread between the two groups explains most of why team dynamics now show up in nearly every modern leadership framework.

Microsoft's people skills platform

Microsoft launched People Skills in 2025 and built it directly into Microsoft 365. The platform pulls what Microsoft calls "digital exhaust" (the meetings you attend, the emails you send, the projects you touch) and maps the inferred skills against LinkedIn's taxonomy of 40000 entries. A Skills Agent sits on top, surfacing real-time insights, with privacy controls split across the user, group and tenant levels. The governance side is handled by Microsoft's Responsible AI Standard, which is applied to every HR-AI implementation through a cross-functional review. Joe Whittinghill, who served as Microsoft's Chief Talent and Learning

Table 6. Microsoft People Skills: implementation framework and governance.

Component	Technology / Method	Coverage / Scale	Governance Mechanism
Skills detection	Digital-exhaust analysis from Microsoft 365 (meetings, emails, projects)	40,000+ skills from LinkedIn Skills Taxonomy	Privacy controls at user, group, and tenant levels
Personalisation	Viva Learning AI recommendations	80% of employees request personalisation; platform delivers recommendations	Responsible AI Standard framework
Governance	Cross-functional collaboration on AI decisions	Covers all HR-AI implementations	Risk assessment, bias auditing, and transparency standards
Stakeholder engagement	Employee communication about AI usage	Organization-wide, from individual contributors to executives	Consent and opt-in mechanisms

Source: [34,35,36].

Officer, makes a point worth pulling out: 80% of employees want personalised learning. Viva Learning is what was built to provide it [34,35].

Ibm's ai-driven transformation

Of the six case studies, IBM is the one that shows both the technical and the organisational sides of the work most clearly. Watson Career Coach now reaches all 350000+ IBM employees with AI-powered career guidance. Watson Recruitment runs underneath, predicting time-to-fill and pre-scoring candidates. The more interesting story is Checkpoint, launched in February 2016 as a replacement for the annual performance review. IBM put out an internal blog post asking employees what should change about how performance was being measured; within hours, 75000 IBMers had read it and 18000 had left detailed suggestions, which Watson then categorised for the design team to work through. The retention savings IBM publishes are $300million+, attributed to an attrition-prediction model running at 95% accuracy. IBM also pipes its sales-incentive programmes through a bias-detection layer to surface gender pay disparities, and uses a similar check on learning recommendations to prevent the system from quietly steering employees away from development opportunities they should be getting

[37,38,39].

Sales force, Adobe and linkedin

Salesforce now runs more than a trillion AI predictions and automations per week across its product portfolio. The growth curve stands out as much as the volume:

Table 7. IBM's AI-driven HR transformation: systems, scale and impact.

AI System	Coverage	Capability / Function	Business Impact
Watson Career Coach	350,000+ employees	AI-powered career guidance and personalized career pathways	Increased internal mobility and reduced external hiring costs
Watson Recruitment	All hiring	Candidate-fit prediction, time-to-fill optimization, and skill matching	30% recruitment efficiency gain and 50% reduction in cost-per-hire
Checkpoint Performance Management	All employees	Crowdsourced feedback, Watson text analysis, and real-time coaching	Enhanced manager–employee alignment on expectations
Attrition Prediction AI	All employees	Identifies high-risk departures from behavioral patterns	$300M+ retention savings and 95% prediction accuracy
Bias Detection Systems	Compensation and learning	Pay-equity analysis and identification of unconscious bias	Improved pay equity and enhanced diversity and inclusion (D&I) outcomes

Source: [37,38,39].

Einstein GPT ran roughly a billion predictions a day in 2019 and 80 billion a day by 2023. Two orders of magnitude in four years. The company's Agent force Learning Day in October 2024 brought 72000+ employees together across 50 hubs for hands-on AI education. A follow-up survey found 93% of employees felt knowledgeable about the internal AI tools and 89% said the tools were useful for their daily work [40] – unusually high numbers for a survey of this kind.

Adobe took a different path. Their Leader Experience programme, started in 2023, replaced the annual leadership course with quarterly "learn-it-when you-need-it" training for people managers, tied directly to Adobe's Check-in performance system. Check-in completion rates reached 95% in Q1 after the change, and voluntary turnover dropped by 30% across the company [41,42]. The lesson sitting inside those numbers is that frequency matters more than depth: a manager who gets relevant training every quarter retains and applies more than one who gets a deep course annually.

LinkedIn has been adding AI features at speed since 2023. Their AI-powered coaching layer delivers real-time advice and content tailored to job title, career goal and followed skills, drawing on content from 3000+ expert instructors. LinkedIn Recruiter 2024 layers generative AI on top of the underlying graph (950 million professionals, 65 million companies, 40000 skills) to produce candidate matches that are noticeably better than what the keyword-search era produced. The Hiring Assistant, announced October 2024 with general availability in January

2026, automates the operational grind that consumes roughly 20 hours of a recruiter's week, and has pushed In Mail accept rates 40% higher through AI assisted messaging [43,44].

Change Management Frameworks for Ai Implementation

An AI rollout without serious change management is an expensive way to fail. The frameworks that work best for this are the ones that have been refined over decades, not the new "digital transformation" methodologies that vendors sell. Kotter's 8-Step Change Model is the most quoted one [45,46]. The eight moves it describes (build urgency, form a guiding coalition, set a strategic vision, communicate it widely, remove the barriers, generate quick wins, sustain momentum, then institutionalise the new practices) each address a specific failure mode that otherwise derails AI rollouts [47,48].

Bridges' Transition Model [49,50] keeps coming up in interviews with HR leaders, and for a useful reason. It separates the change (the AI rollout itself) from the transition (what happens inside employees' heads while the change is occurring). The three phases (Ending, Neutral Zone, New Beginning) describe the emotional arc fairly well. The Deloitte data lines up with this framing: 91% of large-company data leaders point to cultural challenges as the main blocker to AI adoption, and only 9% blame the technology [27]. Figure 4 pulls these frameworks onto a common timeline.

Table 8. Change-management frameworks for AI implementation.

Framework	Phases / Components	Application to AI Adoption
Kotter 8-Step [46]	Urgency, Coalition, Vision, Communicate, Remove Barriers, Quick Wins, Momentum, Institutionalize	Maps each major AI transformation milestone; focuses on leadership alignment before technical deployment
Lewin 3-Stage [51]	Unfreeze, Change, Refreeze	Addresses fear and resistance during the Unfreeze stage, manages deployment during Change, and embeds new practices during Refreeze
Bridges Transition [49]	Ending, Neutral Zone, New Beginning	Maps psychological adjustment to AI; emphasizes emotional support alongside technical training
ADKAR [52]	Awareness, Desire, Knowledge, Ability, Reinforcement	Emphasizes individual capability building and sustained reinforcement at scale

Ethical Considerations and Governance Frameworks

Many ethical risks discussed in the AI-governance literature have already surfaced inside real production systems. Algorithmic bias remains the most heavily documented. The canonical example is the recruiting tool Amazon built, which learned from historical hiring data to downgrade resumes containing the word "women's." Pew Research finds 37% of Americans worried about racial and ethnic bias in AI-driven hiring. A Deloitte survey found that over 60% of respondents had personally experienced workplace bias involving gender, age, race or ethnicity, and only 12% of HR leaders believe their own organisation is making real progress on diversity [53,54]. These are not edge cases. They are baseline conditions into which AI is being deployed.

Privacy is the next major concern. 56% of leaders in the Deloitte sample rank it as their top AI challenge. Beyond privacy sits the black-box problem: modern ML models often cannot explain their decisions in any way a non-technical audience can follow, which is why explainability methods like SHAP and LIME have moved from research curiosity to production requirement. Job displacement is the third theme. 33% of business leaders predicted that ChatGPT would lead to layoffs, and 30% of employees report discomfort with being managed by an AI [39]. Regulation is catching up. The EU AI Act requires compliance by mid-2027 for high-risk applications, with employment decisions explicitly in scope, and New York City has already implemented rules limiting AI in hiring [55].

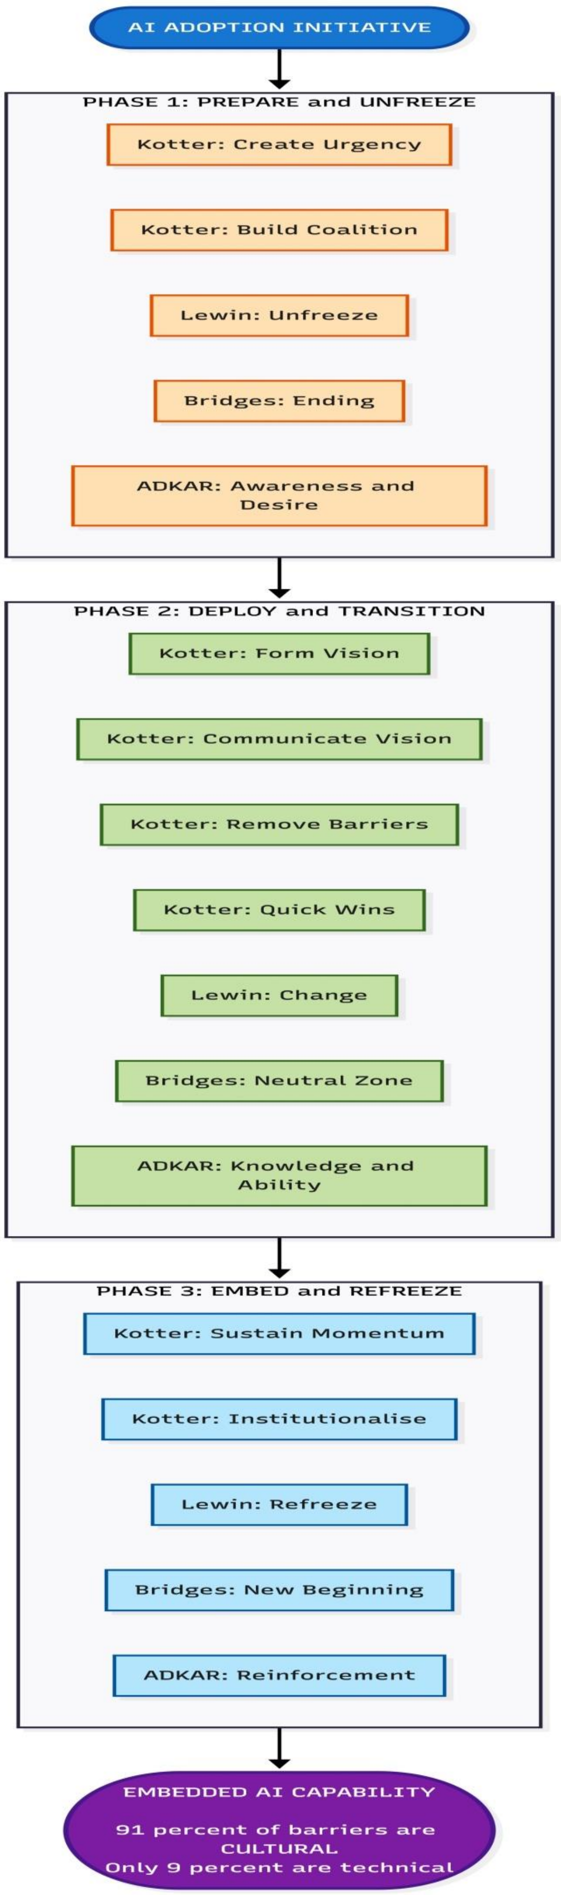

Figure 4. Integrated change-management frameworks (Kotter, Lewin, Bridges, ADKAR) mapped onto the AI-adoption timeline.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue V May 2026

Table 9. Ethical AI governance framework for leadership-assessment systems.

Governance Element	Key Practices	Risk Addressed	Measurement / Verification
Bias auditing	Third-party algorithmic audits, demographic parity testing, disparate-impact analysis	Algorithmic discrimination in hiring and promotion	Quarterly audits; <5% demographic parity gap target
Explainability	SHAP/LIME techniques, decision-justification reports, transparent scoring criteria	Black-box opacity undermining trust	Employee surveys on AI decision clarity; explanation adequacy ratings
Human oversight	Human-in-the-loop for termination and major promotion decisions; final hiring-decision review	Over-reliance on AI for high-stakes decisions	Override rates; appeals processed within 5 business days
Privacy protection	GDPR/CCPA compliance, data minimization, consent, access controls	Unauthorized data use and surveillance concerns	Annual compliance audits; zero unauthorized access incidents
Transparency communication	Clear AI disclosure, employee education, consent mechanisms	Employee distrust and resistance	≥80% employee understanding of AI usage; opt-out availability

Source: [55,53,27,34].

DISCUSSION

Pulling the evidence together yields three findings worth sitting with. None of them is especially flattering to how the industry currently talks about AI in HR.

The first is that the technology works. The 90%+ accuracy figures from MLbased leadership assessment are real. So is the 80% improvement in prediction over traditional methods, and the 40–75% time savings in hiring. These figures hold up across the six case studies, and they matter most where human judgement runs out of bandwidth. Picking the highest-potential people from a pool of 50000 employees is not a task a human team performs well. It is one AI performs well, given clean data.

The second is that the technology working is not sufficient. The Edelman finding (a 21-point gap in AI comfort between high-grievance and low-grievance employees) is the result to dwell on. It inverts the usual order of operations. Trust does not appear because an AI rollout went well; the AI rollout goes well because trust was already in place. Firms that lead with the technology in the hope that trust will follow run into resistance that no amount of model accuracy or explainability tooling fixes. The reverse holds equally: in high-trust environments, the same tools deploy faster and get used more, because employees begin from curiosity rather than from suspicion.

The third is that maturity-stage data shows where most firms are stuck. Only the 38% of organisations in Stages 3–4 outperform their industries financially. The remaining 62% have invested in AI without yet seeing the return. That is the figure CFOs and boards should be tracking, and the route out of Stage 2 is not another tool purchase. It is the cultural shift away from command-and-control management and toward coaching and development. The technology procurement is the easy part. Changing how leaders lead is the part most firms underinvest in.

Strategic Recommendations for Organisations

Six practices consistently separate successful AI rollouts from the ones that quietly fail. Most of these practices are unglamorous, which probably explains why they get skipped.

Begin with baseline metrics, before any vendor conversation. Current costper-hire. Time-to-fill. Leadership-effectiveness scores. Engagement. Retention. AI's return cannot be calculated without a credible before-and-after comparison, and the right first problem for AI to solve cannot be chosen without knowing where your weakest existing data already sits.

Run a trust assessment next. The Great Place to Work Trust Index and the Gallup Employee Engagement Survey are both valid instruments. If the resulting score lands below the 50th percentile for your industry, the AI rollout is the wrong priority. The trust deficit is. Fix that first.

Phase the rollout to match your actual maturity stage, rather than the stage you wish you were in. Stage 1 organisations attempting Stage 3 rollouts produce the most expensive failure mode documented in this literature.

Stand up a cross-functional governance council with HR, IT, Legal, Compliance and Ethics in the room from the start, not after the first incident. The well-publicised AI failures in HR consistently trace back to organisations where one of those five was absent from the design conversation.

Budget for change management properly. The data is consistent: 91% of the barriers to AI adoption are cultural rather than technical, which means 30–40% of the AI project budget going to change-management work is justified rather than extravagant.

Keep humans in the loop for the high-stakes calls. Hiring. Promotions. Terminations. Compensation changes above a meaningful threshold. AI works well as a second opinion on these decisions, but it is not yet good enough to make them alone, and the legal exposure of pretending otherwise is rising fast.

CONCLUSION

The short version of what this paper found: the case for AI in leadership and trust work holds, but only under specific conditions, and those conditions are not the ones vendors typically discuss.

The headline numbers are no longer in dispute. Across the six case studies, the published outcomes line up. Recruitment costs down 30–50%. Hiring time down 40–75%. Attrition predicted at 95% accuracy. Leadership-performance scores 14.9% higher. Google, Microsoft, IBM, Salesforce, Adobe and LinkedIn have all reached these outcomes in production deployments at meaningful scale, not in marketing pilots. The business case for the underlying technology has been settled.

What has not been settled, and what we have tried to make visible in this paper, is the role of trust. Organisations entering an AI rollout with strong pre-existing trust are 21 percentage points more comfortable adopting it, and 2.6× more likely to deploy it successfully [27]. Employees whose existing trust in their employer is high are 20 percentage points more open to AI at work [17]. These are not soft findings dressed up as data. They are data saying that the order of operations matters more than the order in which most companies do this work. Invest in leadership development, psychological safety, transparent communication and fair decision-making first, and the AI rollout becomes the easier project. Try to skip ahead, and the technology will not rescue you.

The maturity-curve research lands in the same place from a different direction.

The gap between Stage 2 and Stage 3 is where most of the financial return on AI sits, and crossing that gap is not a technology project. It is a redefinition of what senior leaders are for: less directive decision-making, more coaching and development. That redefinition requires explicit change management. Most organisations underinvest in this layer because it is harder to budget for than a software contract.

The future this paper points toward is one where AI surfaces high-potential talent earlier, where continuous feedback systems keep trust relationships healthier than annual reviews ever could, where personalised development moves faster, and where predictive analytics inform succession planning and retention. None of this will happen because of AI alone. It will happen, or fail to happen, based on whether the firms involved do the slower, less visible work alongside the technology rollout: building trust, leading the change, governing the algorithms honestly, and keeping humans inside the decisions that matter. The organisations that get this right will define what good leadership looks like for the next decade. The ones that get it wrong will have spent a great deal of money and made themselves materially less trusted in the process.

Declaration. We have taken permission from competent authorities to use the images/data as given in the paper. In case of any dispute in the future, we shall be wholly responsible.

Acknowledgements. The authors thank the editors and anonymous reviewers of ICICC-2026 for their constructive feedback.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

REFERENCES

1. Zenger, J., Folkman, J.: 360-degree leadership assessment. Zenger Folkman Technical Report (2024). https://zengerfolkman.com/, last accessed 2026/05/11
2. Josh Bersin Associates: AI in talent management: A market landscape report. Industry Research Report (2024)
3. Agility at Scale: Proving ROI – Measuring the business value of enterprise AI (2024). https://agility-at-scale.com/implementing/roi-of-enterprise-ai/, last accessed 2026/05/11
4. SHRM: Building trust in AI: Insights from Deloitte & Edelman research (2024). https://www.shrm.org/, last accessed 2026/05/11
5. Bass, B.M.: Leadership and Performance Beyond Expectations. Free Press, New York (1985)
6. Burns, J.M.: Leadership. Harper and Row, New York (1978)
7. Estherita, J., Shanmugam, M.: Influence of artificial intelligence on transformational leadership. ResearchGate Preprint (2024). https://www.researchgate.net/ publication/381145966, last accessed 2026/05/11
8. Luthans, F., Avolio, B.J.: Authentic leadership development. In: Cameron, K.S., Dutton, J.E., Quinn, R.E. (eds.) Positive Organizational Scholarship. BerrettKoehler, San Francisco (2003) Walumbwa, F.O., Avolio, B.J., Gardner, W.L., Wernsing, T.S., Peterson, S.J.:
9. Authentic leadership: Development and validation of a theory-based measure. Journal of Management 34(1), 89–126 (2008)
10. ResearchGate: Leveraging natural language processing to analyze employee feedback for enhanced HR insights (2024). https://www.researchgate.net/, last accessed 2026/05/11
11. Greenleaf, R.K.: Servant Leadership: A Journey into the Nature of Legitimate Power and Greatness. Paulist Press, New York (1977)
12. Boise State University: Distinguishing servant leadership from transactional and transformational leadership (2024). https://scholarworks.boisestate.edu/ipt_ facpubs/, last accessed 2026/05/11
13. Bay Path University: Is AI the ultimate adaptive leadership challenge? (2024). https://www.baypath.edu/news/, last accessed 2026/05/11
14. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An integrative model of organizational trust. Academy of Management Review 20(3), 709–734 (1995)
15. McAllister, D.J.: Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. Academy of Management Journal 38(1), 24–59 (1995)
16. Great Place to Work Institute: The business case for a high-trust culture (2024). https://www.greatplacetowork.com/our-model, last accessed 2026/05/11
17. Edelman Trust Institute: 2024 Edelman Trust Barometer: Special report – Trust at Work (2024). https://www.edelman.com/trust/2024/trust-barometer/, last accessed 2026/05/11

18. Psico-Smart: Utilizing AI and machine learning in leadership potential evaluation (2024). https://psico-smart.com/, last accessed 2026/05/11
19. ACM: Advanced machine learning algorithms for HR analytics. Communications of the ACM (2025). https://cacm.acm.org/, last accessed 2026/05/11
20. Center for Creative Leadership: Benchmarks & measure with a 360 leadership assessment (2024). https://www.ccl.org/leadership-solutions/, last accessed 2026/05/11
21. AIIR Analytics: 360-degree leadership assessment: How executive feedback systems drive competitive advantage (2024). https://www.aiiranalytics.com/blog/, last accessed 2026/05/11
22. Qualtrics: 360 feedback for leadership development (2024). https://www.qualtrics. com/, last accessed 2026/05/11
23. Leadership Circle: The Leadership Circle Profile: 360-degree leadership assessment (2024). https://leadershipcircle.com/, last accessed 2026/05/11
24. Frontiers in AI: Measuring trust in artificial intelligence: Validation of an established scale and its short form (2025). https://www.frontiersin.org/journals/ artificial-intelligence/, last accessed 2026/05/11
25. National Institute of Standards and Technology: Measuring trust in AI systems. NIST Technical Note (2024). https://www.nist.gov/, last accessed 2026/05/11
26. LumApps: Employee sentiment: Understanding and enhancing workplace dynamics (2024). https://www.lumapps.com/, last accessed 2026/05/11
27. Deloitte: Building trust in AI: Insights from Deloitte & Edelman research. SHRM publication (2024). https://www.shrm.org/topics-tools/flagships/ ai-hi/building-trust-in-ai, last accessed 2026/05/11
28. Gallup: How to build trust in the workplace (2024). https://www.gallup.com/ workplace/393401/trust-decline-rebuild.aspx, last accessed 2026/05/11
29. Weill, P., Woerner, S.L., Sebastian, I.M.: What's your company's AI maturity level? MIT Sloan Management Review and MIT CISR (2024). https://mitsloan.mit. edu/ideas-made-to-matter/, last accessed 2026/05/11
30. MITRE Corporation: The MITRE AI maturity model and organizational assessment tool guide. MITRE Technical Report (2025). https://www.mitre.org/, last accessed 2026/05/11
31. Google re:Work: Following the data: The research behind great managers at Google (2024). https://rework.withgoogle.com/intl/en/guides/ managers-research-behind-great-managers, last accessed 2026/05/11
32. BetterUp: Project Oxygen: An inside look at what makes a good manager (2024). https://www.betterup.com/blog/project-oxygen, last accessed 2026/05/11
33. New Age Leadership: Google's Project Oxygen: How Google identifies potential leaders (2024). https://newageleadership.com/, last accessed 2026/05/11
34. Microsoft Careers: Microsoft launches People Skills in Copilot (2025). https:// careers.microsoft.com/, last accessed 2026/05/11
35. HR Executive: 3 key learnings from Microsoft's journey to adopt AI in HR (2024). https://hrexecutive.com/, last accessed 2026/05/11
36. Deel: Employee development at Microsoft: Nurturing talent in the AI era (2024). https://www.deel.com/blog/employee-development-at-microsoft/, last accessed 2026/05/11
37. IBM Think: Artificial intelligence for human resources (2024). https://www.ibm.com/think/topics/ai-in-hr, last accessed 2026/05/11
38. TMI: Understanding IBM's cutting-edge HR strategy (2024). https://www.tmi. org/, last accessed 2026/05/11
39. Workday: IBM Watson Talent: The business case for AI in HR (2024). https://forms.workday.com/, last accessed 2026/05/11
40. Salesforce Careers: Everyone's an Einstein: How Salesforce is investing in employee growth with AI upskilling (2025). https://careers.salesforce.com/, last accessed 2026/05/11
41. Adobe Blog: Announcing a new chapter in fostering a culture of learning at Adobe (2023). https://blog.adobe.com/en/publish/2023/06/22/announcing-new-chapter-fostering-culture-learning-adobe, last accessed 2026/05/11
42. Mokahr: Talent and culture strategy at Adobe (2024). https://www.mokahr.io/, last accessed 2026/05/11
LinkedIn Careers: Reimagining hiring and learning with the power of AI (2024). https://www.linkedin.com/business/talent/, last accessed 2026/05/11

43. Brandon Hall Group: LinkedIn's talent solution: AI-powered hiring assistant transforms recruitment (2024). https://brandonhall.com/, last accessed 2026/05/11
44. Kotter, J.P.: Leading Change. Harvard Business Review Press, Boston (1996)
45. Kotter, J.P.: Accelerate: Building Strategic Agility for a Faster-Moving World. Harvard Business Review Press, Boston (2014)
46. Splunk: Kotter's 8 steps for leading change in organizations (2024). https://www. splunk.com/, last accessed 2026/05/11
47. VisualSP: How to use the Kotter change-management model for organizational transformation (2024). https://www.visualsp.com/, last accessed 2026/05/11
48. Bridges, W.: Managing Transitions: Making the Most of Change. Addison-Wesley, Reading, MA (1991)
49. William Bridges Associates: What is transition? (2024). https://wmbridges.com/ about/what-is-transition/, last accessed 2026/05/11
50. Lewin, K.: Frontiers in group dynamics: Concept, method and reality in social science. Human Relations 1(1), 5–41 (1947)
51. Hiatt, J.M.: ADKAR: A Model for Change in Business, Government and Our Community. Prosci Learning Center Publications, Loveland (2006)
52. TMI: Ethical AI in HR: Challenges, risks and best practices (2024). https://www.tmi.org/blogs/, last accessed 2026/05/11
53. Gartner: 2024 diversity and inclusion benchmark report. Gartner Research (2024) 55. ISACA: Using DTEF to achieve trustworthy AI. ISACA White Paper (2024). https://www.isaca.org/resources/white-papers/2024/, last accessed 2026/05/11
54. Solanki, P., et al.: AI-powered leadership: Transforming management practices with NLP-based sentiment analysis and graph neural networks. Springer (2025).https://link.springer.com/, last accessed 2026/05/11
55. Taylor & Francis Online: Revisiting the classic ABI model of trustworthiness: Interactive effects of trustworthiness components on trust (2024). https://www.tandfonline.com/, last accessed 2026/05/11
56. Blanchard LeaderChat: New research identifies how trust impacts employee work passion (2017). https://leaderchat.org/2017/02/16/ new-research-identifies-how-trust-impacts-employee-work-passion/, last accessed 2026/05/11