

Middle Managers and AI-Assisted Performance Appraisal in Chinese Enterprises: Explainability, AI Literacy, and the Case for Human Oversight

Peng Zhang, Rozaini Binti Rosli*

School of Business & Management, Lincoln University College, Petaling Jaya 47301, Selangor, Malaysia

*Corresponding author: Rozaini Binti Rosli, rozaini@lincoln.edu.my

Copyright: 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY-NC 4.0), permitting distribution and reproduction in any medium, provided the original author and source are credited, and explicitly prohibiting its use for commercial purposes.

Abstract: As AI-based appraisal tools move from pilot projects into daily HR routines, middle managers face a practical dilemma. They are asked to use algorithmic scores and model-generated recommendations in rating, feedback, promotion, and compensation decisions, yet they remain accountable for explaining those decisions to employees. This article offers a structured review of research on AI-assisted performance appraisal, algorithmic management, explainability, AI literacy, and procedural fairness. Rather than treating managerial confidence as a simple matter of technological acceptance, the article argues that what matters is calibrated reliance: the ability to use AI outputs seriously without surrendering judgment to them. Three conditions recur across the literature. First, explainability reduces procedural opacity and makes appraisal outcomes easier to defend. Second, AI literacy equips managers to interpret outputs, detect limitations, and resist both blind trust and reflexive rejection. Third, human oversight preserves accountability, contextual correction, and respectful treatment in employee-facing decisions. Building on these themes, the article develops an integrative framework linking explainability, AI literacy, and oversight to middle managers' reliance on AI-assisted appraisal. The discussion then considers practical implications for Chinese enterprises, where performance evaluation often carries strong consequences for pay, promotion, and internal mobility. The paper contributes to the AI-HRM literature by shifting attention from adoption alone to the managerial conditions under which AI-supported appraisal becomes usable, legitimate, and organizationally defensible.

Keywords: AI-Assisted Performance Appraisal; Middle Managers; Explainability; AI Literacy; Human Oversight; Algorithmic Fairness; Chinese Enterprises

Published: Apr 27, 2026

DOI: <https://doi.org/10.62177/apemr.v3i2.1342>

1. Introduction

Artificial intelligence has moved from the margins of human resource management into core managerial routines. What began with recruitment screening and workforce analytics is now extending into performance appraisal, where AI systems can summarize behavioral traces, detect patterns across large datasets, flag performance deviations, and recommend ratings or follow-up actions^[1]. For organizations, the attraction is obvious: standardization, scale, and speed. For middle managers, the situation is less straightforward. They are often the people who must interpret a score, turn it into feedback, and defend its consequences to both employees and senior leadership.

That position matters. Performance appraisal is not a purely technical exercise. It is also a social and political one. An

appraisal affects pay progression, promotion, developmental opportunities, and informal standing inside the firm. Even when an algorithm appears to improve analytical consistency, the final decision still unfolds in a relationship between manager and employee. This is why AI in managerial work is better understood as a problem of human-AI symbiosis than as simple substitution ^[2]. The practical question is not whether algorithms can process performance data, but under what conditions managers can rely on algorithmic outputs without abandoning accountability.

Recent research has made this tension harder to ignore. Studies of algorithmic management show that algorithms do not merely support work; they can also redistribute control, discretion, and authority in contested ways ^[3]. Parallel reviews in AI and HRM have drawn attention to the promise of AI-enabled decision systems, but they also note persistent problems of opacity, fairness, ethical responsibility, and managerial capability ^[4-6]. These concerns become especially sharp in performance appraisal because the decision is both consequential and visible. Unlike a back-office optimization model, an appraisal score must be explained to a human being who is personally affected by it.

The present article argues that middle managers' use of AI-assisted appraisal is shaped less by abstract enthusiasm for technology than by three linked conditions: explainability, AI literacy, and human oversight. Explainability matters because opaque outputs are difficult to justify, difficult to challenge, and difficult to communicate. AI literacy matters because managers need enough conceptual understanding to read model outputs critically rather than ritualistically. Human oversight matters because appraisal decisions rarely fit the clean boundaries of structured data; they involve context, exception handling, and interpersonal legitimacy.

This paper offers a structured review of recent scholarship and develops a framework of calibrated reliance. Calibrated reliance refers to a mode of use in which managers neither reject AI by default nor defer to it automatically. Instead, they treat AI as a serious but revisable input. The article contributes to the literature in three ways. First, it brings together streams that are often discussed separately: AI in HRM, algorithmic fairness, explainability, AI literacy, and human oversight. Second, it recenters middle managers as the key interpretive layer in AI-assisted appraisal. Third, it translates these insights into implications for Chinese enterprises, where digital HR systems are expanding while appraisal decisions remain tightly tied to managerial authority and employee relations.

2. Review Design and Analytical Lens

This article is a structured review rather than a meta-analysis. It draws mainly on peer-reviewed work published between 2018 and 2025, the period in which AI-assisted HRM, algorithmic management, explainable AI, and AI literacy became a more coherent stream of management and information-systems research. The review was guided by keyword combinations such as 'AI-assisted performance appraisal', 'algorithmic management', 'algorithmic fairness', 'explainability', 'AI literacy', 'human oversight', and 'employee performance evaluation'. Priority was given to studies that speak directly to managerial judgment, employee reactions, organizational design, or the use of AI in people management.

The review deliberately excludes purely technical papers that optimize model performance while saying little about how managers use, communicate, contest, or take responsibility for algorithmic outputs. That choice follows the central premise of this paper: in performance appraisal, the organizational problem is not only predictive accuracy. It is also interpretability, legitimacy, and the translation of machine output into managerial action.

Three analytical questions organize the discussion. First, what makes an AI-supported appraisal understandable enough for a manager to explain and defend? Second, what kind of literacy allows managers to use AI outputs critically rather than passively? Third, what forms of human oversight preserve accountability without reducing oversight to a symbolic gesture? These questions are closely aligned with the literature on opacity in algorithm-based HRM, perceptions of procedural fairness, and the design of human-AI decision systems ^[7,10-12].

Table 1. Analytical themes guiding the review

Theme	Core question	Managerial relevance
Explainability and opacity	Can managers understand, justify, and communicate how the system reached its recommendation?	Without some level of intelligibility, managers struggle to defend ratings and employees are more likely to question the process ^[7,9,11] .

Theme	Core question	Managerial relevance
Fairness and employee reactions	How do employees interpret AI-based evaluation relative to human judgment?	Appraisal acceptance depends on perceived fairness, respectful treatment, and whether the process looks contestable rather than arbitrary ^[10,19,20] .
AI literacy	What do managers need to know in order to use AI outputs well?	Useful literacy is not programming skill alone; it includes recognizing data limits, model assumptions, and appropriate use conditions ^[16-18] .
Human oversight	When should managers follow, revise, or override the system?	Oversight matters when structured indicators miss context, when exceptional cases arise, and when responsibility must remain attributable ^[5,6,8] .
Reliance dynamics	Why do some users distrust algorithms while others over-trust them?	Managerial use sits between algorithm aversion and algorithm appreciation, making calibration more important than simple acceptance ^[12-15] .

3. Explainability and the Problem of Opaque Appraisal

Opacity is one of the defining problems of algorithm-based HRM. Langer and Konig describe opacity not as a purely technical defect but as a multi-stakeholder issue involving users, affected employees, deployers, developers, and regulators^[7]. In appraisal settings, opacity becomes immediately practical. A middle manager who cannot say why the system produced a score is left with two weak options: either repeat the output as if it were self-validating, or quietly replace it with intuition. Neither option produces robust organizational legitimacy.

Research on algorithmic decisions helps explain why this matters. Lee's work on perceptions of algorithmic decisions shows that algorithmic judgments are often seen as less fair and less trustworthy than human ones, especially when the decision affects people in recognizable social contexts^[9]. The problem is not only that the system may be wrong. It is also that people have difficulty locating intention, reason, and recourse in an algorithmic decision. In performance appraisal, that difficulty has immediate consequences because employees do not simply receive a score; they interpret what the score says about their standing, effort, and future prospects.

The design literature makes a similar point from the organizational side. Robert and colleagues argue that fair AI in employee management must be designed around distributive, procedural, and interactional justice, rather than around statistical performance alone^[10]. This is particularly relevant for appraisal. A system may appear distributively efficient by generating consistent scores, yet still fail procedurally if managers cannot explain the pathway from data to outcome, or interactionally if the decision is delivered without recognition of employee context. Explainability, then, should not be understood as a technical add-on. It is part of procedural fairness.

Shin's work strengthens this argument by showing that explainability and causability shape trust, perception, and acceptance^[11]. In managerial practice, explainability does at least three things. First, it gives managers a basis for checking whether the output is coherent with the case at hand. Second, it helps them translate system logic into a form that can be discussed with employees. Third, it makes disagreement possible. That last point is easy to overlook. A fully opaque system does not merely reduce understanding; it also weakens the conditions for meaningful challenge.

At the same time, explainability should not be romanticized. Not every model can be rendered transparent in a way that is both accurate and managerially useful. And more explanation is not always better if it overwhelms the user with technical detail. For middle managers, what matters is actionable explanation: enough clarity to identify the main drivers of a recommendation, enough traceability to see what data mattered, and enough procedural documentation to justify following or modifying the recommendation. That threshold is lower than full model transparency but higher than mere faith in the vendor or system owner.

Seen this way, explainability is less about turning every manager into a data scientist than about preserving the chain of managerial reasoning. If a manager cannot reconstruct why a recommendation deserves weight, the system may still calculate, but it cannot genuinely support judgment. In a domain as consequential as performance appraisal, that weakness is likely to undermine decision confidence, employee acceptance, or both.

4. AI Literacy as Managerial Interpretive Capacity

Explainability alone does not solve the managerial problem. Even a relatively transparent system can be misread, over-trusted, or rejected if the user lacks the conceptual tools to interpret it. This is where AI literacy enters the discussion. In many organizations, literacy is still treated too narrowly, as if it meant coding expertise or technical familiarity with machine learning pipelines. For appraisal practice, that is the wrong benchmark. The relevant question is whether managers know enough to use the system with discernment.

Behavioral research helps clarify why this matters. Users do not approach algorithms neutrally. Some display algorithm aversion, especially after seeing an algorithm make a mistake^[13]. Others can be induced to rely more on algorithms when they retain even a small degree of control over the final output^[14]. Still other studies document algorithm appreciation, showing that people sometimes prefer algorithmic advice to human advice when they believe the model is competent and objective^[15]. Together, these findings suggest that managerial responses to AI are unstable rather than fixed. Managers may oscillate between skepticism and deference depending on task design, error experience, and perceived agency.

AI literacy provides a way of stabilizing that response. Long and Magerko define AI literacy in terms of the competencies people need to interact with and critically evaluate AI systems^[16]. Ng and colleagues extend the idea by framing AI literacy as a multidimensional construct that includes conceptual understanding, critical interpretation, and awareness of social and ethical implications^[17]. Applied to middle managers, this means literacy should include at least four practical capacities: understanding what kind of data feed the model, recognizing that outputs are probabilistic rather than self-evident truths, identifying when important contextual information is missing, and knowing when escalation or override is warranted.

This matters because appraisal often involves mixed evidence. Some elements of performance are highly structured: sales volume, task completion speed, attendance regularity, response times, or customer ratings. Other elements are not. Mentoring effort, crisis handling, team stabilization, political judgment, or the repair of damaged client relationships often leave weaker digital traces. A manager with low AI literacy may confuse data richness with evaluative completeness. A manager with stronger literacy is more likely to treat the model as informative but partial.

Recent work at the executive level is useful here. Pinski, Hofmann, and Benlian show that AI literacy in top management is related to AI orientation and implementation ability^[18]. The same logic applies one level down. Middle managers do not need to build the system, but they do need enough literacy to use it responsibly. In appraisal settings, literacy is therefore a form of interpretive capacity. It helps managers distinguish signal from measurement convenience, and confidence from unwarranted certainty.

This point also reframes training. Many firms introduce AI tools with short operational briefings focused on dashboards, interfaces, and compliance. That is not enough. Managerial training should include model boundaries, common sources of error, bias risks, and case-based exercises on when to question the system. Without that layer, explainability may exist on paper while interpretive competence remains thin in practice.

In short, AI literacy helps middle managers avoid two symmetrical errors: refusing useful analytical support simply because it is algorithmic, and surrendering judgment simply because the output looks precise. Performance appraisal demands a position between these extremes. That position is learned, not automatic.

5. Why Human Oversight Still Matters

The third condition is human oversight. The phrase is used so widely in AI governance that it risks becoming ceremonial. In practice, oversight only matters if humans have both the authority and the competence to question, revise, or contextualize the output. If the manager is expected to rubber-stamp the score, 'human in the loop' becomes a legitimacy device rather than a real safeguard.

The ethics literature on AI and people management makes this point clearly. Varma and colleagues argue that AI in people management raises enduring questions of responsibility, fairness, and moral legitimacy that cannot be solved by efficiency claims alone^[5]. Rodgers and coauthors similarly emphasize the need to think about ethical decision architectures in HR processes rather than treating AI as a neutral computational layer^[6]. These arguments are directly relevant to appraisal

because appraisal decisions are among the most consequential decisions a manager communicates face to face.

Empirical studies reinforce the point. Tong and colleagues show that AI feedback can improve employee performance, but they also find that how AI is deployed and disclosed matters for employee response^[8]. This suggests that performance gains do not remove the communicative and relational dimension of feedback. A system may optimize recommendation quality while still creating friction if employees feel reduced to a score or monitored without meaningful explanation.

Even more directly, Qin and coauthors find that employees may perceive AI as fairer and more accurate than the average human manager in structured, data-intensive evaluation tasks^[19]. Yet their study also shows why human managers remain central: when employees are evaluated by human managers, perceived fairness plays a first-order role in shaping outcomes. In other words, managers do not retain importance because they are always more accurate. They retain importance because fairness in organizational life is not exhausted by calculative precision.

Chun and colleagues push this further by showing what algorithmic evaluation often fails to provide: respectful treatment and individualized consideration^[20]. That finding is critical for performance appraisal. Employees do not judge an appraisal only by whether the number appears unbiased. They also care whether the decision process acknowledges effort, circumstance, growth trajectory, and the feeling of having been treated as a person rather than a datapoint. This does not mean human managers are inherently fairer. It means that relational legitimacy still matters, and algorithms do not automatically supply it.

For middle managers, then, oversight has at least three substantive functions. First, it is an error-correction mechanism: managers can spot omissions, misclassifications, or unusual cases. Second, it is a contextualization mechanism: they can interpret results in light of organizational realities that are weakly captured in the data. Third, it is an accountability mechanism: they can explain the basis of the decision and bear responsibility for its use. These functions are especially important when appraisal outcomes affect salary, bonus allocation, promotion, or disciplinary action.

A useful oversight model is therefore neither full managerial discretion nor full algorithmic autonomy. It is a disciplined hybrid. Managers should be able to override the system, but overrides should require written reasons. Likewise, following the system should also be documented when the case is consequential. That two-sided documentation matters because it prevents oversight from becoming selective or cosmetic. It also produces an audit trail that can support learning, review, and later system improvement.

6. A Framework of Calibrated Reliance

The literature reviewed above suggests that middle managers' use of AI-assisted appraisal is best understood through the lens of calibrated reliance. Calibrated reliance is not simple trust. It is a practical equilibrium in which the manager assigns weight to AI output in proportion to its intelligibility, task fit, and evidentiary adequacy, while still preserving room for contextual judgment. This concept helps move the discussion beyond the familiar opposition between acceptance and resistance.

The framework proposed here links three inputs to calibrated reliance. The first is explainability. When the system offers intelligible reasons, salient input features, and some trace of how the recommendation was formed, managers are better able to judge whether the output deserves confidence. Explainability also supports communication with employees and creates conditions for contestability^[7,9,11].

The second input is AI literacy. Managers need enough knowledge to understand what the output means, what it does not mean, and where the model's blind spots are likely to be. Literacy therefore moderates the practical value of explainability. A transparent output is of limited use if the manager cannot interpret it, and a complex output is especially risky when users mistake statistical confidence for situational adequacy^[13-18].

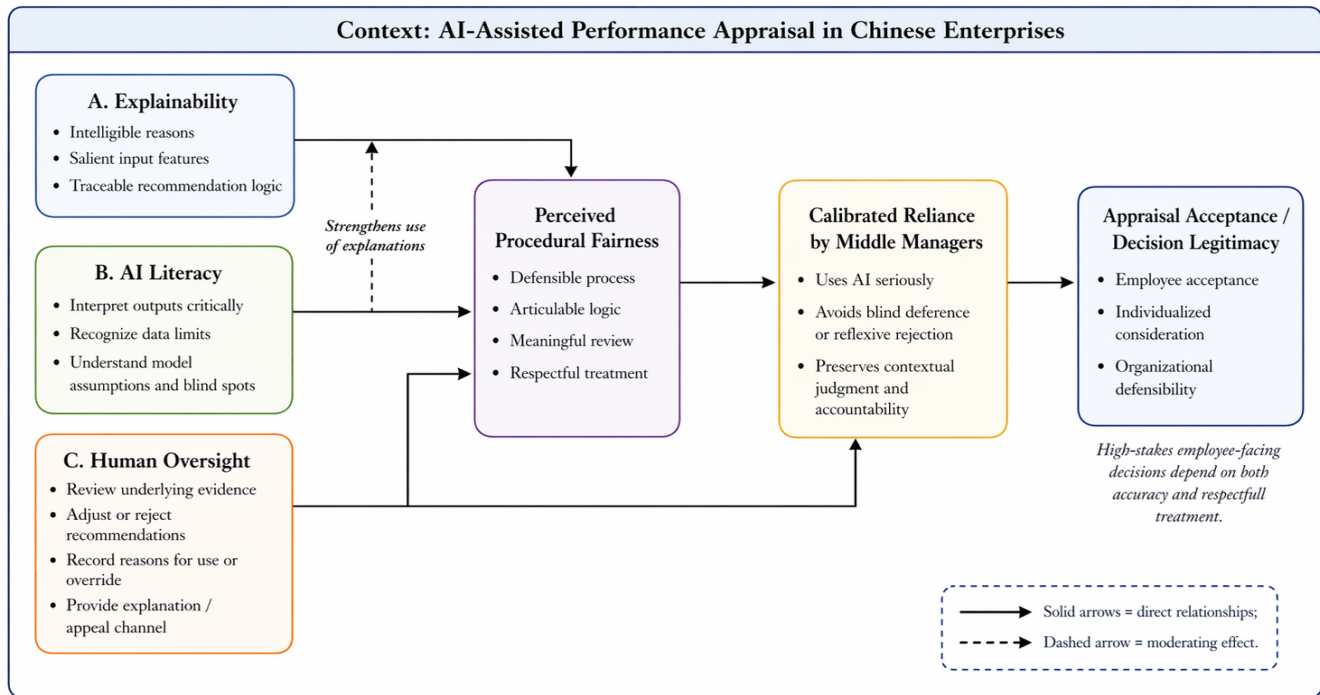
The third input is human oversight. Oversight converts AI from an authority that closes judgment into a tool that supports it. In the framework proposed here, oversight includes the ability to review underlying evidence, adjust or reject recommendations, record reasons, and provide employees with a credible avenue for explanation or appeal. This is what turns managerial involvement into a real governance mechanism rather than a symbolic presence^[5,6,19,20].

These three conditions converge through perceived procedural fairness. Managers are more likely to rely on AI outputs when the process looks defensible, when the logic can be articulated, and when meaningful review remains possible. Employees are more likely to accept AI-supported appraisal when they see that the manager has neither hidden behind the system nor

ignored its analytical value. Calibrated reliance, then, is relational as well as cognitive. It concerns the manager's confidence in using the system and the employee's willingness to treat the decision as legitimate.

On this basis, four propositions can guide future empirical work. First, greater explainability is likely to increase middle managers' reliance on AI-assisted appraisal by improving perceived procedural fairness. Second, AI literacy is likely to strengthen the positive effect of explainability because literate managers can make better use of the information provided by the system. Third, meaningful human oversight is likely to increase calibrated reliance by preserving accountability and contestability. Fourth, in employee-facing appraisal contexts, the effect of AI outputs on decision acceptance is likely to depend not only on perceived accuracy but also on whether the process leaves room for respectful treatment and individualized consideration.

Figure 1. Framework of calibrated reliance in AI-assisted performance appraisal



7. Implications for Chinese Enterprises

The framework has particular relevance for Chinese enterprises, many of which are expanding digital HR infrastructures while also operating in organizational environments where appraisal carries substantial consequences for career progression. In such settings, middle managers sit at a sensitive junction. They are expected to use new systems, demonstrate alignment with data-driven management, and at the same time maintain team stability and interpersonal credibility. That combination makes blind reliance on AI especially risky.

The first implication is design-related. Chinese enterprises should resist the temptation to introduce appraisal systems that produce highly visible scores without adequate explanation layers. A dashboard that displays ranking precision but cannot show the main drivers of the score will place managers in a weak communicative position. For high-stakes decisions, systems should at least display salient indicators, data sources, confidence cues, and a record of whether human adjustment occurred. Managers do not need exhaustive model documentation in daily use, but they do need enough information to defend a recommendation in plain organizational language.

The second implication concerns managerial development. Training should not stop at operational adoption. It should cultivate AI literacy as a managerial competence. Case-based training is especially important: managers should practice reading system outputs, identifying questionable recommendations, and explaining a decision to an employee whose self-assessment differs from the model's conclusion. This kind of training is more demanding than software onboarding, but it is far more relevant to appraisal quality.

The third implication is governance. Firms should formalize two-way accountability around overrides. When managers

depart from AI recommendations, they should state why. When they follow them in consequential cases, they should also record the basis for doing so. This creates a disciplined review process, reduces arbitrary deviations, and produces valuable organizational learning about when the model performs well and when it does not.

The fourth implication is employee-facing procedure. Organizations should offer a clear explanation and appeal channel for AI-supported evaluations. This does not mean every employee should litigate every score. It means there should be an intelligible path for questioning clearly anomalous results or missing contextual evidence. Such channels support legitimacy even when the original recommendation stands.

Finally, firms should be careful about the scope of automation. AI appears most defensible where performance indicators are relatively structured and repeatedly observed. Its role should be more limited where evaluation turns heavily on emergent contribution, tacit coordination, developmental potential, or crisis judgment. A prudent design principle is therefore asymmetry: the more punitive or career-shaping the decision, the stronger the case for deeper human review.

8. Conclusion

AI-assisted performance appraisal is not simply a technical upgrade of an existing HR routine. It changes how performance is seen, measured, explained, and justified inside the organization. For that reason, the key issue is not whether managers are generally 'for' or 'against' AI. The more useful question is whether organizational conditions allow them to rely on AI outputs in a measured and defensible way.

This review has argued that three conditions are central to that outcome: explainability, AI literacy, and human oversight. Explainability reduces opacity and supports justification. AI literacy enables managers to read outputs critically rather than ritualistically. Human oversight preserves accountability, contextual judgment, and respectful treatment. Taken together, these conditions support calibrated reliance: a mode of use in which AI informs appraisal without replacing managerial responsibility.

The argument matters for research and practice alike. For research, it suggests that future studies should move beyond broad measures of adoption or intention to use and examine how reliance is shaped by explanation quality, literacy levels, and oversight design. For practice, it suggests that organizations should treat appraisal not as a suitable site for full automation, but as a domain in which analytical systems and managerial judgment must be carefully assembled. That lesson is particularly relevant for Chinese enterprises as they deepen digital transformation while trying to preserve fairness, legitimacy, and organizational cohesion.

Funding

No

Conflict of Interests

The authors declare that there is no conflict of interest regarding the publication of this paper.

Reference

- [1] Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), 15-42. <https://doi.org/10.1177/0008125619867910>
- [2] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577-586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- [3] Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410. <https://doi.org/10.5465/annals.2018.0174>
- [4] Pan, Y., & Froese, F. J. (2023). An interdisciplinary review of AI and HRM: Challenges and future directions. *Human Resource Management Review*, 33(1), 100924. <https://doi.org/10.1016/j.hrmr.2022.100924>
- [5] Varma, A., Dawkins, C., & Chaudhuri, K. (2023). Artificial intelligence and people management: A critical assessment through the ethical lens. *Human Resource Management Review*, 33(1), 100923. <https://doi.org/10.1016/j.hrmr.2022.100923>

- [6] Rodgers, W., Murray, J. M., Stefanidis, A., Degbey, W. Y., & Tarba, S. Y. (2023). An artificial intelligence algorithmic approach to ethical decision-making in human resource management processes. *Human Resource Management Review*, 33(1), 100925. <https://doi.org/10.1016/j.hrmr.2022.100925>
- [7] Langer, M., & Konig, C. J. (2023). Introducing a multi-stakeholder perspective on opacity, transparency and strategies to reduce opacity in algorithm-based human resource management. *Human Resource Management Review*, 33(1), 100881. <https://doi.org/10.1016/j.hrmr.2021.100881>
- [8] Tong, S., Jia, N., Luo, X., & Fang, Z. (2021). The Janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance. *Strategic Management Journal*, 42(9), 1600-1631. <https://doi.org/10.1002/smj.3322>
- [9] Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5(1), 1-16. <https://doi.org/10.1177/2053951718756684>
- [10] Robert, L. P., Pierce, C., Marquis, L., Kim, S., & Alahmad, R. (2020). Designing fair AI for managing employees in organizations: A review, critique, and design agenda. *Human-Computer Interaction*, 35(5-6), 545-575. <https://doi.org/10.1080/07370024.2020.1735391>
- [11] Shin, D. (2021). The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. *International Journal of Human-Computer Studies*, 146, 102551. <https://doi.org/10.1016/j.ijhcs.2020.102551>
- [12] Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627-660. <https://doi.org/10.5465/annals.2018.0057>
- [13] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114-126. <https://doi.org/10.1037/xge0000033>
- [14] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3), 1155-1170. <https://doi.org/10.1287/mnsc.2016.2643>
- [15] Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90-103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- [16] Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- [17] Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- [18] Pinski, M., Hofmann, T., & Benlian, A. (2024). AI literacy for the top management: An upper echelons perspective on corporate AI orientation and implementation ability. *Electronic Markets*, 34(1), 1-23. <https://doi.org/10.1007/s12525-024-00707-1>
- [19] Qin, S., Jia, N., Luo, X., Liao, C., & Huang, Z. (2023). Perceived fairness of human managers compared with artificial intelligence in employee performance evaluation. *Journal of Management Information Systems*, 40(4), 1039-1070. <https://doi.org/10.1080/07421222.2023.2267316>
- [20] Chun, J. S., De Cremer, D., Oh, E. J., & Kim, Y. (2024). What algorithmic evaluation fails to deliver: Respectful treatment and individualized consideration. *Scientific Reports*, 14(1), 25996. <https://doi.org/10.1038/s41598-024-76320-1>