Trends in Social Sciences and Humanities Research
Print ISSN: 2959-992X
Online ISSN: 2959-9938
DOI: https://doi.org/10.61784/tsshr3223

THE ALIENATION OF BUSINESS ETHICS IN THE AGE OF ALGORITHMS: THE CASE OF BYTEDANCE'S AI RECRUITMENT SYSTEM

YiHan Ma
School of Business, Xi'an International Studies University, Xi'an 710128, Shaanxi, China.

Abstract: Amid the digital wave driven by the pursuit of extreme efficiency, ByteDance has achieved millisecond-level responses in human resources screening through its proprietary AI recruitment system; however, this has also triggered a profound crisis in business ethics. This study focuses on the operational mechanisms and controversies surrounding this system, conducting a penetrating analysis through a three-dimensional framework of "Justice-Dignity-Responsibility". The research finds that, while the system's core logic of "fitting historical optimality" has reduced recruitment costs, it has, at the technical level, solidified and amplified historical biases regarding dimensions such as educational background and region, thereby deviating from Rawls's "Principle of Difference". In terms of dignity, the system reduces job seekers to feature vectors, and its panopticon-style surveillance and evaluation trigger widespread "surveillance anxiety", leading to the dissolution of job seekers' agency. Regarding responsibility, the coupling of organizational incentive mechanisms with algorithmic black boxes causes a systemic breakdown in the chain of accountability, creating a vacuum where "no one is held responsible". This study reveals the deep-seated paradox of ethical compromise under efficiency-driven systems and argues that algorithmic governance must urgently shift from purely commercial logic toward a sustainable balance between efficiency and justice.
Keywords: AI recruitment; Business ethics; Age of algorithms; Human resources

1 INTRODUCTION

As the global digital economy deepens, human resource management is undergoing a governance-level restructuring from human-driven experience to algorithmic decision-making. ByteDance, a technology company with an extremely high concentration of talent, faces the challenges of processing massive volumes of resumes and high-frequency hiring demands. It has taken the lead in embedding an AI system based on the Transformer architecture into the entire recruitment process, aiming to address bottlenecks such as response delays, high costs, and fluctuating subjective standards inherent in manual processes. However, empirical data shows that over 70% of job seekers hold a negative view of the system, with core concerns centered on "algorithmic black box" anxiety stemming from opaque decision-making and inexplicable outcomes. While previous research has made significant strides in addressing algorithmic bias and regulatory mechanisms, it has often focused on either macro-level institutional frameworks or micro-level individual perspectives. The innovation of this section lies in combining the deconstruction of technical patents with social empirical data to construct a "Justice-Dignity-Responsibility" penetrative framework, which seeks to penetrate algorithmic logic and reach the root causes of commercial ethics. This research approach analyzes the technical logic of ByteDance's core patents and integrates social media sentiment data to systematically diagnose the alienation of algorithm-driven human resources governance across three dimensions: outcome fairness, procedural respect, and institutional accountability [1,2].

2 CASE PRESENTATION: OPERATION AND CONTROVERSIES OF BYTEDANCE'S AI RECRUITMENT SYSTEM

2.1 Overview of ByteDance's AI Recruitment System

2.1.1 System design purpose

As one of the technology companies with the highest talent density globally, ByteDance faces complex challenges in talent recruitment, including high frequency, large volume, fragmentation, and cross-region. Driven by the "go all out to achieve miracles" culture that pursues extreme efficiency, the traditional manual screening model can no longer support the speed and scale of its business expansion. To this end, ByteDance independently developed an AI recruitment system, aiming to systematically address three core bottlenecks:

First, delayed response leading to talent loss. Research shows that the recruitment window for top technical talents usually does not exceed 10 days. However, traditional manual processing of a mountain of resumes (an average of over 300 applications per position) often takes 5-7 working days, resulting in nearly 30% of outstanding candidates accepting other offers during the waiting period. After introducing the AI initial screening system, the resume processing time has been reduced to the second level, increasing the talent response speed by nearly 100 times, which significantly improves the candidate experience and corporate competitiveness. Second, mismatching between labor costs and value creation. According to analysis from the human resources industry, junior recruiters spend approximately 70%-80% of their

working time on repetitive resume screening and information verification. Based on ByteDance's handling of millions of applications each year, the initial screening link alone may consume hundreds of thousands of labor hours. Automating this link through AI not only significantly reduces operational costs but also frees the HR team from transactional work to focus on more strategically valuable organizational assessment and talent development work. Third, standard fluctuations caused by subjective evaluations. A study by Harvard Business School found that the consistency of evaluations of the same candidate by different interviewers is usually less than 60%, and personal emotions, fatigue levels, and even interview time slots may affect judgments. ByteDance's algorithmic model has increased the consistency of screening standards to over 85% through a unified evaluation framework, providing a reusable "measurement standard" for cross-departmental and cross-regional talent benchmarking, and reducing unfairness and talent misjudgment caused by subjective biases [3,4].

Based on the above challenges, ByteDance has built an AI recruitment system with deep learning and natural language processing as the core technologies. This is not only a tool upgrade but also the deepening practice of its "instrumental management" concept in the field of human resources—realizing full-link automation from talent identification, evaluation to onboarding through algorithm-driven methods, while improving operational efficiency, it also redefines the underlying logic of large-scale organizational talent screening.

2.1.2 Technical logic

CN111242334A patent technical diagram is shown in Figure 1.

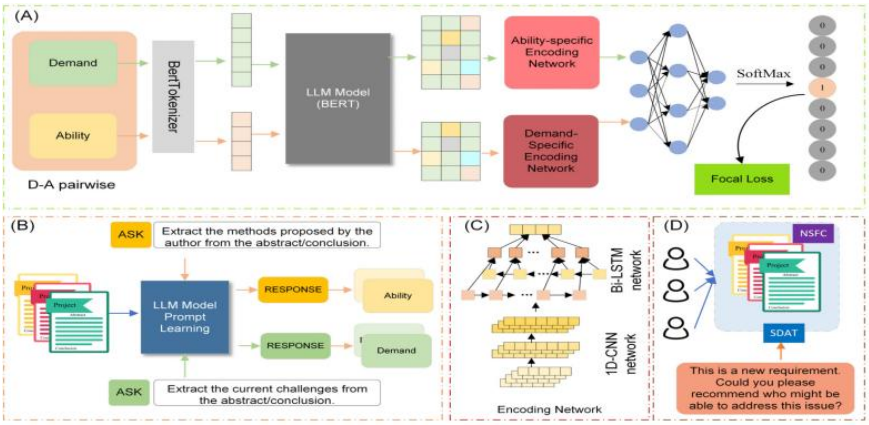

Figure 1 CN111242334A Patent Technical Diagram

Through in-depth decomposition of ByteDance's key patents (such as CN111242334A "A Talent Evaluation Model Training Method and Talent Evaluation Method"), the operating logic of the system can be summarized into three levels: First, the Feature Extraction layer: The system uses a bidirectional encoder (such as a BERT-like structure) to convert unstructured resume content into high-dimensional feature vectors. Explicit features include academic background (weighted by QS ranking), company tenure (calculated by job hopping frequency), and professional matching degree. Implicit features include "stability prediction" and "growth assessment" of candidates extracted through semantic correlation analysis. For example, the algorithm will analyze the promotion frequency of past positions, and even infer their execution ability through the rigor of resume formatting. Second, predicted score output and loss function. The patent flow chart shows that the system directly outputs the predicted score from feature calculation. The core goal of its model training is to reduce the residual between the "system predicted score" and the "historical actual hiring decision". This means that the system is essentially a "decision simulator", and its evolution direction is to continuously approach the established preferences of internal HR in the company. If the historical hiring data prefers certain "famous school and big factory" labels, the algorithm will automatically set these labels as high weights, forming a closed loop of bias. Finally, the multi-modal video interview evaluation level [5]. In the AI interview link, the system introduces computer vision (CV) and speech recognition technologies. The micro-expression monitoring system will capture the facial muscle movements of candidates when answering specific questions (such as eye tremors, muscle micro-tremors) and match them with the "honesty" or "stress resistance" model. In addition to text semantics, the speech feature analysis system will also extract physical features such as fundamental frequency and energy distribution in the speech. Sudden changes in speech rate, excessive pauses, or unsteady intonation may be judged by the algorithm as "lack of self-confidence".

2.1.3 Application scale and industry penetration

Currently, the system has been deeply embedded in ByteDance's entire Human Resource Management (HRM) system. According to internal data, the coverage rate of its AI screening has reached more than 80% of the total recruitment volume. In some highly standardized functional positions (such as sales, operations, customer service), AI even has a "veto power". This large-scale application has made ByteDance a de facto benchmark in the field of "algorithmic recruitment" in China, and also made it the center of ethical controversies [6,7].

2.2 Collation of Ethical Conflict Incidents

2.2.1 Analysis of social media public opinion data

To empirically investigate the real reactions of job seekers to the system, this study collected and analyzed more than 800 valid job seeker feedback data from platforms such as Maimai, Zhihu, and Xiaohongshu during the period from October 2024 to September 2025.

ByteDance AI interview social media public opinion analysis is shown in Figure 2.

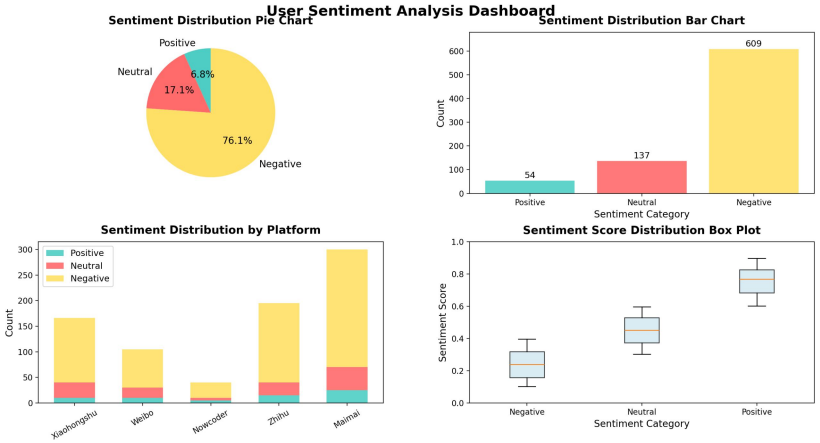

Figure 2 ByteDance AI Interview Social Media Public Opinion Analysis

The data clearly reveals the general negative sentiment [8]. Statistics show that as many as 76.1% of respondents expressed explicit negative attitudes. Their descriptions are full of emotional tension, such as "the process is as cold as playing against a machine", "felt a sense of humiliation being examined by algorithms", and "the interview is no longer a communication but a carefully calculated expression performance". These intuitive feelings of being examined and objectified are highly consistent with the "applicant reactions" highlighted in international academic research (WOS burst term), pointing to a core issue: algorithms are stripping away the human touch and basic respect in recruitment. Among them, the "panopticon"-style surveillance anxiety is particularly prominent. On the Maimai platform, negative voices account for the highest proportion (37.2%) in related discussions. Many job seekers described that the mandatory full-time camera monitoring and real-time micro-expression analysis in AI interviews have plunged them into persistent "performance anxiety", feeling that their most private emotional space has been ruthlessly invaded and evaluated by technology. This has exceeded the scope of efficiency tools and become a soft coercion on individual boundaries and psychological safety.

High-frequency words in negative comments is shown in Figure 3.

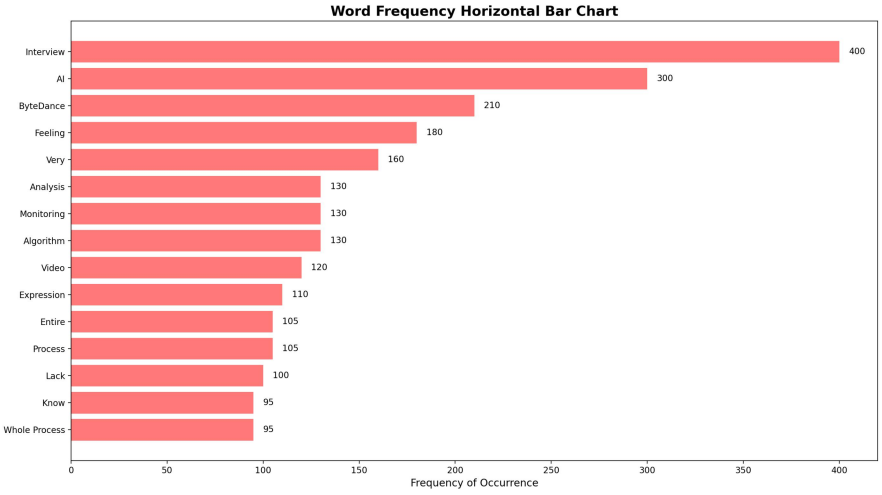

Figure 3 High-frequency Words in Negative Comments

Through word cloud and content analysis of the feedback texts, two core pain points have emerged: first, the widespread doubt about "algorithmic bias", which was directly mentioned 126 times. Candidates strongly suspect that the system has implicit discrimination and unfair downgrading against candidates with "non-first-tier university" backgrounds, older practitioners, or female job seekers. This widespread social suspicion is exactly the real focus reflected by the CNKI core cluster "#2 Algorithmic Discrimination" and the high burst term "Employment Discrimination" (intensity 3.14).

Second, the more profound "technical deprivation of dignity", with relevant expressions appearing 174 times. Job seekers frequently use words such as "feeling like a commodity", "being disassembled into data points", and "the score

determines everything about me", feeling that their rich personality, experience, and potential are brutally flattened in the face of algorithmic scoring. This process of datafying humans into feature vectors directly echoes the ethical dilemma of "personality assessment" studied by the WOS cluster "#11 Big Five", constituting a systematic dissolution of individual subjectivity [9,10].

The public opinion data not only confirms that the system is widely questioned in terms of fairness, i.e., the justice dimension, but also profoundly reveals the universal harm it causes to human dignity, i.e., the dignity dimension. When the interview is alienated from a two-way interactive "intersubjective communication" into a one-way passive "data dissection", its ethical conflict has escalated from technical inaccuracy to a value crisis.

2.2.2 Core ethical controversial points

(1) According to the technical patent flow chart, the path from feature extraction to score output is highly automated, lacking necessary nodes for "manual intervention review" or "ethical bias correction". This leads to when the algorithm filters candidates based on biased features, such as school labels, HR is often in an "unconscious" state and cannot effectively supervise and intervene in the machine's decision. Recruitment decisions thus degenerate from professional judgments that require justification into unexplainable automated administrative rulings, overriding the transparency and accountability principles required by procedural justice.

(2) The optimization goal of the algorithm is set to "reproduce historical hiring decisions as accurately as possible". This design logic essentially embeds the risk mechanism of "algorithmic discrimination" (CNKI cluster #2) into the core of the system. If the enterprise's past recruitment data implicitly contains preferences for a certain gender, age, or academic background, whether reasonable or not, the algorithm will identify it as a "successful model" to learn and strengthen. This means that AI systems not only cannot correct historical biases but may also wrap those discriminatory employment practices in the technical cloak of "data-driven" and "scientifically objective", making discriminatory behaviors more hidden and harder to challenge, forming a vicious circle.

(3) In the algorithm's evaluation framework, vivid and three-dimensional job seekers are reduced to a set of computable and comparable feature vectors. The interview, which should be a social interaction based on two-way communication and in-depth understanding, is thus alienated into a one-way, cold data dissection. This not only seriously erodes individual dignity, reducing them from "ends" to "means", fundamentally violating the basic principles of Kantian ethics, but also may lead to systematic deviations in selection results. That is, what the algorithm ultimately screens out may be "people who are best at playing the role expected by the algorithm in front of the camera", not necessarily "people who have the core competencies and potential required for the position". This fundamentally shakes the validity and legitimacy of talent selection.

2.3 Current Situation Review

2.3.1 Limitations of the enterprise's existing algorithmic governance measures

(1) The algorithm filing has the characteristics of "selective transparency". Although the enterprise has completed the algorithm filing in accordance with regulations, the filing content usually only includes stating the use of technical frameworks such as "deep learning" and "natural language processing". The public and regulatory authorities have no way of knowing the core that determines the fate of job seekers—such as how the weights of specific features (such as educational background, work experience) are allocated and how the decision threshold is set. This "black box" state makes it difficult to verify or falsify allegations of "algorithmic discrimination", and also makes the governance problem pointed out by the CNKI burst term "algorithmic black box" (intensity 2.87) a reality.

ByteDance AI recruitment algorithm filing is shown in Figure 4.

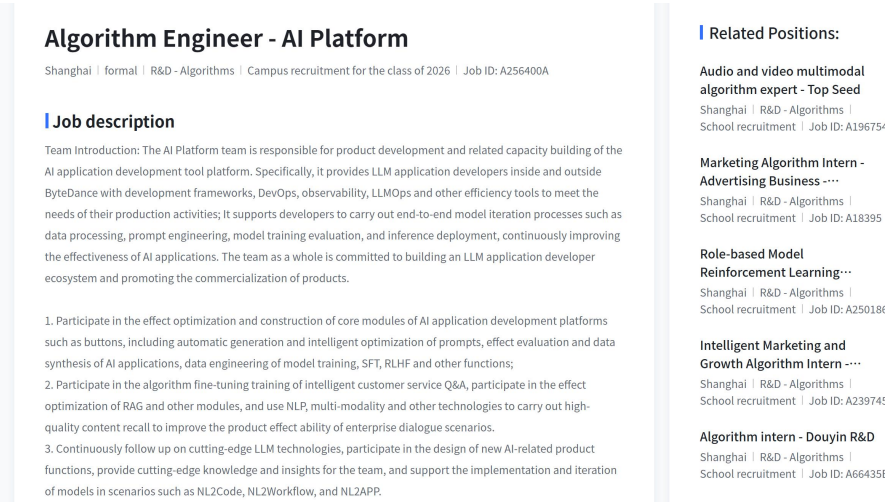

Figure 4 ByteDance AI Recruitment Algorithm Filing

(2) Unlike international enterprises such as Microsoft that set "fairness assessment" as a mandatory link in the algorithm pipeline, ByteDance's system lacks effective manual review and appeal intervention points. This is not a guess but can

be inferred from the "terminal experience" of its recruitment process. In the frontline feedback of many job seekers, common pain points include lengthy processes (the phenomenon of "being rejected after five interviews"), arbitrary transfer between different departments, and vague final rejection reasons (such as "not very matching" and "rejected due to ranking"). When the algorithm deviates or causes controversies, there is a lack of a clear, smooth "manual valve" with the power to interrupt the automated process, which directly leads to the breakdown of the responsibility chain, highlighting the practical urgency of establishing an "algorithmic accountability" mechanism (CNKI burst term, intensity 2.53).

ByteDance recruiter feedback situation is shown in Figure 5.

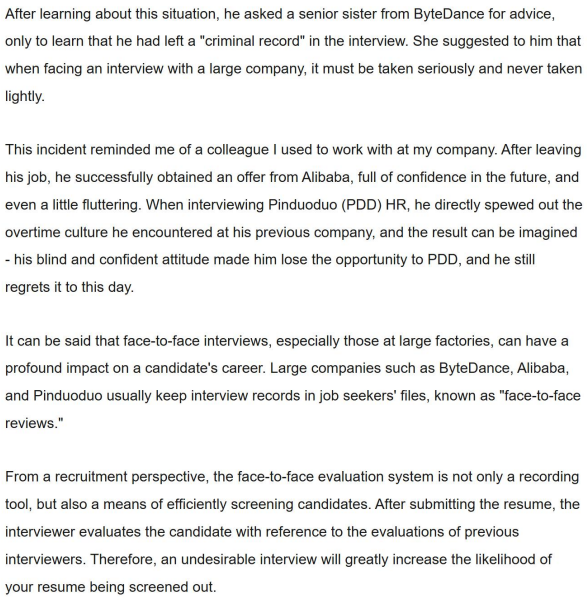
After learning about this situation, he asked a senior sister from ByteDance for advice, only to learn that he had left a "criminal record" in the interview. She suggested to him that when facing an interview with a large company, it must be taken seriously and never taken lightly.

This incident reminded me of a colleague I used to work with at my company. After leaving his job, he successfully obtained an offer from Alibaba, full of confidence in the future, and even a little fluttering. When interviewing Pinduoduo (PDD) HR, he directly spewed out the overtime culture he encountered at his previous company, and the result can be imagined - his blind and confident attitude made him lose the opportunity to PDD, and he still regrets it to this day.

It can be said that face-to-face interviews, especially those at large factories, can have a profound impact on a candidate's career. Large companies such as ByteDance, Alibaba, and Pinduoduo usually keep interview records in job seekers' files, known as "face-to-face reviews."

From a recruitment perspective, the face-to-face evaluation system is not only a recording tool, but also a means of efficiently screening candidates. After submitting the resume, the interviewer evaluates the candidate with reference to the evaluations of previous interviewers. Therefore, an undesirable interview will greatly increase the likelihood of your resume being screened out.

Figure 5 ByteDance Recruiter Feedback Situation

(3) The virtualization of feedback and appeal mechanisms: Currently, job seekers rejected by the system face a situation of "no way to complain". Enterprises have not established a standard "explainable rejection" feedback mechanism, nor do they provide formal channels for dispute resolution. This essentially deprives job seekers of their procedural remedy rights, reducing the interaction that should be responsible to both parties to a one-way technical output, which is contrary to the accountability required by the "responsibility dimension".

2.3.2 Disconnection between ethical statements in esg reports and reality

In its "2024 Corporate Social Responsibility Report", ByteDance does list "Technology for Good" as a core vision and elaborates on its efforts in promoting employment (such as live streaming for job placement) and technological inclusion. However, the report is almost "silent" on the specific ethical risks that have emerged in its core business, such as AI recruitment.

The gap between macro vision and micro risks is shown in Table 1.

Table 1 The Gap between Macro Vision and Micro Risks

Dimension	Macro statements in ESG reports (commitments)	Micro reality in AI recruitment systems (deficiencies)	Core gap
Algorithmic transparency	Emphasize "responsible trust" and build a healthy ecosystem.	Decision logic is a black box, feature weights are not public, and rejection reasons are unexplainable.	Gap in transparency between commitments and practices
Manual intervention and appeals	Focus on "user experience and services" and optimize the rights and interests of all parties.	Lack of mandatory manual review nodes and effective appeal and feedback channels.	Suspension of rights protection mechanisms
Fairness assessment	Advocate an "inclusive and fair healthy business ecosystem".	No public, regular algorithmic fairness audit reports, and no mention of targeted governance of "algorithmic discrimination".	Separation of value propositions and risk control measures
Respect for people	Regard "people" as the purpose and stimulate creativity.	Job seekers are datafied into feature vectors, and the interview process may cause "surveillance anxiety" and "sense of dignity deprivation".	Dissolution of humanistic care in technical processes

2.3.3 Conclusion—ethical concession under efficiency orientation

Compared with international leading AI ethical standards, such as Microsoft's Fairness Assessment, enterprises have intentionally or unintentionally conceded the guarantee of basic ethical principles such as procedural justice, result explainability, and human dignity in system design to maximize recruitment speed and scale benefits. This is not only a problem of a single enterprise. The recent concern of its founder Zhang Yiming about AI talent "overfitting", that is, the model is too close to existing data and lacks innovation and generalization capabilities, exactly reflects the deep dilemma that the entire industry may fall into when pursuing technical efficiency: when the system is optimized to accurately screen "people who are most like excellent employees in the past", it may systematically exclude "future talents" with different backgrounds, thinking paths, and potential creativity. This pursuit of efficiency at the cost of sacrificing diversity and innovation potential constitutes the most hidden and profound paradox in business ethics in the algorithm era, and also makes the ethical review based on the "justice-dignity-responsibility" three-dimensional framework particularly necessary and urgent.

3. IN-DEPTH ANALYSIS OF "THREE-DIMENSIONAL PENETRATING" BUSINESS ETHICS

Based on the bibliometric analysis in Chapter 2, academic research on the ethics of algorithmic recruitment mainly focuses on three clusters: result fairness, process dignity, and institutional responsibility. This exactly maps to the three major controversial points of ByteDance's system revealed in Chapter 3—systematic discrimination, human objectification, and accountability failure. Therefore, this chapter integrates academic concerns and case reality to construct a "justice-dignity-responsibility" three-dimensional analysis framework for in-depth analysis of the case.

3.1 Justice Dimension

Combined with the local pain points such as "employment discrimination" and "school discrimination" revealed by the CNKI cluster #2 Algorithmic Discrimination, and reviewed based on Rawls' "difference principle"—which advocates that social systems should maximize the improvement of the status of the least advantaged members—it can be found that AI recruitment systems represented by ByteDance have a profound conflict between their internal logic and this justice requirement.

The core optimization goal of the system, as shown in Patent CN111242334A, is to "minimize differences from historical hiring decisions". This means that the algorithm does not pursue fairness in a vacuum but efficiently learns and replicates the implicit biases that may exist in the enterprise's past recruitment data. Just as the algorithmic bias solidification mechanism revealed by the University of Washington Research Institute, in ByteDance's practice, it is directly reflected in the systematic strengthening of historical preferences such as "famous school background". The public opinion data in Section 3.2.1 of this report provides evidence for this: in job seeker feedback, doubts about discrimination against "non-first-tier university" backgrounds were directly mentioned 126 times. This is not accidental but an inevitable result of the algorithm automatically replicating the "successful models" in historical recruitment. For example, in a large-scale campus recruitment activity of ByteDance in 2024, a large number of candidates from ordinary universities reported being quickly eliminated in the AI initial screening stage without obtaining specific reasons, which is directly related to the algorithm learning the implicit "famous school preference" in historical data. Therefore, this technical path essentially deviates from the "difference principle". In the pursuit of prediction accuracy, that is, fitting the past, the system will systematically sacrifice the opportunities of groups that are already in a disadvantaged position in historical data (such as "non-first-tier university" graduates and job seekers from specific regions) to improve overall efficiency. The root cause lies in a key ethical choice jointly faced by ByteDance's algorithm team when setting model optimization goals (loss functions) and product managers when determining the "algorithm evaluation coverage rate" KPI: whether to regard "fairness" as a core goal that must be guaranteed or only as a compromise constraint condition. The current choice clearly shows that the priority of commercial efficiency overrides the fair compensation for the least advantaged. This makes technology that should promote equal opportunities alienated into a tool that solidifies and even exacerbates existing social inequalities.

3.2 Dignity Dimension

Comparison chart of human compliance rates with AI recruitment recommendations under different AI bias levels (UW 2025) is shown in Figure 6.

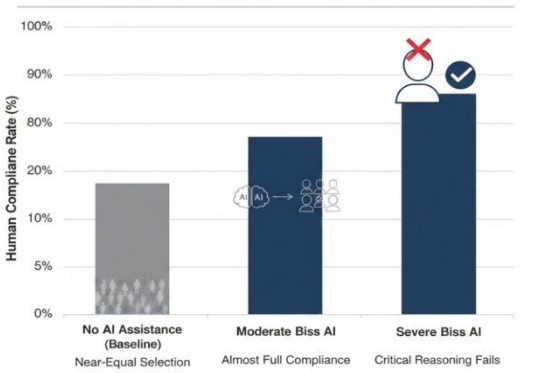

Figure 6 Comparison Chart of Human Compliance Rates with AI Recruitment Recommendations under Different AI Bias Levels (UW 2025)

Combined with the WOS burst term applicant reactions and cluster #11 Big Five, this study examines ByteDance's AI recruitment system based on Kant's fundamental ethical imperative that "human beings are ends in themselves, not means to an end". The system presents a profound ethical paradox in operation: technology that should be people-oriented has actually instrumentalized job seekers in practical applications. By compressing vivid individuals into a set of computable and comparable "feature vectors" through algorithms, the system essentially constitutes a dissolution of human subjectivity and inherent dignity.

In ByteDance's AI interview scenario, this "personality datafication" process is particularly specific and disturbing. Job seekers are no longer facing human interviewers with empathy and communication, but a continuously operating camera and algorithm program. The system will real-time capture and analyze candidates' micro-expressions, speech pauses, and word choices, and reduce them to scores on dimensions such as "extraversion", "conscientiousness", or "emotional stability". For example, a brief silence of a candidate due to in-depth thinking when answering open-ended questions may be judged by the algorithm as "insufficient adaptability"; while a well-rehearsed, fluent but lacking sincerity performative answer may instead receive a higher score. This evaluation mechanism alienates an interaction that should be based on understanding and dialogue into a one-way, gazed technical anatomy. As described in public opinion feedback, job seekers feel that they are "being tried in a black box" or "performing for a machine". In this process, job seekers as "ends" with rich emotions, unique experiences, and independent will are systematically reduced to "means" that can be optimized and screened by algorithms. More alarmingly, this technical logic has begun to erode the professional ethics of recruiters themselves. UW research points out that even if HR realizes that AI recommendations may have obvious biases, they still choose to follow the algorithm's judgment in about 90% of cases. This is far more than "automation bias", but a failure of professional ethics in human resource management. When HR abandons independent judgment and blindly follows algorithms, they essentially deviate from the core ethics of human resource management—"fair evaluation" and "professional prudence". This abandonment may stem from blind trust in "data-driven" decisions or pressure from internal efficiency KPIs. Regardless of the reason, the result is to transform HR, who should assume the role of "gatekeeper", into passive executors of algorithm outputs, reducing recruitment—a professional activity that requires humanistic insight and moral judgment—to a technocratic process operation.

Therefore, the crisis in the dignity dimension lies not only in the "objectification" of individuals but also in the dual decline of human subjectivity in the entire recruitment system, including both job seekers and recruiters. The extreme pursuit of quantifiable efficiency by the system is excluding those qualities that are difficult to be easily captured by algorithms but crucial: sincerity, enthusiasm, unique life experiences, and resilience shown in adversity. When the technical system defaults that "what cannot be measured does not exist", what it maintains is no longer human dignity, but the absolute authority of its own operating logic. This warns us that if enterprises only regard ethics as an external constraint of the technical system, rather than a core value that reshapes its internal logic and professional practice, then any commitment of "Technology for Good" will be empty in practice.

3.3 Responsibility Dimension

Combined with the core concerns revealed by the CNKI burst term "algorithmic accountability", and in accordance with the basic requirements of responsibility ethics, examining ByteDance's AI recruitment system, the core crux exposed by ByteDance's AI recruitment system is far from the superficial problem of technical "black box", but the systematic dissolution of ethical responsibility by its organizational structure and incentive mechanism. The ambiguity and dissipation of responsibility are not inevitable results of technical complexity, but products of conscious or unconscious design by the organization under the dominance of the "efficiency first" culture.

In this system, a dangerous rupture has occurred between individual professional responsibility and organizational ethical responsibility. The KPI of algorithm engineers is defined as improving the "prediction accuracy" of the model (i.e., the fitting degree to historical hiring decisions), not "fairness and explainability". This is not a lack of personal morality of engineers, but a clear value ordering conveyed by the organization through the performance system: placing "efficiency" and "accuracy" above "fairness". Similarly, recruitment managers (HR) face pressure on "processing throughput" and "process timeliness", and their professional success is quantified as "filling positions quickly". When

algorithms can provide seemingly objective and efficient recommendations, following algorithms becomes a choice consistent with individual performance rationality, even if there are ethical doubts in the heart. The "increased human compliance with biased algorithms" revealed in the University of Washington research is exactly the inevitable manifestation of the isomorphism between this individual professional rationality and organizational distorted incentives—systematically rewarding compliance rather than checks and balances. However, it is one-sided to attribute the responsibility only to frontline engineers or HR, which is precisely the mechanism by which organizational responsibility escapes. The real ethical deficiency lies in the institutional dereliction of duty by management: they designed this process with single-point efficiency as the core, but failed to embed checks and balances nodes with equal weight. For example, in model review meetings, there is a lack of a veto power for legal or ethical experts; in the recruitment process, no mandatory, recorded manual review links are set. This design choice essentially regards the externalization of ethical risks as an acceptable cost. Therefore, when a disabled job seeker is eliminated due to the algorithm model's misjudgment of "non-standard" behavioral patterns, he is faced with a system without a clear responsible subject: engineers can claim that "the model has achieved the preset accuracy KPI", HR can defend that "the company's specified automated process was followed", and management hides behind "technical autonomy" and "process standardization". The chain of responsibility is not broken, but intentionally disassembled and virtualized at the beginning of design.

Therefore, the questioning of responsibility ethics must penetrate into the heart of the organization: the company's top management. ByteDance advocates "responsible technology" in its 2024 ESG report, but its practice in AI recruitment shows that the company has not established a rigid governance structure and resource investment matching this commitment. Has an independent algorithmic ethics audit committee been established? Have fairness indicators been included in the assessment of the technical vice president? Has sufficient budget and staffing been reserved for handling appeals and manual reviews? These silent answers collectively point to a conclusion: the current responsibility vacuum is mainly due to the strategic trade-off made by the organization's senior management between commercial efficiency and comprehensive ethical governance. Rebuilding responsibility first requires management to transform its ethical commitments into binding institutional structures and resource allocation, thereby providing support rather than suppression for the professional conscience of every individual.

4 CONCLUSIONS

Based on the analysis of ByteDance's AI recruitment system, this study reveals a profound paradox between efficiency logic and ethical values: while the system accelerates matching speed, it sacrifices workplace diversity, individual dignity, and the effectiveness of accountability. This indicates that current governance gaps are not accidental technical errors, but the inevitable product of an "efficiency-first" organizational culture.

Looking forward, the governance of algorithmic recruitment must be corrected and reconstructed from three dimensions: first, the ethicalization of technical paths, promoting a shift from "black box" algorithms to "Explainable AI" and embedding fairness constraints to break the automated replication of historical biases; second, the return of management functions, re-establishing HR as the "ethical gatekeeper" and building "human-in-the-loop" review mechanisms to prevent the excessive encroachment of technical rationality; third, the substantiation of institutional accountability, requiring enterprises to translate ESG visions into concrete governance actions, such as algorithmic audits and appeal channels, to uphold the baseline of fairness and justice while pursuing commercial efficiency. Only in this way can algorithmic technology truly return to its essence of assisting decision-making, achieving the sustainable development of business logic and the humanistic spirit.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

REFERENCES

[1] Albaroudi E, Mansouri T, Hatamleh M, et al. Addressing intersectional bias in AI recruitment using HITHIRE model: a fair, ethical, green AI and transparent hiring solution for Saudi Arabia's diverse workforce in line with vision 2030. AI and Ethics, 2025, 6(1): 57-57.

[2] Odunlami O, Solano R J, Wadi O E. Evaluating AI Influence on Candidate Authenticity: Risks, Ethics and Solutions for Modern Recruitment. Asian Journal of Economics, Business and Accounting, 2025, 25(12): 329-339.

[3] Leszczyński G, Gaczek P, Ławrynowicz M. The bright side of AI in hiring: Collaborating with algorithms supports ethical decision-making. Journal of Business Research, 2026, 203: 115812-115812.

[4] Herold M, Roedenbeck R M. AI-Driven Research in the Recruitment and Selection Process: Application of an AI Taxonomy With a Systematic Literature Review. SAGE Open, 2025, 15(3).

[5] Balakrishnan S, Thongprayoon C, Wathanavasin W, et al. Evaluating artificial intelligence bias in nephrology: the role of diversity, equity, and inclusion in AI-driven decision-making and ethical regulation. Frontiers in Artificial Intelligence, 2025, 8: 1525937-1525937.

[6] Ouakili E O. The Impact of Artificial Intelligence (AI) on Recruitment Process. Open Journal of Business and Management, 2025, 13(02): 749-762.

[7] Zuzana S, Dana H, Ondřej D, et al. Incorporating artificial intelligence (AI) into recruitment processes: ethical considerations. Vilakshan - XIMB Journal of Management, 2024, 21(2): 293-307.
[8] Norman D B, Ian M. Exploring the Ethical Dynamics of the Use of Artificial Intelligence (AI) in Hiring in Healthcare Organizations. Land Forces Academy Review, 2023, 28(4): 309-321.
[9] Serge P da Motta Veiga, Maria A F, Clark B B. Seeming Ethical Makes You Attractive: Unraveling How Ethical Perceptions of AI in Hiring Impacts Organizational Innovativeness and Attractiveness. Journal of Business Ethics, 2023, 186(1): 199-216.
[10] Lena A H, Alexander K. Is AI recruiting (un)ethical? A human rights perspective on the use of AI for hiring. AI and ethics, 2022, 3(1): 11-15.