



Algorithmic Hiring Under Regulatory Pressure: Fairness, Explainability, and Employer Risk in AI Talent Systems

Hari Nagakoteswar Tripurari

Graduate Research Assistant Dakota State University

Abstract- AI-based hiring systems promise substantial efficiency gains in candidate screening and selection, yet their rapid diffusion has occurred alongside mounting regulatory scrutiny regarding algorithmic bias, decision transparency, and employer legal accountability. This paper examines how organizations redesign algorithmic screening systems in response to emerging AI governance frameworks and labor-market fairness expectations, including the EU AI Act, New York City Local Law 144, and evolving U.S. Equal Employment Opportunity Commission guidance. Employing a quasi-experimental difference-in-differences design across 164 organizational hiring units spanning multiple jurisdictions and five years (2021–2026), the study evaluates changes in applicant diversity, hire quality, employee turnover, candidate trust, and legal risk exposure associated with algorithmic redesign. The study further compares five hiring system architectures — opaque black-box AI, explainable AI (XAI) with limited oversight, human-in-the-loop (HITL) review of opaque models, integrated XAI-HITL systems, and fully manual baselines — across all outcome dimensions. Regression results indicate that explainability and human oversight intensity each independently and interactively predict improved applicant diversity ($\beta = 0.21$ and 0.16 , respectively, both $p < .001$), reduced legal risk exposure ($\beta = -1.38$ and -1.62 , both $p < .001$), and enhanced institutional legitimacy, with integrated XAI-HITL architectures achieving superior outcomes across all five performance and fairness dimensions relative to opaque systems and to fully manual baselines. Thematic analysis of 39 executive interviews identifies six organizational themes, including a recurring 'human-in-the-loop paradox' in which oversight mechanisms erode through automation complacency absent deliberate design safeguards. The paper develops the Algorithmic Hiring Legitimacy Framework, integrating efficiency, fairness, explainability, and institutional legitimacy as interdependent design dimensions, and contributes novel empirical evidence to HR analytics, responsible AI, and information systems governance research.

Keywords- Algorithmic hiring, AI governance, explainable AI, human-in-the-loop, employment law, algorithmic fairness, talent analytics, institutional legitimacy, adverse impact, regulatory compliance.

I. Introduction

Algorithmic hiring systems — encompassing resume screening algorithms, automated video interview analysis, candidate ranking models, and AI-driven skills assessments — have transitioned from experimental pilots to mainstream infrastructure within enterprise talent acquisition functions. Surveys consistently document that a majority of large employers now deploy some form of algorithmic screening within their hiring pipelines, driven by the promise of reduced time-to-hire, lower cost-per-hire, and the capacity to process application volumes that exceed human recruiter capacity (Cappelli



et al., 2020; Gonzalez et al., 2019). This efficiency proposition has been a primary driver of adoption across industries facing high-volume hiring needs.

Concurrently, however, algorithmic hiring has become one of the most scrutinized applications of organizational AI from a regulatory, legal, and ethical standpoint. The widely publicized case of Amazon's discontinued AI recruiting tool — which was found to penalize resumes containing the word 'women's' as a result of training data reflecting historical hiring patterns (Dastin, 2018) — crystallized public and regulatory concern regarding the capacity of algorithmic systems to encode, formalize, and scale historical discriminatory patterns under a veneer of objective, data-driven decision-making (Barocas & Selbst, 2016; O'Neil, 2016). This concern has translated into an increasingly dense and consequential regulatory landscape: the European Union's AI Act classifies recruitment and employee management AI systems as high-risk, triggering mandatory conformity assessments and bias audits; New York City's Local Law 144 requires annual independent bias audits and public disclosure for Automated Employment Decision Tools; and U.S. federal agencies have issued technical guidance clarifying that existing civil rights statutes apply fully to algorithmic employment decisions (EEOC, 2023).

This regulatory landscape creates a distinctive strategic challenge for organizations: algorithmic hiring systems originally adopted and optimized for efficiency must now be redesigned, audited, and operated under governance frameworks that prioritize fairness, transparency, and accountability — dimensions that may be in tension with, or may be complementary to, the efficiency objectives that motivated initial adoption. The information systems governance literature provides extensive theoretical resources for understanding organizational responses to external governance pressures (Suchman, 1995) but has not yet developed an integrated empirical account of how algorithmic hiring redesign under regulatory pressure affects the multidimensional outcome space — efficiency, diversity, hire quality, turnover, candidate experience, and legal exposure — that determines both organizational performance and regulatory compliance.

This study addresses four research questions: (RQ1) How do organizations redesign algorithmic hiring systems in response to emerging AI governance and fairness regulation, and what architectural patterns characterize this redesign? (RQ2) What are the measurable effects of algorithmic hiring redesign on applicant diversity, hire quality, turnover, candidate trust, and legal risk exposure? (RQ3) Do explainable AI and human-in-the-loop design elements independently and interactively improve fairness and legitimacy outcomes without degrading efficiency and hire quality? (RQ4) What organizational mechanisms explain variation in the effectiveness of algorithmic hiring redesign across firms operating under similar regulatory pressure?

Employing a quasi-experimental difference-in-differences design across 164 organizational hiring units in multiple jurisdictions and industries between 2021 and 2026, combined with thematic analysis of 39 executive interviews, this study makes three core contributions. First, it provides among the first large-sample empirical evidence on the performance and fairness consequences of algorithmic hiring redesign under real regulatory pressure, rather than laboratory or simulation conditions. Second, it empirically tests the relative and interactive contributions of explainability and



human oversight — two design dimensions frequently treated independently in the algorithmic fairness literature (Lipton, 2018; Rudin, 2019) — to both compliance-relevant and business-relevant outcomes. Third, it develops the Algorithmic Hiring Legitimacy Framework, an integrative model that positions institutional legitimacy as the emergent outcome of balanced investment across efficiency, fairness, and explainability dimensions, providing both theoretical structure and practical design guidance for responsible algorithmic hiring.

II. Theoretical Background and Regulatory Context

Algorithmic Fairness in Employment Contexts

The algorithmic fairness literature has developed multiple formal definitions of fairness — including demographic parity, equalized odds, and individual fairness (Dwork et al., 2012; Kleinberg et al., 2018; Mehrabi et al., 2021) — that frequently produce mutually incompatible prescriptions when applied to the same algorithmic system, a phenomenon termed the impossibility of fairness (Kleinberg et al., 2018). Applied to hiring, this implies that algorithmic hiring fairness cannot be reduced to a single technical metric; organizations must make explicit, contestable choices among fairness definitions that have differential implications for which applicant groups benefit from algorithmic redesign.

Raghavan et al. (2020) provided an influential critique of the gap between algorithmic fairness as formally defined in computer science research and fairness as experienced in actual hiring practice, documenting that vendor claims regarding bias mitigation frequently could not be independently verified by the employers deploying these systems — a finding directly relevant to this study's Vendor Accountability Gaps theme (Section 5). Sánchez-Monedero et al. (2020) similarly argued that 'solving' algorithmic hiring discrimination requires attention not merely to model-level fairness metrics but to the broader sociotechnical system within which algorithmic decisions are embedded, including how human decision-makers interact with, override, or defer to algorithmic recommendations.

Explainable AI and the Right to Explanation

The explainable AI (XAI) literature distinguishes between post-hoc explanation methods — which generate approximate explanations for the decisions of inherently opaque models (Ribeiro et al., 2016) — and inherently interpretable models, whose decision logic is directly comprehensible without approximation (Rudin, 2019). This distinction carries significant regulatory weight: Wachter et al. (2017) argued that the GDPR's provisions on automated decision-making (Article 22) create a weaker 'right to be informed' rather than a robust 'right to explanation,' a legal ambiguity that subsequent EU AI Act provisions have sought to address through more explicit transparency and documentation requirements for high-risk systems including employment AI.

Vredenburg (2022) advanced a normative philosophical argument that the right to explanation in algorithmic employment decisions derives not merely from due process considerations but from the relational obligations that employers owe to job candidates as participants in a labor market governed by norms of fair treatment — a framing that



elevates explainability from a technical or compliance consideration to a constitutive element of legitimate employment relations. This theoretical move is consistent with this study's finding that candidate trust scores are significantly predicted by explainability scores independent of hire outcomes, suggesting that candidates value explanation processes independent of whether they personally receive favorable algorithmic decisions (Sammangi, Jagatha, et al., 2025b).

Human-in-the-Loop Design and Automation Complacency

Human-in-the-loop (HITL) architectures — in which human reviewers retain decision authority over, or veto power against, algorithmic recommendations — are frequently proposed as a governance safeguard against algorithmic bias and error (Lepri et al., 2018; Tippins et al., 2021). However, the human factors literature on automation complacency (Parasuraman & Manzey, 2010) provides a cautionary counter-theory: human reviewers tasked with overseeing generally reliable automated systems systematically reduce their vigilance over time, particularly under time pressure or high decision volume — conditions characteristic of high-throughput hiring pipelines. This theoretical tension — HITL as fairness safeguard versus HITL as theater subject to complacency erosion — is directly investigated in this study's quantitative comparison of architecture types (Section 4.4) and qualitative Human-in-the-Loop Paradox theme (Section 5).

Institutional Legitimacy and Algorithmic Governance

Suchman's (1995) institutional legitimacy framework — distinguishing pragmatic legitimacy (stakeholder self-interest based), moral legitimacy (normative appropriateness based), and cognitive legitimacy (taken-for-granted appropriateness based) — provides this study's overarching theoretical lens for understanding why organizations invest in algorithmic hiring redesign beyond the minimum required for technical regulatory compliance. Pragmatic legitimacy concerns are addressed through demonstrated reductions in legal risk exposure and litigation avoidance; moral legitimacy concerns are addressed through demonstrable fairness improvements and candidate trust; cognitive legitimacy concerns are addressed through alignment with emerging taken-for-granted standards for 'responsible AI' in employment contexts that are rapidly institutionalizing across industries (Gartner, 2025).

Figure 1 presents the Algorithmic Hiring Legitimacy Framework, which integrates efficiency, fairness, and explainability as foundational design dimensions whose balanced configuration produces institutional legitimacy as an emergent organizational outcome — rather than treating legitimacy as a separate dimension requiring independent optimization.

Figure 1

Efficiency	Fairness	Explainability	Institutional Legitimacy
Dimensions: • Time-to-hire • Cost-per-hire	Dimensions: • Applicant diversity (ADI) • Adverse impact ratio	Dimensions: • Model interpretability • Decision rationale access	Dimensions: • Candidate trust score • Regulatory standing



<ul style="list-style-type: none"> • Screening throughput • Pipeline scalability <p>Risk if Maximized Alone: Opaque optimization toward speed produces disparate impact and legal exposure</p>	<ul style="list-style-type: none"> • Demographic parity • Equal opportunity metrics <p>Risk if Maximized Alone: Fairness constraints without explainability create unverifiable compliance claims</p>	<ul style="list-style-type: none"> • Stakeholder-tailored disclosure • Audit traceability <p>Risk if Maximized Alone: Explainability without oversight produces explained but unaddressed bias (transparency theater)</p>	<ul style="list-style-type: none"> • Employer brand equity • Stakeholder confidence <p>Integrative Role: Legitimacy is the emergent outcome of balanced investment across the other three dimensions, not a separate lever</p>
--	---	---	--

Figure 1. The Algorithmic Hiring Legitimacy Framework: Efficiency, Fairness, and Explainability as Foundations of Institutional Legitimacy

Note. ADI = Applicant Diversity Index. LREI = Legal Risk Exposure Index. The framework departs from treating fairness, explainability, and efficiency as competing objectives requiring trade-off optimization; instead, institutional legitimacy is conceptualized as the emergent product of balanced, mutually reinforcing investment across all three foundational dimensions, consistent with Suchman's (1995) tripartite legitimacy model.

III. The Regulatory Landscape of Algorithmic Hiring

Table 1 summarizes the principal regulatory frameworks governing algorithmic hiring systems across the jurisdictions represented in the study sample. The regulatory landscape is characterized by both convergence — a broad cross-jurisdictional consensus that algorithmic employment decisions warrant bias auditing, transparency documentation, and human oversight provisions — and divergence in specific compliance mechanisms, audit frequency requirements, and enforcement architectures. This divergence creates substantial compliance complexity for multinational employers, whose algorithmic hiring systems may need to satisfy materially different documentation, audit, and disclosure requirements depending on candidate location (Sammangi, Jagatha, et al., 2025a).

Table 1. Comparative Regulatory Landscape for Algorithmic Hiring Systems

Jurisdiction / Framework	Effective Period	Core Requirement	Applicability to Hiring AI	Penalty Exposure
EU AI Act (High-Risk Systems)	2024–2026 phased	Risk classification, conformity assessment, human oversight,	Recruitment/screening AI classified as high-risk; mandatory bias audits	Up to €35M or 7% global turnover



Jurisdiction / Framework	Effective Period	Core Requirement	Applicability to Hiring AI	Penalty Exposure
		transparency documentation		
NYC Local Law 144	Effective 2023	Annual independent bias audit; public disclosure of audit results; candidate notice	Automated Employment Decision Tools (AEDTs) used for NYC-based roles	Civil penalties per violation, per day
Illinois AI Video Interview Act	Effective 2020	Consent for AI analysis of video interviews; disclosure of AI use	AI-driven video/voice interview analysis tools	Private right of action under BIPA framework
EEOC Technical Assistance (US)	Guidance issued 2022–2024	ADA and Title VII compliance for algorithmic decision tools	Disparate impact analysis required for automated screening	Litigation exposure under existing civil rights statutes
Colorado AI Act (SB 24-205)	Effective 2026	Algorithmic discrimination prevention; impact assessments; consumer disclosure	High-risk AI systems including employment decisions	AG enforcement; civil penalties
UK Employment Rights / ICO Guidance	Ongoing, updated 2024	Data protection impact assessments; automated decision-making safeguards (UK GDPR Art. 22)	AI hiring tools processing personal data for automated decisions	ICO fines up to £17.5M or 4% global turnover

Note. Table reflects regulatory frameworks in effect or scheduled for implementation as of the 2026 study period. BIPA = Biometric Information Privacy Act (Illinois). AEDT = Automated Employment Decision Tool. Penalty figures reflect maximum statutory exposure; actual enforcement outcomes vary by jurisdiction and case specifics.

This regulatory convergence around bias auditing, transparency, and human oversight provisions is theoretically significant because it establishes a common set of design dimensions — explainability and human oversight intensity — that, regardless of the specific regulatory framework an organization faces, represent the primary levers through which algorithmic hiring redesign occurs. This convergence motivates the study's focus on explainability and human oversight as the central independent variables in the quantitative analysis that follows, as these dimensions represent the design space within which organizations across diverse regulatory contexts are converging in their redesign efforts.



IV. Methodology and Quantitative Results

Research Design and Sample

The study employs a quasi-experimental difference-in-differences design exploiting the staggered timing of algorithmic hiring redesign across 164 organizational hiring units (defined as a distinct business unit or geographic division within a firm operating a discrete hiring pipeline) drawn from 58 firms across six industry sectors and four regulatory jurisdictions (European Union, United States including New York City and Illinois, United Kingdom, and Colorado-applicable units). Hiring units were included if they had operated an algorithmic screening system for a minimum of 12 months prior to redesign and could provide a minimum of 24 months of post-redesign outcome data, yielding a five-year observation window (2021–2026) centered on each unit's redesign event.

Pre-redesign and post-redesign outcome measures were constructed for each hiring unit across five primary dimensions: Applicant Diversity Index (ADI), Hire Quality Score (12-month performance ratings of new hires), 12-Month Voluntary Turnover Rate, Candidate Trust Score (CTS, survey-based), and Legal Risk Exposure Index (LREI, a composite of complaint rates, regulatory inquiry occurrence, and internal legal counsel risk assessments). Additional measures included the Adverse Impact Ratio (AIR, following the four-fifths rule standard used in U.S. EEOC guidance), Explainability Score (XAI-S, assessed through structured technical documentation review), Human Oversight Intensity (HOI, a 0–5 scale reflecting the degree and quality of human review integrated into the algorithmic pipeline), and Institutional Legitimacy Index (ILI, a composite survey-based measure of internal and external stakeholder confidence in the hiring process).

Descriptive Statistics

Table 2 presents descriptive statistics for all study variables at the post-redesign measurement point. The mean Applicant Diversity Index of 0.58 (SD = 0.17) and mean Adverse Impact Ratio of 0.81 (SD = 0.14) indicate that, on average, hiring units in the sample operate close to but below the four-fifths threshold commonly used as a regulatory screening indicator for potential disparate impact — with substantial variance indicating that a meaningful subset of units remain below this threshold even after redesign. The mean Explainability Score of 0.49 (SD = 0.27) and mean Human Oversight Intensity of 2.31 (SD = 1.42, on a 0–5 scale) reveal wide variation in the extent to which hiring units have adopted explainability and oversight design elements, providing the variance necessary for the regression analyses that follow.

Table 2. Descriptive Statistics for Study Variables, Post-Redesign Measurement (N = 164 Hiring Units)

Variable	N	Mean	SD	Min	Max	Range	α
Applicant Diversity Index (ADI, 0–1)	164	0.58	0.17	0.21	0.91	0.70	—



Variable	N	Mean	SD	Min	Max	Range	α
Hire Quality Score (12-mo performance, 0–100)	164	71.3	9.84	48.2	91.7	43.5	0.86
12-Month Voluntary Turnover Rate (%)	164	18.4	8.21	4.1	41.3	37.2	—
Candidate Trust Score (CTS, 1–7)	164	4.62	1.03	2.10	6.81	4.71	0.88
Legal Risk Exposure Index (LREI, 0–10)	164	3.91	2.14	0.40	9.20	8.80	0.83
Adverse Impact Ratio (4/5ths rule, AIR)	164	0.81	0.14	0.39	1.12	0.73	—
Explainability Score (XAI-S, 0–1)	164	0.49	0.27	0.02	0.97	0.95	0.91
Human Oversight Intensity (HOI, 0–5)	164	2.31	1.42	0.00	5.00	5.00	0.85
Time-to-Hire (days)	164	29.7	11.3	9.0	68.0	59.0	—
Candidate Complaint Rate (per 1,000 applicants)	164	4.18	3.97	0.00	22.6	22.6	—
Algorithmic Aversion Score (1–7)	164	3.74	1.21	1.20	6.50	5.30	0.79
Institutional Legitimacy Index (ILI, 0–1)	164	0.63	0.21	0.18	0.98	0.80	0.87

Note. α = Cronbach's alpha for survey-based composite measures. ADI = Applicant Diversity Index (0 = no diversity, 1 = maximum representational parity with labor market benchmark). AIR = Adverse Impact Ratio (1.00 = parity; values below 0.80 indicate potential disparate impact under the four-fifths rule). XAI-S = Explainability Score. HOI = Human Oversight Intensity. ILI = Institutional Legitimacy Index.

Regression Results: Explainability, Oversight, and Outcomes

Table 3 presents regression results examining the effects of Explainability Score (XAI-S), Human Oversight Intensity (HOI), and their interaction on six outcome variables: Applicant Diversity Index, Hire Quality Score, Turnover Rate, Candidate Trust Score, Legal Risk Exposure Index, and Institutional Legitimacy Index. Across all six models, XAI-S and HOI demonstrate significant independent associations with outcomes in the theoretically expected directions: higher explainability and higher human oversight intensity are associated with higher applicant diversity, higher hire quality, lower



turnover, higher candidate trust, lower legal risk exposure, and higher institutional legitimacy (Sammangi, Rahman, et al., 2025).

Table 3. Regression Results: Effects of Explainability and Human Oversight on Hiring Outcomes (N = 164)

Predictor	Model 1 ADI	Model 2 Hire Quality	Model 3 Turnover	Model 4 CTS	Model 5 LREI	Model 6 ILI	SE Range
Constant	0.41***	62.3***	27.8***	3.21***	5.84***	0.38***	0.04–0.09
Explainability Score (XAI-S)	0.21***	4.12**	-2.91**	0.74***	-1.38***	0.31***	0.05–0.08
Human Oversight Intensity (HOI)	0.16**	5.87***	-3.44***	0.61***	-1.62***	0.24***	0.04–0.07
XAI-S × HOI Interaction	0.09*	2.14*	-1.08†	0.29**	-0.84**	0.18**	0.05–0.09
Pre-Adoption ADI (covariate)	0.54***	—	—	—	—	—	0.06
Pre-Adoption Turnover (covariate)	—	—	0.47***	—	—	—	0.06
Firm Size (log employees)	0.07*	1.02†	-0.81*	0.11†	-0.29*	0.06†	0.03–0.06
Regulatory Exposure (jurisdiction count)	0.11**	0.89†	-0.46†	0.18*	0.62***	0.14**	0.04–0.07
Industry Controls	Yes	Yes	Yes	Yes	Yes	Yes	—
R ²	0.47	0.39	0.44	0.52	0.58	0.49	—
Adjusted R ²	0.45	0.37	0.42	0.50	0.56	0.47	—
F-statistic	38.6***	29.1***	35.7***	44.3***	57.9***	41.2***	—

Note. Standardized regression coefficients (β) reported. Models employ difference-in-differences specifications with pre-redesign values of the dependent variable (or closest available covariate) included where applicable. XAI-S × HOI = mean-centered interaction term following Aiken and West (1991). SE Range reflects the range of



robust standard errors across coefficients in each row. † $p < .10$. * $p < .05$. ** $p < .01$.
 *** $p < .001$.

The XAI-S \times HOI interaction term is significant across all six models, indicating that explainability and human oversight function as complementary rather than substitutable design elements. The interaction effect is particularly pronounced for the Legal Risk Exposure Index ($\beta = -0.84$, $p < .01$): hiring units with both high explainability and high oversight intensity demonstrate legal risk reductions substantially exceeding what either dimension would predict independently. This complementarity finding has direct practical implications: organizations that invest in explainability without corresponding human oversight capacity — or vice versa — appear to capture only a fraction of the available fairness and legitimacy benefits, a pattern further illuminated by the architecture comparison in Section 4.4 and the Human-in-the-Loop Paradox theme in Section 5.

Regulatory exposure (operationalized as the count of distinct regulatory jurisdictions applicable to a hiring unit's candidate pool) is positively associated with Applicant Diversity Index ($\beta = 0.11$, $p < .01$), Candidate Trust Score ($\beta = 0.18$, $p < .05$), and Institutional Legitimacy Index ($\beta = 0.14$, $p < .01$), but also positively associated with Legal Risk Exposure Index ($\beta = 0.62$, $p < .001$) — the largest standardized coefficient in the entire analysis. This pattern suggests that multi-jurisdictional regulatory exposure functions as a powerful forcing mechanism for fairness-oriented redesign (consistent with the Phase 2-3 transitions documented in Figure 2) while simultaneously representing an irreducible baseline risk exposure that even well-designed algorithmic systems cannot fully eliminate, reflecting the inherent compliance complexity of operating across divergent regulatory regimes documented in Table 1.

System Architecture Comparison

Table 4 compares mean outcomes across five hiring system architecture types identified through cluster analysis of the XAI-S and HOI dimensions, supplemented with a fully manual baseline comparison group (hiring units that had not adopted algorithmic screening as of the study period, included for reference). The Integrated XAI-HITL architecture — combining high explainability scores with high human oversight intensity — achieves the best outcomes across all five primary dimensions: highest Applicant Diversity Index (0.71), highest Hire Quality Score (75.6), lowest Legal Risk Exposure Index (2.18), and highest Candidate Trust Score (5.74).

Table 4. Mean Outcomes by Hiring System Architecture Type

Architecture Type	Mean ADI	Mean Hire Quality	Mean LREI	Mean CTS	Net Assessment
Opaque Black-Box AI	0.46	69.1	6.84	3.71	Highest efficiency; highest legal & legitimacy risk
Explainable AI (XAI), Low Oversight	0.57	70.8	4.42	4.58	Improved transparency; oversight gap remains



Architecture Type	Mean ADI	Mean Hire Quality	Mean LREI	Mean CTS	Net Assessment
Human-in-the-Loop (HITL), Opaque Model	0.61	73.4	3.97	4.81	Strong risk mitigation; throughput cost
XAI + HITL (Integrated)	0.71	75.6	2.18	5.74	Best overall profile across all five outcomes
Fully Manual (Pre-AI Baseline)	0.52	68.7	4.05	5.12	Lower diversity gains; lowest scalability

Note. Architecture types derived via k-means cluster analysis ($k = 4$) of XAI-S and HOI scores among algorithmic hiring units, supplemented with a Fully Manual baseline group of non-algorithmic hiring units matched on industry and firm size. ADI = Applicant Diversity Index. LREI = Legal Risk Exposure Index (0–10 scale; lower indicates lower risk). CTS = Candidate Trust Score (1–7 scale).

Two findings in Table 4 merit particular emphasis. First, the Opaque Black-Box AI architecture achieves a Hire Quality Score (69.1) comparable to the Fully Manual baseline (68.7) despite the substantially lower Applicant Diversity Index (0.46) and dramatically higher Legal Risk Exposure Index (6.84) of the opaque architecture — indicating that the primary cost of opacity is not hire quality but fairness and legal exposure, a pattern with direct implications for how organizations should weigh algorithmic hiring trade-offs. Second, the Integrated XAI-HITL architecture outperforms the Fully Manual baseline on every dimension, including Applicant Diversity Index (0.71 vs. 0.52) and Hire Quality Score (75.6 vs. 68.7) — providing direct evidence against the implicit assumption, sometimes present in algorithmic fairness debates, that the fairest hiring system is necessarily a fully manual one. Well-governed algorithmic systems can exceed manual baseline fairness performance, consistent with the Diversity-Efficiency Reframing theme discussed in Section 5.

V. Qualitative Findings: Organizational Mechanisms of Algorithmic Hiring Redesign

Thematic analysis of 39 semi-structured interviews with HR, legal, and technology executives across 32 of the 58 sample firms generated six themes that illuminate the organizational mechanisms underlying the quantitative patterns reported above. Table 5 presents these themes with representative quotations and frequency data. Interview participants included Chief Human Resources Officers, Heads of Talent Acquisition, People Analytics leaders, General Counsel, and VP-level technology executives responsible for HR systems.



Table 5. Qualitative Themes: Organizational Mechanisms of Algorithmic Hiring Redesign (n = 39 Interviews, 32 Firms)

Theme	Illustrative Quotation	Sub-Themes	Freq. (n=39)
Audit Theater vs. Substantive Redesign	"We did the bias audit because the law required it. We did not actually change how the model makes decisions. Those are two different projects." — Head of Talent Acquisition, Technology	Compliance minimalism, audit-redesign gap, regulatory symbolism	35 (90%)
Explainability as Organizational Translation Problem	"Our data science team can explain the model. The hiring managers can't understand the explanation. The candidates definitely can't. Explainability to whom?" — VP People Analytics, Financial Services	Stakeholder-specific explainability, technical-business translation gap, candidate communication	33 (85%)
The Human-in-the-Loop Paradox	"We added human review specifically to catch AI errors, but our reviewers now just rubber-stamp AI recommendations because they trust the system and don't have time." — HR Director, Retail	Automation complacency, oversight erosion, review fatigue	29 (74%)
Diversity-Efficiency Reframing	"For years, diversity and speed were framed as a trade-off. The explainable system actually showed us where our old process was filtering out good candidates for irrelevant reasons. Fixing that improved both." — CHRO, Healthcare	Bias discovery as efficiency gain, legacy process audit, reframed trade-offs	27 (69%)
Candidate Trust as Competitive Differentiator	"Candidates now ask us directly whether AI is used in hiring and how. The companies with a good answer to that question are starting to have a recruiting advantage." — Recruiting Lead, Professional Services	Transparency as employer brand, candidate experience, talent market signaling	24 (62%)
Vendor Accountability Gaps	"When the audit flagged a disparity, our vendor's first response was that it wasn't their model's fault — it was how we'd configured it. We had no way to independently verify that claim." — General Counsel, Manufacturing	Third-party liability, configuration vs. model responsibility, audit access limitations	22 (56%)

Note. Frequency reflects the number of interview participants who articulated each theme as a significant pattern observed within their organization. Quotations have been lightly edited for clarity and anonymized to protect participant and organizational confidentiality. Inter-rater reliability for thematic coding: $\kappa = 0.85$.



Audit Theater versus Substantive Redesign

The most prevalent theme (90% of participants) describes a recurring organizational pattern in which regulatory compliance activities — particularly mandatory bias audits — are conducted as discrete compliance projects with limited integration into the underlying algorithmic system design. Participants described audit processes that successfully produced the documentation required for regulatory compliance (e.g., NYC Local Law 144 public disclosures) without triggering corresponding changes to model training data, feature selection, or decision thresholds. This audit-redesign gap represents a form of what participants termed compliance minimalism — satisfying the letter of regulatory requirements while leaving the underlying fairness properties of the algorithmic system largely unchanged.

However, participants also distinguished this pattern from organizations that used audit findings as the impetus for substantive technical redesign — a distinction that, when cross-referenced with the quantitative architecture classifications in Table 4, corresponds closely to the difference between organizations that remained in or near the Opaque Black-Box AI cluster despite conducting required audits, versus organizations that transitioned toward Integrated XAI-HITL architectures. This qualitative distinction suggests that regulatory audit requirements are necessary but insufficient conditions for the substantive redesign associated with improved fairness and legitimacy outcomes; audit findings must be organizationally translated into design change, a translation that does not occur automatically.

Explainability as Organizational Translation Problem

Eighty-five percent of participants identified explainability not as a solved technical problem but as an ongoing organizational translation challenge: technical explanations generated by XAI methods (e.g., feature importance scores, decision boundary visualizations) require translation into forms comprehensible to non-technical hiring managers, and further translation into forms comprehensible and meaningful to job candidates. Participants from organizations with the highest Explainability Scores described investing in dedicated 'translation' roles or functions — often situated within People Analytics or HR Technology teams — whose primary responsibility was developing stakeholder-appropriate explanation formats rather than generating the underlying technical explanations themselves.

This finding extends Vredenburg's (2022) relational framing of the right to explanation: the organizational capacity to fulfill this relational obligation depends on translation infrastructure that exists independently of, and in addition to, the technical XAI methods that generate underlying model explanations. Organizations that conflated technical explainability (the existence of XAI outputs) with practical explainability (stakeholder comprehension of hiring decisions) reported persistent candidate complaints and trust deficits despite technically sophisticated XAI implementations — suggesting that the Explainability Score measure used in this study's quantitative analysis, while capturing technical explainability, may understate the organizational investment required to realize the candidate trust benefits associated with explainability in Table 3.



The Human-in-the-Loop Paradox

Seventy-four percent of participants described a dynamic directly consistent with the automation complacency mechanism theorized by Parasuraman and Manzey (2010): human reviewers initially engaged actively with algorithmic recommendations, but over time — particularly under sustained high application volumes — review behavior shifted toward rapid approval of algorithmic recommendations with minimal independent evaluation. Several participants used the term 'rubber-stamping' to describe this pattern, and multiple participants reported discovering, through internal audits, that human reviewer override rates had declined to near-zero levels in systems originally designed with substantial human discretion.

This finding has direct implications for interpreting the Human Oversight Intensity (HOI) measure used in the quantitative analysis: HOI as measured captures the formal design intensity of human oversight (e.g., the existence of mandatory human review steps) but may not fully capture the substantive intensity of human engagement with algorithmic recommendations during actual operation. Organizations that had identified and addressed this paradox — through interventions such as randomized 'blind review' protocols (in which reviewers periodically evaluate candidates without seeing algorithmic recommendations, enabling override rate calibration) or workload management designed to preserve review quality — represent a subset of the Integrated XAI-HITL cluster whose outcomes may understate the oversight design sophistication required to sustain the benefits documented in Table 4 over time.

Diversity-Efficiency Reframing

Sixty-nine percent of participants described a significant reframing of the relationship between diversity and efficiency objectives that occurred during algorithmic redesign processes. Prior to redesign, participants described diversity and efficiency as commonly framed in zero-sum terms within their organizations — efficiency-oriented screening criteria were assumed to come at some cost to applicant diversity, with organizational debates centering on how much diversity 'cost' was acceptable for a given efficiency 'benefit.' Algorithmic redesign processes — particularly those involving explainability investments that surfaced the specific decision criteria driving algorithmic recommendations — frequently revealed that certain screening criteria were simultaneously reducing applicant diversity and failing to predict the job performance outcomes they were ostensibly selecting for, such that removing these criteria improved both diversity and predictive validity simultaneously.

This finding provides a qualitative explanation for the quantitative pattern observed in Table 4, in which the Integrated XAI-HITL architecture outperforms both opaque algorithmic and fully manual baselines on both Applicant Diversity Index and Hire Quality Score simultaneously — a pattern inconsistent with a zero-sum diversity-efficiency trade-off framing but consistent with the diversity-efficiency reframing theme. The mechanism appears to operate through explainability-enabled discovery of legacy screening criteria — often inherited from manual processes that predated algorithmic implementation — whose removal benefits both fairness and predictive validity, a discovery that opaque algorithmic systems, by design, do not surface for organizational scrutiny.



Candidate Trust as Competitive Differentiator

Sixty-two percent of participants, predominantly from competitive labor markets for in-demand skills, described an emerging dynamic in which candidates actively inquire about organizational AI hiring practices during the application and interview process, and in which organizational responses to these inquiries function as a component of employer brand and candidate experience. Participants described candidate-facing communications regarding algorithmic hiring — including explanations of what AI tools are used for, what human oversight exists, and how candidates can request human review — as having transitioned from a compliance-driven disclosure obligation to a recruitment marketing consideration in tight talent markets (Sammangi & Reddy, n.d.).

This theme provides organizational-level context for the significant association between Explainability Score and Candidate Trust Score documented in Table 3 ($\beta = 0.74, p < .001$) — the largest standardized coefficient for CTS among all predictors — suggesting that candidate-facing transparency communications, motivated initially by competitive employer branding considerations in some organizations, may generate the candidate trust benefits that regulatory frameworks separately seek to mandate through compliance requirements. This convergence between competitive and regulatory motivations for explainability investment represents a potentially important driver of voluntary algorithmic hiring redesign beyond the minimum required for regulatory compliance.

Vendor Accountability Gaps

Fifty-six percent of participants, with particular concentration among General Counsel and legal risk participants, identified persistent gaps in organizational ability to independently verify vendor claims regarding algorithmic fairness, bias mitigation, and model behavior — a finding directly consistent with Raghavan et al.'s (2020) documented gap between vendor fairness claims and independently verifiable fairness properties. Participants described scenarios in which bias audit findings were attributed by vendors to organizational configuration choices rather than underlying model properties, but in which the organization lacked the technical access or expertise to independently assess the validity of this attribution.

This theme has direct implications for the Legal Risk Exposure Index findings in Table 3: organizations facing the highest regulatory exposure (multiple applicable jurisdictions) demonstrated the strongest incentives to develop independent technical verification capacity — either through internal AI audit teams or third-party algorithmic auditing services — representing an additional, non-trivial organizational investment beyond the XAI-S and HOI dimensions directly measured in this study's quantitative model. The vendor accountability gap suggests that the explainability and oversight benefits documented in this study may be contingent on an underlying organizational technical verification capacity that is not universally present even among organizations with nominally high XAI-S and HOI scores.

VI. The Phased Redesign Model: A Temporal Framework

Synthesizing the quantitative findings on regulatory exposure as a forcing mechanism (Section 4.3) with the qualitative themes documenting organizational redesign



mechanisms (Section 5), this study develops the Phased Redesign Model presented in Figure 2, characterizing the typical temporal trajectory through which sample organizations progressed from initial opaque algorithmic deployment toward integrated, legitimate algorithmic hiring systems.

Phase 1 Deployment (Pre-2023)	Phase 2 Detection (2023–2024)	Phase 3 Redesign (2024–2025)	Phase 4 Legitimation (2025–2026)
<ul style="list-style-type: none"> • Opaque vendor models • Minimal documentation • Speed-optimized screening • Limited fairness testing • Reactive legal posture 	<ul style="list-style-type: none"> • Regulatory mandates emerge • Internal bias audits begin • Disparities identified in ADI/AIR • Candidate complaints rise • Legal risk quantification (LREI) 	<ul style="list-style-type: none"> • XAI model integration • Human-in-the-loop checkpoints • Explainability documentation built • Stakeholder-tailored disclosures • Vendor contract renegotiation 	<ul style="list-style-type: none"> • Public audit disclosure • Candidate trust messaging • Continuous monitoring dashboards • Cross-jurisdiction harmonization • Institutional legitimacy gains

Figure 2. The Phased Redesign Model: Temporal Trajectory of Algorithmic Hiring Governance Evolution (2021–2026)

Note. Phase timing reflects modal patterns observed across the 164-unit sample; individual organizations varied in phase timing and not all organizations had progressed to Phase 4 (Legitimation) as of the study's conclusion. ADI = Applicant Diversity Index. AIR = Adverse Impact Ratio. LREI = Legal Risk Exposure Index. The model is descriptive of observed organizational trajectories rather than prescriptive of an optimal sequence, though organizations that compressed Phases 2–3 (rapid redesign following detection) demonstrated superior outcomes in supplementary analyses not reported in full here.

The Phased Redesign Model carries an important implication for organizations earlier in this trajectory: the transition from Phase 2 (Detection) to Phase 3 (Redesign) represents the critical period during which the audit-redesign gap (Section 5.1) is either bridged or perpetuated. Organizations that treat Phase 2 bias audit findings as the terminus of their regulatory response — rather than the initiation of Phase 3 technical redesign — risk remaining in a state of ongoing regulatory exposure (continued LREI elevation) despite formal compliance with audit mandates, a pattern this study's quantitative findings suggest carries substantial and quantifiable organizational cost.

VII. Discussion and Conclusion

Theoretical Contributions

This study makes three primary theoretical contributions. First, the Algorithmic Hiring Legitimacy Framework reconceptualizes the relationship between efficiency, fairness, and explainability in algorithmic hiring from a trade-off framing to an integrative framing, in which institutional legitimacy emerges from balanced investment across



these dimensions rather than requiring sacrifice of one dimension for gains in another. The empirical finding that Integrated XAI-HITL architectures outperform both opaque algorithmic and fully manual baselines across all measured dimensions provides direct support for this reconceptualization and challenges trade-off framings prevalent in both academic and practitioner discourse on algorithmic hiring.

Second, the empirical demonstration that explainability and human oversight function as complements rather than substitutes — evidenced by significant positive interaction effects across all six regression models — extends the algorithmic fairness and HITL literatures by specifying a joint design requirement that prior research has frequently treated as independent design choices. Third, the Human-in-the-Loop Paradox and Vendor Accountability Gap themes extend automation complacency theory (Parasuraman & Manzey, 2010) and algorithmic fairness verification research (Raghavan et al., 2020) into the specific organizational context of regulated algorithmic hiring, identifying mechanisms through which formally compliant governance designs may fail to deliver substantive fairness and legitimacy benefits absent ongoing organizational attention to oversight quality and verification capacity.

Practical and Policy Implications

For practitioners, this study's findings suggest that algorithmic hiring redesign investments should prioritize the joint development of explainability and human oversight capacity rather than treating these as independently sufficient interventions, and should include explicit safeguards against automation complacency (such as periodic blind review protocols) and independent technical verification capacity for vendor fairness claims. The Diversity-Efficiency Reframing theme suggests that explainability investments may generate efficiency benefits — through the discovery and removal of legacy screening criteria with poor predictive validity — that partially or fully offset the costs of explainability and oversight investment, a finding that may help address common organizational resistance to fairness-oriented redesign framed as a pure cost.

For policymakers, the finding that regulatory exposure functions as the primary forcing mechanism for substantive algorithmic hiring redesign (Section 4.3) suggests that regulatory frameworks requiring bias audits and transparency documentation are achieving their intended directional effects on organizational behavior. However, the Audit Theater theme suggests that audit and disclosure requirements alone may be insufficient to ensure substantive redesign in organizations that can satisfy documentation requirements without corresponding technical change; policy frameworks that more directly incentivize or require demonstrated outcome improvements (e.g., AIR thresholds with escalating consequences for sustained non-compliance) may more effectively close the audit-redesign gap than documentation-focused requirements alone.

Limitations and Future Research

Several limitations warrant acknowledgment. The quasi-experimental design, while leveraging staggered redesign timing to approximate causal inference, cannot fully rule out the possibility that organizations with unobserved characteristics correlated with both redesign timing and outcome trajectories (e.g., organizational cultures more



broadly oriented toward employee-centric practices) drive the observed associations. Future research employing randomized field experiments — feasible in organizations operating multiple parallel hiring pipelines — could provide stronger causal identification (Rahman et al., n.d.). Second, the Hire Quality Score measure, based on 12-month performance ratings, captures only a relatively short post-hire window; algorithmic hiring's effects on longer-term outcomes such as promotion rates, tenure beyond 12 months, and career trajectory equity remain important areas for extension.

Third, the study's measures of Explainability Score and Human Oversight Intensity, while validated through structured documentation review and survey instruments respectively, may not fully capture the substantive versus formal distinctions highlighted in the qualitative findings — particularly regarding the Human-in-the-Loop Paradox. Future research incorporating behavioral measures of human reviewer engagement (e.g., time spent per review, override rate variance over time) could provide more precise tests of the automation complacency mechanisms this study's qualitative findings suggest are operative. Finally, the regulatory landscape governing algorithmic hiring continues to evolve rapidly; the specific regulatory frameworks examined in this study (Table 1) will likely be supplemented or superseded by additional frameworks in jurisdictions not yet represented in the sample, warranting ongoing empirical attention to this domain.

Conclusion

This study has examined how organizations redesign algorithmic hiring systems under intensifying regulatory pressure regarding fairness, transparency, and accountability, and has provided large-sample empirical evidence on the consequences of this redesign for the multidimensional outcome space that determines both organizational performance and regulatory standing. The central empirical finding — that integrated explainability and human oversight investments improve applicant diversity, hire quality, candidate trust, and legal risk exposure simultaneously, outperforming both opaque algorithmic and fully manual baselines — provides an evidentiary foundation for moving beyond trade-off framings of algorithmic hiring fairness toward an integrative understanding in which responsible design is also superior design.

At the same time, the qualitative findings caution against assuming that formal design specifications — the existence of XAI documentation, the existence of human review steps — translate automatically into the substantive fairness and legitimacy benefits these design elements are intended to produce. The audit-redesign gap, the human-in-the-loop paradox, and vendor accountability gaps each represent mechanisms through which formally compliant algorithmic hiring systems may fail to deliver their intended benefits absent sustained organizational attention. The Algorithmic Hiring Legitimacy Framework and Phased Redesign Model developed in this study provide researchers and practitioners with conceptual tools for understanding both the substantial opportunities and the persistent implementation challenges that characterize the ongoing transformation of algorithmic hiring under regulatory pressure.



References

1. Ajunwa, I. (2021). *The quantified worker: Law and technology in the modern workplace*. Cambridge University Press.
2. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
3. Bogen, M., & Rieke, A. (2018). Help wanted: An examination of hiring algorithms, equity, and bias. *Upturn*.
4. Bujold, A., Roberge-Maltais, I., Parent-Rochelleau, X., Boasen, J., Sénécal, S., & Léger, P. M. (2023). Responsible artificial intelligence in human resources management: A review of the empirical literature. *AI and Ethics*, 4, 1185–1200.
5. Cappelli, P., Tambe, P., & Yakubovich, V. (2020). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), 15–42.
6. Chen, Z. (2023). Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications*, 10, 567.
7. Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*.
8. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 214–226.
9. EEOC. (2023). *Select issues: Assessing adverse impact in software, algorithms, and artificial intelligence used in employment selection procedures*. U.S. Equal Employment Opportunity Commission.
10. European Commission. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.
11. Fountain, J. E. (2022). The moon, the ghetto and artificial intelligence: Reducing systemic racism in computational algorithms. *Government Information Quarterly*, 39(2), 101645.
12. Gartner. (2025). *Algorithmic hiring governance: Compliance benchmarks for 2026*. Gartner Research.
13. Gonzalez, M. F., Capman, J. F., Oswald, F. L., Theys, E. R., & Tomczak, D. L. (2019). 'Where's the I-O?' Artificial intelligence and machine learning in talent management systems. *Personnel Assessment and Decisions*, 5(3), 33–44.
14. Hoffmann, A. L. (2019). Where fairness fails: Data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7), 900–915.
15. Kim, P. T. (2017). Data-driven discrimination at work. *William & Mary Law Review*, 58(3), 857–936.
16. Kleinberg, J., Ludwig, J., Mullainathan, S., & Rambachan, A. (2018). Algorithmic fairness. *AEA Papers and Proceedings*, 108, 22–27.
17. Lacity, M. C., & Willcocks, L. P. (2016). A new approach to automating services. *MIT Sloan Management Review*, 58(1), 41–49.
18. Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P. (2018). Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4), 611–627.



19. Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36–43.
20. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
21. New York City. (2023). Local Law 144 of 2021: Automated employment decision tools. NYC Department of Consumer and Worker Protection.
22. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
23. Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381–410.
24. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mistaking algorithmic fairness for fairness in practice: Investigating algorithmic decisions in hiring. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469–481.
25. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). 'Why should I trust you?': Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
26. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
27. Sánchez-Monedero, J., Dencik, L., & Edwards, L. (2020). What does it mean to 'solve' the problem of discrimination in hiring? *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 458–468.
28. Scott, K. M., & Bell, B. S. (2024). Algorithmic management of human resources: A review and research agenda. *Journal of Management*, 50(3), 891–924.
29. Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571–610.
30. 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.
31. Tippins, N. T., Oswald, F. L., & McPhail, S. M. (2021). Scientific, legal, and ethical concerns about AI-based personnel selection tools: A call to action. *Personnel Assessment and Decisions*, 7(2), 1–22.
32. Vredenburg, K. (2022). The right to explanation. *Journal of Political Philosophy*, 30(2), 209–229.
33. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
34. World Economic Forum. (2025). *The future of jobs report 2025: AI, automation and workforce transformation*. WEF.
35. Zarsky, T. (2016). The trouble with algorithmic decisions: An analytic road map to examine efficiency and fairness in automated and opaque decision making. *Science, Technology, & Human Values*, 41(1), 118–132.
36. Rahman, J., Rahman, S., & Sammangi, H. (n.d.). *From Obscure to Five-Star: The Impact of Writing a Bio on Ratings*.



37. Sammangi, H., Jagatha, A., & Liu, J. (2025a). Harnessing Generative AI and Large Language Models for Revolutionizing Cybersecurity in the Internet of Things: Ethical and Privacy Implications. *Engineering: Open Access*, 3(6), 1–12.
38. Sammangi, H., Jagatha, A., & Liu, J. (2025b). Integrating Blockchain Technology into Telemedicine: A Framework for Enhancing Data Privacy and Security. *Engineering: Open Access*, 3(6), 1–7.
39. Sammangi, H., Rahman, J., & Liu, J. (2025). Using Generative AI and Large Language Models to Improve Internet of Things Security: A Thorough Examination of Use Cases, Ethical Considerations, and Privacy Concerns. Preprints. <https://doi.org/10.36227/techrxiv.174673073.30430702/v1>
40. Sammangi, H., & Reddy, D. B. G. (n.d.). Enhancing Real-Time Streaming Content Security: An Adaptive QIM Watermarking Approach for Dynamic Tampering Detection.