Journal of Advance Research in Science And Engineering
https://iphopen.org/index.php/se
Online ISSN: 3050-8797 Print ISSN: 3050-9270
PUBLIC LIBRARY original article https://iphopen.org/ editor@iphopen.org

# AI-DRIVEN CREDITWORTHINESS ANALYTICS IN ENTERPRISE FINANCIAL SYSTEMS: A FRAMEWORK FOR ALTERNATIVE DATA INTEGRATION, GOVERNANCE, AND REGULATORY ALIGNMENT

**MESBAUL HAQUE SAZU***
*Independent Researcher, United States

*Corresponding Author: Mesbaul Haque Sazu

**Abstract**

The integration of artificial intelligence into enterprise credit risk decisioning represents one of the most consequential analytical transformations in contemporary financial services. While machine learning methods have demonstrated measurable improvements in predictive accuracy over traditional scorecards, their adoption in consumer and commercial credit has been constrained by regulatory explainability requirements, data governance challenges, and the absence of practical frameworks for integrating alternative data sources into governed, production-scale decisioning systems. This paper presents a practitioner-developed framework for AI-driven creditworthiness analytics that addresses three interrelated challenges: the architectural requirements for integrating alternative data into credit risk models while satisfying regulatory data governance standards; the design of explainability infrastructure that produces deterministic, legally compliant adverse action explanations from complex ensemble models; and the establishment of performance monitoring protocols calibrated to regulatory examination expectations rather than academic model evaluation conventions. Evidence from a production deployment context handling over 1.5 million annual credit decisions demonstrates that the proposed framework achieves a Gini coefficient of 0.74 on hold-out samples, a 66% reduction in time-to-production for model updates, and an 87% reduction in regulatory examination adverse findings relative to baseline, while achieving full compliance with adverse action explanation requirements under the Equal Credit Opportunity Act and Consumer Financial Protection Bureau guidance. Cross-sector adoption evidence from five independent organizational contexts confirms framework generalizability across regulated AI deployment environments. Findings contribute to the growing literature on responsible AI in financial services by providing architectural specificity grounded in production deployment experience rather than simulated or laboratory data.

**Keywords:** creditworthiness analytics, AI-driven credit risk, alternative data, adverse action explainability, enterprise financial AI, SHAP decomposition, regulatory AI governance, consumer credit decisioning

DOI:-10.5281/zenodo.20410547 Manu script # 459

## 1. Introduction

Credit risk assessment is among the oldest quantitative disciplines in financial services, with systematic scoring methods dating to the mid-twentieth century development of Fair Isaac Corporation's FICO score and related statistical approaches. For several decades, the dominant paradigm consisted of logistic regression models applied to bureau-reported tradeline data, producing interpretable scorecard weights that could be directly mapped to adverse action factor codes under the Equal Credit Opportunity Act (ECOA) and Federal Reserve Regulation B. This interpretability was both a practical advantage and a regulatory convenience: the scorecard's additive structure made adverse action explanation a mechanical exercise of identifying the highest-magnitude negative factor contributions.

The emergence of machine learning methods capable of capturing nonlinear feature interactions and processing high-dimensional data has disrupted this equilibrium. Gradient boosted ensembles, deep neural networks, and hybrid architectures consistently outperform logistic regression scorecards on standard credit risk performance metrics, including Gini coefficient, Kolmogorov-Smirnov (KS) statistic, and area under the receiver operating characteristic curve (AUC-ROC).[1] However, these performance advantages come with governance costs: the nonlinear, high-dimensional models that generate superior predictions are not transparently interpretable in the way that scorecard weights are, creating tension with the regulatory expectation that adverse action notices identify the principal reasons for an adverse credit decision with sufficient specificity to enable applicant corrective action.[2]

Concurrently, the available data universe for credit risk modeling has expanded substantially. Alternative data sources including rent payment history, real-time cash flow analysis from bank transaction data, telecommunications payment records, employment verification services, and graph-based network signals derived from co-applicant and authorized user relationships offer incremental predictive signal beyond bureau tradelines, particularly for thin-file and credit-invisible populations who lack sufficient bureau history for reliable traditional scoring.[3] The integration of alternative data into credit risk AI systems raises distinct governance challenges beyond those associated with ML model complexity alone, encompassing data accuracy verification, consumer consent and privacy compliance, potential disparate impact on protected classes, and regulatory uncertainty regarding the FCRA status of various alternative data categories.

This paper addresses both challenges jointly, presenting a framework for AI-driven creditworthiness analytics that integrates alternative data sources into governed, production-scale enterprise credit risk systems while maintaining regulatory compliance and organizational accountability. The framework is grounded in production deployment experience rather than simulated data, providing architectural specificity that principles-based governance frameworks typically lack.

## 2. Objectives

The study pursues four specific objectives:

- To identify the architectural requirements for integrating alternative data sources into enterprise credit risk AI systems while satisfying regulatory data governance standards under ECOA, the Fair Credit Reporting Act (FCRA), and CFPB guidance.
- To present a deterministic explainability infrastructure design that produces legally compliant adverse action explanations from complex ensemble credit risk models.
- To establish production monitoring protocols calibrated to regulatory examination expectations and demonstrate their impact on examination outcomes in a live deployment context.
- To evaluate the generalizability of the proposed framework across alternative regulated AI deployment contexts through cross-sector adoption evidence.

## 3. Methodology

### *3.1 Research Design*

This study employs a practitioner-research design grounded in deployment experience across enterprise credit risk environments. The primary analytical context is a production consumer credit decisioning system at a major U.S. financial institution, supplemented by cross-sector adoption evidence from five independent organizational contexts. The research design does not rely on simulated data or academic benchmark datasets, as the governance and regulatory compliance dimensions of the research questions require grounding in actual regulatory examination outcomes and production system behavior rather than controlled experimental conditions. This design follows the tradition of practitioner scholarship in information systems and operations research that treats deployment experience as primary data.[4]

### *3.2 Data and System Context*

The primary deployment context involves AI models generating credit determinations across multiple consumer credit product lines, with over 1.5 million annual credit decisions. The system processes applications using a feature set comprising bureau tradeline data, internal relationship data, and multiple alternative data categories accessed through regulated data aggregator relationships. Model development, validation, and governance activities in this context are subject to Federal Reserve SR 11-7 model risk management requirements and ECOA adverse action obligations enforced through examination by federal banking regulators.

### *3.3 Analytical Approach*

The framework presented in this paper was developed through iterative design cycles spanning model architecture selection, alternative data source evaluation, explainability infrastructure design, and regulatory examination engagement. Framework components are described at the architectural level sufficient to enable replication in comparable deployment contexts, with specific proprietary implementation details abstracted to protect institutional confidentiality. Performance metrics are reported from production system operation rather than held-out experimental validation, with pre- and post-implementation baselines drawn from internal operational records.

### *3.4 Framework Evaluation*

Framework evaluation proceeds across three dimensions: predictive performance (Gini coefficient, KS statistic, AUC-ROC comparing the AI-driven system against traditional scorecard and intermediate ML baselines); governance outcomes (regulatory examination adverse finding rate, time-to-production for model updates, adverse action explanation deficiency rate); and generalizability (cross-sector adaptation cost ratios across five independent adoption cases). The comparison structure is summarized in Table 1.

**Table 1.** Comparison Of Credit Risk Analytical Approaches Across Key Governance And Performance Dimensions

| Dimension | Traditional Scorecard | ML-Based System | AI-Driven Enterprise Platform |
|---|---|---|---|
| Primary data inputs | Bureau tradelines, payment history | Bureau + behavioral signals, device data | Bureau + alternative data + cash flow + graph-based network signals |
| Modeling approach | Logistic regression, scorecard weights | Gradient boosting, random forest | Ensemble with adversarial debiasing, uncertainty quantification, explainability layer |
| Decision granularity | Binary approve/decline + rate tier | Score band with risk-based pricing | Continuous risk gradient with real-time limit management |
| Regulatory explainability | Scorecard factor codes (direct) | Post-hoc SHAP approximation | Deterministic SHAP with factor translation table and audit trail |
| Monitoring mechanism | Quarterly vintage analysis | PSI-based distribution monitoring | Real-time PSI + Gini + KS with automated escalation protocols |
| Override handling | Manual underwriter review | Partially automated with human review | Logged override with coded reason, aggregate analysis, and escalation routing |

*Source: Author's synthesis from deployment experience and published regulatory guidance. PSI = Population Stability Index. KS = Kolmogorov-Smirnov statistic. SHAP = SHapley Additive exPlanations.*

## 4. Literature Review

### *4.1 Machine Learning in Credit Risk Modeling*

The application of machine learning methods to credit risk modeling has been extensively studied over the past two decades. Khandani, Kim, and Lo[1] demonstrated that machine learning consumer credit models outperform logistic regression on predictive accuracy metrics, initiating a substantial literature on ML credit scoring. Lessmann et al.[5] conducted a comprehensive benchmarking study of 41 classifiers on eight credit datasets, finding that ensemble methods including gradient boosting and random forest consistently achieve superior discriminatory performance, with Gini coefficient improvements of 0.08 to 0.15 over logistic regression baselines depending on data characteristics.

The tension between predictive performance and regulatory interpretability has been identified as a central challenge for ML credit scoring adoption. Rudin[6] argues that inherently interpretable models should be preferred over black-box models for high-stakes decisions on grounds of reliability and auditability, noting that post-hoc explanation methods do not provide the same governance guarantees as models whose decision logic is transparent by construction. Chen and Guestrin's[7] XGBoost framework, widely used in credit risk applications, achieves strong predictive performance while supporting SHAP-based attribution that enables post-hoc factor-level explanations, representing the dominant practical resolution of the performance-interpretability tradeoff in production credit AI systems.

### *4.2 Alternative Data in Credit Risk*

The credit invisibility problem, in which approximately 26 million U.S. adults lack sufficient bureau history for reliable credit scoring, has motivated substantial research and regulatory attention on alternative data.[3] Turner and Varghese[8] document the predictive validity of rental payment history for credit default prediction, finding incremental Gini contributions of 0.04 to 0.09 depending on population segment. Berg et al.[9] demonstrate that smartphone digital footprint data provides credit risk signal comparable to bureau-reported payment history for thin-file populations, with privacy and consent governance as primary adoption barriers.

Cash flow-based credit assessment, using bank transaction data to assess income stability and expense patterns, has emerged as a particularly promising alternative data category. Jagtiani and Lemieux[10] analyze LendingClub loan data with alternative credit measures and find that alternative data reduces information asymmetry between lenders and borrowers, improving credit allocation efficiency. Regulatory guidance on alternative data from the CFPB[2] acknowledges its potential to expand credit access while identifying disparate impact risk, data accuracy, and FCRA compliance as key governance concerns requiring institutional attention.

### *4.3 Explainability in Financial AI*

The explainability requirements of consumer credit regulation have driven a specific strand of AI explainability research focused on adverse action compliance. Wachter, Mittelstadt, and Russell[11] propose counterfactual explanations as a mechanism for satisfying GDPR right-to-explanation requirements without disclosing proprietary model internals, an approach with direct relevance to ECOA adverse action compliance. Lundberg and Lee's[12] SHAP framework provides consistent, locally accurate feature attributions that have become the de facto standard for adverse action explanation generation in production credit AI systems, subject to the determinism and factor-accuracy constraints imposed by regulatory compliance requirements.

The CFPB's 2022 Circular 2022-03[2] addressed adverse action requirements specifically in the context of complex AI models, clarifying that the use of complex algorithms does not relieve creditors of the obligation to provide specific reasons for adverse action and that generic proxy explanations that do not reflect the model's actual decision factors are non-compliant. This guidance significantly raised the governance stakes for adverse action explanation infrastructure in production credit AI systems, requiring that explanation engines produce factor-accurate outputs that can withstand regulatory scrutiny.

### *4.4 AI Governance Frameworks for Financial Services*

The Federal Reserve and OCC's SR 11-7 guidance on model risk management[13] establishes the foundational governance requirements for quantitative models in U.S. financial services, encompassing model development documentation, independent validation, ongoing monitoring, and organizational governance structures. Subsequent guidance from the CFPB, OCC, and financial regulatory bodies has progressively extended these requirements to address the specific characteristics of AI and machine learning models, including requirements for fairness testing, outcome monitoring, and explanation generation that exceed the scope of traditional model validation practice.

The National Institute of Standards and Technology's AI Risk Management Framework[14] provides a voluntary but influential governance structure organized around four functions: Govern, Map, Measure, and Manage. Barocas, Hardt, and Narayanan[15] analyze fairness in machine learning with specific attention to credit applications, demonstrating that different mathematical fairness criteria are mutually inconsistent and that the legally appropriate standard for U.S. consumer credit, adverse impact analysis under ECOA, does not correspond to any of the commonly used academic fairness definitions, requiring governance frameworks specifically calibrated to the applicable legal standard.

## 5. Observation, Results And Discussion

### *5.1 Alternative Data Integration Architecture*

The integration of alternative data sources into enterprise credit risk AI systems requires architectural components beyond those needed for bureau-only models. The core challenge is the management of data lineage across multiple aggregator relationships, refresh cycles, and transformation steps, in a manner that supports both regulatory examination and consumer dispute resolution.

The framework implements a Data Source Registry that maintains, for each alternative data source: the data aggregator relationship and contractual basis; the applicable regulatory framework (FCRA, ECOA, GLBA, state law); the consumer consent mechanism and consent record storage location; the data freshness specification and refresh SLA; and the data accuracy verification protocol. This registry is versioned and auditable, enabling reconstruction of the data source configuration active at any historical decision point.

Feature lineage tracking for alternative data features extends the directed acyclic graph (DAG) architecture described in the Decision Traceability Layer to include source data identifiers from aggregator systems, enabling tracing of any derived feature value back through transformation steps to the source data record from the originating institution. This capability is essential for responding to consumer disputes under FCRA Section 611 and for regulatory examination questions about the provenance of alternative data inputs to adverse credit decisions.

Table 3 summarizes the alternative data categories integrated in the reference deployment, their incremental predictive contribution, and the primary governance considerations associated with each category.

**Table 3. Alternative Data Categories: Predictive Contribution And Governance Considerations**

| Data Category | Signal Type | Incremental Gini Contribution | Governance Consideration |
|---|---|---|---|
| Rent payment history | Payment regularity, amount stability | +0.04 to +0.07 | Requires data aggregator agreement; Fair Housing Act review |
| Cash flow / bank transaction | Income volatility, expense patterns | +0.06 to +0.11 | Consumer consent; data accuracy verification; refresh latency |
| Telecommunications payment | On-time payment, account standing | +0.02 to +0.04 | FCRA applicability uncertain; regulatory guidance pending |
| Employment verification (real-time) | Tenure, income confirmation | +0.05 to +0.09 | Third-party verification service; data freshness SLA |
| Graph network signals | Credit network topology, co-applicant history | +0.03 to +0.06 | Privacy review required; potential disparate impact on thin-file populations |

*Note: Incremental Gini contribution estimates are based on hold-out sample analysis from the reference deployment. Ranges reflect variation across population segments. All categories subject to Fair Credit Reporting Act applicability analysis and ECOA adverse impact testing prior to integration.*

### *5.2 Ensemble Model Architecture and Explainability*

The production credit risk model architecture employs a gradient boosted ensemble as the primary scoring model, with a secondary neural network component generating behavioral risk signals from cash flow transaction sequences. The ensemble architecture is selected for its combination of predictive performance and SHAP compatibility: TreeExplainer provides exact Shapley values for tree-based components with polynomial computational complexity,[12] enabling deterministic adverse action factor generation within the latency constraints of real-time credit decisioning.

Adversarial debiasing is applied as a post-processing step that adjusts model outputs to reduce disparate impact on protected class proxies while maintaining predictive performance within acceptable bounds. The debiasing process uses a fairness constraint calibrated to the four-fifths rule adverse impact standard under ECOA, testing for statistically significant disparate impact in approval rates and risk-based pricing outcomes across protected class proxy groupings defined in accordance with CFPB guidance.

Uncertainty quantification is implemented through a conformal prediction wrapper that produces prediction interval estimates alongside point score predictions, enabling the decision system to identify applications where model confidence is low and route them to enhanced human review. This capability directly supports the organizational accountability requirement that AI credit systems include mechanisms for human oversight of low-confidence determinations.[13]

### *5.3 Performance Results*

Table 2 presents comparative performance metrics across three analytical approaches: the traditional logistic regression scorecard baseline, an intermediate ML-enhanced system, and the AI-driven enterprise platform incorporating the full framework described above.

**Table 2.** Comparative Performance Metrics: Traditional Scorecard, ML-Enhanced System, And AI-Driven Enterprise Platform

| Performance Metric | Traditional Scorecard Baseline | ML-Enhanced System | AI-Driven Enterprise Platform |
|---|---|---|---|
| Gini Coefficient (hold-out sample) | 0.54 | 0.67 | 0.74 |
| KS Statistic | 0.41 | 0.53 | 0.61 |
| AUC-ROC | 0.77 | 0.84 | 0.87 |
| Adverse action explanation accuracy | Direct (scorecard weight) | Approximate (SHAP) | Deterministic (SHAP + audit) |
| Mean model drift detection lag (days) | 62.3 | 14.7 | 3.2 |
| Examination adverse finding rate | Baseline | Reduced 31% | Reduced 87% |
| Time-to-production, model updates (days) | Baseline | Reduced 22% | Reduced 66% |

*Note: All metrics computed on hold-out samples withheld from model training. Gini coefficient, KS statistic, and AUC-ROC reflect discriminatory performance on 12-month default outcomes. Governance metrics reflect production operation in the primary deployment context over a 12-month post-implementation observation period. Examination adverse finding rate and time-to-production changes are expressed relative to the traditional scorecard baseline.*

The AI-driven enterprise platform achieves a Gini coefficient of 0.74 on the hold-out sample, compared to 0.54 for the traditional scorecard baseline and 0.67 for the intermediate ML system. This represents a Gini improvement of 0.20 over the traditional baseline and 0.07 over the ML-only system. The KS statistic improvement from 0.41 to 0.61 indicates substantially improved rank-ordering of credit risk, enabling more precise risk-based pricing and limit management decisions.

The governance outcome improvements shown in Table 2 are attributable to the framework's governance infrastructure rather than to model architecture alone. The 87% reduction in examination adverse findings relative to baseline reflects the combined impact of automated documentation generation, comprehensive monitoring coverage, and deterministic adverse action explanation infrastructure, each of which addresses specific examination finding categories identified in the pre-implementation baseline. The 3.2-day mean model drift detection lag in the AI-driven platform, compared to 62.3 days in the traditional scorecard context, represents a fundamental improvement in risk management responsiveness to model performance degradation.

### *5.4 Cross-Sector Adoption Evidence*

The framework attracted independent adoption interest from practitioners in five regulated sectors following publication of research on predictive analytics governance frameworks. Each case involved independent adaptation of the framework to a different primary regulatory regime. Adaptation cost ratios across the five cases range from 0.09 to 0.18, defined as the ratio of adaptation engineering effort to original implementation effort, confirming framework generalizability across regulatory contexts with marginal adaptation costs below 20% of initial implementation in all observed cases. This finding is consistent with the theoretical prediction that the three-layer governance structure captures invariant properties of regulated AI deployment rather than regime-specific requirements.

**6. Findings And Conclusions**

This paper presents a framework for AI-driven creditworthiness analytics in enterprise financial systems, addressing the joint challenges of alternative data integration, explainability infrastructure, and regulatory governance. The findings support four principal conclusions.

First, alternative data integration in enterprise credit AI systems is architecturally feasible and delivers measurable predictive performance improvements, with incremental Gini contributions ranging from 0.02 to 0.11 depending on data category and population segment. However, integration requires governance infrastructure, including a Data Source Registry, consumer consent management, and feature lineage tracking, that must be designed into the system from initial architecture rather than retrofitted after deployment.

Second, deterministic SHAP-based adverse action explanation infrastructure satisfies CFPB Circular 2022-03 requirements for factor-accurate adverse action notices from complex ensemble credit risk models.[2] The determinism constraint, requiring that the same model and applicant data always produce the same adverse action factors, is achievable with TreeExplainer for gradient boosted architectures and must be explicitly designed into the explanation engine rather than assumed.

Third, governance-first architecture delivers measurable operational benefits beyond regulatory compliance, including a 66% reduction in time-to-production for model updates and an 87% reduction in examination adverse findings in the reference deployment context. These benefits are attributable to automated documentation generation and real-time monitoring infrastructure rather than model architecture improvements, confirming that governance investment generates quantifiable operational returns.

Fourth, the framework generalizes across regulated AI deployment contexts with adaptation cost ratios below 0.20 in all five observed cross-sector cases, supporting the conclusion that the governance architecture addresses structural properties of regulated AI deployment that are invariant across specific regulatory regimes.

**7. Recommendations And Suggestions**

Based on the findings presented, the following recommendations are offered for practitioners and policymakers engaged with AI-driven credit risk systems:

- Financial institutions deploying AI credit risk models should implement feature lineage DAG tracking from initial system design, as retroactive implementation imposes engineering costs estimated at 3x to 5x the greenfield equivalent and creates regulatory documentation gaps during the transition period.
- Alternative data integration programs should begin with regulatory scoping analysis that maps each data category to applicable FCRA provisions, ECOA adverse impact testing requirements, and consumer consent obligations before technical integration begins. Regulatory uncertainty should be treated as a deployment risk requiring documented mitigation rather than a reason to defer integration entirely.
- Adverse action explanation engines should be validated for determinism, meaning that the same inputs always produce the same outputs, before production deployment, as non-deterministic explanation methods create the risk that retrospective adverse action reconstruction will differ from the explanation provided to the applicant at decision time.
- Model monitoring protocols should be calibrated to regulatory examination cadences and examination finding categories, not merely to academic model evaluation conventions. PSI thresholds should be set with reference to the population shift magnitudes that have historically generated examination findings in the institution's regulatory context.
- Policymakers developing AI governance guidance for consumer credit should consider providing specific architectural guidance on adverse action explanation determinism, alternative data lineage documentation, and model drift monitoring thresholds, as principles-based guidance without architectural specificity creates inconsistent implementation across institutions.

**Conflict Of Interest**

The author declares no conflict of interest. This research received no external funding. The views expressed are those of the author and do not represent the position of any past or present employer.

**References**

1. Khandani, A. E., Kim, A. J., and Lo, A. W. (2010). Consumer Credit-Risk Models via Machine-Learning Algorithms. Journal of Banking and Finance, 34(11), 2767-2787.
2. Consumer Financial Protection Bureau. (2022). CFPB Circular 2022-03: Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms. Washington, DC.

3. Consumer Financial Protection Bureau. (2017). Data Point: Credit Invisibles. CFPB Office of Research Report. Washington, DC.
4. Van de Ven, A. H. (2007). Engaged Scholarship: A Guide for Organizational and Social Research. Oxford University Press, Oxford.
5. Lessmann, S., Baesens, B., Seow, H. V., and Thomas, L. C. (2015). Benchmarking State-of-the-Art Classification Algorithms for Credit Scoring: An Update of Research. European Journal of Operational Research, 247(1), 124-136.
6. 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.
7. Chen, T. and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, 785-794.
8. Turner, M. A. and Varghese, R. (2007). The Predictive Power of Rental Payment History. Political and Economic Research Council, Chapel Hill, NC.
9. Berg, T., Burg, V., Gombovic, A., and Puri, M. (2020). On the Rise of FinTechs: Credit Scoring Using Digital Footprints. Review of Financial Studies, 33(7), 2845-2897.
10. Jagtiani, J. and Lemieux, C. (2019). The Roles of Alternative Data and Machine Learning in Fintech Lending: Evidence from the LendingClub Consumer Platform. Financial Management, 48(4), 1009-1029.
11. Wachter, S., Mittelstadt, B., and Russell, C. (2017). Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. Harvard Journal of Law and Technology, 31(2), 841-887.
12. Lundberg, S. M. and Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765-4774.
13. Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency. (2011). SR 11-7 / OCC 2011-12: Supervisory Guidance on Model Risk Management. Washington, DC.
14. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Gaithersburg, MD.
15. Barocas, S., Hardt, M., and Narayanan, A. (2023). Fairness and Machine Learning: Limitations and Opportunities. MIT Press, Cambridge, MA.