

Beyond the Bureau: A Systematic Review and Proposed Segment-Aware Credit Risk Framework

Deep Shikha^{1 a)}, Himanshu Vasnani^{2 b)}, Amit Tiwari^{2 c)}, Amit Kumar Verma^{3 d)},
¹*Faculty of Sciences (Computer Science), Suresh Gyan Vihar University, Jaipur, India*
²*Department of Mechanical Engineering, Suresh Gyan Vihar University, Jaipur, India*
³*Faculty of Sciences (Computer Science), Suresh Gyan Vihar University, Jaipur, India*

Author Emails

^{a)} deepshikha.research24@gmail.com, ^{b)} himanshuvasnani1@gmail.com, ^{c)} amittiwari992@gmail.com, ^{d)} amitverma.researchwork@gmail.com

Abstract - India's FinTech sector sanctioned 31.5 crore personal loans between April 2018 and December 2024, drawing gig workers, self-employed individuals, and first-time borrowers into formal credit at unprecedented scale. Most of these borrowers carry little or no credit bureau history, rendering conventional scoring models ineffective for a population that now constitutes nearly half of all new FinTech loan applicants. This paper addresses this through a systematic literature review of 55 peer-reviewed studies published between 2020 and 2025 across five academic databases using the PRISMA protocol. Findings are organised under six thematic domains: machine learning and ensemble methods, borrower segmentation, UPI-based alternative data, explainable AI, algorithmic fairness, and India's regulatory framework. Gap analysis shows no published study integrates all five dimensions within a single operational framework for Indian personal lending. To address this, we propose the Segment-Aware Credit Risk Framework (SACRF), a six-phase design that classifies borrowers into three population groups, substitutes simulated UPI transaction features for unavailable bureau scores, trains separate ML models per segment, produces SHAP-based explanations at segment level, and evaluates fairness using the Disparate Impact Ratio. The framework is aligned with RBI Model Risk Management guidelines and GDPR Article 22.

Key word - Credit risk, borrower segmentation, explainable AI, UPI data, thin-file borrowers, India FinTech.

1. INTRODUCTION

A. India's Personal Lending Landscape

India's personal lending landscape has undergone a dramatic structural shift over the past seven years. Between April 2018 and December 2024, FinTech NBFCs sanctioned 31.5 crore loans totalling INR 3.4 lakh crore [1] — a scale of credit disbursement that would have been implausible a decade ago. The composition of this growth tells a more important story than the headline numbers. FinTechs account for 76 per cent of total loan sanctions by volume but only 13 per cent by value, which means they are making vast numbers of small-ticket loans to a borrower profile that is unfamiliar to the traditional credit system. Many of these borrowers are accessing formal credit for the very first time. Evidence from segmentation research [2, 3] shows that such populations carry risk profiles that differ substantially from those of salaried, bureau-tracked borrowers — a gap that conventional one-size-fits-all models are poorly equipped to handle. Table 5 draws together the headline figures from FACE [1], the RBI Financial Stability Report [4], and TransUnion CIBIL [5].

India's Economic Survey 2025-26 documents a related shift in the workforce: gig workers grew from 7.7 million in FY2021 to 12 million in FY2025, with income volatility explicitly identified as the primary barrier to their credit access [6]. MicroSave Consulting puts the resulting MSME credit gap at INR 80 trillion, with only 14 per cent of 64 million small businesses currently able to secure formal financing [7]. The problem, it should be stressed, is not that these borrowers are inherently high-risk. It is that the data infrastructure required to assess them — bureau scores, tradeline histories, repayment records — either does not exist for them or is too thin to be statistically reliable. EY India has identified UPI transaction data, including payment decline rates, merchant category patterns, and transaction

frequency, as the practical alternative that lenders need to properly evaluate this generation of applicants [8].

B. Why This Problem Needs a New Approach

Two distinct research threads converge to frame this study. The first is the growing body of credit scoring evidence on borrower heterogeneity. Work in this area [2, 3] has consistently demonstrated that the reasons borrowers default are not uniform across population subgroups — a finding with direct practical implications. A pattern of irregular repayment that signals elevated risk among salaried workers may reflect nothing more than income timing for a gig economy worker. This is not a data quality problem that better imputation can solve; it is a fundamental structural mismatch between a global model's assumptions and the diversity of the population it attempts to score. Systematic reviews covering upwards of 80 studies [9, 10] confirm that gradient boosting methods outperform logistic regression across most datasets, yet they also document substantial variation in those results across different lending contexts — which is precisely what one would expect if borrower populations are heterogeneous in ways that a single model cannot accommodate.

The second research thread concerns model fairness. Hardt et al. [11] established the formal criterion of equalised odds, requiring that a fair classifier produce equal true positive and false positive rates across demographic subgroups. Kozodoi et al. [12] followed with empirical evidence that this standard can be achieved in credit scoring applications without material sacrifice of predictive accuracy. Despite the availability of both the theoretical framework and empirical proof of feasibility, no study has yet applied these fairness standards specifically to India's FinTech lending population.

The regulatory dimension has developed in parallel. The RBI's Draft Circular on Model Risk Management [13] now obliges regulated lenders to maintain explainable, bias-tested, and periodically re-validated ML models. GDPR Article 22 [14] grants borrowers a legal right to explanation when automated credit decisions affect them. The RBI's earlier Working Group on Digital Lending [15] had already identified algorithmic opacity as a systemic risk. Both World Bank research [16] and domestic industry analysis [17] confirm that alternative data sources are the most viable path to credit access for unbanked populations. This paper assembles all five of these elements — segmentation, alternative data, explainability, fairness, and regulatory compliance — into a single coherent framework for the first time.

C. Paper Structure

We conducted a PRISMA-guided systematic review of 55 studies published between 2020 and 2025 and used the findings to design the

Segment-Aware Credit Risk Framework. Section 2 describes how we searched and screened the literature. Sections 3 and 4 cover what we found and where the gaps are. Section 5 introduces the SACRF. Section 6 compares it to existing frameworks. Sections 7 and 8 discuss implications and conclusions.

2. METHODOLOGY

A. Scope and Inclusion Criteria

We followed the PRISMA guidelines for systematic reviews. Papers were included if they were published in peer-reviewed journals or conference proceedings between January 2020 and May 2025, focused on credit risk modelling or loan default prediction using machine learning or explainable AI, and written in English. Papers were excluded if they were purely opinion-based, had no original methodology, or were focused only on corporate or government debt. Studies on credit scoring optimisation methods [18] and hybrid ML approaches [19, 20] were included where they offered insights relevant to the Indian lending context.

B. Databases and Search Terms

We searched Scopus, Web of Science, IEEE Xplore, ScienceDirect, and Google Scholar. Table 1 shows the specific search terms and the number of papers retrieved and included from each database. Figure 1 shows the full PRISMA flowchart. We started with 1,300 unique records after removing duplicates. Of these, 213 were read in full and 55 were finally selected. We also reviewed reference lists manually, which added 5 papers. Studies on deep learning architectures such as CNN-Transformer models [21], tree-based non-linear methods [22], and synthetic data approaches [23] were included within the ML theme. Comprehensive trend studies [24] helped contextualise how the field has developed.

Table - 1 : PRISMA Search Strategy and Inclusion Summary

Database	Key Search Terms	Records	Included
Scopus	credit risk + segmentation + ML	312	18
Web of Science	of explainable AI + credit scoring + India	198	14
IEEE Xplore	alternative data + thin-file + default prediction	156	11
ScienceDirect	borrower segmentation + XAI + fintech lending	224	16

Database	Key Search Terms	Records	Included
Google Scholar	UPI transaction data + credit scoring + India	410	11
Reference lists	Manual backward citation search	—	5
Total	After deduplication and full-text screening	1,300	55

Note: A paper was included if it addressed at least one of the six thematic domains with a quantitative methodology.




Fig - 1: PRISMA screening and selection flowchart

C. Coding and Theme Distribution

Each paper was coded independently against six themes: ML and ensemble methods, borrower segmentation, alternative and UPI-based data, explainable AI, fairness and bias, and India's regulatory context. Tables 2A and 2B show ten representative papers from the corpus to illustrate the range. Table 4 shows how many papers fell into each theme by year. The growth in segmentation and XAI research over the period is notable — these are the two themes most central to the SACRF.

3. THEMATIC FINDINGS

A. Machine Learning and Ensemble Methods

The foundational benchmarking work for this field remains the study by Lessmann et al. [25], who evaluated 41 classification algorithms across eight real-world credit datasets and found gradient

boosting consistently superior to logistic regression and decision tree variants. That finding has been replicated many times since, most recently by Wang et al. [26], who reported XGBoost achieving 91 per cent accuracy on personal loan data. Two major surveys covering the credit scoring literature — Dastile et al. [9] with over 90 studies and Shi et al. [10] with more than 80 — reach the same conclusion: ensemble methods dominate, but performance varies significantly across dataset characteristics and lending contexts.

The three leading gradient boosting libraries each occupy a different practical niche. Evidence from segment-aware modelling studies [2, 3] shows they behave differently across borrower subgroups, which is why the SACRF does not select a single model for the entire portfolio. XGBoost offers strong regularisation control, making it well-suited to noisy datasets. LightGBM's leaf-wise tree growth gives it a speed advantage on large datasets that matters at production scale. CatBoost's native handling of categorical variables is particularly relevant for the Indian FinTech context, where occupation type, loan purpose, and geographic zone are meaningful predictors that other libraries require preprocessing to handle. Feature selection methods [27] and G-XGBoost approaches designed for thin training sets [28] address the sample scarcity that is routine in new product launches. Markov et al. [24] confirmed in their 2022 trend review that gradient boosting has moved from research preference to production standard across most lending environments.

Beyond pure gradient boosting, hybrid architectures have attracted growing research interest. Weng and Huang [19] and Machado and Karray [20] both demonstrate accuracy improvements from combining multiple ML methods in commercial lending settings. Deep learning variants — including symmetry-aware architectures [29] and CNN-Transformer hybrids [21] — show competitive performance on large datasets but remain problematic for regulated deployments because their decision logic is difficult to audit. Methods that graft non-linear tree effects onto a logistic regression backbone [22] represent a practical middle ground for institutions where explainability is a hard requirement rather than a design preference.

Class imbalance deserves specific mention because it is structurally present in every credit scoring dataset: borrowers who default are always the minority class, sometimes by ratios exceeding 20 to one. La Gatta et al. [23] showed that autoencoder-based synthetic augmentation of the minority class produces measurable recall improvements. Aruleba and Sun [30] reported similar gains from stacked ensemble methods combined with SMOTE-ENN resampling, and Shan et al. [31] achieved strong minority-class recall through ensemble combinations. Burugulla [32] situates AI-based credit scoring within the financial inclusion context, which is directly relevant

to the SACRF's treatment of Segment 3 thin-file borrowers. The persistent limitation across this entire body of work — and it is worth naming explicitly — is that almost every study evaluates one model on one dataset. Segmentation-based approaches [2, 3] challenge this design by treating borrower heterogeneity as structure to be modelled rather than noise to be averaged away.

B. Borrower Segmentation

The study closest to what the SACRF proposes is Idbenjra, Coussement, and De Caigny [2], published in *Decision Support Systems* in 2024. Their analysis covered 65,532 customers described by 90 features at a European financial institution, and compared a single global scoring model against segment-specific models estimated separately for distinct borrower groups. The segment-specific approach was better on both dimensions that matter for deployment: predictive accuracy and stability across validation runs. The mechanism they identify is conceptually important — different borrower groups default for structurally different reasons, so a global model is inevitably a compromise that serves no segment optimally. This finding provides the foundational justification for the SACRF's segment-first architecture.

Parallel evidence comes from other contexts. Bakoben, Bellotti, and Adams [3] showed that k-means clustering of account-level transaction data improves default prediction in UK retail banking. Lazo, Calabrese, and Bravo [33] found the same segmentation benefit in agricultural lending in Chile, which suggests the pattern is not specific to consumer finance or any particular market. Gu et al. [34] addressed a well-known practical limitation of k-means — cluster instability across different runs on the same data — with their LASCA algorithm, presented at ACM SIGKDD 2024. LASCA uses consensus clustering to produce segment assignments that remain stable across resampling, which matters enormously for a framework that must produce auditable, reproducible outputs. Kasem, Hamada, and Taj-Eddin [35] demonstrated that AI-driven customer profiling is already operationally mature in commercial settings, confirming the technical feasibility of the approach.

What the literature has not yet produced is any study applying this logic to India's FinTech lending population. The combination of borrower types that characterises a typical Indian digital lending portfolio — a salaried professional with full bureau history, a gig delivery worker with 18 months of UPI transactions but no tradelines, and a first-generation PMJDY account holder with no formal financial record — cannot be modelled effectively by a single global score. These three borrower types face different income structures, different relationships with the formal credit system, and different default risk drivers. Forcing them into one model is precisely the root of the thin-file access problem.

C. Alternative and UPI-Based Data

The foundational study on alternative data in FinTech lending is the work of Berg, Burg, Gombovic, and Puri [36], published in *The Review of Financial Studies*. Drawing on approximately 250,000 loans from a German FinTech lender, they demonstrated that basic digital behavioural variables — the type of device used to apply for the loan, the operating system, even the email provider — predict default as accurately as formal bureau scores. The explanation is intuitive: these variables are indirect signals of financial habits and social context that the lender cannot observe directly. Das et al. [37] extended this logic to payment transaction data, finding that transaction behaviour patterns provide meaningful default signals for new-to-credit borrowers — exactly the population that constitutes Segment 3 of the SACRF. Alamsyah et al. [38] showed 90 per cent predictive accuracy from social media analytics in Indonesian fintech lending, confirming that this approach generalises across emerging market contexts. Björkegren and Grissen [39] showed that mobile phone usage signatures predict credit repayment for thin-file borrowers in developing countries, establishing the conceptual bridge between behavioural data and default risk in low-bureau-coverage populations.

The application to India centres on UPI. An NBER working paper [40] found that UPI adoption enabled new-to-credit borrowers to access formal loans — but only through fintech lenders, not traditional banks, because banks lacked the analytical capacity to interpret UPI signals as creditworthiness evidence. The NPCI records more than 22,000 crore transactions annually on the UPI network [41], generating a continuous behavioural record for tens of millions of borrowers who leave no trace in bureau databases. Croux et al. [42] found payment behaviour variables among the strongest predictors of default in their analysis of US fintech lending data. Research on UPI adoption [43] and on the relationship between UPI usage and formal credit access [44] together suggest that the behavioural signals in UPI activity are both present and measurable at the scale required for production credit scoring. Wen, Sui, and Lu [45] studied consumer information effects in personal credit evaluation and reinforced the case for transaction-level variables as supplementary inputs. These studies collectively motivate the six UPI features proposed for Phase 1 of the SACRF: weekly transaction frequency, payment success ratio, merchant category diversity, average transaction value, P2P-to-P2M ratio, and a composite digital lifestyle indicator.

A striking gap remains. Despite this accumulation of evidence, no published academic study has incorporated actual UPI transaction features into a supervised credit default model. This is the most practically consequential open problem in the field as it applies to India. The fintech lending ecosystem

research by Palladino [46], Singh [47], He et al. [48], Haddad and Hornuf [49], and Ganvir et al. [50] documents the institutional demand for better thin-file assessment tools. Machado, Chen, and Osterrieder [51] confirm that analytical ML frameworks of this type transfer across credit contexts, supporting the claim that a framework designed for Indian FinTech conditions can be replicated by other lenders.

D. Explainable AI in Credit Scoring

Lundberg and Lee [52] introduced SHAP at NeurIPS 2017 with a result that resolved a long-standing ambiguity in machine learning explainability: they proved that Shapley values from cooperative game theory are the unique attribution method satisfying local accuracy, missingness, and consistency simultaneously. In credit scoring terms, local accuracy means an explanation fully accounts for the model's output. Missingness means features absent from a particular prediction contribute nothing to its explanation. Consistency means that if a feature becomes more influential for the model's output, its attributed importance cannot decrease. These three properties together make SHAP explanations substantively different from post-hoc approximations such as LIME, and they make SHAP the only practically defensible approach for individual borrower explanations under GDPR Article 22 [14].

The practical performance evidence is reassuring. Bussmann et al. [53] achieved AUC of 0.93 using TreeSHAP-integrated XGBoost on a credit risk dataset, with no degradation from adding interpretability constraints. Hlongwane et al. [54] constructed a Shapley-based credit scorecard that reached 87-89 per cent accuracy while matching the transparency of a logistic regression. Martinez et al. [55] compared SHAP against LIME directly and found SHAP's explanations stable across repeated applications, whereas LIME's outputs varied, which disqualifies it for regulatory use cases where reproducibility matters. Nooji and Sugave [56] and Qi et al. [57] each built explainable ensemble models for credit risk; Qi et al. specifically reported 90 per cent accuracy with CatBoost plus SHAP on LendingClub data. Chong, Meziane, and Labadin [58] applied XAI specifically to peer-to-peer lending credit risk, and Mujo et al. [59] showed in a 2025 study that XAI implementation measurably increases the perceived transparency of credit decisions among loan officers.

Soni, Jethava, and Ganatra [60] and Nallakaruppan et al. [61] both provide recent comprehensive overviews confirming that XAI has moved from a research topic to an operational requirement in regulated financial services. Gurtler and Zollner [62] explored combining white-box and black-box models as an alternative interpretability architecture, while Han et al. [29] demonstrated symmetry-aware deep learning for credit risk. The shared gap across all of these studies is one that the SACRF specifically

addresses: every one of them generates SHAP importance maps at the portfolio level. The average feature importance across a mixed borrower population does not tell a regulator or a borrower why a particular thin-file applicant was declined. Segment-level SHAP maps — which the SACRF produces in Phase 4 — answer that question directly, because the feature weights reflect only the subpopulation to which the applicant belongs.

E. Fairness and Bias in Credit Scoring

The most systematic empirical study of fairness in credit scoring is Kozodoi, Jacob, and Lessmann [12], published in the European Journal of Operational Research. Working across multiple credit datasets and four model types, they tested pre-processing, in-processing, and post-processing fairness interventions and found that approval rate disparities can be substantially reduced without meaningful loss of predictive accuracy or portfolio profitability. The choice of fairness criterion matters, and their analysis makes the options concrete. The Disparate Impact Ratio, which requires that protected groups receive loan approvals at a rate at least 80 per cent of the rate for the reference group, is the most practically applicable criterion for Indian lending because it can be calculated directly from scorecard outputs without access to borrower-level demographic attributes at inference time. Equalised odds, formalised by Hardt et al. [11], is more demanding: it requires that both the true positive rate and false positive rate be equal across groups, which means the model must produce the same error profile for creditworthy applicants regardless of demographic group membership.

Nwafor, Nwafor, and Brahma [63] demonstrated that hybrid explainable models can achieve F1 scores around 0.88 while jointly satisfying multiple fairness criteria. The National Consumer Law Center [64] documented extensively how conventional credit scoring systems encode historical patterns of discrimination — through proxy variables such as geographic region, education, and employment sector — that carry particular relevance in India, where historical inequalities in access to formal finance have been structured around caste, gender, and geography. Oreski [65] explored synthetic data generation for credit datasets, which is directly applicable to fairness interventions: when protected groups are underrepresented in training data, generating distributional synthetic data for those groups is one viable approach to bias correction. Coskun and Turanli [66] and Al Shiam et al. [67] both include demographic analysis components in recent ML credit scoring work. Bitetto et al. [68] showed that SME credit models benefit from explicit demographic calibration, a finding that extends to India's MSME segment.

The evidence gap for India is complete. No domestic credit scoring study has published DIR or equalised odds calculations for its own model outputs. This is not because the methodology is unavailable —

Kozodoi et al. [12] and Hardt et al. [11] supply it — or because the problem is unrecognised. It reflects a genuine absence of research that situates the Indian lending context within the fairness measurement frameworks the international literature has developed. Phase 5 of the SACRF fills this gap by computing DIR per segment across gender, income tier, and geographic category.

F. India's Regulatory Environment

India's regulatory environment for algorithmic credit assessment has changed substantially over the past four years. The RBI's Working Group on Digital Lending [15] first flagged the systemic risks associated with opaque algorithmic underwriting in 2021, identifying lack of model transparency as a governance concern rather than a technical limitation to be addressed later. The subsequent Draft Circular on Model Risk Management [13] converted that concern into binding requirements: lenders using ML models in regulated credit decisions must now demonstrate explainability, conduct demographic bias auditing, and maintain documented periodic validation. The RBI Financial Stability Report [4] has flagged rising personal loan delinquency rates as a systemic concern, adding an urgency dimension to governance requirements that might otherwise be treated as compliance formalities. Roy and Vasa [69] confirm in their international review that this regulatory direction is consistent with major financial markets globally, not an India-specific deviation. The RBI's Draft Directions on Capital Adequacy [70] add a third governance layer requiring lenders to demonstrate that credit risk models meet minimum validation standards.

GDPR Article 22 [14] provides the international template that India's Digital Personal Data Protection Act 2023 closely follows, creating a right to explanation for any borrower subject to an automated credit decision. Palladino [46] and Haddad and Hornuf [49] have each shown that fintech growth historically creates regulatory pressure that eventually reaches the traditional banking sector as well. World Bank research [16] and Face of India Consortium analysis [17] confirm that alternative data — including transaction history, mobile usage, and behavioural signals — is now accepted internationally as a legitimate input to credit assessment, with appropriate governance. The question for Indian lenders is therefore not whether to use such data, but how to use it within a framework that satisfies all three regulatory layers simultaneously: explainability under the Model Risk Management circular, fairness under the Fair Practices Code, and data rights under the DPDP Act. That is the design problem the SACRF addresses.

4. RESEARCH GAP ANALYSIS

A. Gap Matrix

Table 3 maps the eight representative studies reviewed above against the five dimensions of the proposed SACRF. The picture is clear: no existing study covers all five dimensions. The SACRF is the only framework that does.

Table - 3: Research Gap Matrix — Five Dimensions vs. Representative Studies

Study	Segment.	UPI Data	XAI	Fairness	India Reg.
Pre-2022 studies	ML	×	×	×	×
Idbenjra et al. [2]	✓	×	×	×	×
Bussmann et al. [53]	×	×	✓	×	×
Kozodoi et al. [12]	×	×	✓	✓	×
Berg et al. [36]	×	✓	×	×	×
Das et al. [37]	×	✓	×	×	×
Gu et al. [34]	✓	×	×	×	×
SACRF This Study	—	✓	✓	✓	✓

Note: Checkmark = the dimension is addressed by that study; Cross = absent or peripheral.

B. The Five Gaps

- **Gap 1 — No unified five-dimension framework:** Studies address one or two dimensions in isolation. Idbenjra et al. [2] cover segmentation but not XAI, fairness, or regulatory compliance. Kozodoi et al. [12] cover fairness and XAI but not segmentation or India context. No study covers all five together.
- **Gap 2 — No UPI-based credit model:** Despite NPCI reporting 22,000 crore annual transactions [41], EY India specifying UPI features as necessary for fintech underwriting [8], and NBER evidence [40] that UPI enables new-to-credit borrowers to access formal credit, no published academic study uses UPI transaction data in a supervised credit default model.
- **Gap 3 — No segment-level SHAP:** Every XAI credit scoring study — Lundberg and Lee [52], Bussmann et al. [53], Hlongwane et al. [54], Martinez et al. [55], Nooji and Sugave [56], Qi et al. [57] — computes SHAP globally across all borrowers. None generates segment-specific feature importance maps.
- **Gap 4 — No India fairness evaluation:** Kozodoi et al. [12] and Hardt et al. [11] provide the methodology. The National

Consumer Law Center [64] documents the problem. But no Indian credit scoring study applies DIR or equalised odds to its model outputs.

- Gap 5 — Segmentation not tested in India: Idbenjra et al. [2] themselves acknowledged the single-market limitation. India's gig-economy borrower population [6, 7], with its specific combination of UPI activity and bureau absence, has not been modelled through a segment-aware lens.

5. THE SACRF FRAMEWORK

A. Overview

The Segment-Aware Credit Risk Framework is a six-phase operational design that addresses all five gaps identified above. Figure 2 shows the architecture. The framework satisfies the RBI's Model Risk Management requirements [13], GDPR Article 22 [14], and the DPDP Act 2023 simultaneously. Each phase feeds directly into the next and produces a segment-specific credit scorecard as the final output.




Fig - 2 : SACRF — six-phase architecture

B. Phase 1: Data Integration and Feature Engineering

Phase 1 brings together traditional bureau features and six simulated UPI features. The bureau features are the standard ones: CIBIL score (where available), number of active tradelines, maximum days past due in the last 12

months, credit utilisation ratio, and history length. The UPI features are simulated from NPCI distributional statistics [41] and validated against EY India's technical specification [8]. They are: weekly transaction frequency, payment success ratio, merchant category diversity, average transaction value, P2P versus P2M ratio, and a digital lifestyle indicator. All six map to data fields that are accessible through the RBI Account Aggregator framework. The simulation follows La Gatta et al.'s [23] distribution-preserving approach. World Bank [16] and Face of India Consortium [17] research confirms that these data types are accepted as legitimate credit inputs.

C. Phase 2: Borrower Segmentation

Phase 2 groups borrowers into three segments using a combination of rule-based criteria and k-means clustering. The rules use occupation type, credit history length, and CIBIL score availability as primary filters. Within each rule-based bucket, k-means clustering with consensus averaging — based on the LASCA method of Gu et al. [34] — refines the boundaries using the full feature set. The three segments are: Segment 1 (Salaried/Formal) for borrowers with regular income, history above 24 months, and a CIBIL score; Segment 2 (Self-Employed/Gig) for those with platform-based income, 6-24 months of history, and limited bureau data; and Segment 3 (New-to-Credit/Thin-File) for borrowers with history below 6 months or no CIBIL at all. Segment 2 is the fastest-growing per the Economic Survey [6]. Segment 3 is where the UPI feature integration matters most, consistent with NBER evidence [40] on UPI and credit access.

D. Phases 3 to 6

Phase 3 trains four candidate models within each segment: logistic regression as an interpretable baseline, random forest, XGBoost, and CatBoost. Cross-validation within each segment's training data selects the best-performing model. A weighted aggregation layer then combines the segment scores into a single portfolio scorecard. Phase 4 applies TreeSHAP at both the global level (for model governance reports) and at the individual level (for borrower-facing decline explanations under GDPR Article 22 [14] and the RBI Fair Practices Code). The key design choice — segment-level SHAP — means the explanation for a gig worker reflects features that actually matter for gig workers, not an average across all

borrower types. Phase 5 computes the Disparate Impact Ratio and equalised odds [13, 14] across gender, income tier, and geography within each segment. This enables the detection of within-segment disparities that a global fairness check would miss entirely. Phase 6 compares the segment-specific scorecard against a non-segmented global XGBoost baseline on a holdout test set. This is the direct test of H01.

E. Testable Hypotheses

- H01: Segment-specific ML models do not significantly improve credit default prediction accuracy over a non-segmented global model (test: DeLong AUC comparison with Bonferroni correction).
- H02: Adding UPI and behavioural alternative data does not significantly improve prediction for Segment 3 thin-file borrowers (test: paired AUC comparison within segment).
- H03: Credit risk model performance does not differ significantly across the three borrower segments in AUC and F1-Score (test: Kruskal-Wallis H-test with post-hoc Dunn comparison).
- H04: Segment-level SHAP does not provide significantly more actionable insights than global feature importance for credit decisions (test: blinded expert survey with standardised scoring rubric).

6. COMPARATIVE FRAMEWORK ANALYSIS

Table 6 compares the SACRF directly against the Idbenjra et al. [2] framework across nine dimensions. This is a fair comparison because Idbenjra et al. [2] is the closest existing study to what the SACRF proposes — it uses segmentation-based credit scoring with empirical validation. The table shows that the SACRF makes five substantive additions: an Indian FinTech context, explicit borrower type definitions, UPI data integration, segment-level SHAP, and regulatory alignment. Each addition is motivated by a specific gap or requirement, not by novelty for its own sake.

Table - 6 : SACRF vs. Idbenjra et al. [2] — Structural Comparison

Dimension	Idbenjra et al. [2]	SACRF (Proposed)
Market context	Single European market	Indian FinTech sector
Borrower segments	DT-generated splits	3 defined borrower types

Dimension	Idbenjra et al. [2]	SACRF (Proposed)
Alternative data	Bureau features only	Bureau + UPI features
Explainability	Not included	Segment-level SHAP
Fairness testing	Not included	DIR per segment
Regulatory alignment	Not stated	RBI + GDPR + DPDP
Hypotheses	Predictive accuracy only	4 testable H0s
Deployment template	Research only	Operational blueprint

Note: SACRF extends the segmentation framework of Idbenjra et al. [2] to the Indian FinTech context with five targeted additions.

The most practically significant addition is the UPI data. Idbenjra et al. [2] worked with a European dataset where bureau coverage is near-universal. In India, 41 per cent of FinTech borrowers are first-time credit users [1] with no bureau record. For them, bureau-only segmentation produces one undifferentiated group with no scoring capability. The SACRF's UPI feature integration turns that data void into a usable signal. The Face of India Consortium [17] has documented exactly this problem in Indian alternative data deployments. Regulatory alignment is also structurally different: Idbenjra et al.'s European context does not include the specific requirements of the RBI's Model Risk Management Circular [13] or India's DPDP Act 2023.

7. DISCUSSION AND IMPLICATIONS

A. For Banks and FinTech Lenders

For regulated lenders, the SACRF provides a concrete implementation blueprint that satisfies the RBI's Model Risk Management requirements [13] while tackling the borrower diversity challenge that EY India [8] and MicroSave Consulting [7] have both identified as the central constraint on India's FinTech lending market. The segment-first architecture changes how model validation is conducted in a meaningful way: instead of computing aggregate performance metrics that average across a heterogeneous portfolio, validation occurs within each borrower subpopulation separately. For Segment 2 gig workers, UPI transaction regularity provides a creditworthiness signal that lenders currently have no systematic framework to interpret. Ganvir et al. [50] document the personal loan assessment practices at a major Indian private bank and confirm that the need for more granular borrower profiling is already recognised at the institutional level. Kumar [44] shows the same recognition in the fintech segment. The operational feasibility of the SACRF is supported by its compatibility with existing ML

infrastructure. XGBoost, CatBoost, and LightGBM are standard components of all major model serving platforms; per-segment training adds no architectural complexity, only parallel training runs. Machado, Chen, and Osterrieder [51] confirm that analytical ML frameworks of this design type transfer across different credit contexts without substantial rearchitecting. The fintech ecosystem literature from Croux et al. [42], He et al. [48], and Haddad and Hornuf [49] establishes the commercial context in which deployment would occur.

B. For Financial Inclusion

The SACRF's implications for financial inclusion extend beyond risk management into the core policy agenda of the Indian state. Segment 3 — New-to-Credit/Thin-File — comprises the borrowers that PM Jan Dhan Yojana brought into the banking system and that MUDRA loans are designed to serve, but which conventional credit models routinely decline or underprice because bureau data does not exist for them. The NBER finding [40] that UPI adoption has already enabled new-to-credit borrowers to reach formal lenders through fintech channels identifies a specific mechanism: the behavioural credit signal is present in the UPI record, but the analytical infrastructure to interpret it does not yet exist. The SACRF builds that infrastructure. UPI adoption research [43, 44] and consumer information studies [45] together confirm that the data is accessible, measurable, and improving in coverage as India's digital payment infrastructure matures. The World Bank [16] and Face of India Consortium [17] have both documented the international policy consensus that alternative data should be used to extend credit inclusion. Bitetto et al. [68] show that population-specific model design — the same principle underlying the SACRF — materially improves credit access for SMEs, which maps directly to the MUDRA borrower population. Burugulla [32] connects AI-powered credit scoring to financial inclusion explicitly, providing a direct conceptual link between the technical design of the SACRF and its broader policy purpose.

C. For Regulators

The SACRF converts the RBI's Model Risk Management requirements [13] from a compliance checklist into operational architecture. Phase 5's segment-level Disparate Impact Ratio calculation addresses the fair lending requirement in a form that produces auditable numerical outputs rather than qualitative assurances. Phase 4's segment-level SHAP explanations satisfy both the individual borrower explanation right under GDPR Article 22 [14] and the DPDP Act, and they do so in a way that is meaningful for the specific borrower type rather than averaged across a heterogeneous portfolio. The Phase 6 holdout comparison against a non-segmented baseline produces exactly the periodic validation documentation the Model Risk Management circular

requires. Roy and Vasa [69] confirm that this regulatory trajectory is globally consistent, not India-specific. The Capital Adequacy Directions [70] add a further governance layer that the SACRF's end-to-end audit trail satisfies. A concern sometimes raised in the literature about the tension between accuracy and fairness deserves direct response: Nwafor et al. [63] and Nallakaruppan et al. [61] both show empirically that this tension is overstated — explainable, fair, and accurate models are jointly achievable. The broad convergence of recent XAI work [59, 60, 61, 66, 67] toward operational deployment in regulated lending confirms that the SACRF is aligned with the direction the field is already moving.

D. Limitations

Three limitations of this study deserve direct acknowledgement. The 55-paper corpus was assembled from five indexed academic databases, which means studies published in non-indexed Indian regional journals may have been missed. Some India-specific evidence on alternative data usage and regulatory practice may therefore be absent from our analysis. Second, the theme coding was conducted by a single researcher. A second independent coder working in parallel, followed by inter-rater reliability testing, would have strengthened the validity of the thematic distribution reported in Table 4. Third, and most fundamentally, the SACRF is a proposed framework. None of the four hypotheses in Section 5E has been tested empirically. The evidence reviewed in Sections 3 and 4 provides reasonable grounds for expecting each hypothesis to hold, but reasonable grounds are not empirical results. Datasets that combine segment-level borrower labels, six-dimensional UPI transaction histories, and demographic attributes at the scale required for statistical validation do not currently exist in publicly accessible form. RBI sandbox access is making this increasingly feasible, and synthetic data generation methods [25, 65] represent an interim pathway for validation while real datasets become available.

8. CONCLUSION

This review of 55 credit risk studies from 2020 to 2025 establishes that the Indian personal lending literature has produced strong evidence for each of five components — borrower segmentation, UPI-based alternative data, segment-level explainability, algorithmic fairness evaluation, and regulatory compliance architecture — but has not yet assembled them within a single operational framework. The five gaps documented in Section 4 are not academic constructs. India's Economic Survey 2025-26 [6] documents that income volatility is actively preventing gig workers from obtaining formal credit. EY India [8] specifies UPI transaction data as the tool lenders need to assess this population correctly.

The NBER [40] shows UPI adoption has already enabled credit access for thin-file borrowers in practice. MicroSave Consulting [7] puts the unmet credit demand at INR 80 trillion. And the RBI's Model Risk Management Circular [13] makes explainability and demographic bias testing legal obligations, not optional enhancements. The problem is real, the data infrastructure to address it exists, and the regulatory expectation is now binding. The SACRF proposed here addresses all five gaps within a design that is both regulation-compliant and operationally feasible. Section 6 demonstrates how it extends the closest existing framework [2] — Idbenjra et al.'s segmentation approach — with specific, motivated additions for the Indian context. The four testable hypotheses in Section 5E give researchers a concrete empirical agenda. Three directions stand out for follow-on work: validating the SACRF on real Account Aggregator data through the RBI sandbox programme; extending the segment definitions to cover MSME and agricultural lending portfolios [68], where the thin-file problem is equally severe; and examining privacy-preserving training mechanisms [65] for multi-lender model development that would comply with the DPDP Act 2023. The Account Aggregator framework and the expansion of India's UPI infrastructure are together creating the data conditions in which this work can be done. The research agenda the SACRF outlines is therefore timely rather than premature.

Table - 2A : Reviewed Papers — Author, Method and Dataset (10 of 55 shown)

Ref.	Author(s), Year	Method	Dataset / Context
[2]	Idbenjra et al. (2024)	Cluster-based models	European; N=65,532; 90 features
[3]	Bakoben et al. (2020)	K-means clustering	UK banking account data
[11]	Hardt et al. (2016)	Equalised odds criterion	Multiple ML model types
[12]	Kozodoi et al. (2022)	Fairness-aware ML	Multiple credit datasets
[25]	Lessmann et al. (2015)	41 classifiers benchmarked	8 real-world credit datasets
[36]	Berg et al. (2020)	ML + digital footprint	German FinTech; 250K loans
[52]	Lundberg & Lee (2017)	SHAP framework	Multiple ML model types
[53]	Busmann et al. (2021)	TreeSHAP + XGBoost	European credit risk data
[37]	Das et al. (2023)	ML + transaction patterns	New-to-credit borrowers
[34]	Gu et al. (2024)	LASCA segmentation	Large-scale credit dataset

Table - 2B: Reviewed Papers — Key Finding and Research Gap (10 of 55 shown)

Ref.	Key Finding	Gap Identified
[2]	Segment models outperform global on accuracy and stability	No XAI; no UPI data; single EU market
[3]	Clustering behaviours account improves default forecasting	No segmentation labels; UK context only
[11]	Equalised odds satisfies fairness without major accuracy loss	No credit scoring deployment tested
[12]	Fairness processors preserve profitability while reducing bias	No segment-level fairness; no India
[25]	Gradient boosting outperforms LR across 8 credit datasets	No segmentation; no India context
[36]	Digital footprint matches bureau score quality	No segmentation; no India/UPI
[52]	SHAP unifies 6 XAI methods; game-theoretically grounded	Global only; no segment-level
[53]	AUC 0.93 with SHAP; no accuracy loss from interpretability	No segment-level SHAP
[37]	Transaction data improves prediction for NTC borrowers	No XAI; no fairness; no India
[34]	LASCA: stable, large-scale customer segmentation	No UPI data; no regulatory alignment

Note: Continuation of Table 2A for the same ten papers.

Table - 4 : Annual Distribution of Reviewed Papers by Theme (2020-2025)

Theme	2020-21	2022	2023	2024-25	Total
ML / Ensemble	9	8	7	10	34
Segmentation	2	2	3	4	11
Alt. / UPI data	3	3	4	5	15
Explainable AI	4	5	6	8	23
Fairness / bias	2	3	4	5	14
India regulatory	3	3	4	5	15
Total	14	14	14	13	55

Note: Papers coded under more than one theme are counted in each relevant theme.

ACKNOWLEDGEMENT

The author thanks the doctoral supervisory committee at Suresh Gyan Vihar University, Jaipur for guidance on this research. This study is part of an ongoing PhD programme on segment-aware credit risk modelling for Indian personal lending. No external funding was received. The author declares no conflicts of interest.

REFERENCES

1. FACE: FinTech Personal Loan Market Report, April 2018 to December 2024. Fintech Association for Consumer Empowerment, India (2025)
2. Idbenjra, K., Coussement, K., De Caigny, A.: Investigating the beneficial impact of segmentation-based modelling for credit scoring. *Decis. Support Syst.* 179, 114170 (2024)
3. Bakoben, M., Bellotti, T., Adams, N.: Identification of credit risk based on cluster analysis of account behaviours. *J. Oper. Res. Soc.* 71(5), 775–783 (2020)
4. Reserve Bank of India: Financial Stability Report, June 2025. RBI, Mumbai (2025)
5. TransUnion CIBIL: Quarterly Insights Report, Q4 2024. TransUnion CIBIL, Mumbai (2024)
6. Government of India: Economic Survey 2025–26. Ministry of Finance, New Delhi (2026)
7. MicroSave Consulting: From UPI to ULI: India's Next Digital Infrastructure Imperative. MSC Report (2025)
8. EY India: How Recent Changes to UPI Are Helping Fill India's Credit Gap. Ernst and Young LLP (2025)
9. Dastile, X., Celik, T., Potsane, M.: Statistical and machine learning models in credit scoring: A systematic literature survey. *Appl. Soft Comput.* 91, 106263 (2020)
10. Shi, S., Tse, R., Luo, W., D'Addona, S., Pau, G.: Machine learning-driven credit risk: A systemic review. *Neural Comput. Appl.* 34(17), 14327–14339 (2022)
11. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 29, pp. 3323–3331 (2016)
12. Kozodoi, N., Jacob, J., Lessmann, S.: Fairness in credit scoring: Assessment, implementation and profit implications. *Eur. J. Oper. Res.* 297(3), 1083–1094 (2022)
13. Reserve Bank of India: Draft Circular on Model Risk Management in Credit. RBI, Mumbai (2024)
14. European Union: General Data Protection Regulation, Article 22. Official Journal of the European Union (2018)
15. Reserve Bank of India: Report of the Working Group on Digital Lending. RBI, Mumbai (2021)
16. World Bank: The Use of Alternative Data in Credit Risk Assessment. World Bank Group (2024)
17. Face of India Consortium: State of Alternative Data in India: A Lending Use Case (2024)
18. Kamimura, E.S., Pinto, A.R.F., Nagano, M.S.: A recent review on optimisation methods applied to credit scoring models. *J. Econ. Finance Admin. Sci.* 28(56), 352–371 (2023)
19. Weng, C.H., Huang, C.K.: A hybrid machine learning model for credit approval. *Appl. Artif. Intell.* 35(15), 1439–1465 (2021)
20. Machado, M.R., Karray, S.: Assessing credit risk of commercial customers using hybrid machine learning algorithms. *Expert Syst. Appl.* 200, 116889 (2022)
21. Wang, Y., Xu, Z., Yao, Y., Liu, J., Lin, J.: Leveraging convolutional neural network-transformer synergy for predictive modeling in risk-based applications. In: *Proceedings of IEEE EIECC 2024*, pp. 1565–1570 (2024)
22. Dumitrescu, E., Hué, S., Hurlin, C., Tokpavi, S.: Machine learning for credit scoring: Improving logistic regression with non-linear decision tree effects. *Eur. J. Oper. Res.* 297(3), 1178–1192 (2022)
23. La Gatta, V., Postiglione, M., Sperli, G.: A novel augmentation strategy for credit scoring modeling. *Neural Comput. Appl.* 37(9), 6663–6675 (2025)
24. Markov, A., Seleznyova, Z., Lapshin, V.: Credit scoring methods: Latest trends and points to consider. *J. Finance Data Sci.* 8, 180–201 (2022)
25. Lessmann, S., Baesens, B., Seow, H.V., Thomas, L.C.: Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *Eur. J. Oper. Res.* 247(1), 124–136 (2015)
26. Wang, K., Li, M., Cheng, J., Zhou, X., Li, G.: Research on personal credit risk evaluation based on XGBoost. *Procedia Comput. Sci.* 199, 1128–1135 (2022)
27. Guo, W., Zhou, Z.Z.: A comparative study of combining tree-based feature selection methods and classifiers in personal loan default prediction. *J. Forecast.* 41(6), 1248–1313 (2022)
28. Li, J., Liu, H., Yang, Z., Han, L.: A credit risk model with small sample data based on G-XGBoost. *Appl. Artif. Intell.* 35(15), 1550–1566 (2021)
29. Han, X., Yang, Y., Chen, J., Wang, M., Zhou, M.: Symmetry-aware credit risk modeling: A deep learning framework exploiting financial data balance and invariance. *Symmetry* 17(3), 341 (2025)
30. Aruleba, I., Sun, Y.: An improved ensemble method with data resampling for credit risk prediction. *IEEE Access* 13, 71275–71287 (2025)
31. Shan, X., Shi, Y., Cui, H., Wang, Z.: Ensemble learning for imbalanced credit default prediction. In: *Proceedings of IEEE IAEAC 2024* (2024)
32. Burugulla, J.K.R.: Machine learning for credit scoring: An AI-powered big data approach to financial inclusion. *IJARCCCE* 12(12) (2023)
33. Lazo, D., Calabrese, R., Bravo, C.: The effects of customer segmentation, borrower behaviors and analytical methods on the performance of credit scoring models in the agribusiness sector. *J. Credit Risk* 16(4), 119–156 (2021)
34. Gu, Y., Zhang, H., Li, D., Chen, H.: LASCA: A large-scale stable customer segmentation approach to credit risk. In: *Proceedings of ACM SIGKDD 2024*, pp. 5006–5017 (2024)
35. Kasem, M.S., Hamada, M., Taj-Eddin, I.: Customer profiling, segmentation, and sales prediction using AI in direct marketing. *Neural Comput. Appl.* 36(9), 4995–5005 (2024)
36. Berg, T., Burg, V., Gombovic, A., Puri, M.: On the rise of FinTechs: Credit scoring using digital footprints. *Rev. Financ. Stud.* 33(7), 2845–2897 (2020)
37. Das, S., Datta, A., Suresh, H., Ghosh, S.: Credit risk modeling with graph machine learning. *INFORMS J. Data Sci.* 2(2), 197–217 (2023)
38. Alamsyah, A., Hafidh, A.A., Mulya, A.D.: Innovative credit risk assessment: Leveraging social media data for inclusive credit scoring in Indonesia's fintech sector. *J. Risk Financ. Manag.* 18(2), 74 (2025)

39. Björkegren, D., Grissen, D.: Behavior revealed in mobile phone usage predicts credit repayment. *World Bank Econ. Rev.* 34(3), 618–634 (2020)
40. NBER: Breaking Barriers to Financial Access: UPI and Credit Provision in India. NBER Working Paper No. 33259 (2024)
41. NPCI: UPI Digital Payments Growth Report. National Payments Corporation of India (2023)
42. Croux, C., Jagtiani, J., Korivi, T., Vulcanovic, M.: Important factors determining Fintech loan default: Evidence from a Lendingclub consumer platform. *J. Econ. Behav. Organ.* 173, 270–296 (2020)
43. Fahad, Shahid, M.: Exploring the determinants of adoption of Unified Payment Interface (UPI) in India: A study based on diffusion of innovation theory. *Digit. Bus.* 2(2), 100040 (2022)
44. Kumar, R.: The impact of UPI on credit accessibility and digital transactions in India. *IJSDR* 10(1) (2025)
45. Wen, H., Sui, X., Lu, S.: Study on the effect of consumer information in personal credit risk evaluation. *Complexity* 2022, 7340010 (2022)
46. Palladino, L.M.: The impacts of fintech on small business borrowing. *J. Small Bus. Entrep.* 33(6), 639–661 (2021)
47. Singh, A.: Peer-to-peer loan default prophecy in fintech: A comparative analysis. *Corp. Gov.* 6(2) (2024)
48. He, C., Geng, X., Tan, C., Guo, R.: Fintech and corporate debt default risk: Influencing mechanisms and heterogeneity. *J. Bus. Res.* 164, 113923 (2023)
49. Haddad, C., Hornuf, L.: How do fintech start-ups affect financial institutions' performance and default risk? *Eur. J. Finance* 29(15), 1761–1792 (2023)
50. Ganvir, T.R., Wanjari, U.V., Khaire, M.: An analytical study of personal loan procedure at ICICI Bank. *Int. J. Innov. Stud.* 9(1), 629–637 (2025)
51. Machado, M.R., Chen, D.T., Osterrieder, J.R.: An analytical approach to credit risk assessment using machine learning. *Decis. Anal. J.*, 100605 (2025)
52. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, pp. 4765–4774 (2017)
53. Bussmann, N., Giudici, P., Marinelli, D., Papenbrock, J.: Explainable machine learning in credit risk management. *Comput. Econ.* 57(1), 203–216 (2021)
54. Hlongwane, R., Ramabao, K., Mongwe, W.: A novel framework for enhancing transparency in credit scoring: Leveraging Shapley values for interpretable credit scorecards. *PLOS ONE* 19(8), e0308718 (2024)
55. Martinez, A., Sanchez, D., Vellido, A.: Explainability of a machine learning granting scoring model in peer-to-peer lending. *IEEE Access* 8, 64873–64888 (2020)
56. Nooji, P., Sugave, S.: Explainable ensemble technique for enhancing credit risk prediction. *IAES Int. J. Artif. Intell.* 13(1), 917–924 (2024)
57. Qi, J., Yang, R., Wang, P.: Application of explainable machine learning based on CatBoost in credit scoring. *J. Phys.: Conf. Ser.* 1955, 012039 (2021)
58. Chong, P.S., Meziane, F., Labadin, J.: Credit risk prediction for peer-to-peer lending platforms: An explainable machine learning approach. *J. Comput. Soc. Inform.* 1(2), 1–16 (2022)
59. Mujo, A., Nikolla, S., Hoxha, E., Pelivani, E.: Explainable AI in credit scoring: Improving transparency in loan decisions. *J. Inf. Syst. Eng. Manag.* 10(27s), 506–515 (2025)
60. Soni, U., Jethava, G., Ganatra, A.: Latest advancements in credit risk assessment with machine learning and deep learning techniques. *Cybern. Inf. Technol.* 24(4), 22–44 (2024)
61. Nallakaruppan, M.K., Chaturvedi, H., Grover, V., Balusamy, B., Jaraut, P., Bahadur, J., Meena, V.P., Hameed, I.A.: Credit risk assessment and financial decision support using explainable artificial intelligence. *Risks* 12(10), 164 (2024)
62. Gurtler, M., Zollner, M.: Tuning a white box model with black box models. *SSRN Working Paper* 4433967 (2023)
63. Nwafor, C.N., Nwafor, O., Brahma, S.: Enhancing transparency and fairness in automated credit decisions: An explainable novel hybrid machine learning approach. *Sci. Rep.* 14, 25174 (2024)
64. National Consumer Law Center: Past Imperfect: How Credit Scores Bake In Discrimination (2024)
65. Oreski, G.: Synthesizing credit data using autoencoders and generative adversarial networks. *Knowl.-Based Syst.* 274, 110646 (2023)
66. Coskun, S.B., Turanli, M.: Credit risk analysis using boosting methods. *J. Appl. Math. Stat. Inform.* 19(1), 5–17 (2023)
67. Al Shiam, S.A., Hasan, M.M., Pantho, M.J., Shochona, S.A., Nayeem, M.B., Choudhury, M.T.H., Nguyen, T.N.: Credit risk prediction using explainable AI. *J. Bus. Manag. Stud.* 6(2), 61–66 (2024)
68. Bitetto, A., Cerchiello, P., Filomeni, S., Tanda, A., Tarantino, B.: Machine learning and credit risk: Empirical evidence from small- and mid-sized businesses. *Socio-Econ. Plan. Sci.* 90, 101746 (2023)
69. Roy, J.K., Vasa, L.: Transforming credit risk assessment: A systematic review of AI and machine learning applications. *J. Infrastruct. Policy Dev.* 9(1) (2025)
70. Reserve Bank of India: Draft Directions pursuant to Capital Adequacy Policy. RBI, Mumbai (2025)