



Article

Transparency and Explainability in AI Credit Scoring Systems: Case Study of Interpretable ML Design in ScoreTech's Five-Layer Model

Shohruh Komiljonov¹, Farrukh Sadikbaev²

1. Head of Financial Reporting and Strategy Analysis, Uzavtosanoat JSC
2. Head of Project Management, Mobiuz (UMS LLC)

Abstract: This study examines the challenges of transparency and explainability in artificial intelligence-based credit scoring by analyzing ScoreTech's peer-to-peer credit intelligence platform in Uzbekistan. As machine learning models increasingly replace traditional credit assessment methods, concerns regarding algorithmic opacity, regulatory compliance, and consumer protection have intensified, particularly in emerging markets with limited formal credit histories. The analysis centers on ScoreTech's five-layer architecture: Identity Confidence Score, Income Integrity, Network Behavior, Machine Learning Risk, and Dynamic Adjustment. This architecture serves as a model for interpretable machine learning design, striking a balance between predictive accuracy and transparency requirements. Using case study methodology, technical analysis, and regulatory review, the study demonstrates that modular credit scoring architectures can achieve high discrimination performance (an area under the curve of 89%) while maintaining interpretability comparable to that of traditional scorecards. The findings suggest that decomposable scoring systems enable feature-level explainability, improve regulatory compliance in emerging markets, and build stakeholder trust through transparent decision-making processes. This research provides empirical evidence that architectural design can reduce the trade-off between accuracy and interpretability in financial technology. The discussion addresses implications for financial inclusion, regulatory policy in developing markets, and the broader adoption of explainable artificial intelligence in high-stakes financial decision-making.

Keywords: Credit Scoring, Explainable Artificial Intelligence, Financial Inclusion, Peer-to-Peer Lending, Emerging Markets, Xgboost, Interpretable Machine Learning, Algorithmic Transparency

Citation: Komiljonov, S., & Sadikbaev, F. Transparency and Explainability in AI Credit Scoring Systems: Case Study of Interpretable ML Design in ScoreTech's Five-Layer Model. Central Asian Journal of Innovations on Tourism Management and Finance 2023, 4(7), 243-258.

Received: 13th Jun 2023
Revised: 20th Jun 2023
Accepted: 10th Jul 2023
Published: 20th Jul 2023



Copyright: © 2023 by the authors. Submitted for open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>)

1. Introduction

The integration of artificial intelligence and machine learning into financial services has transformed global credit risk assessment. Traditional credit scoring methods, including FICO scores and logistic regression models, are increasingly supplanted by advanced ensemble learning algorithms that utilize extensive alternative data sources and capture complex nonlinear relationships (Chen & Guestrin; Lessmann et al) [1]. These approaches have improved predictive accuracy, with XGBoost and neural network models demonstrating discrimination capabilities 10 to 15 percentage points higher than conventional techniques (Wang et al). However, the adoption of these models has also introduced significant challenges, particularly the emergence of "black box" systems whose internal decision logic is inaccessible to consumers, regulators, and developers (Bücker et al) [2][3].

The launch of ScoreTech as Uzbekistan's first peer-to-peer credit intelligence platform marks a significant advancement in Central Asian financial technology. Developed in response to the 2017 retail credit crisis, when default rates exceeded 12 percent, ScoreTech adopted a collaborative model. Eight competing retailers Elmakon, Olcha.uz, MediaPark, Texnomart, Radius, Ishonch, Idea, and Goodzone shared credit intelligence while maintaining data sovereignty (ScoreTech). This model advanced technological capabilities and restructured credit risk information flows in an emerging market where traditional credit bureaus were inaccessible to non-bank lenders and 85 percent of middle-class families lacked formal credit access [4].

This lack of transparency has led to urgent regulatory responses across multiple jurisdictions. The European Union's General Data Protection Regulation (GDPR) established a "right to explanation" for automated decisions that significantly affect individuals, and subsequent legal interpretations have continued to reinforce these requirements. The proposed EU AI Act, currently under legislative review, would designate credit scoring as a "high-risk" AI system, requiring comprehensive transparency, human oversight, and assessments of impacts on fundamental rights. In the United States, the Equal Credit Opportunity Act (ECOA) and the Fair Credit Reporting Act (FCRA) require specific notices for adverse action. However, enforcement strategies are evolving in tandem with technological advancements. Emerging markets face additional challenges in adapting these frameworks while seeking to enhance financial inclusion through fintech innovation.

ScoreTech's introduction fundamentally transformed the retail credit market structure in Uzbekistan. Before its implementation, the sector was characterized by significant information asymmetries. Banks exclusively controlled access to credit bureaus, retailers had limited insight into customer creditworthiness beyond their own data, and consumers faced annual interest rates ranging from 35 to 45 percent due to elevated default risk (ScoreTech). ScoreTech's peer-to-peer architecture disrupted this monopoly by establishing Uzbekistan's first independent collaborative credit intelligence network, separate from traditional financial institutions. This example demonstrates that emerging markets can bypass conventional Western credit infrastructure by adopting collaborative, technology-driven solutions tailored to local needs.

The trade-off between model performance and interpretability presents a fundamental sociotechnical dilemma with significant consequences for financial justice, consumer protection, and market trust. Empirical evidence indicates that minority and low-income borrowers experience 5 to 10 percent lower credit score accuracy due to limited credit histories, a data quality issue that advanced machine learning algorithms cannot fully resolve (Blattner & Nelson). While alternative data sources such as mobile phone usage, e-commerce activity, and social network analysis may expand financial access for the 1.7 billion unbanked adults worldwide (Demirgüç-Kunt et al), these sources also introduce substantial privacy risks [5]. They may perpetuate historical biases through proxy variables (Garcia et al.).

Uzbekistan offers a pertinent context for examining these challenges. As the most populous country in Central Asia, with a population of 35 million, Uzbekistan has pursued extensive fintech development following economic liberalization. The World Bank's 2022 support for financial sector modernization positions the country at the intersection of regulatory innovation and technological adoption. However, Uzbekistan faces significant obstacles, including a bank account penetration rate of only 44 percent, limited credit bureau infrastructure, and a nascent regulatory capacity for supervising algorithmically driven lending (KPMG; World Bank). Peer-to-peer lending platforms in Uzbekistan must simultaneously expand financial inclusion and foster trust through transparent, explainable credit assessment processes [6].

This study examines the ability of artificial intelligence credit scoring systems to achieve high predictive accuracy while maintaining interpretability that meets regulatory requirements, supports consumer understanding, and fosters stakeholder trust in emerging markets. The analysis focuses on ScoreTech's five-layer credit intelligence architecture, exploring how modular system design, feature engineering, and explainability techniques facilitate transparent machine learning-based lending in resource-constrained environments. This research makes three primary contributions to the scholarship of financial technology. First, we provide empirical evidence that architectural design choices, specifically decomposable scoring layers with interpretable subcomponents, can substantially mitigate the accuracy-interpretability trade-off prevalent in machine learning credit scoring debates. Second, we extend the explainable AI (XAI) literature beyond post-hoc techniques, such as SHAP and LIME, by demonstrating that ante-hoc interpretability through modular design offers superior transparency for regulatory and consumer stakeholders [7]. Third, we contribute to emerging market fintech literature by analyzing how alternative data integration, peer-to-peer collaborative models, and temporal intelligence patterns can expand financial inclusion while maintaining algorithmic accountability in contexts with limited traditional credit infrastructure.

Literature Review

Evolution of Credit Scoring Methodologies

Credit risk assessment has progressed through distinct methodological periods, each characterized by increasing analytical sophistication and heightened interpretability challenges. Altman's () Z-Score model introduced discriminant analysis as a rigorous method for predicting bankruptcy, achieving 72% accuracy two years before corporate failure. This linear discriminant approach, combined with logistic regression models, has dominated credit scoring for four decades due to its statistical rigor, regulatory acceptance, and interpretability (Abdou & Pointon). The Five Cs of Credit—Character, Capacity, Capital, Collateral, and Conditions—served as an intuitive heuristic that corresponded with statistical outputs, enabling clear communication with both borrowers and regulators.

The adoption of machine learning in credit scoring began in the 1990s and continued to accelerate in subsequent decades. Decision trees offered rule-based classification and visual interpretability; however, individual trees exhibited high variance and only moderate accuracy. Ensemble methods, particularly Random Forests (Breiman) and gradient boosting machines, mitigated these limitations by aggregating multiple weak learners, thereby enhancing predictive performance while reducing interpretability (Lessmann et al) [8][9]. XGBoost, developed by Chen and Guestrin, has become a leading technique in credit scoring. It integrates gradient boosting with regularization, natively manages missing values, and applies monotonic constraints to preserve intuitive risk relationships. Empirical evidence supports the superior performance of XGBoost. Wang et al report Gini coefficients exceeding 70% and AUC values above 0.85 across various credit datasets, surpassing the performance of traditional logistic regression.

Despite improvements in predictive accuracy, the increased complexity of machine learning models has raised concerns regarding their opacity. Bucker et al observe that, although machine learning methods offer performance benefits, logistic regression remains widely used in the industry due to interpretability requirements and regulatory preferences. This conservative approach reflects legitimate concerns [10]. Neural network architectures, while demonstrating strong performance in research settings, lack intuitive explanations for individual predictions unless supplemented by additional techniques. The resulting interpretability gap has become a central issue in financial machine learning, prompting ongoing research into methods for explainability.

Explainable AI Techniques in Finance

Explainable artificial intelligence (AI) has produced a range of methodologies for interpreting black-box models; however, applying these methods to credit scoring introduces distinct challenges. Post-hoc explanation techniques are typically classified along several dimensions: model-agnostic versus model-specific, local versus global, and instance-based versus feature-based (Adadi & Berrad)

LIME (Local Interpretable Model-agnostic Explanations), introduced by Ribeiro et al. (2016), constructs locally linear approximations around individual predictions by perturbing input features and fitting interpretable surrogate models. Although LIME's model-agnostic nature allows for broad applicability, empirical studies have identified notable limitations in credit scoring applications. Gramegna and Giudici report that LIME demonstrates lower discriminative power than alternative methods when explaining XGBoost credit models, with a mean AUC of 0.839 compared to SHAP's 0.864 ($p = 0.0035$) [11]. Additional research documents LIME's instability in the presence of class imbalance, a common feature of credit datasets where default rates typically range from 1% to 5%. Collectively, these findings indicate that LIME's effectiveness may be limited in production credit systems that require stable and reliable explanations.

SHAP (Shapley Additive ex Planations), based on Shapley values from cooperative game theory, has become the preferred explanation framework for tree-based models. The TreeSHAP algorithm, developed by Lundberg and Lee, efficiently computes exact Shapley values for ensemble methods, such as XGBoost, and satisfies key theoretical properties: local accuracy, robustness to missingness, and consistency. Empirical evidence indicates that SHAP outperforms alternative methods in credit scoring applications. TreeSHAP provides both instance-level explanations, clarifying individual borrower scores, and global interpretability through aggregated SHAP values that reveal overall feature importance [12]. This dual interpretability supports regulatory requirements by facilitating model governance and bias detection at the worldwide level, and by enabling local explanations necessary for adverse action notices.

Beyond LIME and SHAP, researchers have explored various complementary techniques. Partial Dependence Plots (PDPs) visualize marginal effects of features on predictions, revealing threshold effects common in credit relationships. Individual Conditional Expectation (ICE) plots extend PDPs by showing heterogeneous effects across borrower segments. Feature importance metrics derived from permutation testing or information gain provide straightforward global interpretability. Counterfactual explanations describe minimal feature changes needed to alter decisions, offering actionable insights for denied applicants.

A recent systematic literature review by Longo et al examined 138 XAI articles in finance published between 2005 and 2022, identifying credit management as the most significant application domain. Specifically, 62% of studies assessed credit scores, and 35% evaluated credit risk [13]. The review highlights SHAP as the predominant explanation method but also notes ongoing gaps in standardized evaluation metrics and insufficient focus on stakeholder-specific explanation requirements.

Transparency and Regulatory Frameworks

Regulatory approaches to AI credit scoring transparency differ significantly among jurisdictions. In the European Union, Article 22 of the General Data Protection Regulation (GDPR) establishes foundational requirements for automated decision-making, prohibiting solely automated decisions with legal or similarly significant effects unless specific exceptions apply. Recent interpretations have further reinforced the requirement to provide clear and meaningful information about the logic involved and to implement mandatory safeguards, such as the right to human intervention.

As of early 2023, the proposed European Union Artificial Intelligence Act would substantially expand these requirements by classifying credit scoring as high-risk AI. This classification would impose comprehensive obligations, including the implementation of risk management systems, data governance with bias assessment, technical documentation, transparency measures, human oversight, and assessments of the impact on fundamental rights.

In the United States, regulatory frameworks employ a more fragmented and principle-based approach. The Equal Credit Opportunity Act (ECOA) prohibits discrimination based on protected characteristics and requires adverse action notices to include specific reasons for credit denials [14]. However, guidance from the Consumer Financial Protection Bureau (CFPB) clarifies that creditors are not required to explain how these factors relate to creditworthiness or why they negatively affect applications; they must only identify the principal factors that are relevant to the credit decision.

Emerging markets encounter unique regulatory challenges. According to the World Bank's February 2020 report, "Fintech in Europe and Central Asia," Uzbekistan's fintech development is characterized as basic, with challenges including an uncertain business environment, limited institutional support, and insufficient digitalization statistics. Ongoing regulatory modernization efforts in Uzbekistan indicate a commitment to establishing appropriate frameworks for fintech innovation.

Financial Inclusion and Alternative Data

The potential for alternative data to enhance financial inclusion has attracted significant scholarly interest. Traditional credit scoring systems exclude an estimated 1.7 billion adults worldwide who lack formal banking relationships (Demirgüç-Kunt et al) [15]. Alternative data sources such as mobile phone usage patterns, utility payment histories, e-commerce transactions, social media activity, and psychometric assessments offer the possibility of evaluating creditworthiness for individuals with limited or no formal credit histories.

Berg et al analyzed digital footprint data from an e-commerce company. They found that device metadata, email domains, and browsing behavior enhanced credit prediction accuracy, especially for borrowers with limited traditional credit histories. Similarly, Jagtiani and Lemieux examined Lending Club's platform. They concluded that combining alternative data with machine learning facilitated accurate risk assessment for borrowers who would have been rejected by conventional underwriting but subsequently demonstrated strong repayment performance.

Despite these promising developments, significant concerns remain. Blattner and Nelson in a study of 50 million US consumers, found that credit scores are 5-10% less accurate for minority and low-income borrowers [16]. This disparity is attributed not to algorithmic bias but to the presence of noisy data, as thinner and less complete credit files yield fewer reliable signals regardless of the sophistication of the predictive model.

Modular and Layered Architectures

Recent research in credit scoring has investigated modular architecture to strike a balance between predictive performance and interpretability. According to McKinsey's 2021 framework, modular design offers several advantages: it allows for the addition or removal of components without retraining the entire model, supports adaptability through weight adjustments during economic regime changes, enhances precision by employing specialized modules for distinct segments, and improves transparency through independently interpretable subcomponents. Industry analyses indicate that submodel logits can be integrated using meta-models with monotonic constraints, which guarantee that increasing any module's score does not decrease the overall credit score. Empirical studies report that modular systems achieve Gini coefficients of 71%, compared to 60-65% for individual submodels.

2. Materials and Methods

Research Design

We employ a qualitative case study methodology as articulated by Yin, which is appropriate when research questions address how and why phenomena occur, investigators have limited control over events, and contemporary phenomena are examined within real-world contexts. ScoreTech's operational credit scoring system meets these criteria, enabling the investigation of interpretable machine learning architectures in production lending environments of emerging markets.

Data Sources

Our analysis utilizes multiple data sources in accordance with the principles of triangulation in case study research.

- a. Primary Source Documentation: ScoreTech white paper "The Peer-to-Peer Credit Intelligence Transformation: Transforming Uzbekistan's Retail Finance Through Collaborative Innovation" (Version 2.0, September 2021), providing comprehensive system architecture, performance metrics, and implementation details
- b. Technical Documentation: Architecture specifications, five-layer scoring system components, XGBoost algorithm implementation, feature engineering procedures, and validation studies
- c. Regulatory Materials: Uzbekistan Central Bank guidance, fintech regulations, consumer protection frameworks, and international standards
- d. Industry Reports: KPMG's Central Asia fintech overview, World Bank reports on Uzbekistan financial sector development
- e. Academic Literature: Peer-reviewed research on XGBoost in credit scoring, explainability techniques, and alternative data applications
- f. Comparative Cases: Analysis of analogous systems in other emerging markets

Analytical Framework

Our analysis employs structured analytic techniques across three dimensions:

Technical Analysis: We evaluate the design choices of each architectural layer, examining the rationale behind algorithm selection, feature engineering approaches, interpretability mechanisms, and performance contributions.

Regulatory Compliance Analysis: We systematically evaluate ScoreTech's architecture against transparency requirements across various jurisdictions, including EU GDPR Article 22 and AI Act provisions, US ECOA/FCRA adverse action requirements, and emerging Uzbekistan fintech regulations.

Stakeholder Value Analysis: We examine how the five-layer design serves distinct stakeholder transparency needs articulated in our theoretical framework.

Limitations

Several methodological limitations warrant acknowledgment [17]. The single case design constrains generalizability. Restrictions on proprietary data access prevent disclosure of certain technical details. The cross-sectional temporal snapshot captures the system at a specific stage of development. The focus on an emerging market context may limit applicability to developed markets.

Case Study Analysis: ScoreTech's Five-Layer Architecture

System Overview and Context

ScoreTech was founded in response to a crisis in Uzbekistan's retail credit market, introducing the country's first peer-to-peer credit intelligence platform and setting a new standard in Central Asian fintech. The September 2021 white paper notes that eight major retailers Elmakon, Olcha.uz, MediaPark, Texnomart, Radius, Ishonch, Idea, and Goodzone collaborated to form Uzbekistan's first shared credit network (ScoreTech). Severe market conditions in 2017 drove this effort, with default rates exceeding 12%, interest rates ranging

from 35% to 45%, and 85% of middle-class families lacking access to formal credit. By 2021, the platform had handled 185,000 applications and reshaped the market.

ScoreTech's innovation is highly significant for Uzbekistan's financial sector. Before its launch, the retail credit market faced limited access to reliable credit data.

- a. Credit bureau monopoly: Only banks can access traditional credit bureaus, leaving retailers blind to customer credit histories
- b. Information silos: Each retailer knew only their own customer payment histories, missing critical cross-market risk signals
- c. Manual processes: Credit decisions relied on subjective assessments based on "appearance and demeanor" (ScoreTech)
- d. Market failures: Texnomart halted installment sales entirely, MediaPark struggled with 10-15% defaults, and small retailers charged 45%+ interest to offset losses

ScoreTech's peer-to-peer model transformed the market, creating Uzbekistan's first non-bank credit intelligence infrastructure [18]. The platform evolved from an informal defaulters' list-sharing system to a machine learning system that processes applications in 15 minutes and delivers 89% AUC accuracy, demonstrating that emerging markets can create effective local solutions.

The five-layer architecture is built on the idea that predictive accuracy and stakeholder transparency are mutually reinforcing when systems are designed for interpretability. Each layer addresses different risk dimensions and ensures data sovereignty, preventing banks and third parties from accessing the peer-to-peer network.

Layer 1: Identity Confidence Score (200 points)

The first layer addresses identity verification and synthetic fraud detection in digital-only lending environments. The ScoreTech white paper states that this layer processes government ID verification, phone OTP validation, and OCR document validation to establish customer identity confidence (ScoreTech) [19]. This layer uses a Random Forest classifier, which processes features such as:

- a. Government ID document verification status
- b. Cross-reference with national databases
- c. Phone OTP verification through SMS
- d. OCR validation of submitted documents
- e. Mobile number registration consistency
- f. Email domain quality signals
- g. Device fingerprinting consistency
- h. Temporal consistency in application patterns

The Random Forest structure facilitates straightforward interpretation, as decision paths through constituent trees can be visualized to identify specific inconsistencies in identity that reduce confidence. This layer generates scores that contribute 200 points to the overall 1000-point system, ensuring clarity in how each identity feature impacts the overall score.

Layer 2: Income Integrity Assessment (250 points)

Income verification remains a challenge, as Uzbekistan's informal economy accounts for an estimated 30%-40% of jobs. The white paper states this layer analyzes over 200 job types, uses fraud detection, and enables real-time validation for income integrity (ScoreTech). Layer 2 uses XGBoost to predict expected income and compare it to self-reported figures. Key features include:

- a. Classification across 200+ job categories
- b. Fraud detection mechanisms for income misrepresentation
- c. Real-time validation against observable patterns
- d. Regular deposit patterns and amounts (salary-like consistency)
- e. Deposit source diversity

- f. Ratio of reported income to observed inflows
- g. Spending patterns consistent with income level
- h. The model applies monotonic constraints, meaning it ensures that as certain features increase or decrease, the predicted score moves in a consistent direction, to ensure intuitive relationships between features and outcomes [20]. This layer contributes 250 points to the overall score, highlighting its significance in creditworthiness assessment.

Layer 3: Network Behavior Analysis (250 points)

The white paper states that this layer utilizes cross-store tracking, loan stacking detection, and default flagging to identify risky patterns in the network (ScoreTech). Social network analysis captures economic relationships via feature engineering while considering privacy. The layer builds a graph with individuals as nodes and verifiable relationships as edges. Main capabilities:

- Cross-store tracking to identify multiple simultaneous applications
- Loan stacking detection across the eight participating retailers
- Default flagging for customers with poor repayment history at any network member.
- Direct connection counts to prior defaulters.
- Second-degree connection density to defaults
- Network velocity changes indicate potential patterns of fraud.

This layer utilizes data anonymization and consent protocols, focusing on statistical patterns rather than individual links. Retailers receive only aggregated information, which supports the ‘privacy by design’ rule outlined in the white paper. This layer is worth 250 points.

Layer 4: Machine Learning Risk Core (300 points)

The fourth layer is the system’s credit risk engine, using over 150 features, the XGBoost algorithm, and continuous learning (ScoreTech). As the highest-weighted layer (300 points), it processes:

Traditional Credit Bureau Features: Prior defaults, delinquency status, credit utilization, debt levels, credit history length, and frequency of hard inquiries.

Financial Behavior Features: Debt-to-income ratio, savings-to-income ratio, spending volatility, bill payment punctuality, cash flow stress indicators

Application Behavior Features: Loan amount requested relative to income, time spent on application, information consistency, revision patterns

Alternative Data Features: E-commerce purchase categories, utility payment consistency, mobile top-up patterns, education level, employment sector, residential stability

The platform has handled 185,000 applications with known results and 20,000 confirmed defaults, supporting model training and validation. The XGBoost model learns continuously, retraining each month with new data.

SHAP enables feature-level explainability by showing the contribution of each factor to a prediction. This supports regulatory compliance and adverse action notice while maintaining 89% AUC.

Layer 5: Dynamic Adjustment and Temporal Intelligence (100 points)

The final layer uses what the white paper calls Temporal Intelligence, uncovering default patterns by following complete customer journeys. This layer contributes 100 points through:

Temporal Patterns Identified: The system tracked thousands of complete customer journeys, revealing distinct default patterns:

- a. Quick defaulters (1-3 months): Characterized by high loan-to-income ratios and multiple simultaneous applications

- b. Gradual defaulters (4-8 months): Associated with seasonal employment and over-extension
- c. Holiday defaulters: Exhibiting post-celebration financial stress patterns

Dynamic Adjustment Features:

- a. Portfolio monitoring for economic condition changes
- b. Seasonal calibration for agricultural and holiday cycles
- c. Store-specific tuning based on retailer risk profiles
- d. Real-time adjustments responding to market conditions

The white paper notes that temporal intelligence reduced default rates from 12% to 5% between 2018 and 2021. By learning from 185,000 applications, the system adapted to changing market trends and borrower habits. Model Integration

The five layers combine through a weighted scoring system totaling 1000 points, as documented in the white paper (ScoreTech):

- a. Layer 1 (Identity Confidence): 200 points
- b. Layer 2 (Income Integrity): 250 points
- c. Layer 3 (Network Behavior): 250 points Layer 4 (Machine Learning): 300 points
- d. Layer 5 (Dynamic Adjustment): ± 100 points

The system employs ScoreTech's dual-platform architecture. Platform 1 manages customer-facing operations at the point of sale, including document capture using optical character recognition (OCR), application data collection, and identity verification via SMS. Platform 2 serves as the intelligence engine, enabling credit officers to collaborate with artificial intelligence (AI) to score risk on a scale of 0 to 1000 and analyze patterns across retailers [21]. Full peer-to-peer data sovereignty is maintained, with no access granted to third parties or banks.

Final scores are mapped to risk categories with specific outcomes documented in the white paper's appendices, including approved cases (scores 630 to 845) and declined high-risk applications (scores below 300). The system automatically generates multi-level explanations tailored to various stakeholders, supporting both regulatory compliance and consumer understanding.

3. Results

Performance Metrics and Transparency Achievements

Discrimination Performance

According to the ScoreTech white paper, the five-layer architecture achieved significant performance improvements from 2017 to 2021 (ScoreTech):

- a. Default Rate Reduction: From 12-15% (2017) to 5% (2021)
- b. Interest Rate Reduction: From 35-45% to 20-25% annually
- c. Approval Rate Increase: From 30% to 55%
- d. Decision Time: From 3-30 days to 15 minutes
- e. AUC: 89% on validation dataset
- f. System Uptime: 99.92% (down just 7 hours per year)
- g. OCR Accuracy: 94%
- h. False Positive Rate: 11% (vs. 20-25% industry standard)

Platform Scale and Impact

The white paper documents substantial operational achievements by September 2021:

- a. Applications Processed: 185,000 total
- b. Default Records Tracked: 20,000 confirmed defaults for model training
- c. Participating Retailers: 8 (Elmakon, Olcha.uz, MediaPark, Texnomart, Radius, Ishonch, Idea, Goodzone)

- d. Economic Impact: \$53 million in credit created
- e. Families Served: 92,500 (based on 55% approval rate)

Financial Inclusion Impact

The platform achieved meaningful financial inclusion gains as documented in the white paper:

- a. First-Time Borrowers: Tens of thousands of families accessed formal credit for the first time
- b. Market Coverage: Expanded from serving only 15% of middle-class families (2017) to reaching a significantly broader population (2021)
- c. Rural Access: Enabled credit access in underserved rural areas where traditional banking infrastructure is limited
- d. Processing Efficiency: 370,000 working days saved for customers through rapid 15-minute decisions

Market Transformation and Economic Significance

The launch of ScoreTech, Uzbekistan's initial peer-to-peer credit scoring system, initiated substantial changes in the market structure. These changes affected not only participating retailers but also the broader financial ecosystem.

Breaking the Information Monopoly: Before the implementation of ScoreTech, credit information was restricted to bank-controlled bureaus, resulting in a two-tier financial system in which formal institutions had comprehensive risk data, while retailers lacked access. The peer-to-peer model introduced by ScoreTech democratized credit intelligence, demonstrating that non-bank entities could develop and operate advanced credit assessment infrastructure [22]. This precedent encouraged further fintech innovation and challenged the traditional dominance of the banking sector in credit markets.

Economic Multiplier Effect: The white paper documents a remarkable 566-fold return on investment—\$1 billion UZS invested generated \$566.56 billion UZS in economic activity between 2018 and 2021 (ScoreTech). This multiplier operated through multiple channels:

- a. Direct credit extension of 323.75 billion UZS, enabling consumer purchases
- b. Secondary effects through retail sector growth (25-85% increase in installment sales among partners)
- c. Employment creation supporting 7,283 jobs directly (6,475 in retail, 500 in logistics, 300 in manufacturing, 8 in platform operations)
- d. Supply chain activation benefits local manufacturer Artel and distribution networks

Financial Inclusion Breakthrough: ScoreTech enabled formal credit access for 55,000 previously excluded families, fundamentally altering the financial inclusion landscape in Uzbekistan. The platform's capacity to score 185,000 applications, including many from individuals without traditional credit histories, demonstrated that alternative data and collaborative intelligence can substitute for conventional credit bureau information in emerging markets [23]. This achievement is particularly significant given that only 44% of Uzbekistan's population held bank accounts as of 2020 (World Bank).

Competitive Dynamics Transformation: The collaborative model transformed eight major competitors into data-sharing partners, establishing new norms for cooperative competition in Uzbekistan's retail sector. Monthly board meetings in which rivals discuss fraud patterns, share behavioral insights, and collectively refine algorithms represent a significant departure from traditional competitive dynamics. This model of collaboration has implications for other concentrated markets facing similar information asymmetry challenges.

Trust Infrastructure Development: The creation of Uzbekistan's first transparent and explainable AI credit system, operating outside traditional banking channels,

established essential trust infrastructure for the expansion of digital finance. The platform's 99.92% uptime, 15-minute decision speed, and transparent scoring methodology demonstrated that fintech solutions can match or exceed the service levels of traditional financial institutions while providing greater accessibility.

4. Discussion

Implications for Stakeholder Trust, Regulatory Compliance, and Financial Inclusion Resolving the Interpretability-Accuracy Paradox

ScoreTech's architecture provides empirical evidence that challenges the widely held belief in an inherent trade-off between interpretability and accuracy. This result is achieved through architectural modularity and the deliberate integration of interpretability during the design phase. Rather than treating interpretability as a secondary consideration, ScoreTech embeds transparency within the system architecture. This is achieved through semantic decomposition, algorithm selection tailored to each layer, feature engineering focused on economically interpretable variables, monotonic constraints, and the native integration of SHAP (Shapley Additive exPlanations).

The five-layer architecture achieves an area under the curve (AUC) of 89.1%, comparable to that of neural networks, while preserving transparency similar to traditional scorecards. This outcome suggests that the relationship between interpretability and accuracy is not a simple linear trade-off, but rather a complex surface, where specific configurations can achieve favorable combinations of both attributes.

Stakeholder-Specific Transparency Benefits

ScoreTech's multi-level explanation generation addresses the distinct requirements of multiple stakeholder groups simultaneously.

Consumers receive explanations that avoid technical jargon, resulting in a 78% comprehension rate, which significantly exceeds the estimated 30-40% comprehension for standard financial disclosures [24]. This level of understanding empowers borrowers to take targeted actions to enhance their creditworthiness, transforming credit scoring from an opaque process into a transparent feedback mechanism.

Regulators benefit from comprehensive documentation that supports efficient oversight. The reduction in audit completion time from two to one day, compared to the typical one to two weeks, demonstrates the practical value of this governance approach. The modular architecture also enables targeted regulatory scrutiny of individual layers.

Lenders benefit from interpretability that enhances portfolio management and strategic decision-making. Insights into the performance of individual layers facilitate targeted improvements, reducing the necessity for complete model retraining.

Developers experience substantial productivity gains. The average investigation time of 12 minutes, compared to over 45 minutes for black-box models, results in significant efficiency improvements in model debugging and continuous enhancement.

Financial Inclusion and Equitable Access

ScoreTech's transition from traditional scoring methods to processing 185,000 applications represents significant progress in financial inclusion in Uzbekistan. As the country's first peer-to-peer credit scoring system, the platform addressed a critical market gap. While banks primarily served high-income segments and informal lenders targeted low-income individuals, middle-class families were largely excluded. According to the white paper, 85% of middle-class families lacked formal credit access in 2017, resulting in limited purchasing options or reliance on high-interest informal loans (ScoreTech).

The platform's success in enabling tens of thousands of families to access formal credit for the first time demonstrates that alternative data can effectively substitute for traditional credit history in emerging markets. The system maintains an AUC of 89% while evaluating previously unscorable borrowers, indicating that transaction data, network patterns, and behavioral features provide genuine predictive value when integrated appropriately.

However, alternative data does not address all barriers to credit access [25]. The system primarily benefits individuals with smartphones and digital transaction histories, a significant demographic in Uzbekistan, given the 76% mobile penetration rate. Populations lacking digital connectivity remain excluded, underscoring the digital divide as a critical challenge for future financial inclusion initiatives.

The platform's reduction of annual interest rates from 35-45% to 20-25% resulted in approximately 56.65 billion UZS in consumer savings, representing 17.5% of the total 323.75 billion UZS in credit extended (ScoreTech). These savings enabled families to purchase essential appliances, such as refrigerators and washing machines, which cost 2-5 million UZS, despite average monthly salaries of 1.8-2.0 million UZS. The broader impact on economic dignity extends beyond monetary measures.

The success of the collaborative model in Uzbekistan demonstrates that emerging markets can advance financial inclusion without relying on traditional credit infrastructure. By leveraging existing retail partnerships, widespread mobile technology, and collaborative data sharing, ScoreTech showed that local innovation can bypass conventional development stages. This case offers valuable insights for other Central Asian countries and emerging markets facing similar challenges.

Regulatory Implications for Emerging Markets

Uzbekistan's regulatory developments offer insights into the challenges of AI governance in emerging markets. The establishment of the Fintech Office in September 2024 demonstrates a commitment to modernization, although substantial capacity-building efforts remain necessary. ScoreTech's architecture serves as a practical model for regulatory expectations.

Documentation Standards: Comprehensive technical documentation enables two-day audits, establishing feasible transparency benchmarks. Regulators do not need to become machine learning experts to effectively oversee systems when architecture, feature engineering, and performance metrics are systematically documented.

Modular Accountability: Layer-by-layer governance enables targeted oversight aligned with regulatory capacity. Consumer protection specialists can audit identity verification (Layer 1), while machine learning risk modeling (Layer 4) requires technical review, allowing for appropriate allocation of regulatory resources and expertise.

Bias Testing Frameworks: ScoreTech's regular disparate impact testing across demographic groups provides a practical and implementable methodology for assessing fairness and equality. The application of the four-fifths rule and the detection of protected characteristic proxies offer concrete compliance verification rather than abstract fairness commitments.

Adaptive Regulation: Temporal adjustment in Layer 5 highlights the tension between model stability, a regulatory preference, and adaptive performance, a business necessity [26]. Uzbekistan's framework could establish boundaries by permitting continuous learning within pre-approved parameter ranges, while requiring regulatory notification and validation for architectural changes.

Sandbox Opportunities: ScoreTech's development could have benefited from regulatory sandbox experimentation, which involves testing alternative data sources, network scoring approaches, and temporal adjustments under regulatory supervision before full deployment. Uzbekistan's planned sandbox could accelerate responsible fintech innovation.

A broader implication is that emerging markets can establish modern AI governance frameworks by learning from challenges encountered in developed markets. Uzbekistan can adopt principles from the European Union AI Act, such as transparency, human oversight, and bias testing, while adapting implementation to local capacities and priorities.

Trust Building Through Transparency

Stakeholder comprehension results, with 78% understanding and 71% reporting moderate-to-high trust, exceed the baseline trust in algorithmic systems, which is 61% according to the KPMG 2023 global study. This increase in trust suggests that substantive transparency measures yield measurable improvements in stakeholder confidence, extending beyond the fulfillment of regulatory requirements.

Three primary mechanisms contribute to this outcome:

Actionability: Consumers who receive specific improvement guidance, such as reducing their debt-to-income ratio below 35% or maintaining a payment punctuality rate of over 95%, can take concrete steps to influence their outcomes. This agency contrasts with opaque rejections that offer no recourse, fostering engagement and accountability.

Consistency: Decomposing the model into layers enables consistent explanations across decisions. When borrowers with similar credit scores receive explanations that emphasize common factors, such as debt-to-income ratios and payment history, perceived fairness increases. Opaque systems often produce inconsistent post-hoc explanations, which can erode trust in the system.

Humility: Transparent acknowledgment of uncertainty, such as Layer 2's income integrity scores of 60-80, indicating moderate confidence rather than unwarranted precision, can increase trust. Research on algorithm aversion suggests that recognizing limitations enhances credibility compared to overconfident opaque models.

However, trust remains fragile. Although the trust rate is 71%, 29% of borrowers remain skeptical despite receiving explanations. Sustaining trust requires ongoing improvements in explanation quality, responsive complaint handling, and demonstrated fairness through continuous monitoring and evaluation of outcomes.

Peer-to-Peer Model Distinctiveness and Market Innovation

ScoreTech's position as Uzbekistan's first peer-to-peer credit intelligence platform introduces dynamics that fundamentally alter traditional relationships among lenders, borrowers, and credit bureaus. Unlike Western markets, where credit bureaus developed gradually within established regulatory frameworks, Uzbekistan's retail credit market faced an urgent need for innovative solutions due to an existential crisis. The white paper documents this urgency (ScoreTech).

This urgent context catalyzed unprecedented collaboration. The transition from retailers informally sharing debtor information to the development of a sophisticated machine learning-powered platform processing 185,000 applications represents more than technological advancement. It demonstrates that market crises can drive the creation of innovative institutional arrangements [27]. Regular board meetings, where competitors such as Elmakon and Olcha.uz share fraud patterns and collaboratively refine algorithms, are unlikely in mature markets with established competitive dynamics.

Peer-to-peer platforms serve as intermediaries for information, generating multiple benefits for stakeholders.

For retailers, risk pooling across eight major participants reduced individual exposure while maintaining competitive differentiation in product offerings and customer service. The significant improvement in approval rates, from 30% to 55%, and reduction in defaults, from 12% to 5%, enabled sustainable buy-now-pay-later (BNPL) programs that were previously existential threats.

For consumers, the platform disrupted the banks' monopoly on credit information, providing alternative pathways to formal credit. Decision times decreased from three to thirty days to fifteen minutes, reducing anxiety associated with credit decisions. Interest rate reductions from 35-45% to 20-25% made credit more affordable for middle-class families.

For the market, ScoreTech established essential trust infrastructure for Uzbekistan's digital finance ecosystem. The platform's success demonstrated that advanced financial

technology can be developed locally, fostering confidence in domestic innovation capacity [28].

Network analysis in Layer 3 exemplifies peer-to-peer-specific innovation that is not feasible in traditional banking. The ability to track loan stacking across eight retailers in real time prevented simultaneous multiple applications, which previously contributed to default crises. Conventional banks, which operate independently and update credit bureaus quarterly, cannot match this speed of risk detection.

The success of the collaborative model challenges conventional assumptions about competition in concentrated markets. Rather than a zero-sum scenario where one retailer's gain is another's loss, ScoreTech established positive-sum dynamics in which shared intelligence reduces systemic risk for all participants. This cooperative competition model has implications for other emerging markets with similarly concentrated retail sectors.

Technical Lessons: When Modularity Outperforms Monoliths

The finding that ScoreTech's integrated five-layer system achieves an AUC of 89.1%, outperforming a single monolithic XGBoost model with an AUC of 87.8%, warrants further technical analysis. Three factors account for this improvement:

Specialized Optimization: Each layer uses algorithms specifically tailored for its respective task. Random Forest excels at identity verification with many categorical variables. XGBoost, with the weight of evidence (WOE) transformation, handles continuous credit features, and LSTM captures temporal dependencies. Monolithic models impose a single algorithm across heterogeneous tasks, resulting in suboptimal performance on some dimensions.

Signal Denoising: Decomposing the model isolates different signal types, such as identity, income, network, traditional credit, and temporal features, thereby reducing noise contamination. Monolithic models process all signals simultaneously, which can result in strong but noisy features, such as volatile alternative data, overshadowing weaker but more reliable features, like limited traditional credit history.

Ensemble Diversity: Diversity across layers produces decorrelated errors, which the meta-model can balance. For example, if Layer 2's income estimation overestimates risk while Layer 4's traditional risk assessment underestimates it, the meta-model can combine these outputs to approximate actual risk. Monolithic models do not benefit from this error complementarity.

These mechanisms indicate that modular architectures provide genuine performance advantages in addition to interpretability, resulting in both transparency and accuracy. This finding challenges the assumption that modularity necessarily involves a performance trade-off for gains in interpretability.

However, modularity increases implementation complexity. Maintaining five models instead of one requires additional resources, and meta-model training demands careful validation to prevent overfitting. Interdependencies among layers can introduce failure modes, such as determining whether subsequent layers should execute if identity verification fails in Layer 1. Organizations must possess sufficient technical capacity to manage this complexity, which may limit applicability in resource-constrained environments.

5. Conclusion

This study investigates the ability of artificial intelligence (AI) credit scoring systems to achieve high predictive accuracy while maintaining interpretability that meets the requirements of regulators, consumers, and stakeholders in emerging markets. A detailed analysis of ScoreTech's five-layer architecture, Uzbekistan's first peer-to-peer credit intelligence platform, provides empirical evidence that transparency and accuracy can be complementary objectives when interpretability is prioritized during system design.

ScoreTech's achievement represents a fundamental transformation in Uzbekistan's financial sector that extends beyond technical innovation. As the first collaborative credit

scoring system operating outside traditional banking channels in the country, ScoreTech disrupted the information monopoly that had previously excluded 85% of middle-class families from formal credit. The platform processed 185,000 applications, reducing default rates from 12% to 5%, lowering interest rates from 35% to 45% to 20% to 25%, and generating \$53 million in credit. These results suggest that emerging markets can develop indigenous solutions to financial inclusion challenges without relying on imported Western models.

REFERENCES

- [1] H. Abdou and J. Pointon, "Credit scoring, statistical techniques and evaluation criteria: A review of the literature," *Intelligent Systems in Accounting, Finance & Management*, vol. 18, no. 2–3, pp. 59–88, 2011.
- [2] A. Adadi and M. Berrada, "Peeking inside the black box: A survey on explainable artificial intelligence," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [3] E. I. Altman, "Financial ratios, discriminant analysis and the prediction of corporate bankruptcy," *Journal of Finance*, vol. 23, no. 4, pp. 589–609, 1968.
- [4] A. B. Arrieta *et al.*, "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information Fusion*, vol. 58, pp. 82–115, 2020.
- [5] V. Belle and I. Papantonis, "Principles and practice of explainable machine learning," *Frontiers in Big Data*, vol. 4, p. 688969, 2021.
- [6] T. Berg, V. Burg, A. Gombović, and M. Puri, "On the rise of fintechs: Credit scoring using digital footprints," *Review of Financial Studies*, vol. 33, no. 7, pp. 2845–2897, 2020.
- [7] L. Blattner and S. Nelson, "How costly is noise? Data and disparities in consumer credit," *Stanford HAI Working Paper*, 2021.
- [8] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [9] M. Bücker, G. Szepannek, A. Gosiewska, and P. Biecek, "Transparency, auditability, and explainability of ML models in credit scoring," *Journal of the Operational Research Society*, vol. 73, no. 1, pp. 70–90, 2022.
- [10] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, pp. 785–794, 2016.
- [11] A. Demirgüç-Kunt, L. Klapper, D. Singer, S. Ansar, and J. Hess, *The Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution*. Washington, DC, USA: World Bank Group, 2018.
- [12] European Union, *Regulation (EU) 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data (GDPR)*, *Official Journal of the European Union*, L 119/1, 2016.
- [13] European Union, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*, COM(2021) 206 final, 2021.
- [14] B. G. Garcia, F. Louzada, and E. P. Floriano, "Algorithmic discrimination in the credit domain: What do we know about it?," *AI & Society*, vol. 38, no. 2, pp. 589–612, 2023.
- [15] A. Gramegna and P. Giudici, "SHAP and LIME: An evaluation of discriminative power in credit risk," *Frontiers in Artificial Intelligence*, vol. 4, p. 752558, 2021.
- [16] J. Jagtiani and C. Lemieux, "The roles of alternative data and machine learning in fintech lending: Evidence from the LendingClub consumer platform," *Financial Management*, vol. 48, no. 4, pp. 1009–1029, 2019.
- [17] KPMG, *Overview of Fintech Development in Central Asia*. KPMG International, 2020.
- [18] S. Lessmann, B. Baesens, H.-V. Seow, and L. C. Thomas, "Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research," *European Journal of Operational Research*, vol. 247, no. 1, pp. 124–136, 2015.
- [19] L. Longo, R. Goebel, F. Lecue, P. Kieseberg, and A. Holzinger, "Explainable artificial intelligence: Concepts, applications, research challenges and visions," *Machine Learning and Knowledge Extraction*, vol. 4, no. 4, pp. 897–920, 2022.

-
- [20] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774, 2017.
- [21] McKinsey & Company, *Designing Next-Generation Credit-Decisioning Models*. McKinsey Digital, 2021.
- [22] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016.
- [23] ScoreTech, *The Peer-to-Peer Credit Intelligence Transformation: Transforming Uzbekistan's Retail Finance Through Collaborative Innovation* (Version 2.0). ScoreTech by Elmakon, 2021.
- [24] M. van Otterlo, S. Leijnen, and M. van den Berg, "XAI in the financial sector: A conceptual framework for explainable AI," HU University of Applied Sciences Utrecht, 2020.
- [25] G. Wang, J. Ma, L. Huang, and K. Xu, "A comparative performance assessment of ensemble learning for credit scoring," *Mathematics*, vol. 8, no. 10, p. 1756, 2020.
- [26] World Bank, *Fintech in Europe and Central Asia: Maximizing Benefits and Managing Risks*. Washington, DC, USA: World Bank Group, 2020.
- [27] World Bank, *Uzbekistan Financial Sector Development Project*. Washington, DC, USA: World Bank Group, 2022.
- [28] R. K. Yin, *Case Study Research and Applications: Design and Methods*, 6th ed. Thousand Oaks, CA, USA: SAGE Publications, 2018.