



Explainable AI Frameworks for Regulatory-Compliant Buy-Now-Pay-Later Credit Risk Assessment in Real-Time Cloud Banking Architectures

Siva Prakash*

Bharathidasan University, India

* Corresponding Author Email: emailtosivaprakash@gmail.com - ORCID: 0009-0001-0080-7838

Article Info:

DOI: 10.22399/ijcesen.5107
Received : 02 November 2023
Accepted : 26 December 2023

Keywords

XAI,
BNPL,
GDPR,
Real Time,
Banking Architecture.

Abstract:

The rapid proliferation of Buy-Now-Pay-Later (BNPL) services has transformed digital lending ecosystems, necessitating robust, scalable, and transparent credit risk assessment frameworks. Traditional credit scoring mechanisms are insufficient for BNPL contexts characterized by real-time decisioning, thin credit files, and dynamic consumer behavior. This paper presents a comprehensive academic analysis of Explainable Artificial Intelligence (XAI) frameworks tailored for regulatory-compliant BNPL credit risk assessment within real-time cloud banking architectures. The study synthesizes advances in machine learning, interpretability techniques, regulatory mandates (e.g., GDPR), and cloud-native financial infrastructures. The paper proposes a layered architectural framework integrating explainability, fairness, and compliance into AI-driven credit decision systems. The discussion aligns technological developments with evolving financial regulations and highlights open research challenges.

1. Introduction

The emergence of Buy-Now-Pay-Later (BNPL) services has significantly altered consumer financing by enabling instant credit decisions at the point of sale. Unlike traditional lending, BNPL systems operate under stringent latency constraints and rely heavily on alternative data sources, including transactional behavior and digital footprints (Lupşa-Tătaru, 2023). BNPL systems differ fundamentally from conventional credit models, requiring real-time analytics and adaptive risk modeling, which has led to the increasing integration of AI-driven decision systems in fintech platforms.

The rapid adoption of BNPL has also introduced systemic risks, including borrower overextension and increased regulatory scrutiny. Empirical studies indicate that BNPL users often belong to subprime or thin-file segments, raising concerns about financial stability and consumer protection (Di Maggio et al., 2022).

Simultaneously, artificial intelligence (AI) has emerged as a dominant paradigm for credit risk assessment, offering improved predictive accuracy and automation (Lessmann et al., 2015; Khandani et al., 2010). However, the opacity of complex

models—particularly deep learning—poses challenges for transparency, fairness, and regulatory compliance. The need for explainability has therefore become central to financial AI systems, especially under frameworks such as the “right to explanation” mandated by data protection regulations (Goodman & Flaxman, 2017; Demajo et al., 2020). This paper addresses the intersection of XAI, BNPL credit risk, and cloud-native banking architectures, proposing a unified framework for regulatory-compliant real-time decision-making.

2. Literature Review

2.1 Evolution of AI in Credit Risk Assessment

Machine learning techniques have significantly improved credit risk prediction compared to traditional statistical models, marking a substantial shift in financial analytics between 2018 and 2023. Early empirical studies demonstrated that advanced algorithms such as random forests, support vector machines, and gradient boosting models outperform conventional approaches like logistic regression in predicting default risk, particularly in complex and high-dimensional datasets (Lessmann et al., 2015). These models are capable of capturing non-linear

relationships and intricate feature interactions that traditional methods often fail to identify.

More recent advancements have extended these capabilities through the integration of deep learning and reinforcement learning, which enable continuous learning and dynamic adaptation to evolving financial behaviors and market conditions (Heaton et al., 2017). In addition, hybrid AI models that incorporate alternative data sources—including transactional data, behavioral analytics, and digital footprints—have been shown to enhance predictive performance while simultaneously promoting financial inclusion (Berg et al., 2020).

2.2 Explainable AI in Financial Services

Explainable Artificial Intelligence (XAI) has emerged as a critical response to the “black-box” nature of modern machine learning systems, particularly in high-stakes domains such as financial services. As the complexity of predictive models increases, the need for interpretability and transparency becomes essential for both operational and regulatory purposes (Bussmann et al., 2020).

Techniques such as Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro et al., 2016) and Shapley Additive Explanations (SHAP) (Lundberg & Lee, 2017) are widely adopted to interpret model predictions. LIME provides localized explanations by approximating complex models with simpler, interpretable representations, whereas SHAP employs a game-theoretic approach to quantify the contribution of individual features to a model’s output.

Research demonstrates that these methods improve transparency without significantly compromising predictive accuracy and can bridge the gap between high-performance predictive models and regulatory requirements (Misheva et al., 2021). Bussmann et al. (2020) further emphasize that explainability is essential for fostering trust, ensuring auditability, and achieving regulatory acceptance in fintech risk management systems.

2.3 Regulatory Drivers for Explainability

The increasing adoption of AI in financial decision-making has been accompanied by the introduction of regulatory frameworks that mandate transparency and accountability. Regulations such as the General Data Protection Regulation (GDPR) and the Equal Credit Opportunity Act (ECOA) explicitly require that automated decision-making systems provide clear and understandable explanations for their outcomes (Goodman & Flaxman, 2017).

These regulations are designed to protect consumers from opaque and potentially biased decision processes, ensuring fairness and non-discrimination in credit allocation. Research highlights that compliance with these regulatory requirements necessitates the integration of explainability mechanisms directly into AI systems, rather than treating them as supplementary features (Demajo et al., 2020). Consequently, explainability has evolved from a technical enhancement to a legal obligation within modern banking systems.

2.4 Cloud-Native Banking Architectures

The advancement of cloud computing technologies has played a crucial role in enabling scalable and real-time financial services, particularly in the context of BNPL systems. Cloud-native banking architectures leverage microservices-based designs, distributed data pipelines, and event-driven processing frameworks to support high-throughput and low-latency operations (Dragoni et al., 2017).

These architectures allow financial institutions to process large volumes of transactional data in real time, facilitating instantaneous credit decisioning at the point of sale. Moreover, cloud platforms enable seamless deployment, monitoring, and updating of AI models, thereby enhancing system scalability, resilience, and operational efficiency (Zhang et al., 2018). By integrating AI-driven credit risk models within cloud-native infrastructures, financial institutions can achieve real-time decision-making capabilities while maintaining flexibility and compliance with evolving regulatory requirements.

3. BNPL Credit Risk Assessment: Challenges

3.1 Data Heterogeneity and Alternative Data

Buy-Now-Pay-Later (BNPL) systems rely extensively on non-traditional and heterogeneous data sources to evaluate creditworthiness, which distinguishes them from conventional lending models that primarily depend on structured financial histories. These data sources include transactional histories, behavioral analytics, and e-commerce activity, all of which provide granular insights into consumer behavior and spending patterns. While the incorporation of such alternative data enhances predictive capabilities and enables the inclusion of previously underserved or thin-file customers, it also introduces significant complexity in data preprocessing, feature engineering, and model validation. The diversity and unstructured nature of these data streams require sophisticated data integration techniques and robust validation

frameworks to ensure model reliability and generalizability across different user segments.

3.2 Real-Time Decision Constraints

Unlike traditional credit approval processes, which may involve manual review and extended processing times, BNPL systems operate under strict real-time constraints, often requiring credit decisions to be made within milliseconds at the point of sale. This necessitates the development of low-latency inference pipelines capable of processing incoming data and generating predictions almost instantaneously. In addition, efficient model deployment strategies are essential to ensure that AI models can be scaled and executed seamlessly within high-throughput environments. Achieving this level of performance requires careful optimization of computational resources, streamlined data pipelines, and the use of cloud-native technologies that support rapid and reliable decision-making.

3.3 Regulatory Compliance

Regulatory compliance represents a critical challenge in the deployment of BNPL credit risk systems, particularly given the increasing scrutiny of AI-driven financial decision-making. Key compliance requirements include ensuring transparency and explainability in automated decisions, implementing mechanisms for bias mitigation and fairness, and maintaining comprehensive audit trails for all decision processes. Financial institutions must demonstrate that their models do not discriminate against protected groups and that decisions can be clearly justified to both regulators and customers. This necessitates the integration of explainable AI techniques and governance frameworks that enable continuous monitoring, validation, and reporting of model behavior in accordance with regulatory standards.

3.4 Model Interpretability vs. Accuracy Trade-off

A fundamental challenge in AI-driven credit risk assessment is the trade-off between model interpretability and predictive accuracy. High-performing models, particularly deep neural networks and complex ensemble methods, are capable of capturing intricate patterns in data and delivering superior predictive performance. However, these models often function as “black boxes,” offering limited insight into how decisions are made. This lack of transparency creates tension

between the need for accurate risk prediction and the requirement for interpretability imposed by regulatory frameworks. As a result, financial institutions must carefully balance these competing objectives, often adopting hybrid approaches that combine high-performance models with post-hoc explainability techniques to achieve both accuracy and compliance.

4. Explainable AI Techniques for Credit Risk

4.1 Model-Agnostic Methods

Model-agnostic explainability techniques have gained significant prominence in credit risk assessment due to their ability to interpret a wide range of machine learning models without requiring access to their internal structures. Among the most widely used approaches are Local Interpretable Model-Agnostic Explanations (LIME) and Shapley Additive Explanations (SHAP). LIME provides localized interpretations by approximating the behavior of complex models in the vicinity of a specific prediction using simpler, interpretable models, thereby enabling stakeholders to understand individual decision outcomes. In contrast, SHAP offers a unified framework grounded in cooperative game theory to quantify the contribution of each feature to both individual predictions and overall model behavior. By delivering both local and global interpretability, SHAP has become particularly valuable in financial applications where transparency and accountability are critical. These model-agnostic methods are widely adopted due to their flexibility and compatibility with diverse model types, including ensemble and deep learning models, making them highly suitable for complex credit risk environments.

4.2 Intrinsically Interpretable Models

Intrinsically interpretable models represent an alternative approach to achieving transparency in credit risk assessment by design, rather than through post-hoc analysis. Models such as logistic regression, decision trees, and rule-based systems are inherently transparent, as their decision-making processes can be directly understood and communicated without additional explanation layers. Logistic regression provides clear insights into feature importance through its coefficients, while decision trees offer a visual and intuitive representation of decision paths. Rule-based systems further enhance interpretability by expressing decisions in the form of explicit logical rules. Despite these advantages, intrinsically

interpretable models often exhibit limitations in predictive performance, particularly when dealing with complex, high-dimensional datasets. Compared to advanced ensemble methods or deep learning models, they may struggle to capture intricate patterns and interactions within the data, which can reduce their effectiveness in accurately predicting credit risk in dynamic environments such as BNPL systems.

4.3 Hybrid Approaches

To address the inherent trade-off between interpretability and predictive performance, recent research has increasingly advocated for hybrid approaches that combine high-performance black-box models with post-hoc explainability techniques. In this framework, complex models such as gradient boosting machines or deep neural networks are employed to achieve superior predictive accuracy, while explainability methods like SHAP or LIME are applied to interpret their outputs. This combination enables financial institutions to leverage the strengths of advanced machine learning models while still meeting regulatory requirements for transparency and accountability. By balancing accuracy with interpretability, hybrid approaches provide a practical and scalable solution for credit risk assessment in modern financial systems, particularly in environments that demand both high performance and strict compliance with regulatory standards.

5. Proposed Framework: XAI-Driven BNPL Credit Risk Architecture

5.1 Architectural Overview

The proposed framework for Explainable AI-driven BNPL credit risk assessment is structured as a multi-layered architecture designed to integrate predictive performance, interpretability, and regulatory compliance within real-time financial systems. At the foundational level, the data ingestion layer is responsible for collecting and managing both real-time streaming data and batch data. Real-time inputs include transactional data and behavioral signals generated during user interactions, while batch data consists of historical credit records and previously stored financial information. This combination ensures that the system captures both immediate behavioral patterns and long-term credit histories, enabling a comprehensive understanding of borrower risk profiles. Building upon this, the feature engineering layer transforms raw data into meaningful and

predictive features through various preprocessing and transformation techniques. This layer plays a crucial role in integrating alternative data sources, such as digital activity and behavioral metrics, with traditional financial indicators. By harmonizing heterogeneous data inputs, it enhances the model's ability to detect subtle patterns and relationships that are indicative of credit risk. The effectiveness of this layer directly influences the quality and robustness of subsequent predictive models.

The AI modeling layer constitutes the core analytical component of the framework, where advanced machine learning algorithms are employed to generate credit risk predictions. Ensemble learning models, including techniques such as XGBoost and Random Forest, are utilized to improve predictive accuracy and robustness by combining multiple decision trees. In addition, deep learning components are incorporated to capture complex, non-linear relationships within high-dimensional datasets, particularly those derived from behavioral and transactional data. This combination of models enables the system to achieve high performance in diverse and dynamic financial environments.

To address the critical requirement of transparency, the explainability layer integrates state-of-the-art XAI techniques that provide both global and local interpretability of model outputs. SHAP is employed to deliver comprehensive insights into feature contributions across the entire model as well as for individual predictions, while LIME is used to generate instance-level explanations that help stakeholders understand specific decision outcomes. This layer ensures that the system's predictions are not only accurate but also interpretable and justifiable, thereby supporting regulatory compliance and fostering user trust.

The final layer, compliance and governance, is designed to ensure that the entire system adheres to regulatory standards and ethical guidelines. This layer incorporates mechanisms for model auditing, bias detection, and regulatory reporting, enabling continuous monitoring and validation of model performance and fairness. By maintaining detailed records of decision processes and ensuring transparency in model behavior, this layer plays a vital role in meeting legal requirements and mitigating risks associated with automated decision-making.

5.2 Cloud Deployment Considerations

The deployment of the proposed framework is facilitated through cloud-native technologies that support scalability, flexibility, and real-time processing capabilities. Containerization

technologies such as Docker and orchestration platforms like Kubernetes enable the modular deployment of system components, ensuring efficient resource utilization and ease of maintenance. Serverless inference APIs further enhance the system's ability to handle variable workloads by dynamically allocating computational resources based on demand, thereby reducing latency and operational costs. Additionally, distributed data lakes provide a scalable and centralized repository for storing and managing large volumes of structured and unstructured data, enabling seamless data access and integration across the system. Together, these cloud-native technologies ensure that the framework can operate efficiently in high-throughput environments, delivering real-time credit risk assessments while maintaining robustness and compliance with evolving financial regulations.

6. Regulatory Compliance Integration

6.1 Transparency and Accountability

Transparency and accountability are fundamental requirements in AI-driven credit risk assessment systems, particularly within regulated financial environments such as BNPL services. Explainability mechanisms play a central role in ensuring that credit decisions are interpretable and can be clearly understood by both internal stakeholders and external regulators. By providing insights into how specific features influence model predictions, these mechanisms enable financial institutions to justify automated decisions in a meaningful and accessible manner. Furthermore, the availability of interpretable outputs allows stakeholders, including auditors and compliance officers, to systematically examine and validate model behavior, thereby reinforcing accountability and trust in AI-driven systems.

6.2 Fairness and Bias Mitigation

Fairness and bias mitigation are critical components of regulatory compliance, as financial institutions must ensure that their credit decision systems do not produce discriminatory outcomes. Explainable AI techniques facilitate the identification of potentially discriminatory features and uncover hidden biases embedded within training data or model structures. By analyzing feature contributions and decision patterns, these techniques enable practitioners to detect and address sources of bias that could lead to unfair treatment of certain demographic groups. This capability is particularly important in BNPL

systems, where alternative data sources may inadvertently introduce bias if not carefully managed. Consequently, the integration of XAI supports the development of fair and ethical AI systems that align with regulatory expectations and promote equitable access to financial services.

6.3 Auditability and Reporting

Auditability and reporting are essential for demonstrating compliance with regulatory standards and ensuring the integrity of AI-driven decision-making processes. Regulators require that financial institutions maintain traceable decision logs that document how credit decisions are made, as well as reproducible model outputs that allow for independent verification of results. To meet these requirements, the proposed framework incorporates comprehensive logging and monitoring systems that capture model inputs, outputs, and intermediate processes in a structured and accessible manner. These systems enable continuous tracking of model performance and facilitate the generation of detailed reports for regulatory review. By ensuring that decision processes are transparent, traceable, and reproducible, the framework supports robust governance and strengthens confidence in the deployment of AI technologies within financial services.

7. Real-Time Decisioning in Cloud Banking

Real-time decisioning in BNPL systems necessitates a highly optimized technological infrastructure capable of delivering rapid and reliable credit assessments at the point of transaction. Such systems require low-latency model inference to ensure that predictions are generated within milliseconds, enabling seamless user experiences during checkout processes. In addition to speed, continuous model updates are essential to adapt to evolving user behaviors, market dynamics, and emerging risk patterns. This dynamic updating process allows models to remain relevant and accurate over time, thereby improving decision quality. Furthermore, scalable infrastructure is a critical requirement, as BNPL platforms must handle fluctuating transaction volumes and peak demand scenarios without compromising performance. AI-driven systems play a pivotal role in enabling automated underwriting, significantly reducing processing time while enhancing the accuracy of credit risk assessment, as highlighted in industry analyses such as those from TrustDecision. Cloud-based architectures further strengthen these capabilities by providing elasticity, resilience, and distributed

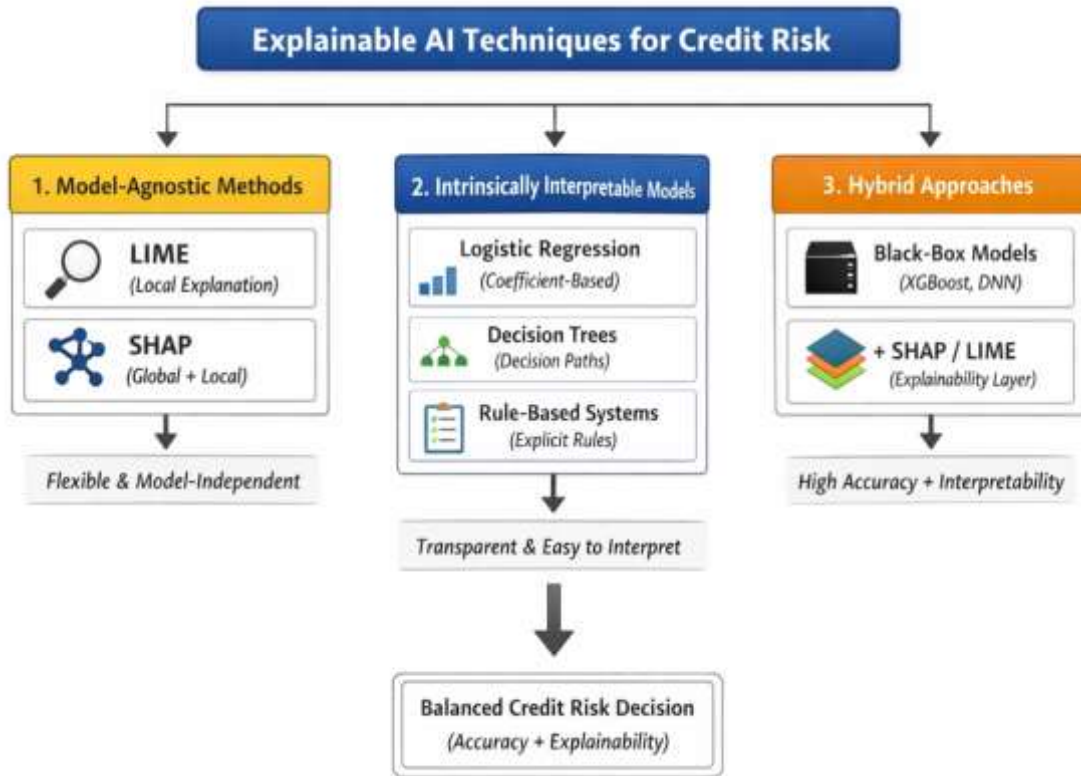


Figure 1: Explainable AI Techniques for Credit Risk

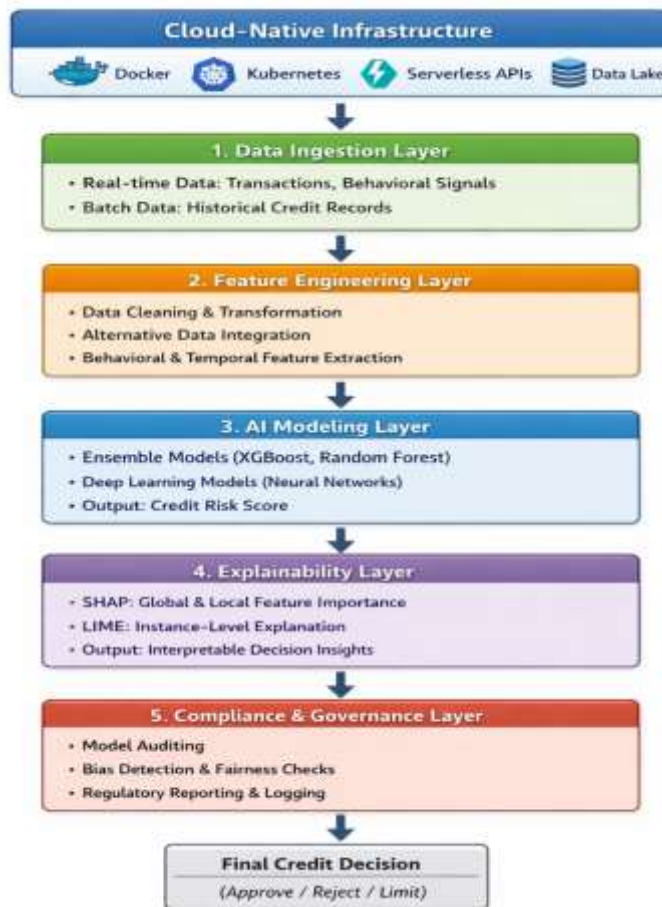


Figure 2: XAI-Driven BNPL Credit Risk Architecture

computing resources, allowing financial institutions to manage high transaction volumes efficiently while maintaining system reliability and performance.

8. Discussion

The integration of Explainable Artificial Intelligence into BNPL credit risk systems represents a significant paradigm shift in the evolution of financial technology, reflecting the convergence of advanced analytics, regulatory requirements, and cloud computing. One of the most critical observations is the emergence of explainability as a regulatory imperative, where transparency is no longer optional but a mandatory requirement for compliance with modern financial regulations. Financial institutions must ensure that automated decisions can be interpreted and justified, thereby reinforcing trust among regulators and consumers. Another important aspect is the management of the trade-off between predictive accuracy and interpretability. Hybrid modeling approaches, which combine high-performance black-box models with post-hoc explainability techniques, offer a practical solution by balancing these competing objectives. Additionally, the adoption of cloud-native architectures provides a distinct advantage by enabling real-time processing capabilities, which are essential for the operational success of BNPL systems.

9. Conclusions

This paper has presented a comprehensive analysis of explainable AI frameworks for BNPL credit risk assessment within the context of real-time cloud banking architectures. The findings highlight the critical role of explainability in ensuring regulatory compliance, fostering trust, and enhancing operational efficiency in AI-driven financial systems. The proposed framework demonstrates how explainable AI can be systematically integrated into modern fintech infrastructures, addressing key challenges related to transparency, fairness, and scalability. By combining advanced machine learning techniques with robust explainability and governance mechanisms, the framework provides a practical solution for deploying reliable and compliant credit risk assessment systems. As the adoption of BNPL services continues to expand, future research must focus on developing standardized metrics for explainability, improving the robustness and adaptability of AI models, and ensuring alignment with evolving regulatory landscapes. These efforts will be essential for advancing the responsible use

of AI in financial services and sustaining innovation in digital lending ecosystems.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

- [1] Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, 33(7), 2845–2897.
- [2] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2020). Explainable AI in fintech risk management. *Frontiers in Artificial Intelligence*, 3, 26.
- [3] Demajo, L. M., et al. (2020). Explainable AI for interpretable credit scoring. *arXiv preprint arXiv:2012.03749*.
- [4] Di Maggio, M., Yao, V., & others. (2022). Fintech borrowing and financial fragility. *Journal of Financial Economics*.
- [5] Dragoni, N., et al. (2017). Microservices: Yesterday, today, and tomorrow. *Present and Ulterior Software Engineering*.
- [6] Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine*, 38(3), 50–57.
- [7] Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep learning for finance: Deep portfolios. *Applied Stochastic Models in Business and Industry*, 33(1), 3–12.
- [8] Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11), 2767–2787.

- [9] Lessmann, S., Baesens, B., Seow, H. V., & Thomas, L. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring. *European Journal of Operational Research*, 247(1), 124–136.
- [10] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*.
- [11] Lupşa-Tătaru, D. A. (2023). Buy Now Pay Later: A perspective. *Economies*, 11(8), 218.
- [12] Misheva, B. H., et al. (2021). Explainable AI in credit risk management. *arXiv preprint arXiv:2103.00949*.
- [13] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” Explaining the predictions of any classifier. *Proceedings of KDD*.
- [14] Zhang, Q., Chen, M., & Li, L. (2018). Cloud computing: State-of-the-art and research challenges. *Journal of Internet Services and Applications*, 9(1).