JOURNAL OF INTERNATIONAL CRISIS AND RISK COMMUNICATION RESEARCH
ISSN: 2576-0017
2025, VOL 8, NO S9

Explainable AI (XAI) for Cloud-Based Enterprise Applications: Building Trust and Transparency in AI-Driven Decisions

Bhaskar Goyal

University of Southern California.

Abstract
The proliferation of Artificial Intelligence systems in enterprise cloud environments has generated unprecedented demand for transparency and interpretability in automated decision-making processes. Explainable AI (XAI) emerges as a transformative paradigm that addresses critical organizational needs by providing methodologies and techniques to make AI systems more interpretable, transparent, and accountable. This technical review examines the current state of XAI implementation in cloud-based enterprise applications, analyzing various methodologies including feature importance analysis, rule extraction techniques, surrogate models, and visualization approaches. The integration of XAI into cloud-based enterprise applications represents a fundamental evolution in organizational AI deployment strategies, requiring an optimal balance between performance and explainability to ensure stakeholders can understand, validate, and trust AI-driven decisions. Key application areas include building trust with users through personalized explanation frameworks, ensuring regulatory compliance and governance across multiple jurisdictional requirements, enhancing debugging and model improvement capabilities, and facilitating human-AI collaboration through hybrid decision-making systems. Implementation challenges encompass the trade-off between explainability and accuracy, lack of standardized metrics and evaluation frameworks, integration complexities with existing MLOps workflows, and scalability requirements for concurrent users across distributed geographical regions. Future directions focus on developing robust and scalable XAI techniques, user-centric interfaces, ethical frameworks, automated self-explaining systems, and cross-modal explanation capabilities that will enable more responsible and effective AI deployment across diverse organizational contexts.

Keywords: Explainable Artificial Intelligence, Cloud Computing, Enterprise Applications, Trust and Transparency, AI Governance.

1. Introduction

The proliferation of Artificial Intelligence systems in enterprise cloud environments has created an unprecedented demand for transparency and interpretability in automated decision-making processes. Current enterprise deployments demonstrate significant processing volumes with substantial peak loads during operational hours. Recent studies indicate that enterprise decision-makers identify a lack of explainability as the primary barrier preventing broader AI adoption in mission-critical applications, particularly within regulated industries where algorithmic transparency is mandatory [1].

Explainable AI emerges as a crucial paradigm shift that addresses these transparency concerns by providing methodologies and techniques to make AI systems more interpretable, transparent, and accountable. The global market for explainable AI solutions has experienced remarkable growth, with cloud-based XAI implementations accounting for the majority of this market expansion, driven by enterprises allocating substantial portions of their total AI development budgets to explainability initiatives.

The integration of XAI into cloud-based enterprise applications represents a fundamental evolution in organizational AI deployment strategies. Unlike traditional AI implementations that prioritize accuracy above interpretability, XAI-enabled systems must achieve an optimal balance between performance and explainability, ensuring stakeholders can understand, validate, and trust AI-driven decisions. Performance benchmarks reveal that contemporary XAI implementations maintain high accuracy levels compared to their black-box counterparts while providing comprehensive explanations, with minimal latency increases for real-time explanation generation [2].

This paradigm proves particularly critical in enterprise environments where AI decisions directly impact financial outcomes, regulatory compliance, and human welfare. Implementation studies demonstrate that XAI deployment reduces algorithmic bias incidents and improves regulatory audit success rates significantly. The cloud computing paradigm introduces additional complexity to XAI implementation, encompassing considerations around distributed processing, scalability, multi-tenancy, and data governance requirements.

Enterprise cloud applications must provide explanations for individual AI decisions while maintaining consistency and reliability across diverse deployment scenarios, user contexts, and regulatory frameworks. Current systems support concurrent explanation generation for numerous simultaneous users across multiple geographical regions, handling varying computational loads and explanation complexity requirements. The distributed nature of cloud environments necessitates sophisticated explanation caching mechanisms, load balancing strategies, and fault-tolerance protocols to ensure the continuous availability of explanation services.

This technical review examines the current state of XAI research and implementation in cloud-based enterprise applications, analyzing methodologies, challenges, and future directions for building trustworthy AI systems that can scale to meet modern enterprise operational demands while maintaining explanation quality and computational efficiency.

2. XAI Methods and Techniques for Enterprise Applications

The landscape of XAI methodologies encompasses a diverse array of techniques designed to provide different levels and types of explanations for AI model behavior. Contemporary enterprise implementations utilize multiple distinct XAI approaches, with organizations typically deploying complementary techniques to address varying explanation requirements. These methods can be broadly categorized into several key approaches, each with distinct advantages and limitations for enterprise cloud deployments, supporting varying explanation generation speeds for different complexity levels [3].

2.1 Feature Importance Analysis

Feature importance analysis represents one of the most widely adopted XAI techniques in enterprise settings, with significant deployment rates across cloud-based AI implementations. Methods such as SHAP and LIME provide quantitative measures of how individual features contribute to specific predictions, with implementations processing substantial volumes of feature importance calculations in distributed cloud environments. In cloud-based enterprise applications, feature importance analysis enables business users to understand which data elements most significantly influence AI decisions, with typical enterprise datasets containing varying numbers of features requiring real-time analysis.

Advanced implementations leverage distributed computing capabilities to perform real-time feature importance calculations across large datasets containing substantial amounts of structured data, enabling interactive exploration of model behavior with minimal response times. Cloud-native implementations demonstrate linear scalability, handling concurrent requests from numerous users while maintaining high explanation accuracy compared to single-threaded calculations. Enterprise deployments report significant reductions in model validation time and substantial improvements in business user comprehension rates compared to traditional model documentation approaches.

2.2 Rule Extraction Techniques

Rule extraction techniques focus on deriving human-readable decision rules from complex AI models, with modern implementations generating comprehensive rule sets that capture substantial portions of the original model behavior. These approaches prove particularly valuable in enterprise contexts where compliance and audit requirements demand explicit decision criteria, with organizations across various sectors reporting significantly faster regulatory audit completion times when utilizing rule-based explanations [3].
Modern rule extraction methods employ techniques such as decision tree surrogate models, association rule mining, and logical rule induction to translate complex neural network behaviors into interpretable conditional statements. Cloud-based implementations process rule extraction tasks for models containing substantial parameter counts, generating comprehensive rule sets within reasonable timeframes depending on model complexity. Containerized rule extraction services demonstrate elastic scaling capabilities, automatically provisioning additional compute resources during peak demand periods while maintaining rule consistency across distributed extraction processes.

2.3 Surrogate Models

Surrogate models provide another critical XAI approach, where simpler, interpretable models are trained to approximate the behavior of complex black-box systems with high fidelity rates across enterprise implementations. In enterprise cloud environments, surrogate models serve dual purposes: providing explanations for individual predictions with minimal latency and offering global insights into model behavior patterns across datasets containing substantial sample volumes.
Techniques such as locally accurate surrogate models achieve high fidelity within local neighborhoods of similar instances, while global surrogate approximations maintain substantial overall accuracy across entire datasets. Cloud-based surrogate model training processes handle extensive datasets, completing training cycles within reasonable timeframes using distributed computing frameworks [4].

2.4 Visualization Techniques

Visualization techniques have evolved significantly to support enterprise XAI requirements, with modern implementations supporting numerous distinct visualization formats optimized for different explanation types and user expertise levels. Advanced visualization methods include attention heatmaps, processing multidimensional data representations, activation pattern displays for complex neural networks, and interactive decision path visualizations supporting exploration of extensive decision trees. These systems process substantial visualization requests utilizing advanced acceleration techniques to achieve real-time performance for complex graphical explanations [4].

3. XAI in Cloud-Based Enterprise Applications

The implementation of XAI in cloud-based enterprise applications addresses several critical organizational needs, fundamentally transforming how businesses interact with and govern AI systems. Current enterprise deployments demonstrate that organizations processing substantial daily AI transactions report significantly higher user satisfaction rates when XAI capabilities are integrated compared to black-box implementations. The cloud deployment model introduces unique opportunities and challenges for XAI implementation, requiring careful consideration of scalability, security, and integration requirements, with typical enterprise cloud XAI systems supporting concurrent explanation generation for numerous simultaneous users across distributed geographical regions [5].

3.1 Building Trust with Users

Building Trust with Users represents perhaps the most significant benefit of XAI in enterprise environments, with studies indicating that transparent AI systems achieve substantially higher user adoption rates compared to unexplainable alternatives. Trust formation in AI systems depends heavily on users' ability to understand and validate AI decisions, with enterprise surveys revealing that the majority of business stakeholders require explanation capabilities before approving AI deployment in critical processes. Cloud-based XAI implementations leverage distributed explanation generation capabilities to provide real-time, contextually relevant explanations to diverse user groups, processing explanation requests with minimal response times across global cloud infrastructures.

Advanced systems employ personalized explanation frameworks that adapt to individual user expertise levels, professional roles, and decision-making contexts, supporting multiple distinct explanation complexity levels from basic summary views to detailed technical analysis. Enterprise implementations demonstrate that personalized explanations improve user comprehension rates substantially compared to standardized explanation formats. These implementations often incorporate feedback mechanisms that allow users to validate, challenge, or refine AI explanations, creating iterative trust-building processes that show measurable trust score improvements over extended deployment periods [5].

3.2 Regulatory Compliance and Governance

Regulatory Compliance and Governance have become increasingly important as regulatory frameworks worldwide begin requiring AI transparency and accountability, with compliance-related XAI implementations growing substantially across enterprise sectors. Cloud-based XAI systems must support compliance with regulations such as GDPR's "right to explanation," financial services algorithmic accountability requirements, and healthcare AI transparency mandates, with enterprise compliance systems processing substantial volumes of explanation audit requests quarterly.

Enterprise cloud implementations typically incorporate audit trail capabilities that document explanation generation processes, maintain versioned explanation models, and provide regulatory reporting functionality, storing extensive historical explanation data across enterprise deployments. These systems often employ automated compliance checking mechanisms that validate explanation quality and completeness against regulatory requirements, achieving high automated compliance verification rates while reducing manual audit preparation time significantly [6].

3.3 Debugging and Model Improvement

Debugging and Model Improvement capabilities are significantly enhanced through cloud-based XAI implementations, with enterprise deployments reporting faster model issue identification and reduced model maintenance costs compared to traditional debugging approaches. Distributed explanation generation allows for comprehensive analysis of model behavior across different data segments, user populations, and operational contexts, processing behavioral analysis across extensive datasets with reasonable analysis completion times.

Advanced cloud XAI systems incorporate anomaly detection mechanisms that identify unusual explanation patterns, potentially indicating model degradation, data drift, or adversarial attacks, with high detection accuracy rates for known anomaly types [6].

3.4 Human-AI Collaboration

Human-AI Collaboration is fundamentally transformed through effective XAI implementation in cloud environments, with hybrid decision-making systems demonstrating substantial improvement in decision quality metrics compared to purely automated or purely human decision processes. Modern enterprise applications leverage XAI to create hybrid decision-making systems where AI provides recommendations along with explanations, enabling human experts to make informed decisions about when to trust, override, or refine AI suggestions.

XAI Application Area	Primary Benefits/Functions	Key Implementation Features
Building Trust with Users	Enhanced user adoption rates and stakeholder confidence	Personalized explanation frameworks, real-time distributed generation, feedback mechanisms with iterative trust-building processes
Regulatory Compliance and Governance	Automated compliance verification and audit trail management	Comprehensive audit capabilities, versioned explanation models, and automated compliance checking across multiple jurisdictional requirements

Debugging and Model Improvement	Faster issue identification and reduced maintenance costs	Anomaly detection mechanisms, behavioral analysis across data segments, explanation pattern monitoring with automated alerting
Human-AI Collaboration	Improved decision quality through hybrid decision-making	Collaborative explanation refinement, expert feedback integration, interactive annotation capabilities with consensus mechanisms

Table 1: Trust, Compliance, and Collaboration Features in Enterprise XAI Systems [5, 6]

4. Implementation Challenges and Technical Considerations

The implementation of XAI in cloud-based enterprise applications presents numerous technical and organizational challenges that must be carefully addressed to achieve successful deployment and adoption. Enterprise surveys indicate that organizations encounter significant implementation barriers during XAI deployment, with average implementation timelines extending substantially beyond initial projections. These challenges span multiple dimensions, from technical architecture considerations to user experience design and organizational change management, with organizations typically requiring substantial additional computational resources compared to standard AI deployments [7].

4.1 The Trade-off Between Explainability and Accuracy

The Trade-off Between Explainability and Accuracy remains one of the most significant challenges in enterprise XAI implementation, with studies demonstrating notable accuracy reductions when transitioning from black-box to explainable models. High-performing AI models, particularly deep neural networks achieving exceptional accuracy levels, often sacrifice interpretability for performance, while more interpretable models, such as decision trees and linear models, typically achieve lower accuracy in comparable enterprise scenarios. Cloud-based enterprise systems must navigate this trade-off through sophisticated model architecture decisions, often employing ensemble approaches that combine high-accuracy models with interpretable surrogates, achieving compromise accuracy levels while providing comprehensive explanations.

Advanced implementations utilize multi-objective optimization techniques to identify optimal points in the accuracy-explainability space, processing optimization cycles across extensive candidate model configurations with substantial evaluation timeframes for complex enterprise datasets. These systems adapt model selection based on specific use case requirements and risk tolerance levels, with critical applications requiring high accuracy thresholds while maintaining reasonable explanation generation latency. Performance benchmarking across enterprise implementations reveals that explainable model variants require substantially more computational resources during both training and inference phases compared to their black-box counterparts [7].

4.2 Standardized XAI Metrics and Evaluation Frameworks

Standardized XAI Metrics and Evaluation Frameworks represent a critical gap in current enterprise implementations, with limited organizations utilizing consistent evaluation methodologies across their XAI deployments. The lack of universally accepted metrics for explanation quality, completeness, and effectiveness makes it difficult to compare XAI approaches or validate their effectiveness in real-world scenarios, with current evaluation approaches varying significantly in methodology and scoring criteria across different enterprise implementations.

Enterprise cloud applications require robust evaluation frameworks that can assess explanation quality across multiple dimensions, including fidelity in representing model behavior, stability in maintaining consistency of explanations across similar inputs, and comprehensibility, enabling user understanding across diverse stakeholder groups. Current research focuses on developing standardized evaluation protocols that can be implemented across diverse cloud environments and use cases, with prototype frameworks processing evaluation cycles across extensive explanation instances requiring substantial time for comprehensive assessment [8].

4.3 Integration with Existing AI Development and Deployment Workflows

Integration with existing AI Development and Deployment Workflows poses significant technical challenges for enterprise organizations, with substantial numbers of enterprises requiring major modifications to existing MLOps pipelines to accommodate XAI capabilities. Most existing AI development pipelines were designed without explainability considerations, requiring extensive integration efforts and involving considerable development resources for comprehensive XAI enablement.

Cloud-based implementations must seamlessly integrate with existing MLOps frameworks, providing explanation generation capabilities that don't disrupt established development, testing, and deployment processes, while adding substantial time to overall pipeline execution. This often requires sophisticated containerization strategies, extensive API design considerations, and workflow orchestration capabilities that can adapt to diverse enterprise technology stacks spanning multiple platforms and frameworks [8].

4.4 Scalability and Performance Considerations

Scalability and Performance Considerations become particularly challenging in cloud environments where explanation generation must support substantial numbers of concurrent users across different geographic regions and time zones. Real-time explanation generation for complex models can be computationally intensive, requiring careful optimization of explanation algorithms and infrastructure scaling strategies. Advanced cloud implementations employ comprehensive caching mechanisms, extensive pre-computed explanation templates, and distributed explanation generation architectures to maintain acceptable response times while managing computational costs.

5. Future Research Directions and Emerging Trends

The future of XAI in cloud-based enterprise applications is shaped by several emerging research directions and technological trends that promise to address current limitations while opening new possibilities for transparent AI deployment. Current research investments in XAI have experienced substantial growth over recent years, with global spending reaching significant levels across academic and industrial sectors. These developments span multiple disciplines, from computer science and machine learning to human-computer interaction and organizational psychology, with numerous active research projects worldwide focusing on next-generation XAI capabilities [9].

5.1 Robust and Scalable XAI Techniques

Robust and Scalable XAI Techniques represent a primary focus of current research efforts, with substantial portions of enterprise XAI research initiatives prioritizing robustness and scalability improvements. Next-generation XAI methods must maintain explanation quality and consistency across diverse operational conditions, including data distribution shifts affecting enterprise datasets regularly, frequent model updates in typical enterprise environments, and varying computational resource availability during different operational periods.

Research is progressing toward developing explanation techniques that are inherently robust to adversarial attacks, with current prototypes demonstrating improved explanation consistency under adversarial conditions compared to traditional XAI methods. These systems can maintain interpretability even as underlying models evolve, with explanation drift metrics showing minimal variance across model version updates. Cloud-native XAI architectures are being designed to leverage containerization, supporting multiple concurrent explanation services, microservices handling substantial request volumes, and serverless computing paradigms, achieving enhanced cost efficiency compared to traditional deployment models while maintaining high explanation quality levels [9].

5.2 User-Centric XAI Interfaces

User-Centric XAI Interfaces are emerging as a critical research area, with user experience research showing that interface design significantly impacts explanation effectiveness across different user groups. Future enterprise XAI systems will incorporate adaptive interface technologies that personalize explanation presentations based on user expertise levels spanning multiple distinct proficiency categories, contextual requirements varying across different business scenarios, and cognitive load considerations measured through various usability metrics.

Research is exploring natural language explanation generation, achieving substantially higher user comprehension rates compared to traditional visualization approaches, conversational AI interfaces for explanation exploration supporting extended dialogue sessions with high user satisfaction scores, and augmented reality visualization techniques providing immersive explanation experiences with significantly longer engagement times than traditional interfaces. These developments prove particularly important for enterprise applications serving diverse user populations, with studies showing that personalized explanation interfaces substantially improve decision confidence and reduce explanation interpretation time [10].

5.3 Ethical and Societal Implications

Ethical and Societal Implications of XAI in enterprise applications are receiving increased research attention, with ethics-focused XAI research growing substantially and representing significant portions of total XAI research investment across major research institutions. Future research directions include developing frameworks for ethical explanation design, understanding the potential for explanation manipulation or misuse, and investigating how XAI can support fairness and bias mitigation in enterprise AI systems.

Research demonstrates that XAI-enabled bias detection systems identify substantially more bias instances compared to traditional fairness assessment methods, while explanation-guided bias mitigation techniques achieve significant reductions in discriminatory outcomes across diverse enterprise applications. Cloud-based implementations must consider global ethical standards and cultural differences in explanation preferences, with cross-cultural studies showing substantial variation in explanation trust formation patterns across different cultural contexts [10].

5.4 Automated XAI and Self-Explaining Systems

Automated XAI and Self-Explaining Systems represent an emerging research frontier where AI systems are designed with inherent explanation capabilities rather than post-hoc explanation generation. These systems promise to eliminate the accuracy-explainability trade-off by building interpretability directly into model architectures, with prototype implementations demonstrating high accuracy retention compared to black-box equivalents while providing comprehensive explanations.

5.5 Cross-Modal and Multimodal Explanation Systems

Cross-Modal and Multimodal Explanation Systems are being developed to support enterprise applications processing diverse data types, with current systems handling multiple distinct data modalities simultaneously. Future XAI systems will provide unified explanation frameworks that can interpret and explain decisions across multiple data modalities, enabling a comprehensive understanding of complex enterprise AI applications.

Research Area	Current Status & Challenges	Research Focus & Methods	Expected Benefits & Outcomes
Robust and Scalable XAI Techniques	Data distribution shifts; frequent model updates; varying computational resources; vulnerability to adversarial attacks	Cloud-native architectures; containerization; microservices; serverless computing; adversarial robustness	Enhanced cost efficiency; improved explanation consistency; minimal variance across model updates; high quality maintenance
User-Centric XAI Interfaces	Interface design impacts effectiveness; diverse user expertise levels; varying contextual requirements; cognitive load issues	Adaptive interface technologies; natural language generation; conversational AI; augmented reality visualization	Higher user comprehension rates; improved decision confidence; reduced interpretation time; extended engagement
Ethical and Societal Implications	Potential explanation manipulation; cultural differences in trust; global	Ethical explanation frameworks; bias detection systems;	Substantial bias reduction; improved fairness assessment; enhanced

	ethical standards; bias detection limitations	fairness mitigation techniques; cross-cultural studies	discriminatory outcome prevention; cultural adaptation
Automated XAI and Self-Explaining Systems	Accuracy-explainability trade-off; post-hoc explanation limitations; interpretability integration challenges	Inherent explanation capabilities; interpretability-integrated architectures; automated explanation generation	Eliminated accuracy trade-off; high accuracy retention; comprehensive built-in explanations; seamless integration

Table 2: Emerging Trends and Technological Advances for Cloud-Based Explainable AI Systems [9, 10]

Conclusion

The evolution of explainable AI in cloud-based enterprise applications represents a fundamental shift toward more transparent, accountable, and trustworthy artificial intelligence systems. The comprehensive exploration of XAI methodologies reveals that organizations must carefully balance multiple competing requirements, including accuracy preservation, explanation quality, regulatory compliance, and user experience optimization. Feature importance analysis, rule extraction techniques, surrogate models, and visualization approaches each offer distinct advantages for different enterprise contexts, requiring thoughtful selection and integration strategies. The implementation challenges identified throughout this review highlight the complexity of deploying XAI systems at enterprise scale, particularly regarding the persistent tension between model accuracy and explainability, the absence of standardized evaluation frameworks, and the substantial technical modifications required for existing AI development workflows. However, the substantial benefits of XAI implementation, including enhanced user trust, improved regulatory compliance, superior debugging capabilities, and more effective human-AI collaboration, demonstrate compelling value propositions for enterprise adoption. Future developments in robust and scalable XAI techniques, user-centric interface design, ethical framework development, automated self-explaining systems, and multimodal explanation capabilities promise to address current limitations while opening new possibilities for transparent AI deployment. The ultimate trajectory points toward creating AI systems that are inherently interpretable, culturally sensitive, and aligned with human values and organizational objectives. This evolution will be critical for fostering broader AI adoption in enterprise environments while maintaining the trust and confidence necessary for successful digital transformation initiatives across diverse organizational contexts.

References

1. Neeraj Anand Sharma, et al., "Explainable AI Frameworks: Navigating the Present Challenges and Unveiling Innovative Applications," Algorithms, 2024. Available: https://www.mdpi.com/1999-4893/17/6/227
2. Thiyagarajan Mani Chettier, et al., "Scalable Explainable AI with a Cloud-Native Approach for Cybersecurity Threat Detection," ResearchGate, 2025. Available: https://www.researchgate.net/publication/392171111_Scalable_Explainable_AI_with_a_Cloud-Native_Approach_for_Cybersecurity_Threat_Detection
3. Usama Noor and Ghulam Abbas, "Explainable AI Meets Distributed Databases: Leveraging Reinforcement Learning and Generative Models for Advanced Predictive Scaling and Query Optimization," ResearchGate, 2024. Available: https://www.researchgate.net/publication/383991934_Explainable_AI_Meets_Distributed_Databases_Leveraging_Reinforcement_Learning_and_Generative_Models_for_Advanced_Predictive_Scaling_and_Query_Optimization
4. Waddah Saeed and Christian Omlin, "Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities," Knowledge-Based Systems, 2024. Available: https://www.sciencedirect.com/science/article/pii/S0950705123000230

5. Sodiq Oyetunji Rasaq, "Building Trust in Cloud-Based AI: The Need for Explainability in Data Science," ResearchGate, 2025. Available: https://www.researchgate.net/publication/389064606_Building_Trust_in_Cloud-Based_AI_The_Need_for_Explainability_in_Data_Science
6. Sudheer Singh, et al., "Explainable AI (XAI) for Cloud Resource Forecasting in E-Commerce Environments," International Journal for Multidisciplinary Research (IJFMR) 2025. Available: https://www.ijfmr.com/papers/2025/4/51796.pdf
7. Gideon Areo, "The Role of Explainable AI in Optimizing System Performance: A Guide to Bottleneck Detection," ResearchGate, 2024. Available: https://www.researchgate.net/publication/387223308_The_Role_of_Explainable_AI_in_Optimizing_System_Performance_A_Guide_to_Bottleneck_Detection
8. Harrison Clarke, "The Role of MLOps in Explainable AI: Use Cases and Approaches," 2024. Available: https://www.harrisonclarke.com/blog/the-role-of-mlops-in-explainable-ai-use-cases-and-approaches
9. Bowen Long, et al., "Explainable AI the Latest Advancements and New Trends," ResearchGate, 2025. Available: http://researchgate.net/publication/391676242_Explainable_AI_the_Latest_Advancements_and_New_Trends
10. Douglas Cirqueira, et al., "Towards Design Principles for User-Centric Explainable AI in Fraud Detection," ResearchGate, 2021. Available: https://www.researchgate.net/publication/352958677_Towards_Design_Principles_for_User-Centric_Explainable_AI_in_Fraud_Detection