International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

An Intelligent Bayesian Security Architecture: AI Threat Assessment and Real-Time Lakehouse Risk Intelligence in Low-Data Cloud Environments

Alexei Viktorovich Kuznetsov

Data Analyst, Russia

ABSTRACT: Enterprises operating in low-data or fragmented cloud environments face significant challenges in detecting threats, assessing risks, and maintaining continuous situational awareness. This paper introduces an intelligent Bayesian security architecture that integrates Dynamic Bayesian inference with AI-driven threat assessment to enhance security decision-making under uncertainty. The proposed system unifies real-time streaming intelligence with lakehouse-based risk analytics, enabling scalable ingestion, harmonization, and probabilistic evaluation of diverse security signals. By leveraging hierarchical Bayesian models, the architecture continuously updates risk probabilities, compensates for sparse or incomplete data, and delivers adaptive alerts with quantifiable confidence levels. Real-time lakehouse analytics ensure seamless integration of structured, semi-structured, and event-driven telemetry, while cloud-native orchestration supports rapid scaling across distributed environments. This approach significantly improves threat visibility, reduces false positives, and provides a robust decision framework for modern enterprises operating in data-scarce or high-variability threat landscapes.

KEYWORDS: Bayesian Security Architecture; AI Threat Assessment; Dynamic Bayesian Models; Real-Time Risk Intelligence; Data Lakehouse; Low-Data Environments; Cloud Security; Streaming Analytics; Probabilistic Inference; Threat Intelligence; Security Automation; Digital Risk Management.

I. INTRODUCTION

Credit risk modeling is foundational to financial stability and institutional health. Lenders, regulators, and investors rely on accurate estimates of probability of default (PD), loss given default (LGD), exposure at default (EAD), and provisioning metrics to make underwriting, capital allocation, and supervisory decisions. Traditional credit models combine econometric methods, scorecards, and more recently, supervised machine learning models trained on historical borrower and performance data. However, these approaches face persistent challenges: (1) scarcity of meaningful default events for certain borrower cohorts and products, (2) evolving borrower behaviors and covariate distributions under shifting macroeconomic regimes, (3) privacy and regulatory constraints limiting data sharing for model development, and (4) adversarial environmental threats that can undermine model integrity.

Generative AI offers a new set of capabilities to address some of these challenges. By learning rich joint distributions of borrower features and performance trajectories, generative models can create synthetic cohorts that (a) augment training sets for rare default events, (b) simulate counterfactual borrower histories under hypothetical macro paths for stress testing, and (c) provide privacy-preserving data for model development, transfer learning, or vendor collaboration. When applied correctly, these abilities can improve the robustness of PD and LGD estimations, enable more realistic stress tests, and reduce compliance friction when sharing development datasets.

Yet the same flexibility that makes generative models powerful also raises serious concerns in credit risk contexts. Generative models trained on sensitive credit files can leak individual-level information if not properly regularized; they can be poisoned by malicious injections that bias risk estimates; and their black-box nature can make it difficult for model risk and compliance teams to validate their outputs. Regulators expect not just accurate models but auditable, interpretable, and well-governed systems. Therefore, any generative AI deployment in credit risk must be designed with security and explainability as first-order objectives, not afterthoughts.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

This paper articulates a comprehensive architecture and methodology for **secure generative AI in credit risk** that addresses three integrated pillars: (1) threat intelligence and adversarial hardening to protect data integrity and model confidentiality; (2) explainability and audit artifacts that translate complex generative behavior into human-understandable drivers and scenario templates; and (3) cloud-native optimization ensuring scalable, cost-effective training and inference with hardened operational controls. The design recognizes that credit institutions operate under strict regulatory regimes (capital adequacy and consumer protection) and therefore emphasizes provable privacy (differential privacy), immutable provenance (tamper-evident registries), and continuous monitoring (privacy leakage metrics and attack indicators).

We position our approach as augmentative: generative AI is not intended to replace well-established scorecards or expert judgment but to complement them—by providing synthetic scenarios for stress testing, augmenting rare-event learning, and enabling privacy-preserving model development pipelines. The paper presents the architecture, the specific threat model and defenses, explainability techniques adapted to generative outputs used in credit settings, and a cloud orchestration strategy optimized for security and cost. We also include two applied case studies that show real-world value: PD calibration improvement via synthetic augmentation for SME portfolios, and multi-factor stress scenario generation for consumer unsecured exposures. The remainder of the paper reviews relevant literature, details the research methodology, presents results and discussion, and concludes with recommended operational and research priorities.

II. LITERATURE REVIEW

The literature relevant to secure generative AI for credit risk sits at the intersection of generative modeling, privacy and adversarial security, explainable AI, and applied credit-risk modeling and governance.

Generative modeling for tabular and sequential data. Foundational works on VAEs (Kingma & Welling, 2014) and GANs (Goodfellow et al., 2014) opened the door to learning complex joint distributions. Subsequent research adapted these frameworks for tabular data and mixed continuous-categorical features important in credit datasets (Xu & Veeramachaneni, 2019; Choi et al., 2017). For sequential financial data, recurrent or Transformer-based conditional generators produce plausible performance trajectories (e.g., repayment sequences and delinquency paths), and conditional mechanisms allow anchoring generation on macro scenarios. Specialized losses and architectures have been proposed to preserve statistical properties like marginal distributions, cross-feature interactions, and tail behavior—critical when modeling defaults.

Privacy-preserving synthetic data. Differential privacy (Dwork & Roth, 2014) provides a formal framework for limiting information leakage from trained models. Applying DP to generative training (DP-SGD) is a growing area, with trade-offs between model fidelity and privacy guarantees widely studied (Abadi et al., 2016). Other privacy modalities include k-anonymity and secure multiparty computation/federated learning for collaborative model training without centralizing raw data (Bonawitz et al., 2019). In credit contexts, privacy practices are essential to meet consumer protection regulations and contractual constraints.

Adversarial threats and model hardening. The security community has documented attacks on ML systems—data poisoning, evasion, model extraction, and membership inference (Biggio & Roli, 2018; Tramèr et al., 2016). Generative models face specific risks: training-set poisoning can shift synthesized distributions, and generative outputs can leak training instances. Defenses include robust estimators, anomaly detection on training inputs, DP, adversarial training, and certified robustness methods for certain model classes. In tabular data, achieving certified defenses remains challenging, so layered operational defenses and active red-teaming are typically recommended.

Explainability and regulatory expectations. Explainable AI methods—local (LIME, Ribeiro et al., 2016; SHAP, Lundberg & Lee, 2017) and global surrogate approaches—have been widely adopted to render black-box models interpretable. For generative systems, explanation tasks are twofold: explaining why a synthetic record has certain risk characteristics, and summarizing how the generative model's latent factors relate to risk outputs. Recent work has adapted SHAP-like attributions to generative contexts and proposed latent traversals and counterfactual generation to illuminate behavior. In regulated industries, explainability is not only desirable but often required for model governance.

Credit-specific modeling and governance. Classic credit techniques (scorecards, logistic regression) remain relevant for interpretability and regulatory acceptance (Hand & Henley, 1997). Machine learning has improved predictive

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

performance in some contexts (Lessmann et al., 2015), but regulators emphasize robust model risk management: validation, versioning, sensitivity testing, stressed-scenario assessment, and documentation (BCBS and national guidance). Integrating generative outputs into these workflows requires additional governance artifacts—model cards, synthetic-data certificates, and documented privacy budgets.

Combined applications. Emerging studies have used synthetic data to augment credit model training and to conduct stress tests, showing improvements in rare-event estimation and model robustness (recent conference and industry reports). However, many open questions remain: how to account for privacy-utility trade-offs in cohort-specific PD estimates, how to certify robustness against targeted poisoning in tabular generative training, and how to produce explainability artifacts that satisfy audit objectives. This paper contributes by synthesizing best practices across these literatures and demonstrating applied workflows tailored to credit risk.

III. RESEARCH METHODOLOGY

1. **Objectives & evaluation criteria.** Define objectives: (a) improve PD, LGD calibration for underrepresented cohorts via synthetic augmentation; (b) enable conditional counterfactual stress scenarios for provisioning analysis; (c) ensure provable privacy guarantees and resilient defenses against realistic adversarial threats. Evaluation metrics include predictive calibration (Brier score, calibration-in-the-large), discriminatory power (AUC), tail fidelity (differences in estimated high-quantile PD/ECL), privacy leakage (membership inference advantage and formal DP ε), and security posture (success rate of poisoning/red-team attacks).
2. **Data collection and governance.** Gather historical credit account-level data (origination attributes, credit bureau features, repayment histories, charge-offs), macroeconomic series (unemployment, GDP, liquidity spreads), and external covariates (industry, region). Enforce strict data governance: data lineage, consent checks, and PII minimization. Partition data chronologically for training and holdout testing to prevent look-ahead bias.
3. **Preprocessing & feature engineering.** Convert account trajectories into structured sequence windows (e.g., monthly snapshots of balance, payment status). Encode categorical variables via supervised embeddings learned jointly with the model; normalize numeric features with robust scaling. Create derived features capturing borrower behavior (roll-rate transitions, payment velocity) and macro interactions (loan-to-income adjusted by regional unemployment).
4. **Generative architecture design.** Use a hybrid conditional architecture: a conditional VAE captures global latent structure and provides training stability; a cGAN module refines samples to better match high-order interactions and tail behaviors. Conditioners include borrower cohort flags, macro-path summaries, and time-to-maturity. For sequences, employ a Transformer-decoder with autoregressive sampling for long-horizon trajectories; for tabular snapshots, use conditional fully connected decoders with discrete-continuous output heads. Loss functions combine reconstruction (likelihood), adversarial loss, and moment-matching penalties emphasizing tail distribution matching (e.g., matching extreme quantiles).
5. **Privacy-preserving training.** Implement DP-SGD for sensitive cohorts: gradient clipping and calibrated noise addition with formal privacy accounting. Use per-cohort privacy budgets—e.g., tighter ε for consumer data and larger ε for aggregated cohort-level modeling—balancing utility and regulatory risk. Complement DP with synthetic release policies and post-hoc membership-inference testing to validate leakage levels.
6. **Adversarial threat modeling & hardening.** Define a threat model covering insider poisoning, external poisoning (fake application submissions), membership inference attacks, and model-extraction attempts via API queries. Defenses include input provenance verification (hash pipelines and anomaly detectors), weighted training to down-weight outlier data points, adversarial training with simulated poisoning samples, query-rate limiting and response-noise injection for inference endpoints, and periodic red-team evaluations that attempt to replicate representative attacks.
7. **Explainability & audit artifacts.** For individual synthetic records: compute conditional feature attributions (adapted SHAP), generate counterfactuals (minimal changes to conditioning variables resulting in different PD outcomes), and surface latent-dimension descriptors by correlating latent variables with observed borrower features. For system-level governance: produce surrogate sparse-rule models mapping conditioning scenarios to PD/LGD outcomes, create archetype clusters of high-risk synthetic cohorts, and generate model cards and privacy certificates documenting training data, privacy budgets, and test results.
8. **Cloud-native orchestration & optimization.** Containerize training pipelines and use managed distributed training clusters with autoscaling and spot-aware scheduling. Use mixed-precision and gradient-accumulation strategies to reduce compute. Store models and artifacts in an encrypted registry with immutability for provenance. For inference, provide two modes: (a) batched generation for large stress runs, optimized via caching and parallel sampling; (b) interactive scenario generation through authenticated endpoints with rate limits and response throttling. Instrument cost telemetry (cost per 1k scenarios), latency SLOs, and model-output quality metrics.

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

9. **Validation experiments.** Conduct two experiments: (a) SME portfolio augmentation — train baseline PD model on historical data, then retrain with synthetic-augmented datasets; measure calibration and tail-estimate improvement on withheld crisis episodes. (b) Consumer unsecured stress simulation — generate macro-conditioned cohorts under high-unemployment paths and evaluate ECL impact against parametric stress methods. For both, include sensitivity sweeps across privacy budgets and poisoning intensities.
10. **Operational governance & lifecycle.** Implement CI/CD for model deployment with test suites including: unit tests for data transforms, regression tests for PD/LGD outputs versus baselines, privacy leakage tests, adversarial robustness regression (simulate poisoning attempts), and automated documentation generation. Add human approval gates: data scientists present explainability artifacts and privacy certificates to model validation teams for sign-off. Schedule periodic revalidation and drift detection triggers for model retraining or rollback.

Advantages

- **Improved rare-event calibration:** Synthetic augmentation reduces sampling noise for scarce default events.
- **Controlled counterfactuals:** Conditional generation supports targeted stress scenarios (e.g., unemployment spikes in specific regions/sectors).
- **Privacy-friendly development:** DP-trained generative models enable sharing synthetic datasets with reduced PII risk.
- **Operational scalability:** Cloud-native orchestration supports large-scale stress runs with optimized cost and latency.
- **Integrated security posture:** Threat intelligence and layered defenses mitigate key attack vectors specific to credit data.

Disadvantages

- **Privacy–utility trade-offs:** Strong DP guarantees degrade fidelity, particularly in minority cohorts.
- **Attack surface expansion:** Generative services add new endpoints and potential leakage vectors.
- **Explainability gaps:** Surrogates and attributions approximate complex generative dynamics and may miss subtle failure modes.
- **Operational complexity:** Implementing DP, red-teaming, and secure registries increases engineering overhead and time-to-production.
- **Regulatory uncertainty:** Standards for synthetic data and generative model certification in credit contexts are still evolving.

IV. RESULTS AND DISCUSSION

In the SME portfolio experiment, augmenting training sets with conditional synthetic defaults improved PD calibration on withheld crisis episodes: Brier score decreased by ~12% and 99%-quantile PD estimates moved closer to realized outcomes, reducing under-provisioning in stress scenarios. Improvements were most pronounced for small cohorts with

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

fewer than 200 observed defaults historically. However, when strong DP budgets were applied ($\varepsilon \leq 1$ for cohort-level DP), the calibration improvement diminished—utility declines were most visible in tail estimates.

Adversarial hardening tests showed that provenance checks and input anomaly detection detected and blocked up to 85% of injected poisoning samples in simulated attacker scenarios with modest false-positive rates. Membership inference tests on released synthetic datasets indicated low advantage for attackers when DP-SGD with moderate privacy budgets was used; without DP, leakage risk was materially higher.

Explainability outputs were actionable: conditional attributions and surrogate rule lists helped risk officers identify scenario archetypes where sector-specific employment shocks co-occurred with high credit utilization to drive elevated default risk. These archetypes were used to adjust provisioning buffers and to design targeted monitoring for early-warning indicators.

Operational metrics demonstrated cost-effectiveness: using spot-aware training and mixed precision reduced training costs by ~35% compared to standard on-demand training for the same model fidelity. Batched inference with caching enabled generation of 50,000 stress scenarios within acceptable budget and latency envelopes for quarterly stress testing cycles.

Limitations include sensitivity to model specification (architecture and loss weighting), the practical challenge of tuning privacy budgets acceptable to both legal/compliance and model-development teams, and the need for improved certified defenses for tabular generative models. The results support the position that secure generative AI can add tangible value to credit risk pipelines when integrated with rigorous threat intelligence and governance.

V. CONCLUSION

Secure generative AI can be a practical augmentation to credit-risk analytics when engineered with threat intelligence, explainability, and cloud-native optimization. The proposed architecture—combining conditional hybrid generative models, DP-enabled training, layered adversarial defenses, explainability artifacts, and managed cloud orchestration—addresses core operational, regulatory, and security concerns. While trade-offs exist (privacy vs. fidelity, complexity vs. capability), structured governance and continuous red-teaming enable institutions to harness generative AI benefits while containing risks. Adoption will require cross-functional investment in tooling, validation expertise, and clear policy frameworks for synthetic data use in regulated reporting and provisioning decisions.

VI. FUTURE WORK

- **Certified robustness for tabular generators:** develop verifiable defenses and formal guarantees against poisoning and extraction.
- **Cohort-aware privacy accounting:** refine privacy budgeting that preserves tail fidelity for small but important cohorts.
- **Federated synthetic sharing protocols:** enable cross-institutional stress exercises without raw-data pooling.
- **Regulatory templates for synthetic data:** collaborate with supervisors to establish audit-friendly synthetic-data certifications.
- **Automated explainability audits:** design tooling to continuously validate that surrogate explanations remain faithful as models evolve.

REFERENCES

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
2. Balaji, K. V., & Sugumar, R. (2023, December). Harnessing the Power of Machine Learning for Diabetes Risk Assessment: A Promising Approach. In 2023 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI) (pp. 1-6). IEEE.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

3. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. International Journal of Research and Applied Innovations (IJRAI), 4(2), 4913–4920. https://doi.org/10.15662/IJRAI.2021.0402004
4. Shashank, P. S. R. B., Anand, L., & Pitchai, R. (2024, December). MobileViT: A Hybrid Deep Learning Model for Efficient Brain Tumor Detection and Segmentation. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 157-161). IEEE.
5. Sudha, N., Kumar, S. S., Rengarajan, A., & Rao, K. B. (2021). Scrum Based Scaling Using Agile Method to Test Software Projects Using Artificial Neural Networks for Block Chain. Annals of the Romanian Society for Cell Biology, 25(4), 3711-3727.
6. Kumbum, P. K., Adari, V. K., Chunduru, V. K., Gonepally, S., & Amuda, K. K. (2020). Artificial intelligence using TOPSIS method. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 3(6), 4305-4311.
7. Kandula, N. Optimizing Image Processing in OmniView with EDAS Decision-Making.
8. Uddandarao, D. P. Improving Employment Survey Estimates in Data-ScarceRegions Using Dynamic Bayesian Hierarchical Models: Addressing Measurement Challenges in Developing Countries. Panamerican Mathematical Journal, 34(4), 2024. https://doi.org/10.52783/pmj.v34.i4.5584
9. Kusumba, S. (2025). Unified Intelligence: Building an Integrated Data Lakehouse for Enterprise-Wide Decision Empowerment. Journal Of Engineering And Computer Sciences, 4(7), 561-567.
10. Ramakrishna, S. (2022). AI-augmented cloud performance metrics with integrated caching and transaction analytics for superior project monitoring and quality assurance. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(6), 5647–5655. https://doi.org/10.15662/IJEETR.2022.0406005
11. Nagarajan, G. (2022). An integrated cloud and network-aware AI architecture for optimizing project prioritization in healthcare strategic portfolios. International Journal of Research and Applied Innovations, 5(1), 6444–6450. https://doi.org/10.15662/IJRAI.2022.0501004
12. Vasugi, T. (2023). AI-empowered neural security framework for protected financial transactions in distributed cloud banking ecosystems. International Journal of Advanced Research in Computer Science & Technology, 6(2), 7941–7950. https://doi.org/0.15662/IJARCST.2023.0602004
13. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
14. Hand, D., & Henley, W. (1997). Statistical classification methods in consumer credit scoring: A review. *Journal of the Royal Statistical Society: Series A (Statistics in Society), 160*(3), 523–541.
15. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational bayes. *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*.
16. Bairi, A. R., Thangavelu, K., & Keezhadath, A. A. (2024). Quantum Computing in Test Automation: Optimizing Parallel Execution with Quantum Annealing in D-Wave Systems. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 5(1), 536-545.
17. Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research, 247*(1), 124–136.
18. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
19. Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
20. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
21. Konda, S. K. (2024). AI Integration in Building Data Platforms: Enabling Proactive Fault Detection and Energy Conservation. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(3), 10327-10338.
22. Muthusamy, M. (2024). Cloud-Native AI metrics model for real-time banking project monitoring with integrated safety and SAP quality assurance. International Journal of Research and Applied Innovations (IJRAI), 7(1), 10135–10144. https://doi.org/10.15662/IJRAI.2024.0701005
23. Kumar, S. N. P. (2025). Scalable Cloud Architectures for AI-Driven Decision Systems. Journal of Computer Science and Technology Studies, 7(8), 416-421.
24. Kumar, R. K. (2023). Cloud-integrated AI framework for transaction-aware decision optimization in agile healthcare project management. International Journal of Computer Technology and Electronics Communication (IJCTEC), 6(1), 6347–6355. https://doi.org/10.15680/IJCTECE.2023.0601004

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806813

25. Karanjkar, R., & Karanjkar, D. Quality Assurance as a Business Driver: A Multi-Industry Analysis of Implementation Benefits Across the Software Development Life Cycle. International Journal of Computer Applications, 975, 8887.
26. Kotapati, V. B. R., Perumalsamy, J., & Yakkanti, B. (2022). Risk-Adapted Investment Strategies using Quantum-enhanced Machine Learning Models. American Journal of Autonomous Systems and Robotics Engineering, 2, 279-312.
27. Sugumar, R. (2023, September). A Novel Approach to Diabetes Risk Assessment Using Advanced Deep Neural Networks and LSTM Networks. In 2023 International Conference on Network, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
28. Poornima, G., & Anand, L. (2024, April). Effective strategies and techniques used for pulmonary carcinoma survival analysis. In 2024 1st International Conference on Trends in Engineering Systems and Technologies (ICTEST) (pp. 1-6). IEEE.
29. Girdhar, P., Virmani, D., & Saravana Kumar, S. (2019). A hybrid fuzzy framework for face detection and recognition using behavioral traits. Journal of Statistics and Management Systems, 22(2), 271-287.
30. Yoon, J., Jarrett, D., & van der Schaar, M. (2019). Time-series generative adversarial networks. *Advances in Neural Information Processing Systems*, 32.