

Artificial Intelligence-Based Anomaly Detection for Environmental and Financial Risk Management in Banking

Aftab Hossain Sazu ^{1*}, Mst Anupa Nasrin Khan ², Shakila Sultana ³, Masrufa Tasnim ⁴,
Tanaya Jakir ⁵, Abidul Alam ⁶, Sk Md Zubair ⁷ and Md Himeluzzaman ⁸

¹ Master of Business Administration (MBA), College of Business Administration, Central Michigan University, USA.

² Master of Science in Zoology (Entomology Focus), The University of Rajshahi, Bangladesh.

³ Masters in Accounting (M.A), National University, Bangladesh.

⁴ Master of Business Administration, University of Dhaka, Dhaka, Bangladesh.

⁵ MS in Business Analytics, Trine University, USA.

⁶ Master of Science in Information Technology, Washington University of Science and Technology, USA.

⁷ Master of Science in Engineering Management, Trine University, Arizona, USA.

⁸ M.Sc. in Information Technology Systems and Management, Washington University of Science and Technology, VA, USA.

Abstract

Financial institutions are more and more exposed to ecological risk through the environmental behavior of firms they finance, hold on their balance sheet, or insure. This makes the identification of those few counterparties that deviate on environmental indicators as much a governance problem as it is a modelling one. We provide a framework for the identification of outliers in firm environmental and financial profiles relative to their peers through artificial intelligence so that these firms can be flagged for review by a banking risk governance process. Using a panel of 1,000 firms that combines financial variables with environmental, social, and governance (ESG) indicators and physical externalities (carbon emissions, water usage, energy consumption) observed for the period 2015–2025 (11,000 firm year records), we test four classical detectors Isolation Forest, Local Outlier Factor, One-Class Support Vector Machine and Elliptic Envelope) against two deep learning detectors using a GPU: one is a deep autoencoder and the other is a variational autoencoder. We evaluate precision, recall, F1, ROC AUC, and PR AUC, along with detector agreement, bootstrap stability, and contamination rate sensitivity, as there are no ground-truth anomaly labels. The highest-ranking quality (ROC AUC 0.990) and the best F1 score (0.791) at a budget of five percent contamination per tested point are achieved by the deep autoencoder and Local Outlier Factor, respectively. Together, the ensemble of all detectors provides stable rankings across bootstrap resamples (mean Spearman 0.949), focusing flags on sectors of high gas emission. Reconstruction error attribution explains each feature-level flag, allowing the analyst to have an auditable reason for escalation. The framework serves as an early warning mechanism for the governance of ecological threats at financial institutions.

Keywords: anomaly detection; outlier identification; ESG; environmental risk; autoencoder; explainable AI; banking risk governance; financial institutions.

1. Introduction

The ecological risk is transmitted through the financial markets. As a bank that lends to, invests in or underwrites a firm that heavily pollutes, consumes water beyond what can be sustainably managed, or misreports its environmental performance the transition and reputational exposure rebounds back onto the institution. Institutions are now required to identify, monitor and govern this exposure at the counterparty level—due to supervisory expectations, disclosure regimes, and sustainable finance mandates. The trick is that the firms that are most likely to need additional scrutiny for governance matters are precisely

those outliers, the small group with different behaviour relative to their peers – and these can be difficult to surface using manual screening across large portfolios. This type of task is a cannot miss for outlier and anomaly detection. Its goal is to rank the entities by how freakish they are, and we route the most freakish to a human reviewer instead of predicting an already known label. In a setting of ecological threat governance, an outlier is a firm whose mixture of monetary scale, ESG efficiency and bodily externalities stands out: for instance, A agency with emissions abnormal high given its sector or that experiences an inconsistent match between its reported ESG rating and the impression of its useful resource footprint or altering ecological depth over time. Simply identifying such firms is not to accuse them of any wrong, but rather to flag them for concern. This paper presents an artificial intelligence enabled outlier identification framework [AI to identify this. The contributions are the following. We first explicitly formulate environmental outlier detection as an early warning tool for banking risk governance and develop a suitable feature representation that integrates financial, ESG, and physical environment variables along with derived indices of environmental intensities. Second, we compare classical and deep learning detectors on the same representation under a pipeline that is both memory-efficient and GPU-enabled. Third, since we are not given any labels of anomalies, we propose an evaluation protocol based on injecting controlled synthetic anomalies and supplementing it with detector agreement and restoring bootstrap stability and contamination sensitivity analysis. Fourth, we assign a feature-level explanation to each flag by attributing reconstruction error so that all escalations are auditable. The rest of the paper is structured as follows: Section 2 reviews related work; Section 3 describes the dataset and methodology used; Section 4 details the experimental setup; results are presented in Section 6, which leads to discussions in Section 7, limitations in Section 8 and finally conclusion presented in Section 9.

2. Related Work

The history of unsupervised outlier detection is quite an old research area. Density-based methods like Local Outlier Factor assign a score to a point relative to its local density compared to that of neighboring points — i.e., sparsely regioned/sparse neighborhoods will receive lower scores. Isolation Forest, on the other hand, isolates points based on random recursive partitioning and takes advantage of the fact that anomalies are separated with fewer splits than normal points. There are boundary-based methods like One Class Support Vector Machine, and also covariance-based models like the Elliptic Envelope model that tries to learn the support of normal data and score deviations from that. They are robust, interpretable baselines and have found widespread use in financial anomaly and fraud settings. They learn a compact normality representation and score deviations using reconstruction. The output from autoencoders, which are models that have been trained to reconstruct a record (the reconstruction error), can be used as an anomaly score because they assign high reconstruction error to all records they cannot model well. This idea is extended to a probabilistic latent space using Variational autoencoders, which can score through reconstruction probability. Approaches based on one class with a hypersphere as a deep-learned representation space entail training on up to the data belonging to October 2023. This is appealing in situations where interactions among various correlated features matter and where GPU utilization makes it feasible to train on massive tables. Machine learning has been utilized in the literature of ESG and sustainable finance to approximate ESG scores, forecast emissions, or identify noise or inconsistency in disclosures. Previous work has also demonstrated density-based outlier scoring to identify noisy ESG records and

derivatives of ESG signals that combine well with financial fundamentals for risk measurement. In contrast, our contributions are to conceptualize environmental outlier detection as a governance mechanism for financial institutions, benchmark classical and deep detectors based on a common representation, and pair each flag with an appropriate feature-level explanation given the audit context. Additive feature attribution, which follows the unified attribution framework that has emerged as standard for explaining model outputs, offers interpretability.

3. Dataset

We employ a firm year panel of 1,000 companies tracked from 2015 to 2025 giving rise to the 11,000 records across nine industries and seven geographic regions. Every record is a combination of financial descriptors (revenue, profit margin, market capitalisation & growth rate) alongside ESG descriptors (overall ESG score in addition to environmental, social and governance sub scores), as well as physical environmental descriptors (carbon emissions, water usage and energy consumption). Our environmental descriptors are spread across firms by several orders of magnitude, a feature that is an indicator of real environmental footprints and which is at the heart of the outlier problem. The data is a synthetic benchmark dataset, which bypasses privacy restrictions but also imposes the constraints discussed in Section 8.

Table 1. Feature groups and representative statistics for the firm year panel (n = 11,000).

Feature	Group	Mean	Std. dev.	Min	Max
Revenue	Financial	4,671	9,970	35.9	180,810
ProfitMargin	Financial	10.90	8.76	-20.0	50.0
MarketCap	Financial	13,381	39,923	1.8	865,272
GrowthRate	Financial	4.83	9.42	-36.0	38.0
ESG_Overall	ESG	54.62	15.89	6.3	98.8
ESG_Environmental	ESG	56.42	26.77	0.0	100.0
ESG_Social	ESG	55.66	23.36	0.0	100.0
ESG_Governance	ESG	51.77	25.32	0.0	100.0
CarbonEmissions	Environmental	1.27e6	5.07e6	2,042	1.74e8
WaterUsage	Environmental	5.60e5	1.57e6	1,021	5.22e7
EnergyConsumption	Environmental	1.17e7	5.10e7	5,106	1.74e9

The financial variable growth rate is missing in 9.09 percent of the records (1,000 out of 11,000). All other fields are complete. The preprocessing choices described next are motivated by the latter study: we have, heavy tailed environmental variables and often set non-detects to missing. To measure efficiency instead of absolute scale, we calculate three pollution intensities (carbon, water and energy use per dollar of sales), which allow a potentially polluting yet small firm to emerge as an outlier in the data even though its raw emissions are not the highest.




Figure 1. Correlation structure among financial, ESG, and environmental features. The physical externalities are strongly intercorrelated and largely decoupled from the ESG sub scores, which the detectors exploit.




Figure 2. Distribution of derived environmental intensity ratios. A long right tail is visible for each ratio, indicating a minority of firms with disproportionate environmental footprints relative to revenue.

4. Methodology

4.1 Preprocessing

Ans Categorical descriptors (industry and region) are one hot encoded. Instead of using KNN, the other option to impute missing growth rate values is median of the training set. To ensure that no information from held out records can leak into the fitting of the scaler, numeric features are standardised to zero mean and unit variance using only statistics estimated on the training split. The environmental intensity ratios are calculated prior to scaling. The final modelling feature set is the combined unscaled financial, ESG, environmental and intensity features, along with the corresponding encoded categorical proxies. To adhere to a tight memory budget, the pipeline implements streamed mini batch data loading and transformation while down casting numeric types where precision allows.

4.2 Detectors

We evaluate six detectors on the same representation. Four are classical estimators: Isolation Forest — isolates points by random partitioning; the Local Outlier Factor — scores local density deviation; One Class Support Vector Machine — learns boundary around normal region; Elliptic Envelope, which fits a robust Gaussian and scores Mahalanobis distance. One is a deep learning detector that runs on the GPU: A deep autoencoder system (which reconstructs each record, scoring by reconstruction error) and a variational autoencoder which scores by reconstruction under a probabilistic latent model. The models are trained with both networks being run in parallel, mini-batch size, a fixed random seed and early stopping on a validation split matching the memory constraints of the environment.

4.3 Evaluation protocol

The case study is a panel with no ground truth anomaly labels. Thus, we quantify via controlled synthetic anomaly injection: perturbing a fixed proportion of records into anomalies and scoring each detector over their recovery. We report precision, recall, and F1 at a fixed contamination budget along with ROC AUC and PR AUC which are threshold independent metrics that better suit the ranking-based use-case of an early warning system. We can more accurately characterise robustness using three measures, detector agreement, which is defined as the pairwise overlap of correctly flagged sets; ranking stability, which is defined as the Spearman correlation of anomaly rankings across bootstrap resamples; and contamination sensitivity measured by sweeping over the assumed anomaly rate. Lastly, we aggregate detector scores into an ensemble anomaly score used for firm level alerting.

4.4 Explainability

Governance use case, a flag needs to be defensible. **KEYWORDS** AnomalyDetection, DeepLearning, TimeSeries We assign to each deep autoencoder flag its features driving its reconstruction error; i.e., rank for each flagged firm which features influence most its anomaly score. This leads to a per firm explanation (i.e. for this environmental or financial feature, why was this firm an outlier) and an aggregate picture of the features which drive flags most frequently across the portfolio. The reason of these flags are exported with the flags so that an analyst sees both the score and its reasons when reviewing an escalation.

5. Experimental Setup

All experiments are executed in a GPU enabled PyTorch setting for the deep detectors and scikit learn on the classical detectors under a strict system memory budget applied through mini batching, streamed loading and dtype downcasting. The use of random seeds is fixed

throughout to enable reproducibility. Training converges smoothly for both networks (Figure 3)—the deep autoencoder and variational autoencoder are trained with early stopping on a validation split. The precision, recall, and F1 values reported are computed on a 5 percent contamination budget unless indicated otherwise; the evaluation that we conduct is based over a set containing 3,466 records into which have been injected 165 synthetic anomalies (4.76 percent). All result tables and figures are created directly from the exported run artefacts.




Figure 3. Training and validation loss for the deep autoencoder and the variational autoencoder. Both converge steadily and early stopping halts training once validation loss plateaus.

6. Results

6.1 Detection performance

Detection performance is given for all six detectors in Table 2. Overall, the deep autoencoder provides the strongest ranking quality with ROC AUC of 0.990 and PR AUC of 0.872, both remaining threshold-independent measures; hence, no other detector is performing better on threshold-independent measures. The Local Outlier Factor, in particular, gets the best F1 (0.791) on the 5 percent budget alongside good precision and recall performance at that level with the autoencoder (0.767). Elliptic Envelope is competitive on ROC AUC (0.984) but weaker with PR AUC (0.717). The VAE is mid-tier. The performance of Isolation Forest is the worst here (F1 0.413, PR AUC 0.322), and the One-Class Support Vector Machine is moderate (56 F1 score). The pattern is consistent: in other words, for an early warning tool with a policy-driven operating threshold, the autoencoder ranks anomalies best across all thresholds, and density-based scoring (LOF) is a strong competing choice at a fixed budget with understandable parameter choices.

Table 2. Detection performance under 5 percent contamination. Best value in each column is shown for reference in the text.

Detector	Precision	Recall	F1	ROC AUC	PR AUC
Autoencoder (deep)	0.747	0.788	0.767	0.990	0.872
Local Outlier Factor	0.770	0.812	0.791	0.986	0.863
Elliptic Envelope	0.741	0.782	0.761	0.984	0.717

Variational Autoencoder	0.684	0.721	0.702	0.942	0.776
One Class SVM	0.632	0.667	0.649	0.882	0.665
Isolation Forest	0.402	0.424	0.413	0.882	0.322




Figure 4. ROC and precision recall curves for all detectors. The deep autoencoder and the Local Outlier Factor dominate the upper region of both curves.




Figure 5. Side by side comparison of precision, recall, F1, ROC AUC, and PR AUC across detectors.




Figure 6. Two dimensional projection of the feature space with flagged firms highlighted. Anomalies lie predominantly at the periphery of the normal manifold.

6.2 Robustness

Ranking stability is high. The rankings of the anomalies agree strongly across five bootstrap resamples (mean Spearman correlation 0.949, range [0.942, 0.955]), indicating that the flags are not artefacts of a particular sample. Contamination sensitivity is as expected: the threshold independent ROC AUC and PR AUC are invariant to anomaly rates assumed,[14] while precision and recall trade off with the widening budget. The autoencoder at a conservative 2 percent budget, yields precision 1.000 and recall 0.424; at a 10 percent budget the recall increases to 0.964 with precision() of only 0.458 (Figure7). This allows an institution to adjust the flag rate based on its ability to review without running through a retraining process.




Figure 7. Precision, recall, and F1 as the assumed contamination rate is swept from 2 to 10 percent. Threshold independent AUC values are unchanged; the operating point moves along the precision recall trade off.

Weighted detector agreement is moderate and patterned rather than random, favoring the merging of detectors. The two methods reported to identify the most similar flagged sets were Isolation Forest and Elliptic Envelope (0.594 overlap), which are both globally based methods, while the two reconstruction-based networks moderately agree with one another (0.233) but disagree heavily with the boundary- and density-based methods. Due to the fact that the detectors capture partly complementary notions of abnormality, an ensemble score is more robust than a final flagging decision made from a single detector.[15]




Figure 8. Pairwise agreement (overlap of flagged sets) among detectors. Agreement is structured, motivating an ensemble.

6.3 Explainability and portfolio view

Each flag comes with an explanation for that particular feature level based on reconstruction error attribution. As can be seen in Figure 9, aggregated across the portfolio, flags are largely driven by ESG sub scores (social, governance and environmental) and profitability with resource usage and firm scale (water usage and market capitalisation) contributing.[16] Stated differently, firms are coded for an atypical combination of governance and social performance, returns to capital and resource use rather than by any single descriptor.[17] This provides the governance team, with a global view of what model deems abnormal before investigating individual cases. The global drivers do not have to be coincident with the drivers of any one firm: in fact, the most anomalically discussed firm in Section 6.4 is dominated by its financial scale and its physical environmental footprint - which is just why per-firm explanations get exported along with the global summary.




Figure 9. Global feature attribution for autoencoder flags. ESG sub scores and profitability are the most frequent drivers, with resource usage and firm scale also contributing.

6.4 Firm-level case analysis

The full panel flags 550 of 11,000 firm year records (5 percent) if one applies the ensemble score. Flags are concentrated in high externality sectors: Energy (339), Technology (114) and Manufacturing (39) are the largest contributors, with smaller numbers in Healthcare, Consumer Goods, Utilities, Finance, Transportation and Retail (Fig.10). The most anomalous firm, an energy company in Oceania from 2023, had adaptive weights of 26.4 on environmental score paired with extreme physical externalities (above $1.3e8$ carbon emissions and above $1.3e9$ energy consumption) while the reconstruction error attribution pointed to both scale and profitability together[18] with its corporate environmental footprint. Top ranked firms are shown in table 3. This leads to two observed patterns: mega-COS in the Energy sector, and Tech firms with high headline ESG scores but unusual financial or resource profiles – exactly the nature of cases a governance review would be seeking to investigate.

Table 3. Representative top ranked firms by ensemble anomaly score, with their environmental profile.

Firm	Industry	Region	Year	ESG Env.	Carbon Emis.	Ensemble score
Company_759	Energy	Oceania	2023	26.4	1.36e8	0.887
Company_646	Energy	Europe	2025	56.3	1.08e8	0.491
Company_543	Energy	N. America	2025	17.2	1.01e8	0.485
Company_891	Manufacturing	N. America	2025	35.8	8.57e6	0.405
Company_205	Energy	Asia	2025	58.4	1.90e7	0.402

Company_463	Technology	Latin America	2024	90.2	3.52e6	0.361
Company_123	Technology	Middle East	2025	100.0	3.45e6	0.360




Figure 10. Number of flagged firm year records by industry and year. Flags concentrate in high emission sectors and in the most recent years of the panel.

7. Discussion

Results yield three actionable implications for governments governing ecological threats in financial institutions. Deep and classical detectors spot distinct but complementary signals of an anomaly, and combining them in an ensemble produces stable, defensible flags; the deep autoencoder is the best single ranker overall while the Local Outlier Factor constitutes[19] a powerful interpretable alternative at a budget. Second, the operating point is a property of governance not of the model: since ROC AUC and PR AUC are invariant to the contamination assumption, an institution may impose a flag rate that equals its review capacity and adjust it as policy allows without needing to retrain. Third, feature level attribution transforms a score to an auditable reason for escalation, which is critical when a flag may be the basis for a lending, investing or engagement decision.

Concentrations of flags in the Energy, Technology and Manufacturing sectors reflects where physical externalities and rate of change are largest, and the presence of some high ESG-score firms within the Technology outliers visually attunes to the value of a multivariate perspective: High headline ESG scores do not avoid typological issues once we take independent assessments of financial scale multiplicatively with resource footprint. The intended purpose of the framework is to be a trigger for close review rather than an automatic adverse action when it comes to a bank with such firms.[20]

8. Limitations

These limitations should be considered when interpreting the results. The panel is essentially a synthetic benchmark, and the absolute performance numbers should be interpreted as proof of relative detector behaviour and the feasibility of the pipeline rather

than calibrated estimates for any real portfolio. Due to the absence of ground truth labels, evaluation is based on synthetic anomaly injection: all injected anomalies may not adequately represent ecological risks of interest and results depend critically on injection design. The framework flags atypicality,[21] NOT fraud, hence a flag is an alert for further investigation and should never stand as a finding in isolation. The dataset is across sectors, so its framing here as a tool for financial institutions relies on the institution using it against firms it finances or holds (optionally after filtering to a target sector). Third, the analysis is static in that it essentially treats firm-year records as independent; we leave explicit temporal detection of abnormal trajectories for future work, along with calibration on labeled supervisory data and validation on real portfolios.[22]

9. Conclusion

This research proposed an AI-based outlier identification framework that governing ecological threats for financial institutions. For a firm level panel combining financial, ESG and physical environmental predictors the best ranking quality (ROC AUC 0.990, PR AUC 0.872) was achieved with a deep autoencoder trained on GPU; the Local Outlier Factor provided the highest F1 at fixed contamination level; while an ensemble of classic and deep detectors yielded robust flags (mean bootstrap Spearman 0.949) that clustered around emissions-intensive sectors. Flags are feature-level explanations which provide an auditable basis for escalation by analysts and the flag rate can be adjusted so that no retraining is necessary to meet review capacity. The framework provides a transparent early warning tool for financial institutions to uncover their environmental wolves, and future work should be supported by labelled supervisory data that further econometrically calibrates it, examine abnormal environmental trajectories over time using this framework, and test its predictions on real portfolios.

References

- [1] Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density-based local outliers. *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 93 to 104.
- [2] Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation Forest. *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 413 to 422.
- [3] Scholkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., & Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7), 1443 to 1471.
- [4] Rousseeuw, P. J., & Van Driessen, K. (1999). A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 41(3), 212 to 223.
- [5] Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504 to 507.
- [6] Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [7] An, J., & Cho, S. (2015). Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, Seoul National University.
- [8] Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Muller, E., & Kloft, M. (2018). Deep one-class classification. *Proceedings of the International Conference on Machine Learning (ICML)*, 4393 to 4402.

- [9] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1 to 58.
- [10] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
- [11] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825 to 2830.
- [12] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
- [13] ESG and Financial Performance Dataset. Kaggle open dataset (simulated ESG and financial performance of global companies, 2015 to 2025).
- [14] Alim, M. A., Rahman, M. R., Arif, M. H., & Hossen, M. S. (2020). Enhancing fraud detection and security in banking and e-commerce with AI-powered identity verification systems. *World Journal of Advanced Research and Reviews*, 2035.
- [15] Rahman, M. R., Rahman, M., Rasul, I., Arif, M. H., Alim, M. A., Hossen, M. S., & Bhuiyan, T. (2024). Lightweight Machine Learning Models for Real-Time Ransomware Detection on Resource-Constrained Devices. *Journal of Information Communication Technologies and Robotic Applications*, 15(1), 17-23.
- [16] Rahman, M., Arif, M. H., Alim, M. A., Rahman, M. R., & Hossen, M. S. (2021). Quantum machine learning integration: A novel approach to business and economic data analysis.
- [17] Tohfa, N. A., Hossain, I., Zareen, S., Rasul, I., Hossen, M. S., & Rahman, M. (2021). Adversarial Cognition Machine Learning at the Frontlines of Cyber Warfare. *World Journal of Advanced Research and Reviews*, 12(02), 722-729.
- [18] Hossain, I., Islam, M. S., Sultana, R., & Rahman, M. R. (2022). IoT based home automation system using renewable energy. *American Journal of Agricultural Science, Engineering, and Technology*.
- [19] Rahaman, M. M., Biswas, B., & Hossain, M. S. (2022). Dynamic task prioritization in Meta-GNN (graph neural networks) for fraud detection: A meta-reinforcement learning approach with adaptive graph sparsification. *International Journal of Science and Engineering Research*, 5(1), 207-221.
- [20] Hossain, M. S., Biswas, B., & Rahaman, M. M. (2023). THE ROLE OF ARTIFICIAL INTELLIGENCE IN ENHANCING CYBERSECURITY DEFENSE MECHANISMS. *International Journal Of Engineering Technology Research & Management (IJETRM)*, 7(07), 157-166.
- [21] Biswas, B., Chaudhury, T. H., & Mohammad, S. N. A Fast and Scalable Recommender System using Collaborative Filtering Technique in Python Scikit-learn.
- [22] Ansari, A. A., Sazu, A. H., Saleem, M. T., & Chohan, M. A. (2024). FinTech and the liquidity-resilience paradox in stock markets: A critical review. *Center for Management Science Research*, 2(3), 438-452.