Early Identification of Deforestation using Anomaly Detection

Nethmi Wijesinghe
Faculty of Information Technology
University of Moratuwa
Moratuwa, Sri Lanka
nethmiw.17@itfac.mrt.ac.lk

Rashmi Perera
Faculty of Information Technology
University of Moratuwa
Moratuwa, Sri Lanka
rashmiperera1997@gmail.com

Nethmee Sellahewa
Faculty of Information Technology
University of Moratuwa
Moratuwa, Sri Lanka
nssellahewa@gmail.com

Priyanga Dilini Talagala
Department of Computational Mathematics
University of Moratuwa
Moratuwa, Sri Lanka
priyangad@uom.lk

2023 8th International Conference on Information Technology Research (ICITR) | 979-8-3503-5950-3/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICITR61062.2023.10382919

***Abstract*—Research involving anomaly detection in image streams has seen growth through the years, given the proliferation of high-quality image data in various applications. One such application that is in urgent need of attention is deforestation. Detecting anomalies in this context, however, remains challenging due to the irregular and low-probability nature of deforestation events. This study introduces two anomaly detection frameworks utilizing machine learning and deep learning for the early detection of deforestation activities in image streams. Furthermore, Explainable AI was used to explain the black box models of the deep learning-based anomaly detection framework. The class imbalance problem, the inter-dependency between the images with time, the lack of available labelled images, a data-driven anomalous threshold, and the trade-off of accuracy while increasing interpretability in the black box optimization methods are some key aspects considered in the model-building process. Our novel framework for anomaly detection in image streams underwent rigorous evaluation using a range of datasets that included synthetic and real-world data, notably datasets related to Amazon's forest coverage. The objective of this evaluation was to detect occurrences of deforestation in the Amazon. Several metrics were used to evaluate the performance of the proposed framework.**

***Index Terms*—Anomaly Detection, Image Time Series, Machine Learning, Deforestation, Explainable AI**

I. Introduction

Deforestation is a pressing global issue driven by various factors such as agriculture, urbanization, and logging, resulting in far-reaching ecological, environmental, and socio-economic consequences [1]. Detecting deforestation in its early stages is critical for effective conservation and mitigation efforts. Traditional methods, however, exhibit limitations in terms of scalability, accuracy, and real-time responsiveness [2]. Timely detection and intervention are essential for sustainable forest conservation and sustainable land management. Recent technological advancements have led to the proliferation of high-resolution satellite imagery and remote sensing capabilities, which have opened up new avenues for early detection and monitoring of deforestation, offering valuable insights into the dynamic nature of forest ecosystems [3]. Nevertheless, monitoring extensive forested areas for signs of deforestation is a challenging task, given the sheer volume of satellite and aerial images generated daily. The timely identification of deforestation events is pivotal to the success of conservation initiatives. In recent years, the integration of remote sensing technologies, coupled with machine learning techniques, has demonstrated promising potential in the monitoring of deforestation activities. This paper introduces a novel framework for the early identification of deforestation using anomaly detection in image streams, addressing several key research gaps within the existing literature. One of the primary challenges in anomaly detection, as highlighted by previous research, is the class imbalance problem [4] [5] [6]. In the context of deforestation monitoring, anomalies (i.e., deforestation events) are often rare compared to normal or typical scenes in image streams. To mitigate this issue, the authors propose a novel strategy: the treatment of anomaly detection as a one-class classification problem. This approach centres on modelling the typical behaviour, i.e., non-deforestation situations, within forested areas. Subsequently, the framework is trained accordingly. This novel approach ensures that the system is better equipped to identify deforestation events accurately, addressing the class imbalance challenge. Moreover, traditional anomaly detection methods often rely on manual thresholding techniques, making them less adaptive and subject to unrealistic assumptions [7]. In contrast, this study introduces a data-driven anomalous threshold based on the Extreme Value Theory (EVT) [8]. This threshold is dynamically adjusted to accommodate the evolving characteristics of the image stream, improving the robustness and accuracy of the anomaly detection system and making it better suited to the dynamic landscape of deforestation monitoring. Acknowledging the temporal interdepen-

RETINA research and innovation project funded by UNESCO and the International Development Research Centre, Ottawa
979-8-3503-5950-3/23/$31.00 ©2023 IEEE

dencies within image streams, especially as forested areas evolve over time, is essential [9]. To tackle this challenge, the authors incorporated time-series forecasting techniques into the framework. The integration enables the capture of temporal patterns and trends associated with deforestation activities. Consequently, the system gains heightened sensitivity to early-stage deforestation events, facilitating timely intervention and mitigation strategies. In addition, the study addresses a prevalent issue in the field of anomaly detection - the dearth of explainability in deep learning models [10]. Such models, including Convolutional Neural Networks (CNNs), often operate as enigmatic "black boxes," complicating the interpretation of their decision-making processes. Recognizing the importance of transparency and interpretability, particularly for stakeholders engaged in deforestation monitoring and management, the authors integrate Explainable AI (XAI) techniques into their framework. This allows for valuable insights into the decision-making process of the framework, making it more accessible and trustworthy for stakeholders involved in deforestation monitoring and management.

The remainder of this paper is organized as follows: Section II provides an overview of the proposed framework and implementation approaches taken in this research, followed by Section III, Results and Discussions which details the evaluations of the research. Finally, Section IV, Conclusion, provides a concise and clear summary.

II. METHODOLOGY

A. Overview

In this study, we characterized a deforestation situation as an anomaly, denoting it as an occurrence that exhibits an exceedingly low likelihood within the projected distribution. The proposed framework is based on two main assumptions: first, that deforestation situations manifest as substantial deviations from the typical behaviour of the monitored ecosystem; and second, that a representative dataset encapsulating the typical behaviour of said ecosystem is available to define a model for the typical behaviour of the image streams generated within the ecosystem under consideration. In this study, two distinct anomaly detection frameworks were developed to identify deforestation, employing machine learning and deep learning methodologies. The primary objective was to conduct a comprehensive comparison of these technologies to determine the optimal approach for early deforestation identification. Furthermore, Explainable AI was integrated into the deep learning-based framework to interpret the reasoning behind the results generated by the opaque deep learning models. The overall proposed framework is demonstrated in Figure 1.

B. Conventional Machine Learning and Deep Learning based Anomaly Detection

In this study, two distinct anomaly detection frameworks were developed to identify deforestation, employing machine learning and deep learning methodologies. The primary objective was to conduct a comprehensive comparison of these technologies to determine the optimal approach for early deforestation identification. Furthermore, Explainable AI was integrated into the deep learning-based framework to interpret the reasoning behind the results generated by the opaque deep learning models. Both conventional machine learning and deep learning approaches for anomaly detection consist of three main components: computer vision, univariate time series forecasting, and an unsupervised anomaly detection component. Each step followed in two frameworks is discussed in detail below.

1) **Pre-processing:** This initial step in the methodology involves foundational activities that play a pivotal role in enhancing data quality, reducing noise, and facilitating data preparation for precise modelling. In both frameworks discussed in this paper, a standardized set of pre-processing tasks, including image cropping, resizing, and normalization, was particularly applied as an essential preparatory step before advancing to subsequent stages of analysis. Following the execution of these pre-processing tasks, the resulting processed image series was subsequently input into the feature extraction component.
2) **Feature extraction using computer vision:** The prime objective of the feature extraction component is to enhance image recognition and facilitate automated detection of prominent features without any human intervention. Within the conventional machine learning approach, a multifaceted feature extraction strategy is employed, encompassing Gabor Wavelet features, edge detection features, first-order statistical features, and Gray-level Co-Occurrence Matrix (GLCM) feature extraction techniques. These techniques are instrumental in capturing the dynamic nature associated with anomalous behaviour. In contrast, within the deep learning framework, the feature extraction process utilizes several pre-trained Convolutional Neural Network (CNN) architectures, specifically VGG 16, ResNet 50, Inception V3, DenseNet 121, and Xception. However, a pivotal adjustment is made by removing the top-level, fully connected layers to eliminate their inherent classification capabilities. This modification ensures that these networks are served exclusively for feature extraction purposes, thus preserving their capacity to uncover and extract prominent and complex patterns within the image data.
3) **Dimensionality reduction:** Following the feature extraction phase, the resultant feature set yields a considerable number of features for every image within the image series dataset. The management of such a high-dimensional feature set raises concerns, notably the Curse of Dimensionality. To mitigate these challenges, Principal Component Analysis (PCA) is employed within the machine learning module. Meanwhile, the Deep Learning module incorporates a Global Max Pooling layer, and subsequently, the average feature vector is computed for each image, resulting in a sub-

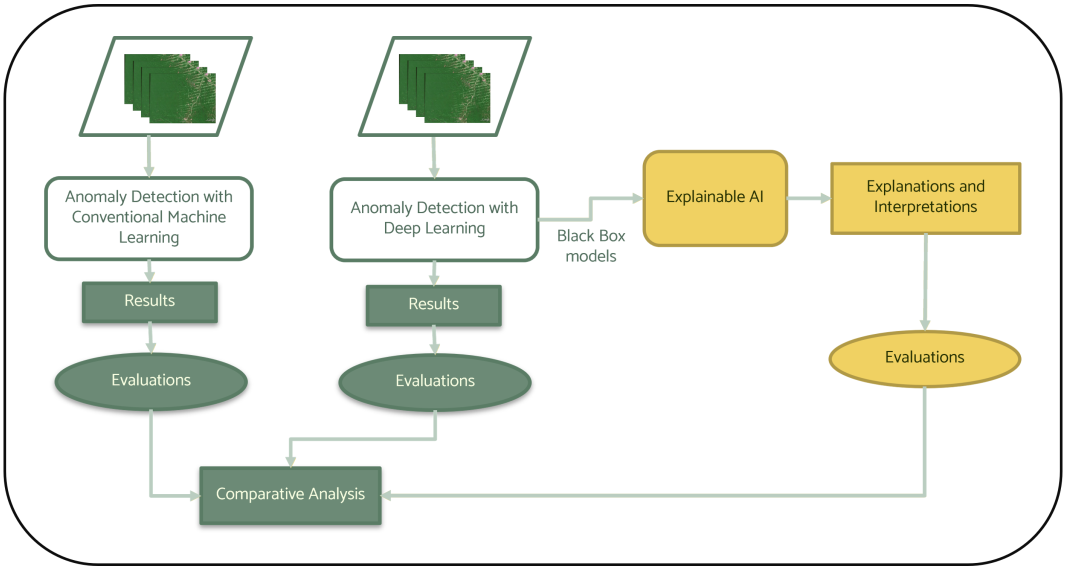

Fig. 1. An overview of the anomaly detection framework architecture.

stantial reduction of dimensionality. These techniques are utilized to obtain a univariate feature vector for each image.

4) **Time series forecasting:** Time-series forecasting is a statistical technique used to predict future values or trends in chronologically ordered data points, i.e., time-series data. The univariate dataset, post-dimensionality reduction, is split into training and test datasets. The training set is employed to train the time-series forecasting model, which in the case of Machine Learning employs the Auto-Autoregressive Integrated Moving Average (ARIMA) model, while the Deep Learning model leverages a Long Short Term Memory (LSTM). The model yields fitted values are residuals for the training set, along with forecast error series from the test dataset. Residuals denote the disparity between actual training values and the fitted values, while the forecast error series represents the deviation between actual test values and forecasted values.
5) **Data driven anomalous threshold calculation:** In the subsequent stage of model development, the residuals obtained from the previous step play a pivotal role in establishing a data-driven threshold, a critical element for classification purposes. The selected approach for threshold computation is firmly grounded in EVT. In contrast to the conventional box plot-based threshold calculation method, which is based on unrealistic assumptions, the EVT-based approach hinges on the fundamental principle that anomalies, resembling rare events, share similarities with extreme values. In this approach, a child distribution for generalized extreme values is plotted, and a threshold calculation is subsequently conducted.
6) **Binary Classification:** The final stage of the methodology includes binary classification, wherein data points within the forecasted series that surpass the upper anomalous threshold or dip below the lower anomalous threshold are categorized as instances of deforestation. Conversely, data points residing within the calculated threshold are designated as representative of typical behaviour (i.e., non-deforestation situations).

C. Explainable AI for Interpretation of Deep Learning-based Framework

Even though the pre-trained CNN-based feature extractors play a significant role in directing the overall anomaly detection module towards its success or failure, it is very much a black-box operation, rendering them less interpretable. This lack of transparency can impede user understanding, trust, and validation, undermining accountability. To address this issue, Explainable AI techniques are leveraged to facilitate the interpretation of predictions generated by the deep learning-based anomaly detection models. In this module, using SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME) methods, high importance is given to explaining the feature extraction and its contribution to the outcome of the anomaly detection module by evaluating the validity of the importance given to the extracted features by the anomaly detection model.

III. Results and Discussions

Both the proposed anomaly detection systems were evaluated for their prediction accuracy using various metrics, including accuracy, F1-score, sensitivity, and specificity. These metrics were used in the comparative analysis between the machine learning-based anomaly detection system and the deep learning-based anomaly detection system. The machine learning-based anomaly detection framework performed fairly well for the small dataset and displayed comparatively low

TABLE I
PERFORMANCE COMPARISON OF THE MACHINE LEARNING-BASED ANOMALY DETECTION FRAMEWORK ON FOREST COVERAGE DATASETS.

	Small Dataset	**Large Dataset**
	Box Plot Threshold Calculation Technique	
Accuracy	70.00%	70.67%
Sensitivity	0.84	1.0
Specificity	0.45	0.12
F1-Score	0.78	0.82
	EVT Threshold Calculation Technique (with 95% confidence level)	
Accuracy	83.33%	68.8%
Sensitivity	0.79	0.53
Specificity	0.91	1.0
F1-Score	0.86	0.69

TABLE II
PERFORMANCE COMPARISON OF THE DEEP LEARNING-BASED ANOMALY DETECTION FRAMEWORK USING VARIOUS CNN EXTRACTORS ON A SMALL FOREST COVERAGE DATASET.

CNN Architecture	**InceptionV3**	**ResNet50**	**VGG16**	**Xception**	**DenseNet121**
	Box Plot Threshold Calculation Technique				
Accuracy	71.43%	96.43%	92.86%	67.86%	96.43%
Sensitivity	0.6	0.95	0.95	0.55	0.95
Specificity	1.0	1.0	0.875	1.0	1.0
F1-Score	0.750	0.974	0.950	0.710	0.974
	EVT Threshold Calculation Technique (with 95% confidence level)				
Accuracy	71.43%	96.43%	96.43%	85.71%	96.42%
Sensitivity	0.6	0.95	0.95	0.8	0.95
Specificity	1.0	1.0	1.0	1.0	1.0
F1-Score	0.750	0.974	0.974	0.889	0.974

TABLE III
PERFORMANCE COMPARISON OF THE DEEP LEARNING-BASED ANOMALY DETECTION FRAMEWORK USING VARIOUS CNN EXTRACTORS ON A LARGE FOREST COVERAGE DATASET.

CNN Architecture	**InceptionV3**	**ResNet50**	**VGG16**	**Xception**	**DenseNet121**
Box Plot Threshold Calculation Technique					
Accuracy	94.39%	87.58%	67.16%	70.09%	67.02%
Sensitivity	1.0	1.0	1.0	0.958	1.0
Specificity	0.832	0.628	0.016	0.188	0.012
F1-Score	0.960	0.915	0.802	0.810	0.802
EVT Threshold Calculation Technique (with 95% confidence level)					
Accuracy	99.60%	92.92%	79.57%	70.09%	79.57%
Sensitivity	1.0	1.0	1.0	0.958	1.0
Specificity	0.988	0.788	0.388	0.188	0.388
F1-Score	0.997	0.950	0.867	0.810	0.867

performance under both threshold calculation techniques for the large dataset. However, the results obtained from the machine learning framework were less impressive than the performance of the deep learning-based anomaly detection framework, which scored good and consistent accuracy. An overview of the performance results of both frameworks is given in Table I, II, and III.

Both frameworks performed better with the EVT-based threshold than the box plot-based approach. Therefore, out of all the frameworks, the VGG-16-LSTM-based deep learning model and the ResNet-50-LSTM-based deep learning model along with the EVT-based threshold calculation worked best for the small dataset. As for the large dataset, the InceptionV3-LSTM-based deep learning anomaly detection model along with the EVT-based threshold was selected as the best framework. Both final frameworks resulted in good accuracy scores and F1 scores, indicating that the proposed frameworks are befitted for imbalanced datasets. Additionally, it is visible that the deep learning-based anomaly detection framework has a very low time complexity compared to the machine learning-based anomaly detection framework.

The two XAI methods were used to analyze the accuracy of feature extraction and their relevance to the final output of the InceptionV3, ResNet50, and VGG16-based feature extractors. In both SHAP and LIME explanations, InceptionV3 and ResNet50 identified the most appropriate regions as major contributing areas when classifying an image as typical or not, compared to the VGG16 model. Furthermore, SHAP

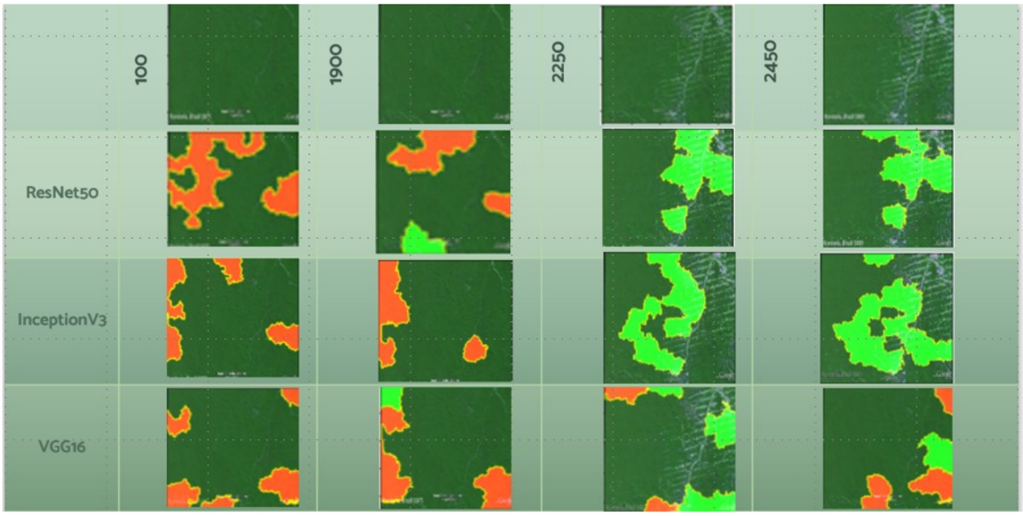

Fig. 2. Performance comparison of the LIME explanations on a large forest coverage dataset.

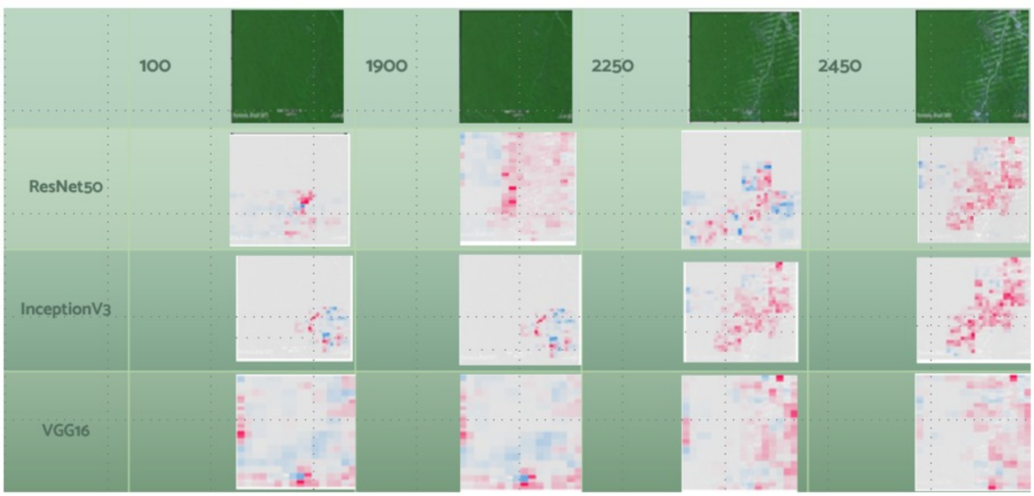

Fig. 3. Performance comparison of the SHAP explanations on a large forest coverage dataset.

predictions also indicated the degree of contribution of each result. An overview of the performance results of XAI methods are given in Figures 2 and 3.

IV. Conclusion

Comparing the different methodologies used in the research, the authors conclude that deep learning works better, the deep learning approach is faster compared to the machine learning approach, and the EVT threshold calculation provides more accurate results than the boxplot-based threshold. The XAI module further demonstrates that the InceptionV3 and ResNet50 models outperform in extracting the most relevant features.

In the research analysis of the performance of machine learning, deep learning, and Explainable AI within this research, the findings lead to the following conclusions; Firstly, in terms of performance, deep learning can be deemed to be performing better since it displayed higher accuracy and significantly faster processing when compared with the machine learning approach. However, it can be noted that the machine learning approach is more interpretable than the deep learning module, which can be identified as an advantage and would work fairly well for small datasets.

When evaluating the threshold calculation techniques, it is evident that the Extreme Value Theory (EVT) approach consistently yields more precise results in contrast to the box plot-based threshold method. As anomalies have a low probability of happening, extreme values can be considered to represent anomalies, and EVT was able to demonstrate its competency when it comes to identifying and addressing extreme values within the data. Additionally, the XAI analysis shows the exceptional performance of the InceptionV3 and ResNet50 models in feature extraction, outperforming other deep learning models in capturing the most important features.

While deep learning models used in this research consisted of black-box models and architectures, such as the CNN architectures, which are known for not being interpretable, it is noteworthy that their interpretability can be improved through the incorporation of explainable AI techniques. The comprehensive evaluations done in this research confirm that the CNN architectures, responsible for feature extraction within the deep learning framework, have performed well in identifying and isolating relevant features pertaining to deforestation.

As for future work, the authors are planning to make several improvements to the frameworks to increase their performance. The machine learning module can sport a more optimized feature extraction and selection procedure that would increase its performance when processing larger datasets. Likewise, the deep learning module can potentially be optimized to accommodate smaller datasets as well, and good hyperparameter tuning methods can be assisted to increase the performance of the framework. The XAI module can be further improved by accommodating standard quantitative evaluation techniques such as BAM (Benchmark Attribution Method) which is highly demanded in the state of art. Furthermore, a back propagation-based XAI technique such as Layer Wise Relevance Propagation (LRP) can be implemented in comparison to the existing perturbation-based XAI techniques used in this study.

Acknowledgment

This work was carried out with the aid of a grant from UNESCO and the International Development Research Centre, Ottawa, Canada. The Views expressed herein do not unnecessarily represent those of UNESCO, IDRC or its Board of Governors.

References

[1] M. Albughdadi, D. Kouamé, G. Rieu, and J.-Y. Tourneret, "Missing data reconstruction and anomaly detection in crop development using agronomic indicators derived from multispectral satellite images," in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 5081–5084, IEEE, 2017.

[2] O. C. Stickler, T. R. Powell, M. R. Coe, M. P. Anderson, K. M. Dalrymple, P. V. Soares-Filho, P. B. Oliveira, A. A. Nogueira, P. L. Alencar, and P. M. Fearnside, "Limitations of traditional methods for detecting deforestation in the amazon rainforest," *Environmental Research Letters*, vol. 11, no. 5, p. 054006, 2016.

[3] A. Sánchez-Azofeifa, B. Rivard, J. Wright, J.-L. Feng, P. Li, M. M. Chong, and S. A. Bohlman, "Estimation of the distribution of tabebuia guayacan (bignoniaceae) using high-resolution remote sensing imagery," *Sensors*, vol. 11, no. 4, pp. 3831–3851, 2011.

[4] S. Rodda and U. S. R. Erothi, "Class imbalance problem in the network intrusion detection systems," in *2016 international conference on electrical, electronics, and optimization techniques (ICEEOT)*, pp. 2685–2688, Ieee, 2016.

[5] J. Kong, W. Kowalczyk, S. Menzel, and T. Bäck, "Improving imbalanced classification by anomaly detection," in *International Conference on Parallel Problem Solving from Nature*, pp. 512–523, Springer, 2020.

[6] D. Devi, S. K. Biswas, and B. Purkayastha, "Learning in presence of class imbalance and class overlapping by using one-class svm and undersampling technique," *Connection Science*, vol. 31, no. 2, pp. 105–142, 2019.

[7] E. Vignotto and S. Engelke, "Extreme value theory for anomaly detection–the gpd classifier," *Extremes*, vol. 23, no. 4, pp. 501–520, 2020.

[8] L. Wilkinson, "Visualizing big data outliers through distributed aggregation," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 256–266, 2017.

[9] S. A. Sumon, R. Goni, N. B. Hashem, T. Shahria, and R. M. Rahman, "Violence detection by pretrained modules with different deep learning approaches," *Vietnam Journal of Computer Science*, vol. 7, no. 01, pp. 19–40, 2020.

[10] R. Marcinkevičs and J. E. Vogt, "Interpretability and explainability: A machine learning zoo mini-tour," *arXiv preprint arXiv:2012.01805*, 2020.