AnoGAN for Tabular Data: A Novel Approach to Anomaly Detection

Pavan Reddy[1] and Aditya Singh[1]

George Washington University, Washington, DC 20052, USA
{pavan.reddy,adityasingh}@gwu.edu

Abstract. Anomaly detection, a critical facet in data analysis, involves identifying patterns that deviate from expected behavior. This research addresses the complexities inherent in anomaly detection, exploring challenges and adapting to sophisticated malicious activities. With applications spanning cybersecurity, healthcare, finance, and surveillance, anomalies often signify critical information or potential threats. Inspired by the success of Anomaly Generative Adversarial Network (AnoGAN) in image domains, our research extends its principles to tabular data. Our contributions include adapting AnoGAN's principles to a new domain and promising advancements in detecting previously undetectable anomalies. This paper delves into the multifaceted nature of anomaly detection, considering the dynamic evolution of normal behavior, context-dependent anomaly definitions, and data-related challenges like noise and imbalances.

Keywords: AnoGAN · Tabular data · Anomaly detetction

1 Introduction

Anomaly detection identifies unexpected patterns in data across various domains, often referred to as anomalies, outliers, or contaminants. It finds applications in cybersecurity, safety-critical system monitoring, fraud detection in finance and healthcare, and military surveillance for non-traditional enemy activities.

Anomaly detection, distinct from noise accommodation and removal [5], addresses unwanted noise in data. Noise, defined as irrelevant data phenomena hindering interpretation, requires elimination before analysis. Conversely, noise accommodation shields statistical model estimates from outlier impacts.

Anomaly detection plays a vital role across various domains by uncovering crucial insights from data anomalies [19, 20, 33, 9, 36, 14, 17]. An unusual network traffic pattern [2] may indicate a security breach, while abnormal MRI scans [31] could signal the presence of tumors. Anomalies in aviation sensors [3] may highlight potential aircraft component issues, and deviations in credit card transactions often signify fraudulent activity. Anomaly detection finds practical

[1] Both authors contributed equally to this research.

applications in manufacturing, finance, and medical imaging, relying on models to identify abnormal patterns amidst regular data. Despite extensive research, managing complex, high-dimensional data remains a challenge. Various communities have developed specialized anomaly detection techniques tailored to specific domains.

Generative Adversarial Networks (GANs), introduced by Ian Goodfellow [10] and colleagues, have emerged as a powerful modeling approach for handling high-dimensional data. Anomaly Generative Adversarial Network (AnoGAN) integrates traditional anomaly detection methods with GAN architecture, enabling it to generate data while simultaneously learning typical data properties for anomaly detection.

Our main contributions in this paper can be summarized as follows:

- How do different reconstruction errors affect the performance of GAN for anomaly detection in tabular data, and what thresholds or criteria can be established for effective detection?
- How does the performance of a GAN-based anomaly detection model with the optimal reconstruction error, trained on single-class data, compare to traditional anomaly detection methods in terms of accuracy and efficiency?

1.1 Challenges in Anomaly Detection

Identifying anomalies in datasets presents several challenges. Firstly, distinguishing abnormal patterns from normal ones requires algorithms that minimize false positives while effectively detecting anomalies. Defining "normal" behavior is complex and evolves over time, necessitating adaptable baseline models. Anomaly detection must also adapt to recognize malicious activities acing as normal behavior, balancing adaptability and resilience. Furthermore, normal behavior changes dynamically, requiring systems that learn and update accordingly. Context-dependent anomaly definitions add another layer of complexity, demanding models that contextualize data based on specific circumstances. Data-related challenges, including handling noisy or imbalanced datasets, addressing blind spots, and implementing preprocessing techniques, further complicate anomaly detection. Finally, managing the general complexity inherent in anomaly detection, such as balancing sensitivity and specificity and adapting to changing conditions, requires expertise in machine learning, data analysis, and domain knowledge.

1.2 Overview of GANs

The introduction of Generative Adversarial Networks (GANs) (Figure 2) by Ian Goodfellow and associates [10] offers a strong modeling approach for addressing the problem of high-dimensional data. Two adversarial networks, a generator (G) and a discriminator (D) are involved in conventional GANs. While D learns to distinguish between actual data and samples provided by G, G is in charge of modeling the data by learning a mapping from latent random variables z (derived from Gaussian or uniform distributions) to the data space. GANs are seeing increasing use in speech and medical imaging, where they have shown empirical success as natural image models.

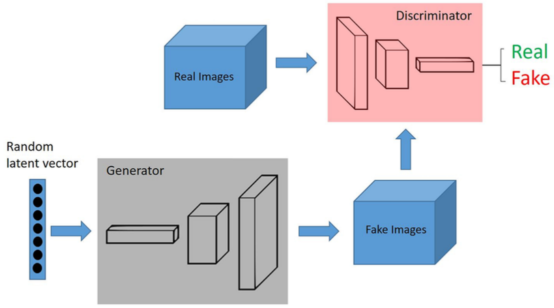

Fig. 1. GANs, or Generative Adversarial Networks, are intricate deep neural network structures consisting of two networks, namely the generator and discriminator. These networks work in opposition to each other, which is why they are called "adversarial." The generator accepts random numbers as input and produces an image. This generated image is then presented to the discriminator, along with a continuous stream of images sourced from the genuine, ground-truth dataset. [**?**]

Anomaly Generative Adversarial Network (Figure 3), or AnoGAN [29], is a cutting-edge anomaly detection method that blends features of conventional anomaly detection approaches with generative adversarial networks (GANs). The GAN architecture is expanded to include anomaly detection with AnoGAN, which enables the model to produce data while concurrently learning the properties of typical data to detect anomalies.

AnoGAN's versatility enables its application across diverse domains for anomaly detection. Particularly in medical imaging, finance, and cybersecurity, it excels in identifying irregular patterns. In finance, it detects fraud by recognizing unusual transaction behaviors, while in cybersecurity, it spots anomalies in network traffic. Its adaptability to time series data facilitates predictive maintenance, anticipating equipment breakdowns. In text analysis, it identifies abnormal language patterns. Overall, AnoGAN offers a flexible solution for anomaly detection in various industries, from healthcare to industrial process monitoring.

2 Related Work

In recent years, numerous approaches have emerged in the field of deep learning-based anomaly detection and object detection [16, 18, 6, 39, 1, 21, 13, 15, 24, 35, 8, 41, 32, 22, 11, 26, 27, 25, 28, 23, 12]. Researchers have [4] comprehensively reviewed deep learning techniques for anomaly detection, emphasizing the promise and adaptability of these models across diverse data types. Their work echoes the sentiment of Zhou and Paffenroth [40], who proposed an anomaly detection method

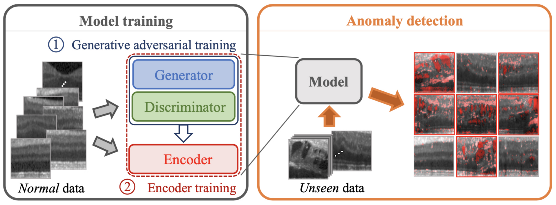

Fig. 2. Framework for identifying anomalies. This framework involves two main stages of model training: generative adversarial training, which results in a trained generator and discriminator, and encoder training, which yields a trained encoder. Both of these training phases are executed using normal or "healthy" data. Subsequently, the framework is used for anomaly detection, where it is applied to both unseen healthy cases and anomalous data. [29]

utilizing autoencoders, this method employs neural network architectures to reconstruct input data. Anomalies are identified when the reconstruction error surpasses a predefined threshold, enabling effective detection of deviations from normal patterns. They demonstrated their efficacy in capturing non-linear transformations and achieving impressive detection rates. Further, Ruff et al. extended the applicability of one-class classification using deep neural networks, shedding light on the strength of deep architectures in isolating normal from anomalous patterns. Their approach extends the scope of one-class classification through the integration of deep neural networks. The model is trained on normal instances, enhancing its ability to discern anomalies by learning complex patterns inherent in the normal data distribution. In the context of Generative Adversarial Networks (GANs), Zenati [37] introduced an adversarially-trained one-class classifier, which proved pivotal in benchmark datasets. They introduced an adversarial element to one-class classification, this method employs a classifier trained to distinguish between normal and anomalous samples. Adversarial training enhances the model's robustness, making it more adept at discerning subtle anomalies.

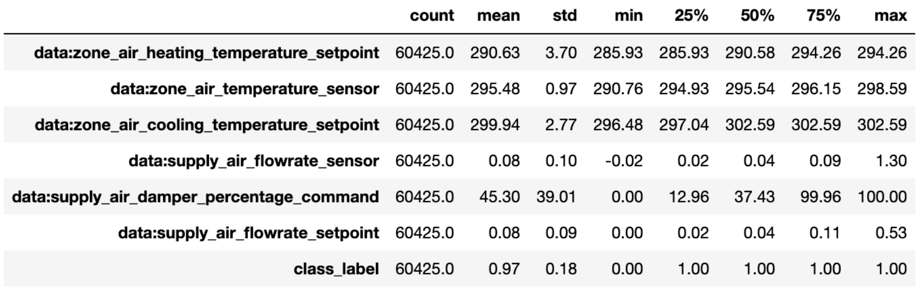

	count	mean	std	min	25%	50%	75%	max
data:zone_air_heating_temperature_setpoint	60425.0	290.63	3.70	285.93	285.93	290.58	294.26	294.26
data:zone_air_temperature_sensor	60425.0	295.48	0.97	290.76	294.93	295.54	296.15	298.59
data:zone_air_cooling_temperature_setpoint	60425.0	299.94	2.77	296.48	297.04	302.59	302.59	302.59
data:supply_air_flowrate_sensor	60425.0	0.08	0.10	-0.02	0.02	0.04	0.09	1.30
data:supply_air_damper_percentage_command	60425.0	45.30	39.01	0.00	12.96	37.43	99.96	100.00
data:supply_air_flowrate_setpoint	60425.0	0.08	0.09	0.00	0.02	0.04	0.11	0.53
class_label	60425.0	0.97	0.18	0.00	1.00	1.00	1.00	1.00

Fig. 3. Statistical values of the data collected from the Google campus for the Variable Air Volume devices [30]

In another critical development, researchers [7] explored the time series implications in anomaly detection, emphasizing the temporal dynamics of data. This method addresses anomalies by considering temporal dependencies. It accounts for patterns evolving over time, enabling the identification of deviations from expected temporal sequences. Moreover, there has been a growing interest [38] into hybrid models, merging traditional statistical methods with deep learning, establishing a bridge between classical and contemporary methodologies.

3 Methodology

The methodology is executed systematically, beginning with rigorous data collection and preprocessing to ensure the quality of the dataset. Our methodology tackles the issue of prediction randomness in CTGAN models, attributed to the Gumbel Softmax activation, which is essential for ensuring stable training. Instead of eliminating the Gumbel Softmax, we mitigate this randomness through the implementation of a hard Gumbel Softmax approach, enhancing the precision and reliability of the model's outputs. AnoGAN methodology is then applied for anomaly detection, involving the optimization of noise vectors through backpropagation using Mean Squared Error (MSE) loss. This process results in an anomaly score that reflects disparities between synthetic and original samples. Further refinement includes the determination of an optimal threshold through Receiver Operating Characteristic (ROC) analysis, enhancing the discriminative capacity of the framework. The investigation of individual feature differences between normal and synthetic samples concludes this methodological phase, offering detailed insights into the detected anomalies within structured datasets.

3.1 Dataset

The Smart Buildings Anomaly Detection dataset [30], covers 14 days from October 8 to 21, 2019, at a Google campus in California. It includes 60,425 (Figure 4) observations from 15 Variable Air Volume (VAV) devices, with 1,921 instances (3.2%) showing anomalies. These devices, designed to regulate air temperature, have two operational modes: a stricter comfort mode on weekdays (6:00 am to 10:00 pm) and a more relaxed eco mode during off-hours. The data is enriched with temporal markers like the day and hour.

3.2 Data preprocessing

In the initial phase of our methodology, the raw dataset is subjected to a thorough preprocessing pipeline to enhance its suitability for subsequent analysis. The dataset, once loaded undergoes a series of transformations for optimal utility. Initially, irrelevant features, namely 'dow' (day of the week) and 'hod' (hour of the day), are removed to streamline the dataset.

Further refinement includes categorizing the data into specific classes, with anomalies identified by the class label '0' being segregated into an exclusive dataframe for testing. Concurrently, data instances classified as normal are compiled into a distinct dataframe, as our training exclusively focuses on normal data.

To standardize and normalize the dataset, the Min-Max scaling technique is applied, rescaling feature values to a specified range (-1 to 1). This transformation is particularly valuable for maintaining the integrity of the anomaly and normal data distributions while preparing them for integration into the subsequent anomaly detection framework.

The multimodal input data undergoes processing via a Gaussian Mixture Model to derive multiple unimodal representations. This methodology enhances the stability of training Generative Adversarial Networks (GANs) by facilitating the learning of a broader array of simpler distributions, as opposed to a singular complex distribution.

$$p(x) = \sum_{i=1}^{M} \pi_i \mathcal{N}(x|\mu_i, \Sigma_i) \rightarrow \text{GMM Transformation}$$

where $p(x)$ is the probability density function of the data x modeled as a mixture of M Gaussian distributions, π_i are the mixing coefficients, μ_i and Σ_i are the mean and covariance of the i-th Gaussian component, respectively.

3.3 CT-GAN Implementation and Randomness Handling

In our research methodology, the utilization of the original CT-GAN (Conditional Tabular Generative Adversarial Network) (Figure 5) [34] implementation plays a pivotal role in generating synthetic samples for anomaly detection. CT-GAN is a specialized GAN variant designed for tabular data, aiming to faithfully reproduce the statistical characteristics of the given dataset. However, the original CT-GAN implementation introduces a layer of unpredictability during the testing phase, manifested in the randomness inherent in the predictions generated by the generator. To address this challenge and enhance the stability of the generated samples, we employ a modified softmax gumbel activation. This modification is instrumental in mitigating the unpredictable nature of the test predictions, ensuring a more consistent and reliable generation of synthetic samples. The application of the softmax gumbel activation contributes to the robustness and effectiveness of our anomaly detection framework by providing a more controlled and deterministic generation process for the synthetic data. Figure 6 shows the KDE for both the real and generated data after applying CT-GAN.

Original Gumbel Softmax

$$\begin{aligned}
G_i &= -\log(-\log(U_i)) && \text{where } U_i \text{ is } Uniform(0,1) \text{ and } G_i \text{ is Gumbel Noise} \\
h_i &= \frac{l_i + G_i}{\tau} && \text{where } \tau \text{ is Temperature Scaling Constant and } l_i \text{ is logits} \\
y_i &= \frac{\exp(h_i)}{\sum_j \exp(h_j)} && \text{Softmax on noised and scaled logits}
\end{aligned}$$

Hard Gumbel Softmax

$$\begin{aligned}
y_i &= \frac{\exp(l_i/\tau)}{\sum_j \exp(l_j/\tau)} && \text{Temperature Scaling and softmax on logits} \\
y_{\text{hard},i} &= \begin{cases} 1 & \text{if } i = \arg\max_j y_j \\ 0 & \text{otherwise} \end{cases} && \text{Selecting the max component} \\
y &= \text{stop\_gradient}(y_{\text{hard}} - y) + y && \text{Straight through Estimator}
\end{aligned}$$

3.4 Optimizing Noise Vector and Anomaly Scoring

In this phase of our methodology, the focus is on refining the generated synthetic samples to closely mirror the characteristics of the original data. For a given sample, we initiate the optimization process of the noise vector, a critical parameter in the generative process. The objective is to adjust the noise vector such that the generator produces synthetic samples that are the most similar to the corresponding original samples.

To quantify the dissimilarity between the synthetic and original samples, we employ the Mean Squared Error (MSE) loss during the backpropagation process to optimize the latent vector. We tried several other loss-measuring functions but MSE worked the best for the data. The MSE loss serves as a measure of the average squared differences between corresponding elements of the synthetic and original samples. This penalizes larger differences and allows us to capture the nuances and intricacies of the data distribution, facilitating a more precise evaluation of the generative process.

The computed MSE loss between the original sample and generated sample serves as the anomaly score, representing the magnitude of deviation between the two samples. This score becomes a crucial metric for distinguishing normal from anomalous samples. Leveraging this anomaly score, we establish a threshold that optimally discriminates between normal and anomalous instances, determined through Area Under Curve-Receiver Operating Characteristic (AUC-ROC) analysis. This meticulous threshold determination enhances the discriminative power of our anomaly detection framework.

Furthermore, delving into individual feature differences between normal and synthetic samples provides insights into the specific aspects contributing to detected anomalies. By scrutinizing these differences, we gain a nuanced understanding of the reasons behind the anomalies, allowing for more informed and

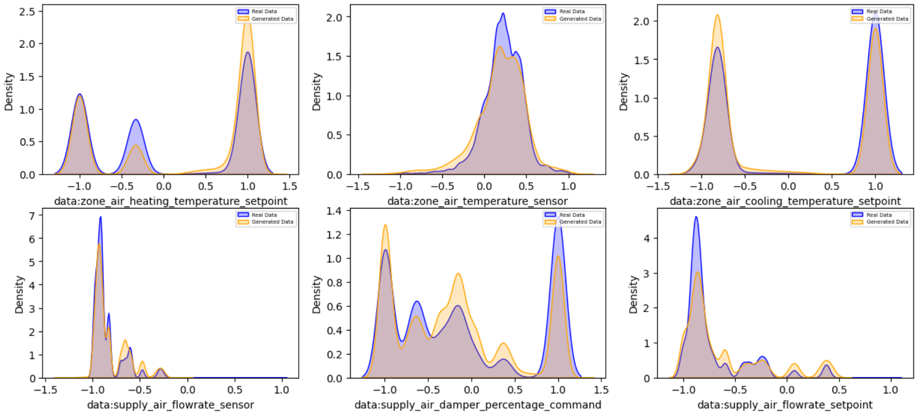

Fig. 4. The figure provides a visual representation of the Kernel Density Estimation (KDE) for both the real and generated datasets. The KDE serves as a smoothed probability density function, offering insights into the underlying distribution of the data. Specifically, the KDE for the real data showcases the probability density across different values, illustrating the distribution's characteristics and patterns. On the other hand, the KDE for the generated data, produced by the Generator component of the Generative Adversarial Network (GAN), offers a comparative view. This comparison enables an assessment of how well the generator has learned to replicate the statistical properties of the real data. By visually inspecting the KDE curves, one can discern the fidelity of the generated data distribution in relation to the authentic dataset, providing a valuable tool for evaluating the performance and quality of the GAN's generative capabilities.

targeted interventions in subsequent stages of analysis. This comprehensive approach to anomaly scoring and threshold determination forms a robust foundation for the effectiveness of our anomaly detection methodology.

4 Results

In this section, we present a comprehensive evaluation of our anomaly detection framework based on AnoGAN, comparing its performance against two alternative methodologies, namely One-Class Support Vector Machine (OCSVM) and k-Nearest Neighbors (KNN). The choice of OCSVM and KNN as comparative methods reflects their widespread usage and effectiveness in anomaly detection tasks. Table 1 captures the accuracy of all the methods implemented. Through a rigorous analysis of the results obtained from these three approaches, we aim to elucidate the strengths and limitations of our AnoGAN-based framework in detecting anomalies within structured datasets.

The progression of Generative Adversarial Network (GAN) losses across epochs is a key aspect of our findings. As the GAN undergoes training, the Generator Loss, representing the ability to generate realistic data, exhibits notable

Methods	AUC-ROC
OCSVM	55.6%
KNN	50%
AnoGAN	72%

Table 1. Comparison of Anomaly Detection Methods based on Accuracy.

fluctuations. Initially high, this loss gradually diminishes as the Generator refines its capacity to produce more authentic samples. Simultaneously, the Discriminator Loss, indicating the Discriminator's accuracy in distinguishing real from generated data, follows a similar trajectory. After a certain point the loss starts to increase but our implementation stops the progression with early stopping. The delicate interplay between these losses is pivotal for achieving equilibrium in GAN training and generating high quality samples.

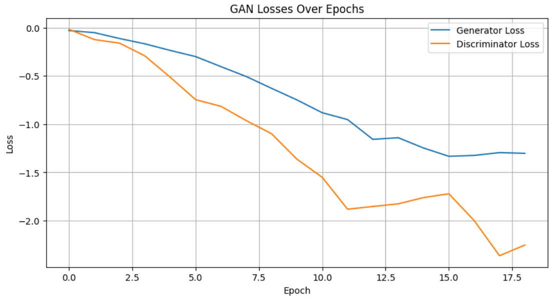

Fig. 5. The figure illustrates the progression of Generator and Discriminator losses over training epochs in the Generative Adversarial Network (GAN), showcasing how the model refines generative and discriminative capabilities during learning.

Our innovative approach involves applying AnoGAN to synthetic data generated by CTGAN and normal data samples, achieving a robust anomaly detection accuracy of approximately 72%. Figure 7 shows the loss over epoch for training and Figure 8 shows the ROC curve for our method. Extended training substantially improves accuracy, reaching around 80%. This outperforms traditional methods like KNN and OCSVM, where accuracy is comparable to random assignment at 50%. The superior performance underscores the efficacy of our methodology in detecting anomalies with higher precision and reliability.

Our methodology demonstrates notable efficacy in scenarios characterized by infrequent data anomalies. The system exhibits exceptional performance, achieving an accuracy exceeding 85% when tested on datasets featuring a mere 50 anomaly samples. The observed variance in accuracy, contingent upon the anomaly count, can be attributed to the inherent challenge of discerning the underlying generative function for anomalies. This task is comparatively more

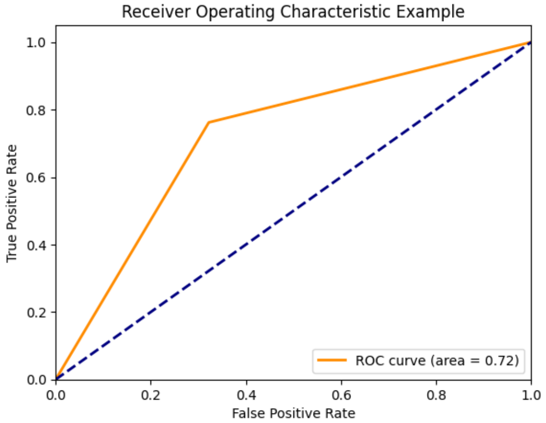

Fig. 6. The ROC curve visually represents the Receiver Operating Characteristic for the Generative Adversarial Network (GAN), offering insights into its discriminative performance and the trade-off between true positive and false positive rates

straightforward with a smaller anomaly set, where distinct patterns are easily recognizable, as opposed to the heightened complexity associated with larger anomaly datasets.

5 Conclusion and Future work

Our meticulously executed methodology encompasses systematic data collection and preprocessing, ensuring dataset quality. Incorporating the original CT-GAN implementation, we address inherent randomness in test predictions through softmax gumbel activation, enhancing the stability of generated samples. Leveraging AnoGAN for anomaly detection, our approach optimizes noise vectors using Mean Squared Error (MSE) loss, resulting in a nuanced anomaly score reflective of synthetic-original disparities. Refinement involves ROC analysis for threshold determination, enhancing the framework's discriminative capacity.

Looking ahead, our research opens avenues for several promising directions. Firstly, exploring the integration of domain-specific knowledge into the anomaly detection process can enhance the model's interpretability and performance. Additionally, investigating the adaptability of our methodology to incorporate categorical variables will broaden its applicability across a wide range of domains. Further research into refining the threshold determination process and extending the framework to handle dynamic datasets with evolving patterns could deepen its practical applicability. The exploration of ensemble techniques and hybrid models, combining the strengths of different anomaly detection methods, presents another intriguing avenue for future investigation. Continuous refinement and adaptation of our methodology will be essential for addressing evolving challenges in anomaly detection across various real-world scenarios.

References

1. Abdelli, K., Cho, J.Y., Azendorf, F., Griesser, H., Tropschug, C., Pachnicke, S.: Machine-learning-based anomaly detection in optical fiber monitoring. Journal of optical communications and networking **14**(5), 365–375 (2022)
2. Ahmed, M., Mahmood, A.N.: Novel approach for network traffic pattern analysis using clustering-based collective anomaly detection. Annals of Data Science **2**(1), 111–130 (2015)
3. Basora, L., Olive, X., Dubot, T.: Recent advances in anomaly detection methods applied to aviation. Aerospace **6**(11), 117 (2019)
4. Chalapathy, R., Chawla, S.: Deep learning for anomaly detection: A survey. arXiv preprint arXiv:1901.03407 (2019)
5. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM computing surveys (CSUR) **41**(3), 1–58 (2009)
6. Chen, Z., Liu, J., Gu, W., Su, Y., Lyu, M.R.: Experience report: Deep learning-based system log analysis for anomaly detection. arXiv preprint arXiv:2107.05908 (2021)
7. De Stefani, J., Le Borgne, Y.A., Caelen, O., Hattab, D., Bontempi, G.: Batch and incremental dynamic factor machine learning for multivariate and multi-step-ahead forecasting. International Journal of Data Science and Analytics **7**(4), 311–329 (2019)
8. Dou, G., Zhou, Z., Qu, X.: Time majority voting, a pc-based eeg classifier for non-expert users. In: International Conference on Human-Computer Interaction. pp. 415–428. Springer (2022)
9. Dunning, T., Friedman, E.: Practical machine learning: a new look at anomaly detection. " O'Reilly Media, Inc." (2014)
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)
11. Gui, S., Song, S., Qin, R., Tang, Y.: Remote sensing object detection in the deep learning era—a review. Remote Sensing **16**(2), 327 (2024)
12. Han, D., Wang, Z., Chen, W., Zhong, Y., Wang, S., Zhang, H., Yang, J., Shi, X., Yin, X.: Deepaid: Interpreting and improving deep learning-based anomaly detection in security applications. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. pp. 3197–3217 (2021)
13. Key, M.L., Mehtiyev, T., Qu, X.: Advancing eeg-based gaze prediction using depthwise separable convolution and enhanced pre-processing. In: International Conference on Human-Computer Interaction. pp. 3–17. Springer (2024)
14. Koome Murungi, N., Pham, M.V., Dai, X., Qu, X.: Trends in machine learning and electroencephalogram (eeg): A review for undergraduate researchers. arXiv e-prints pp. arXiv–2307 (2023)
15. Li, W., Zhou, N., Qu, X.: Enhancing eye-tracking performance through multi-task learning transformer. In: International Conference on Human-Computer Interaction. pp. 31–46. Springer (2024)
16. Luo, Y., Xiao, Y., Cheng, L., Peng, G., Yao, D.: Deep learning-based anomaly detection in cyber-physical systems: Progress and opportunities. ACM Computing Surveys (CSUR) **54**(5), 1–36 (2021)
17. Murungi, N.K., Pham, M.V., Dai, X.C., Qu, X.: Empowering computer science students in electroencephalography (eeg) analysis: A review of machine learning algorithms for eeg datasets. SIGKDD (2023)

18. Nayak, R., Pati, U.C., Das, S.K.: A comprehensive review on deep learning-based methods for video anomaly detection. Image and Vision Computing **106**, 104078 (2021)
19. Pang, G., Shen, C., Cao, L., Hengel, A.V.D.: Deep learning for anomaly detection: A review. ACM computing surveys (CSUR) **54**(2), 1–38 (2021)
20. Patcha, A., Park, J.M.: An overview of anomaly detection techniques: Existing solutions and latest technological trends. Computer networks **51**(12), 3448–3470 (2007)
21. Patel, V., Pan, L., Rajasegarar, S.: Graph deep learning based anomaly detection in ethereum blockchain network. In: International conference on network and system security. pp. 132–148. Springer (2020)
22. Qu, X.: Time Continuity Voting for Electroencephalography (EEG) Classification. Ph.D. thesis, Brandeis University (2022)
23. Qu, X., Hall, M., Sun, Y., Sekuler, R., Hickey, T.J.: A personalized reading coach using wearable eeg sensors. CSEDU (2019)
24. Qu, X., Hickey, T.J.: Eeg4home: A human-in-the-loop machine learning model for eeg-based bci. In: Augmented Cognition: 16th International Conference, AC 2022, Held as Part of the 24th HCI International Conference, HCII 2022, Virtual Event, June 26–July 1, 2022, Proceedings. pp. 162–172. Springer (2022)
25. Qu, X., Liu, P., Li, Z., Hickey, T.: Multi-class time continuity voting for eeg classification. In: Brain Function Assessment in Learning: Second International Conference, BFAL 2020, Heraklion, Crete, Greece, October 9–11, 2020, Proceedings 2. pp. 24–33. Springer (2020)
26. Qu, X., Liukasemsarn, S., Tu, J., Higgins, A., Hickey, T.J., Hall, M.H.: Identifying clinically and functionally distinct groups among healthy controls and first episode psychosis patients by clustering on eeg patterns. Frontiers in psychiatry **11**, 541659 (2020)
27. Qu, X., Mei, Q., Liu, P., Hickey, T.: Using eeg to distinguish between writing and typing for the same cognitive task. In: Brain Function Assessment in Learning: Second International Conference, BFAL 2020, Heraklion, Crete, Greece, October 9–11, 2020, Proceedings 2. pp. 66–74. Springer (2020)
28. Qu, X., Sun, Y., Sekuler, R., Hickey, T.: Eeg markers of stem learning. In: 2018 IEEE Frontiers in Education Conference (FIE). pp. 1–9. IEEE (2018)
29. Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G.: Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: International conference on information processing in medical imaging. pp. 146–157. Springer (2017)
30. Sipple, J.: Interpretable, multidimensional, multimodal anomaly detection with negative sampling for detection of device failure. In: International Conference on Machine Learning. pp. 9016–9025. PMLR (2020)
31. Wang, N., Chen, C., Xie, Y., Ma, L.: Brain tumor anomaly detection via latent regularized adversarial network. arXiv preprint arXiv:2007.04734 (2020)
32. Wang, R., Qu, X.: Eeg daydreaming, a machine learning approach to detect daydreaming activities. In: International Conference on Human-Computer Interaction. pp. 202–212. Springer (2022)
33. Wu, R., Keogh, E.J.: Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress. IEEE transactions on knowledge and data engineering **35**(3), 2421–2429 (2021)
34. Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K.: Modeling tabular data using conditional gan. Advances in neural information processing systems **32** (2019)

35. Yi, L., Qu, X.: Attention-based cnn capturing eeg recording's average voltage and local change. In: Artificial Intelligence in HCI: 3rd International Conference, AI-HCI 2022, Held as Part of the 24th HCI International Conference, HCII 2022, Virtual Event, June 26–July 1, 2022, Proceedings. pp. 448–459. Springer (2022)
36. Yunoki, I., Berreby, G., D'Andrea, N., Lu, Y., Qu, X.: Exploring ai music generation: A review of deep learning algorithms and datasets for undergraduate researchers. In: International Conference on Human-Computer Interaction. pp. 102–116. Springer (2023)
37. Zenati, H., Romain, M., Foo, C.S., Lecouat, B., Chandrasekhar, V.: Adversarially learned anomaly detection. In: 2018 IEEE International conference on data mining (ICDM). pp. 727–736. IEEE (2018)
38. Zhang, G.L., Ma, W.L., Liu, R.B.: Cluster correlation expansion for studying decoherence of clock transitions in spin baths. Physical Review B **102**(24), 245303 (2020)
39. Zhang, J., Xie, Y., Li, Y., Shen, C., Xia, Y.: Covid-19 screening on chest x-ray images using deep learning based anomaly detection. arXiv preprint arXiv:2003.12338 **27**(10.48550) (2020)
40. Zhou, C., Paffenroth, R.C.: Anomaly detection with robust deep autoencoders. In: Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 665–674 (2017)
41. Zhou, Z., Dou, G., Qu, X.: Brainactivity1: A framework of eeg data collection and machine learning analysis for college students. In: International Conference on Human-Computer Interaction. pp. 119–127. Springer (2022)