Unsupervised Anomaly Detection in ECG Signals Using Denoising Autoencoders: A Comparative Study

Samuele Russo[1], Pasquale Silvestri[1] and Imad Eddine Tibermacine[1]

[1]Neuroimaging Laboratory, IRCCS Santa Lucia Foundation, Rome, Italy

Abstract

Anomaly detection is essential in various domains, including healthcare, where early identification of irregular patterns in data can significantly impact patient outcomes. This paper presents a novel approach to unsupervised anomaly detection using the ECG5000 dataset, focusing on electrocardiogram (ECG) data. We introduce multiple autoencoder architectures—linear, convolutional, and LSTM-based—reframing the traditionally supervised classification task as an unsupervised anomaly detection problem. By disregarding original labels, we emphasize the models' ability to generalize across different ECG abnormalities. Our extensive experiments reveal that a denoising linear autoencoder outperforms more complex architectures, achieving an accuracy of 97.73%, within 0.7% of the current state-of-the-art. Furthermore, we conduct a comprehensive analysis of the latent space representations, providing insights into the models' feature extraction capabilities. These findings suggest that our approach not only reduces model complexity but also maintains high accuracy, offering a viable solution for real-time anomaly detection in medical settings.

Keywords

ECG, Anomaly Detection, Autoencoders, Deep Learning, Signal Processing

1. Introduction

Anomaly detection is a crucial task in data science and machine learning[1], involving the identification of patterns in data that do not conform to expected behavior[2]. Its applications span a wide range of fields, including finance, cybersecurity, healthcare [3], and manufacturing, where the detection of anomalies can prevent catastrophic failures[4, 5], secure systems against breaches, and identify early signs of disease[6, 7, 8]. In healthcare, particularly in cardiology, the timely detection of anomalies[9] in electrocardiogram (ECG) is vital for diagnosing potentially life-threatening conditions such as arrhythmias[10, 11, 12, 13]. Traditional methods for ECG analysis rely heavily on labeled data for supervised learning; however, obtaining labeled data can be challenging and expensive[14, 15]. This limitation has driven the need for effective unsupervised anomaly detection methods that can operate reliably without labeled data[16, 17].

Recent advancements in anomaly detection within healthcare have progressed from traditional statistical methods, such as Z-Score and Interquartile Range (IQR)[18], to more sophisticated machine learning and deep learning approaches[19, 20], capable of handling complex, high-dimensional data like ECG signals[21, 22, 23, 24]. Traditional methods often struggle with such intricate temporal patterns, while density-based[25] (e.g., LOF, DBSCAN) and distance-based techniques (e.g., k-NN) have shown improved efficacy but tend to be computationally intensive and less scalable for large datasets[26, 27, 28].

In contrast, deep learning methods, particularly autoencoders, offer significant advantages by learning compressed representations that capture the underlying structure of high-dimensional data[29, 30, 31]. Variants such as denoising, contractive, and variational autoencoders have been explored extensively for their robustness across diverse anomaly types[32, 33, 34, 35]. Building on these developments, our work applies autoencoder architectures to the ECG5000 dataset, a benchmark in ECG analysis, reframing the anomaly detection task as an unsupervised learning problem to enhance model generalizability across various ECG abnormalities without relying on labeled data[10, 36, 37].

This study addresses this need by reframing the ECG5000 dataset[38, 39], typically used for classification, as a benchmark for unsupervised anomaly detection. We explore a variety of autoencoder architectures—linear, convolutional, and LSTM-based—to evaluate their effectiveness in identifying anomalies without the guidance of labels[40]. The research focuses on two main variations of autoencoders: denoising and contractive. Our objective is to identify an architecture that balances model complexity with performance, making it suitable for real-time medical applications where computational resources may be limited.

Through rigorous experimentation, we demonstrate that a denoising linear autoencoder achieves near state-of-the-art performance with significantly reduced complexity. Additionally, we perform an in-depth analysis of the latent space representations generated by our mod-

ICYRIME 2025: 10th International Conference of Yearly Reports on Informatics, Mathematics, and Engineering. Czestochowa, January 14-16, 2025

s.russo@hsantalucia.it (S. Russo); d.tibermacine@hsantalucia.it (I. E. Tibermacine)

els, offering insights into how these architectures capture essential features for anomaly detection in ECG data.

2. Methodology

In this section, we outline the methodologies employed to explore the effectiveness of various autoencoder architectures for unsupervised anomaly detection in ECG data. The methodology encompasses the design and testing of multiple autoencoder variants, the evaluation of their performance, and the analysis of the latent space representations.

2.1. Autoencoder Architectures

We investigated three different autoencoder architectures: Linear, Convolutional, and LSTM-based autoencoders[41]. These architectures were selected based on their ability to capture different characteristics of the ECG data—spatial hierarchies in the case of convolutional layers and temporal dependencies for LSTM layers.

- **Linear Autoencoder**: A fully connected (linear) architecture with two layers in both the encoder and decoder. The encoder transforms the input data from 140 to 32 features, and subsequently to an 8-dimensional latent space using ReLU activation. The decoder mirrors this process to reconstruct the original input.
- **Convolutional Autoencoder**: The encoder consists of two convolutional layers with a 1x9 kernel, reducing the input to an 8-dimensional latent space, followed by a dropout layer to prevent overfitting. The decoder uses transposed convolutions to reconstruct the input.
- **LSTM Autoencoder**: This architecture employs three unidirectional LSTM layers in both the encoder and decoder, with the final hidden state forming an 8-dimensional latent vector. This latent vector is then replicated across the sequence length to reconstruct the ECG signal.

2.2. Model Variants

Each of the aforementioned architectures was further developed into different variants to assess their robustness and effectiveness:

- **Contractive Autoencoders**: These models introduce a regularization term based on the Frobenius norm of the Jacobian matrix, which encourages the model to learn stable latent representations that are less sensitive to small input perturbations.
- **Denoising Autoencoders**: Here, Gaussian noise is added to the input data during training. The model is trained to reconstruct the clean input from the noisy version, enhancing its ability to filter out noise and focus on the underlying signal structure.
- **Mixed Models**: Combining both contractive and denoising strategies, these models aim to leverage the strengths of both approaches, though they require careful tuning of hyperparameters to avoid over-regularization.

2.3. Training Pipeline

The training process for these models is illustrated in Figure 1. The autoencoders were trained on normal ECG samples, allowing them to learn the distribution of non-anomalous data. The training involved minimizing the mean squared error (MSE) between the input and reconstructed output. The training process was optimized using the Adam optimizer, with the best model performance achieved using early stopping criteria to prevent overfitting.

Post-training, the models were evaluated using a threshold-based classification approach. The MSE distribution of normal and abnormal samples was analyzed to set a threshold that distinguishes between the two. Specifically, the threshold was calculated by interpolating the MSE distributions of normal and abnormal samples, with one standard deviation added to the mean of the normal sample distribution and subtracted from the mean of the abnormal sample distribution.

During inference, an ambiguous ECG sample is passed through the autoencoder, and its MSE is calculated. This MSE is then compared to the threshold: if the MSE is below the threshold, the sample is classified as normal; otherwise, it is classified as anomalous.

2.4. Latent Space Analysis

A key aspect of our methodology was the analysis of the latent space produced by the autoencoders. Principal Component Analysis (PCA) was applied to reduce the dimensionality of the latent space and visualize the separation between normal and abnormal samples. Additionally, a simple logistic regression discriminator was trained on the latent representations to assess their ability to distinguish between normal and anomalous data.

The effectiveness of the latent space was quantified by measuring the accuracy of the discriminator, with higher accuracy indicating a more distinct and informative latent representation. This analysis was crucial in understanding the models' capacity to encode meaningful features in the latent space, which directly impacts the anomaly detection performance.

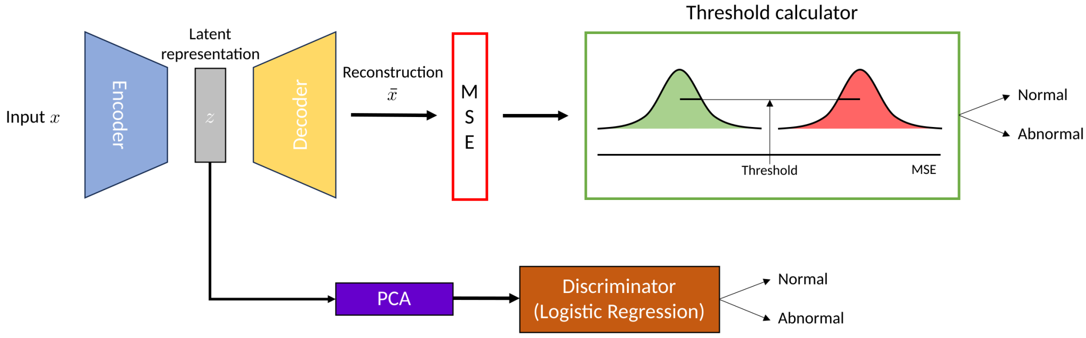

Figure 1: Model pipeline for both anomaly detection and latent space analysis. This pipeline integrates the training, evaluation, and post-processing steps used to assess model performance.

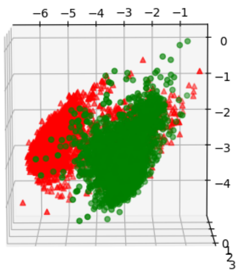

Figure 2: Latent space with true labels of the denoising linear model. The separation between normal and abnormal samples is clearly visible.

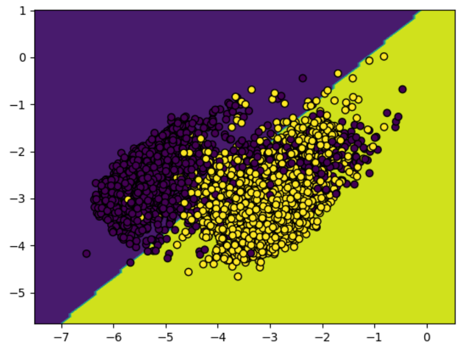

Figure 3: Decision boundary computed by the discriminator of the denoising linear model.

For comparison, Figures 4 and 5 illustrate the latent space and decision boundary for the denoising LSTM model, which shows less distinct separation, leading to lower classification accuracy.

Figures 6 and 7 present the latent space and decision boundary for the contractive linear model, which also demonstrated lower effectiveness in distinguishing between normal and abnormal samples compared to the denoising linear model.

Further analysis was conducted on the mixed linear model (Figures 8 and 9), and the denoising convolutional model (Figures 10 and 11). These results highlighted the strengths and limitations of combining contractive and denoising approaches in the same architecture.

2.5. Training

The training process was repeated multiple times to optimize the hyperparameters for each model variant. The final training was conducted over 15 epochs, with a learning rate of 0.01 using the Adam optimizer, which proved more stable and faster than stochastic gradient descent (SGD). The denoising linear autoencoder achieved the best performance, with training early-stopped at 11 epochs based on validation performance.

The optimal hyperparameters included a contractive lambda of 0.0001 and a noise level of 0.05 (representing the maximum absolute value for the Gaussian noise added to the input). Notably, the contractive variants took approximately 10 to 15 times longer to train than their denoising counterparts, without yielding superior

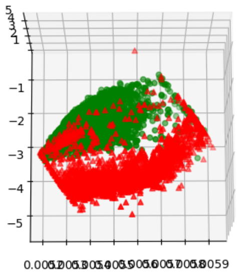

Figure 4: Latent space with true labels of the denoising LSTM model.

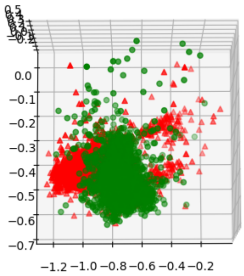

Figure 6: Latent space with true labels of the contractive linear model.

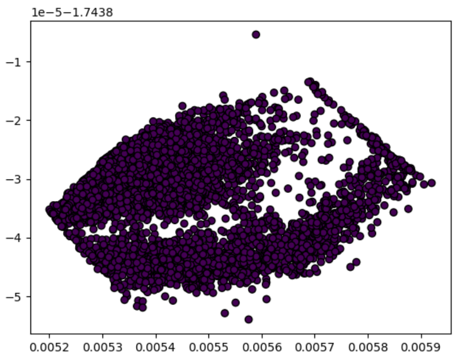

Figure 5: Decision boundary computed by the discriminator of the denoising LSTM model.

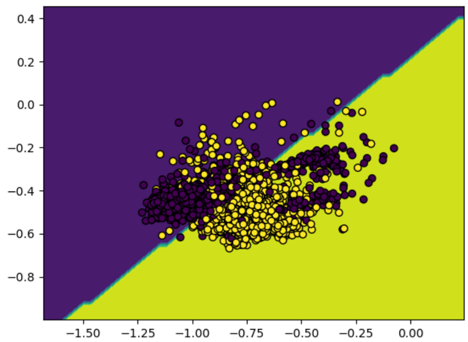

Figure 7: Decision boundary computed by the discriminator of the contractive linear model.

results.

3. Experimental Setup

In this section, we detail the experimental setup used to evaluate the performance of the various autoencoder architectures on the ECG5000 dataset. This includes a description of the dataset, data preprocessing steps, model training configurations, and evaluation metrics.

3.1. Dataset Description

The ECG5000 dataset is a well-known benchmark for anomaly detection tasks involving electrocardiogram (ECG) signals. The dataset consists of 5000 one-dimensional ECG recordings, each containing 140 time steps. The dataset is divided into five classes, with one representing normal heartbeats and the remaining four representing various types of anomalies, including:

- **Class 1**: Normal beats.
- **Class 2**: Premature ventricular contractions.
- **Class 3**: Fusion beats.
- **Class 4**: Unclassifiable beats.
- **Class 5**: Anomalous beats.

For the purposes of unsupervised anomaly detection, the original labels are disregarded during training. Only samples from the normal class (Class 1) are used to train the autoencoders, with the goal of identifying anomalies based on reconstruction error during testing.

3.2. Data Preprocessing

The ECG signals were normalized to have zero mean and unit variance. This step ensures that the models focus on

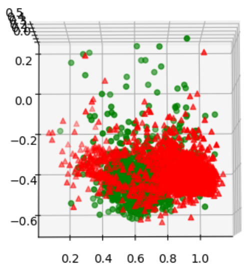

Figure 8: Latent space with true labels of the mixed linear model.

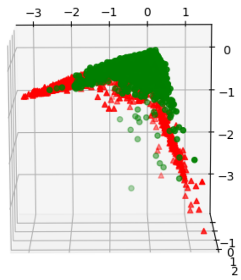

Figure 10: Latent space with true labels of the denoising convolutional model.

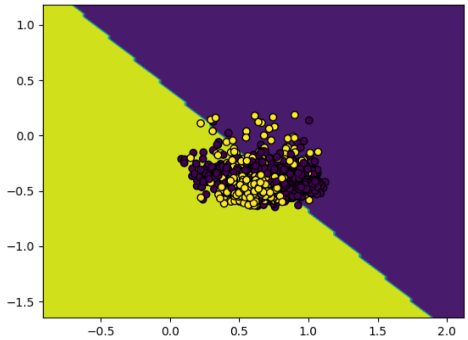

Figure 9: Decision boundary computed by the discriminator of the mixed linear model.

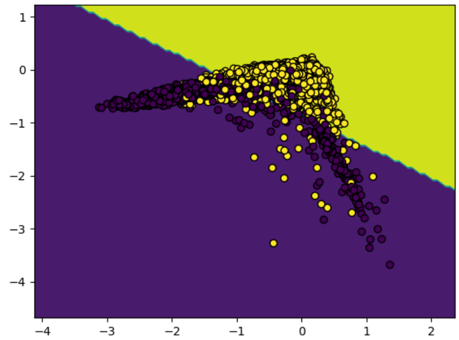

Figure 11: Decision boundary computed by the discriminator of the denoising convolutional model.

the shape of the signal rather than its absolute amplitude, which is crucial for the generalization of the anomaly detection task.

Given the nature of the dataset, no further data augmentation techniques were applied, as the goal was to evaluate the autoencoders' performance on raw, unaltered ECG signals. The dataset was split into training, validation, and test sets with the following proportions:

- **Training Set**: 60% of the normal samples.
- **Validation Set**: 20% of the normal samples.
- **Test Set**: 20% of the normal samples, along with all abnormal samples.

The validation set was used for hyperparameter tuning and early stopping, while the test set was used for final performance evaluation.

3.3. Model Training Configuration

The models were trained using the Adam optimizer with a learning rate of 0.01. The training process was conducted over 15 epochs, with early stopping applied if the validation loss did not improve for three consecutive epochs. The batch size was set to 64, which provided a balance between computational efficiency and gradient estimation accuracy.

- **Contractive Autoencoders**: A contractive loss term with a regularization coefficient (λ) of 0.0001 was added to the standard reconstruction loss. This encouraged the model to learn more stable representations.
- **Denoising Autoencoders**: Gaussian noise with a standard deviation of 0.05 was added to the input during training. The model was trained to reconstruct the original, clean ECG signal from

the noisy input.

- **Mixed Models**: Both contractive loss and denoising noise were applied. These models required careful tuning of the regularization coefficient and noise level to avoid over-regularization.

All models were implemented using PyTorch, leveraging GPU acceleration to expedite the training process. The best-performing model, based on validation performance, was selected for final testing.

3.4. Threshold Selection

The selection of the decision threshold for anomaly detection was a critical aspect of the experimental setup. The threshold was determined by analyzing the MSE distribution for normal and anomalous samples. The final threshold was set as the midpoint between one standard deviation above the mean MSE of the normal samples and one standard deviation below the mean MSE of the anomalous samples.

This approach ensured that the threshold was not overly conservative, allowing the models to generalize better to unseen data, especially in scenarios where the separation between normal and anomalous samples was subtle.

3.5. Computational Resources

All experiments were conducted on a high-performance computing cluster equipped with NVIDIA GPUs. The use of GPU acceleration significantly reduced training times, particularly for the more complex models like the LSTM-based autoencoder and the contractive autoencoders, which require extensive matrix computations.

4. Results

In this section, we present the performance results of the different autoencoder models evaluated on the ECG5000 dataset. The results are organized to highlight the effectiveness of each model in detecting anomalies based on the various metrics discussed in the experimental setup.

4.1. Model Performance

Table 1 summarizes the accuracy of the various autoencoder model configurations. The results show that the denoising linear autoencoder outperformed the other models, achieving the highest accuracy.

4.2. Training and Evaluation Losses

Figures 12-17 provide a visual representation of the training and evaluation losses, as well as the MSE scores for each model. The denoising linear autoencoder demonstrated consistent and stable training, with the lowest overall loss and a clear separation between the reconstruction errors of normal and anomalous samples.

Table 1
Accuracy of the Various Autoencoder Models

	Accuracy (%)		
	Denoising	Contractive	Mixed
Linear	**97.73**	95.69	95.50
Convolutional	94.94	94.51	-
LSTM	93.60	-	-

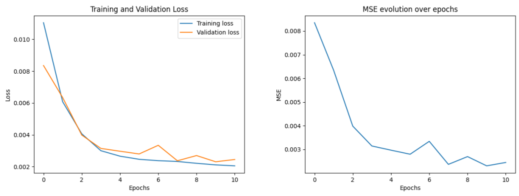

(a) Train and eval losses (b) MSE score

Figure 12: Training graphs for the denoising linear autoencoder

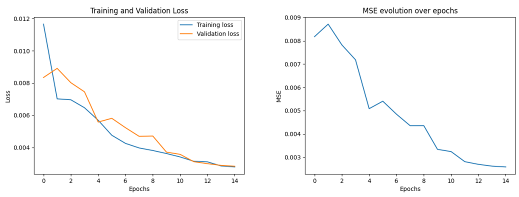

(a) Train and eval losses (b) MSE score

Figure 13: Training graphs for the contractive linear autoencoder

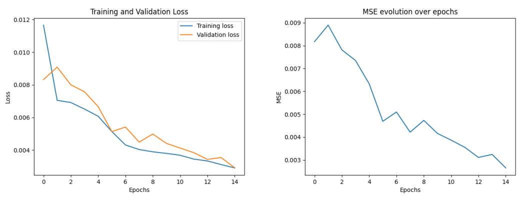

(a) Train and eval losses (b) MSE score

Figure 14: Training graphs for the mixed linear autoencoder

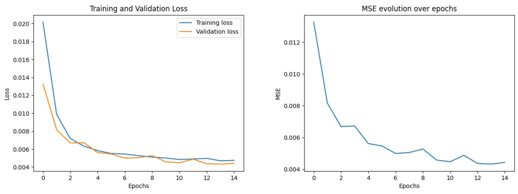

(a) Train and eval losses (b) MSE score

Figure 15: Training graphs for the denoising convolutional autoencoder

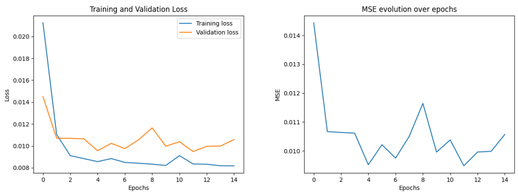

(a) Train and eval losses (b) MSE score

Figure 16: Training graphs for the contractive convolutional autoencoder

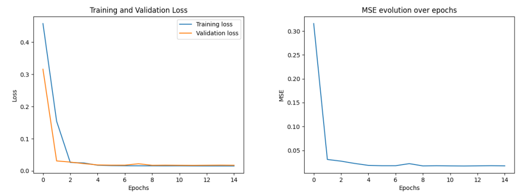

(a) Train and eval losses (b) MSE score

Figure 17: Training graphs for the denoising LSTM autoencoder

4.3. Latent Space Analysis

The latent space analysis revealed significant differences in the representation quality of each model. The denoising linear autoencoder produced a well-separated latent space, allowing for clear discrimination between normal and anomalous samples. This is reflected in the high accuracy of the discriminator applied to the latent space, as shown in Table 2.

Table 2
Discriminator Accuracy for Each Model

	Discriminator Accuracy (%)		
	Denoising	Contractive	Mixed
Linear	**93.10**	87.40	87.60
Convolutional	91.00	58.00	-
LSTM	54.90	-	-

Figures 18-21 illustrate the latent spaces and corresponding decision boundaries for the best-performing models. The denoising linear model shows clear clusters for normal and anomalous data, whereas the contractive and mixed models exhibit more overlapping clusters, indicating less effective separation in the latent space.

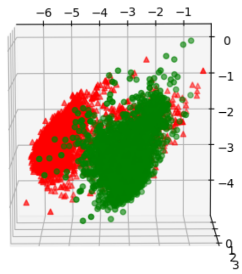

Figure 18: Latent space with true labels of the denoising linear model

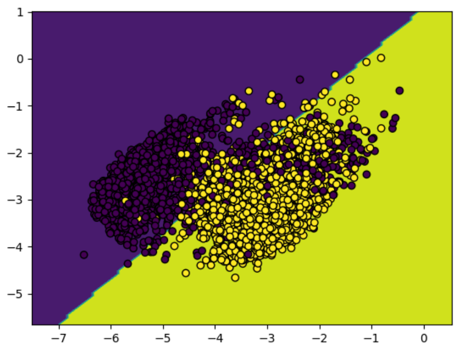

Figure 19: Decision boundary computed by the discriminator of the denoising linear model

The performance of the LSTM-based model was notably lower, as evidenced by the near-random discriminator accuracy and the poor separation in the latent space (Figures 20 and 21). This suggests that the LSTM model struggled to capture relevant features in the latent space for effective anomaly detection.

4.4. Summary of Results

The results demonstrate that the denoising linear autoencoder was the most effective model for anomaly detection

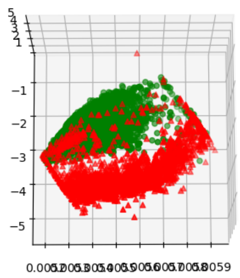

Figure 20: Latent space with true labels of the denoising LSTM model

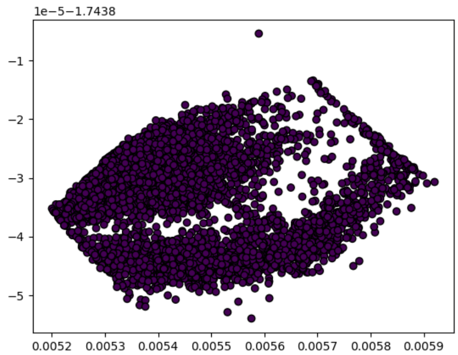

Figure 21: Decision boundary computed by the discriminator of the denoising LSTM model

on the ECG5000 dataset. It achieved the highest accuracy, the most distinct latent space separation, and outperformed the contractive and LSTM-based models. The success of the denoising approach highlights the importance of handling noise in ECG signals and suggests that simpler models with well-tuned noise management can outperform more complex architectures in this context.

5. Discussion

In this section, we discuss the implications of the results obtained from the various autoencoder models tested on the ECG5000 dataset, with a particular focus on the insights gained from the latent space analysis.

5.1. Model Performance

The performance results indicate that the denoising linear autoencoder outperformed other models in terms of accuracy and robustness in anomaly detection. This suggests that for the task of unsupervised anomaly detection on ECG data, simpler models with well-managed noise handling capabilities can provide superior performance compared to more complex architectures such as LSTM or convolutional models.

The denoising approach effectively enhances the autoencoder's ability to learn meaningful representations by forcing the network to reconstruct clean data from noisy inputs. This technique seems to be particularly well-suited for ECG data, where noise is prevalent due to various sources of interference during signal acquisition. The success of the denoising linear autoencoder demonstrates that the simplicity of the architecture, combined with an effective noise reduction strategy, can lead to robust performance in anomaly detection tasks.

5.2. Contractive vs. Denoising Autoencoders

The comparison between contractive and denoising autoencoders highlights distinct differences in how these architectures manage latent space representations. While contractive autoencoders aim to enforce robustness by minimizing the sensitivity of the latent space to small perturbations in the input, they tend to be more computationally expensive and, in this study, did not outperform the denoising models. This suggests that, at least in the context of ECG anomaly detection, denoising strategies are more effective in maintaining the balance between reconstruction accuracy and computational efficiency.

One notable observation is the performance of the contractive convolutional autoencoder, which suffered from latent space collapse—where all inputs were mapped to nearly identical latent representations. This outcome highlights the potential risks of over-regularization in contractive models, particularly when the regularization strength is not carefully tuned.

5.3. Latent Space Analysis

The latent space analysis provided valuable insights into the internal workings of the different autoencoder models. The use of Principal Component Analysis (PCA) and a simple discriminator allowed us to visualize and quantify the quality of the latent space representations.

The denoising linear autoencoder produced a well-separated latent space, as evidenced by the clear clusters corresponding to normal and anomalous data. This well-defined separation is crucial for effective anomaly detection, as it allows for the reliable identification of outliers based on the reconstruction error. The high discriminator accuracy (93.10%) further confirms the effectiveness of the latent space representation in distinguishing between normal and anomalous samples.

In contrast, the latent space produced by the denoising LSTM and contractive convolutional models exhibited significant overlap between normal and anomalous data, resulting in poor discriminator performance. The near-random accuracy of the discriminator (around 50%) in these cases indicates that the latent space failed to capture the essential features needed for effective anomaly detection. This finding suggests that while LSTM-based models are well-suited for capturing temporal dependencies, they may struggle with unsupervised tasks where the latent space must be highly informative for anomaly detection.

The issues observed with the contractive convolutional model, including the latent space collapse, underscore the importance of balancing regularization strength to avoid over-constraining the model. When the contractive loss term is too strong, the model may prioritize minimizing the latent space sensitivity to the extent that it disregards the actual data structure, leading to poor performance.

5.4. Practical Implications

The findings from this study have several practical implications. First, the success of the denoising linear autoencoder suggests that for ECG anomaly detection, model simplicity combined with effective noise management can yield strong performance. This insight is particularly relevant for deployment in resource-constrained environments, where computational efficiency is paramount.

Second, the latent space analysis highlights the importance of selecting appropriate architectures and regularization strategies to ensure that the latent space remains informative and well-structured. The trade-offs between model complexity, regularization strength, and performance must be carefully considered when designing autoencoders for anomaly detection.

Finally, the limitations observed with LSTM-based models in this study suggest that alternative strategies, such as attention mechanisms or hybrid models, may be needed to effectively capture temporal dependencies in unsupervised anomaly detection tasks. Future research could explore these alternatives to improve the performance of sequential models in this context.

5.5. Limitations and Future Work

While the results of this study are promising, several limitations should be noted. The models were tested on a single dataset (ECG5000), and the findings may not generalize to other types of ECG data or different domains. Additionally, the study focused on unsupervised anomaly detection; future work could explore semi-supervised or fully supervised approaches to leverage labeled data for enhanced performance.

Furthermore, the contractive autoencoders showed potential issues with over-regularization, which warrants further investigation. Future work could explore adaptive regularization techniques or alternative forms of regularization to address these challenges. Similarly, the exploration of more sophisticated noise models for the denoising autoencoder could provide insights into further improving robustness and accuracy.

Finally, expanding the latent space analysis to include other dimensionality reduction techniques or incorporating more advanced discriminative models could provide deeper insights into the structure and utility of the learned representations.

6. Conclusion

In conclusion, this study demonstrates that denoising autoencoders, particularly those with simple linear architectures, are highly effective for unsupervised anomaly detection in ECG signals. The ability of these models to generate robust latent representations, coupled with their computational efficiency, makes them strong candidates for deployment in clinical settings where rapid and accurate detection of cardiac anomalies is crucial. The findings also underscore the importance of careful selection and tuning of regularization techniques, as inappropriate combinations can hinder rather than enhance model performance. Future research should continue to explore these themes, with a focus on generalizability, interpretability, and computational efficiency.

7. Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] C. Li, C. Zheng, C. Tai, Detection of ecg characteristic points using wavelet transforms, IEEE Transactions on biomedical Engineering 42 (1995) 21–28.

[2] M. C. Chuah, F. Fu, Ecg anomaly detection via time series analysis, in: Frontiers of High Performance Computing and Networking ISPA 2007 Workshops: ISPA 2007 International Workshops SSDSN, UPWN, WISH, SGC, ParDMCom, HiPCoMB, and IST-AWSN

Niagara Falls, Canada, August 28-September 1, 2007 Proceedings 5, Springer, 2007, pp. 123–135.

[3] M. Woźniak, D. Połap, R. K. Nowicki, C. Napoli, G. Pappalardo, E. Tramontana, Novel approach toward medical signals classifier, in: Proceedings of the International Joint Conference on Neural Networks, volume 2015-September, 2015. doi:10.1109/IJCNN.2015.7280556.

[4] S. eddine Boukredine, E. Mehallel, A. Boualleg, O. Baitiche, A. Rabehi, M. Guermoui, A. Douara, I. E. Tibermacine, Enhanced performance of microstrip antenna arrays through concave modifications and cut-corner techniques, ITEGAM-JETIA 11 (2025) 65–71.

[5] H. Li, P. Boulanger, A survey of heart anomaly detection using ambulatory electrocardiogram (ecg), Sensors 20 (2020) 1461.

[6] G. Sivapalan, K. K. Nundy, S. Dev, B. Cardiff, D. John, Annet: A lightweight neural network for ecg anomaly detection in iot edge sensors, IEEE Transactions on Biomedical Circuits and Systems 16 (2022) 24–35.

[7] A. TIBERMACINE, W. GUETTALA, I. E. TIBERMACINE, Efficient one-stage deep learning for text detection in scene images., Electrotehnica, Electronica, Automatica 72 (2024).

[8] G. D. Clifford, F. Azuaje, P. McSharry, et al., Advanced methods and tools for ECG data analysis, volume 10, Artech house Boston, 2006.

[9] B. Nail, B. Djaidir, I. E. Tibermacine, C. Napoli, N. Haidour, R. Abdelaziz, Gas turbine vibration monitoring based on real data and neuro-fuzzy system, Diagnostyka 25 (2024).

[10] E. Iacobelli, D. Pelella, V. Ponzi, S. Russo, C. Napoli, et al., A fast and accessible neural network based eye-tracking system for real-time psychometric and hci applications, in: CEUR WORKSHOP PROCEEDINGS, volume 3870, CEUR-WS, 2024, pp. 32–41.

[11] G. Zimatore, A. R. Fetoni, G. Paludetti, M. Cavagnaro, M. V. Podda, D. Troiani, Post-processing analysis of transient-evoked otoacoustic emissions to detect 4 khz-notch hearing impairment - a pilot study, Medical Science Monitor 17 (2011) MT41–MT49. doi:10.12659/MSM.881793.

[12] S. Russo, I. E. Tibermacine, A. Tibermacine, D. Chebana, A. Nahili, J. Starczewscki, C. Napoli, Analyzing eeg patterns in young adults exposed to different acrophobia levels: a vr study, Frontiers in Human Neuroscience 18 (2024). doi:10.3389/fnhum.2024.1348154.

[13] G. Zimatore, C. Serantoni, M. C. Gallotta, L. Guidetti, G. Maulucci, M. De Spirito, Automatic detection of aerobic threshold through recurrence quantification analysis of heart rate time series, International Journal of Environmental Research and Public Health 20 (2023). doi:10.3390/ijerph20031998.

[14] I. E. Tibermacine, A. Tibermacine, W. Guettala, C. Napoli, S. Russo, Enhancing sentiment analysis on seed-iv dataset with vision transformers: a comparative study, in: Proceedings of the 2023 11th International Conference on Information Technology: IoT and Smart City, 2023, pp. 238–246.

[15] N. Brandizzi, A. Fanti, R. Gallotta, S. Russo, L. Iocchi, D. Nardi, C. Napoli, Unsupervised pose estimation by means of an innovative vision transformer, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), volume 13589 LNAI, 2023, p. 3 – 20. doi:10.1007/978-3-031-23480-4_1.

[16] E. Iacobelli, V. Ponzi, S. Russo, C. Napoli, Eye-tracking system with low-end hardware: development and evaluation, Information 14 (2023) 644.

[17] C. Napoli, C. Napoli, V. Ponzi, A. Puglisi, S. Russo, I. E. Tibermacine, Exploiting robots as healthcare resources for epidemics management and support caregivers (short paper)., in: AIRO@ AI* IA, 2023, pp. 1–10.

[18] I. Naidji, A. Tibermacine, W. Guettala, I. E. Tibermacine, Semi-mind controlled robots based on reinforcement learning for indoor application., in: ICYRIME, 2023, pp. 51–59.

[19] N. Boutarfaia, S. Russo, A. Tibermacine, I. E. Tibermacine, Deep learning for eeg-based motor imagery classification: Towards enhanced human-machine interaction and assistive robotics, in: CEUR Workshop Proceedings, volume 3695, 2023, p. 68 – 74.

[20] A. Tibermacine, N. Djedi, Neat neural networks to control and simulate virtual creature's locomotion, in: 2014 International Conference on Multimedia Computing and Systems (ICMCS), IEEE, 2014, pp. 9–14.

[21] A. Srivastava, Statistical techniques for anomaly detection, Towards Data Science (2019). URL: https://towardsdatascience.com/statistical-techniques-for-anomaly-detection-6ac89e32d17a.

[22] S. Falciglia, F. Betello, S. Russo, C. Napoli, Learning visual stimulus-evoked eeg manifold for neural image classification, Neurocomputing 588 (2024). doi:10.1016/j.neucom.2024.127654.

[23] G. Medics, Understanding an ecg, https://geekymedics.com/understanding-an-ecg/, 2021.

[24] B. Nail, M. A. Atoussi, S. Saadi, I. E. Tibermacine, C. Napoli, Real-time synchronisation of multiple fractional-order chaotic systems: an application study in secure communication, Fractal and Fractional 8 (2024) 104.

[25] D. Samariya, A. Thakkar, A comprehensive

survey of anomaly detection algorithms, Annals of Data Science 10 (2023) 829–850. URL: https://doi.org/10.1007/s40745-021-00362-9. doi:10.1007/s40745-021-00362-9.

[26] M. Celik, F. Dadaser-Celik, A. S. Dokuz, Anomaly detection in temperature data using dbscan algorithm, in: 2011 International Symposium on Innovations in Intelligent Systems and Applications, 2011, pp. 91–95. doi:10.1109/INISTA.2011.5946052.

[27] E. Iacobelli, V. Ponzi, S. Russo, C. Napoli, Eye-tracking system with low-end hardware: Development and evaluation, Information (Switzerland) 14 (2023). doi:10.3390/info14120644.

[28] A. Tibermacine, D. Akrour, R. Khamar, I. E. Tibermacine, A. Rabehi, Comparative analysis of svm and cnn classifiers for eeg signal classification in response to different auditory stimuli, in: 2024 International Conference on Telecommunications and Intelligent Systems (ICTIS), IEEE, 2024, pp. 1–8.

[29] P. Bergmann, M. Fauser, D. Sattlegger, C. Steger, Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 9584–9592. doi:10.1109/CVPR.2019.00982.

[30] C. Napoli, G. Pappalardo, E. Tramontana, A hybrid neuro-wavelet predictor for qos control and stability, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 8249 LNAI (2013) 527 – 538. doi:10.1007/978-3-319-03524-6_45.

[31] Y. Zou, J. Jeong, L. Pemula, D. Zhang, O. Dabeer, Spot-the-difference self-supervised pre-training for anomaly detection and segmentation, 2022. arXiv:2207.14315.

[32] J. Hyun, S. Kim, G. Jeon, S. H. Kim, K. Bae, B. J. Kang, Reconpatch : Contrastive patch representation learning for industrial anomaly detection, 2023. arXiv:2305.16713.

[33] S. Russo, S. Ahmed, I. E. Tibermacine, C. Napoli, Enhancing eeg signal reconstruction in cross-domain adaptation using cyclegan, in: 2024 International Conference on Telecommunications and Intelligent Systems (ICTIS), IEEE, 2024, pp. 1–8.

[34] R. Brociek, G. D. Magistris, F. Cardia, F. Coppa, S. Russo, Contagion prevention of covid-19 by means of touch detection for retail stores, in: CEUR Workshop Proceedings, volume 3092, 2021, p. 89 – 94.

[35] B. Ladjal, I. E. Tibermacine, M. Bechouat, M. Sedraoui, C. Napoli, A. Rabehi, D. Lalmi, Hybrid models for direct normal irradiance forecasting: a case study of ghardaia zone (algeria), Natural Hazards (2024) 1–23.

[36] K. Batzner, L. Heckler, R. König, Efficientad: Accurate visual anomaly detection at millisecond-level latencies, 2023. arXiv:2303.14535.

[37] S. Bouchelaghem, I. E. Tibermacine, M. Balsi, M. Moroni, C. Napoli, Cross-domain machine learning approaches using hyperspectral imaging for plastics litter detection, in: 2024 IEEE Mediterranean and Middle-East Geoscience and Remote Sensing Symposium (M2GARSS), IEEE, 2024, pp. 36–40.

[38] J. Pereira, M. Silveira, Learning representations from healthcare time series data for unsupervised anomaly detection, in: 2019 IEEE International Conference on Big Data and Smart Computing (BigComp), 2019, pp. 1–7. doi:10.1109/BIGCOMP.2019.8679157.

[39] P. W. Code, Papers with code, 2023. URL: https://paperswithcode.com/sota/outlier-detection-on-ecg5000, accessed on June 12th, 2023.

[40] A. Tibermacine, I. E. Tibermacine, M. Zouai, A. Rabehi, Eeg classification using contrastive learning and riemannian tangent space representations, in: 2024 International Conference on Telecommunications and Intelligent Systems (ICTIS), IEEE, 2024, pp. 1–7.

[41] S. Rifai, P. Vincent, X. Muller, X. Glorot, Y. Bengio, Contractive auto-encoders: Explicit invariance during feature extraction, in: Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11, Omnipress, Madison, WI, USA, 2011, p. 833–840.