



EBAQ: An Entropy-Based Bit Allocation Framework for Lightweight Autoencoder Models

Zaidir¹, Mohammad Diqi²

¹Department of Information Technology, Universitas Respati Yogyakarta, Yogyakarta, Indonesia

²Department of Informatics, Universitas Respati Yogyakarta, Yogyakarta, Indonesia

¹zaidirtan@respati.ac.id, ²dqi@respati.ac.id

Abstract

Autoencoder-based models have shown strong potential for anomaly detection in complex time-series data; however, they often assume equal importance across latent dimensions, resulting in inefficiencies and reduced precision. This study addresses this limitation by introducing the Entropy-Based Bit Allocation Quantizer (EBAQ), a novel quantization framework that adaptively allocates bits to each latent dimension based on its entropy, preserving more precision where information content is highest. The primary objective is to enhance representational efficiency and anomaly detection performance without increasing model complexity or computational cost. EBAQ is implemented as a plug-and-play module within a standard autoencoder architecture, requiring no retraining or architectural modification. The method was evaluated using a publicly available ECG dataset, where reconstruction-based anomaly detection was employed to assess its performance. Results show that EBAQ outperforms the standard autoencoder baseline, achieving higher accuracy (94.9%), precision (99.4%), and recall (91.4%), while also demonstrating more apparent separation between normal and anomalous data in latent space visualizations. These findings confirm that entropy-aware quantization improves both fidelity and interpretability in unsupervised anomaly detection. Overall, this work presents a theoretically grounded and practically efficient solution that bridges information theory and deep learning, offering a human-centered approach to developing more intelligent and efficient AI systems for real-world applications.

Keywords: autoencoder; anomaly detection; edge computing efficiency; entropy-based quantization; latent space representation;

How to Cite: Z. Zaidir and Mohammad Diqi, "EBAQ: An Entropy-Based Bit Allocation Framework for Lightweight Autoencoder Models", *J. RESTI (Rekayasa Sist. Teknol. Inf.)*, vol. 10, no. 1, pp. 171 - 180, Feb. 2026.
Permalink/DOI: <https://doi.org/10.29207/resti.v10i1.6925>

Received: July 9, 2025

Accepted: December 16, 2025

Available Online: February 26, 2026

This is an open-access article under the CC BY 4.0 License

Published by Ikatan Ahli Informatika Indonesia

1. Introduction

In recent years, the growing reliance on real-time data monitoring in critical domains such as healthcare, industrial automation, and cybersecurity has intensified the demand for robust and efficient anomaly detection systems [1]. Clinicians and engineers increasingly depend on continuous data streams to support timely and informed decision-making. Electrocardiogram (ECG) signals provide vital diagnostic information that enables early identification of life-threatening cardiac conditions, including arrhythmias [2]. Healthcare professionals use ECG monitoring to detect subtle physiological changes that may indicate serious risk. However, the complexity and variability of ECG waveforms, combined with the large volume of streaming data, create substantial challenges for accurate and scalable anomaly detection [3].

Traditional machine learning approaches often struggle to generalize effectively in dynamic and high-dimensional environments [4]. Many deep learning models improve predictive performance but require substantial computational resources and offer limited interpretability. These limitations restrict their applicability in real-time and resource-constrained settings. Moreover, most existing models assume equal importance across latent dimensions, ignoring the natural asymmetry in feature-level information contribution [5]. This assumption can produce redundant representations and reduce detection performance in real-world deployments where precision and responsiveness remain critical [6].

Addressing these challenges represents both a technical priority and a human-centered necessity, particularly in clinical decision-support contexts [4]. Accurate and

efficient anomaly detection in streaming biomedical signals can strengthen clinical workflows and reduce diagnostic errors [7]. Healthcare systems can deploy intelligent monitoring tools to assist medical professionals without replacing their expertise. At the same time, minimizing computational overhead remains essential for implementing models on wearable sensors, mobile health platforms, and embedded medical devices [8]. Therefore, researchers must design novel methods that prioritize information-rich features, suppress irrelevant noise, and enable real-time decisions without sacrificing performance [9].

Recent research has explored autoencoder-based models to overcome limitations in latent space representation for anomaly detection. Qien Yu et al. [10] proposed a framework with orthogonal subspaces to enhance structured feature learning. Ghafourian et al. [11] introduced a simple regularization term to stabilize training and improve detection capability. Alaghbari et al. [12] developed an integrated model that combines anomaly detection and feature extraction in IoT networks. Ivanovska and Štruc [13] designed Y-GAN with dual latent spaces to strengthen anomaly discrimination, while Niaz et al. [14] incorporated spatial awareness into 3D convolutional autoencoders for video anomaly detection.

Other studies have investigated vector reconstruction error and probabilistic modeling to enhance detection performance. Torabi et al. [15] proposed using vector reconstruction error as a more sensitive anomaly indicator. Alshameri and Xia [16] evaluated Variational Autoencoders for credit card fraud detection and demonstrated competitive results. Soleimani-Babakamali et al. [17] analyzed dimensionality reduction techniques in structural health monitoring and showed that information loss significantly affects unseen anomaly detection. These findings highlight the importance of preserving meaningful variation within latent representations. They also emphasize the need for principled strategies that manage information distribution effectively.

Advancements in latent vector quantization further reveal unequal information distribution across dimensions in deep learning models. Shin et al. [18] introduced a VQ-VAE framework to improve channel state information feedback efficiency in massive MIMO systems. Kotyan et al. [19] proposed the k* Distribution methodology to analyze class-specific latent sample distributions. Ge et al. [20] developed Contextual Sequential Quantization to improve learned image compression efficiency. Gu et al. [21] presented Entropy Optimized deep Weighted Product Quantization to balance codeword representation, while Chen et al. [22] applied VQ-VAE to cancer subtype identification and improved clustering performance.

Researchers have also explored entropy-based optimization techniques for neural networks and compression tasks. Entropy-constrained implicit neural

representations have improved image compression efficiency [23]. Entropy coding strategies have compressed deep neural network weights without severe accuracy loss [24]. Novel entropy models based on deep Gaussian process regression have enhanced rate-distortion performance in image coding [25]. Pruning and quantization approaches have reduced model complexity while preserving min-entropy evaluation accuracy [26]. Task-aware dynamic optimization methods have allocated network parameters efficiently in multi-task learning scenarios [27].

Despite these advances, a critical gap remains in integrating information-theoretic quantization directly into the latent space of autoencoder-based anomaly detection systems. Existing quantized autoencoders often rely on uniform or fixed-bit schemes that overlook fine-grained differences in feature-level information content [28], [29]. Motivated by this limitation, this study introduces the Entropy-Based Bit Allocation Quantizer (EBAQ), a lightweight and interpretable module that adaptively assigns bit depth based on latent entropy. EBAQ prioritizes information-rich dimensions and compresses redundant ones without requiring architectural modification or retraining. This research aims to develop and evaluate EBAQ as a principled framework that enhances representational efficiency, strengthens anomaly discrimination, and supports reliable deployment in healthcare and edge computing environments [9].

2. Methods

This chapter outlines the methodological framework of the study, describing the systematic procedures and design choices undertaken to develop, implement, and evaluate the proposed anomaly detection approach in a rigorous and human-centered manner.

2.1 Research Methodology Workflow

To provide a clearer overview of the experimental pipeline, the complete research process is illustrated in Figure 1. This diagram outlines the sequential stages from dataset acquisition to model evaluation and conclusion.

The diagram illustrates the end-to-end research workflow for the proposed anomaly detection framework, beginning with ECG dataset collection and systematic preprocessing through normalization and train-test splitting to ensure data consistency and reliability. The processed data are then fed into an autoencoder model consisting of an encoder, latent space, and decoder, forming the foundation for representation learning. At the core of the framework, the EBAQ performs entropy calculation, adaptive bit allocation per latent feature, and latent quantization to enhance representational efficiency. The model is subsequently trained using the Adam optimizer with L1 loss across defined epochs and validated through

reconstruction error analysis to distinguish normal from anomalous signals. Finally, performance is assessed using accuracy, precision, recall, and threshold setting,

followed by comparative analysis, visualization, and concluding insights, ensuring a rigorous, transparent, and human-centered evaluation process.




Figure 1. Research methodology flowchart of the proposed EBAQ-based autoencoder framework

2.2 Dataset and Preprocessing

This study utilizes a publicly available electrocardiogram (ECG) time-series dataset provided by Google TensorFlow and curated specifically for benchmarking anomaly detection algorithms. Each sample contains a single-lead ECG segment with 140 sequential data points that represent the heart's rhythmic electrical activity, and the dataset is labeled into normal and anomalous classes to distinguish physiologically typical patterns from irregular rhythms often associated with arrhythmias. This binary structure supports an unsupervised learning framework in which models train exclusively on normal signals and later evaluate performance on both normal and anomalous data.

As shown in Figure 2, normal ECG traces display smooth and periodic waveforms with clearly defined peaks, whereas anomalous signals exhibit distortions, irregular intervals, or abrupt fluctuations, illustrating both the clinical relevance and the representational complexity of the data. Prior to training, the signals undergo min-max normalization to scale values into the $[0, 1]$ range, the dataset is randomly partitioned into balanced training and testing subsets with only normal samples included in training, and all subsets are converted into PyTorch tensors to enable stable, efficient GPU-based mini-batch processing within a rigorous and human-centered anomaly detection pipeline.




Figure 2. Dataset Sample

2.3 Autoencoder Architecture with EBAQ

This section introduces the integration of the EBAQ into the autoencoder architecture, outlining the structural design and conceptual rationale that guide the proposed framework.

Figure 3 presents the proposed EBAQ integrated into a standard autoencoder framework for anomaly detection.

The model first processes time-series inputs, such as ECG signals, through an encoder that transforms them into a compact latent representation using nonlinear mappings. The EBAQ module then adaptively quantizes this latent vector by assigning different bit depths to each dimension according to its entropy, preserving highly informative features with greater precision while compressing less informative ones. The

decoder reconstructs the signal from the quantized representation, and the system evaluates reconstruction error to identify anomalies. This entropy-guided design balances efficiency and robustness, offering a practical and human-centered solution for real-time health

monitoring and resource-constrained edge environments.

Table 1 summarizes the mathematical notations used in the proposed framework, including the latent matrix, entropy measures, and quantization parameters.




Figure 3. Architecture of the proposed EBAQ

Table 1. Notations and Definitions

Notation	Definition
$\mathbf{Z} \in \mathbb{R}^{N \times D}$	Latent matrix produced by the encoder, with N samples and D dimensions.
$\mathbf{z}^{(d)} \in \mathbb{R}^N$	Vector of latent values for the d -th dimension across all samples.
N	Number of input samples
D	Number of latent dimensions
$p_i^{(d)}$	Estimated probability for the i -th sample in the d -th latent dimension
ϵ	Small constant for numerical stability (e.g., 10^{-8})
$H^{(d)}$	Entropy of the d -th latent dimension
B_{total}	Total number of bits to be allocated across all dimensions
$b^{(d)}$	Number of bits allocated to the d -th latent dimension
$b_{\min}$	Minimum number of bits allowed per dimension (e.g., 1)
$b_{\max}$	Maximum number of bits allowed per dimension (e.g., 4)
$\text{range}^{(d)}$	Value range $[\min(\mathbf{z}^{(d)}), \max(\mathbf{z}^{(d)})]$ used for quantization
$\text{bin}_j^{(d)}$	Bin boundary for the j -th quantization level in dimension d
$\hat{z}_i^{(d)}$	Quantized value of the i -th sample in dimension d
$\hat{\mathbf{Z}} \in \mathbb{R}^{N \times D}$	Quantized latent matrix after applying EBAQ
$\text{clip}(x, a, b)$	Function to limit the value x within the range $[a, b]$

Let $\mathbf{Z} \in \mathbb{R}^{N \times D}$ denote the latent matrix produced by the encoder, where N is the number of samples and D is the number of latent dimensions. For each latent dimension $d \in \{1, 2, \dots, D\}$, we first compute a probability distribution $p^{(d)}$ across its values. This is formulated in Equation 1.

$$p_i^{(d)} = \frac{z_i^{(d)}}{\sum_{j=1}^N z_j^{(d)} + \epsilon}, \quad \text{for } i = 1, \dots, N, \quad (1)$$

ϵ is a small constant introduced for numerical stability.

The entropy of each latent dimension, representing its information richness, is then computed using Shannon entropy as shown in Equation 2.

$$H^{(d)} = -\sum_{i=1}^N p_i^{(d)} \log(p_i^{(d)} + \epsilon) \quad (2)$$

Next, a fixed total bit budget B_{total} is proportionally distributed across dimensions according to their entropy contributions. This allocation strategy is expressed in Equation 3.

$$b^{(d)} = \left\lfloor \frac{H^{(d)}}{\sum_{j=1}^D H^{(j)}} \cdot B_{\text{total}} \right\rfloor \quad (3)$$

subject to the constraint $b_{\min} \leq b^{(d)} \leq b_{\max}$, where $b_{\min}$ and $b_{\max}$ are practical limits (e.g., 1 to 4 bits).

Finally, each latent dimension is quantized into $2^{b^{(d)}}$ bins using uniform bucketization, as defined in Equation 4.

$$\hat{z}_i^{(d)} = \text{bucketize}(z_i^{(d)}, \text{bins}_{b(d)}) \quad (4)$$

The result is a quantized latent matrix $\hat{Z} \in \mathbb{R}^{N \times D}$, which is then forwarded to the decoder for reconstruction. This entropy-aware quantization ensures that information-dense dimensions retain high precision, while redundant features are efficiently compressed.

In a conventional autoencoder, the encoder $f_\theta: \mathbb{R}^n \rightarrow \mathbb{R}^d$ compresses high-dimensional input data $x \in \mathbb{R}^n$ into a latent vector $z = f_\theta(x)$, which is then reconstructed by the decoder $g_\phi: \mathbb{R}^d \rightarrow \mathbb{R}^n$ as $\hat{x} = g_\phi(z)$. However, this continuous latent representation treats all dimensions equally, disregarding their varying informational contributions.

EBAQ addresses this limitation by inserting an intermediate quantization layer $Q: \mathbb{R}^d \rightarrow \mathbb{R}^d$, which discretizes z into $\hat{z} = Q(z)$ according to the entropy-based allocation defined in Equations (1)–(4). The updated forward pass of the model is expressed in Equation 5.

$$\hat{x} = g_\phi(Q(f_\theta(x))) \quad (5)$$

This design requires no retraining or modification of the encoder–decoder parameters, yet it significantly improves the model’s ability to preserve critical information under constrained representational capacity. By aligning the precision of each latent dimension with its relative importance, the EBAQ-augmented autoencoder achieves a principled trade-off between compression and fidelity. Such a strategy is especially valuable in resource-constrained environments, where computational efficiency and diagnostic accuracy must coexist.

2.4 Training Parameters

The training configuration in this study was carefully designed to balance accuracy, stability, and practical feasibility. The autoencoder was optimized using Adam with a learning rate of 1×10^{-3} to effectively handle the non-convex optimization landscape of deep networks. A mini-batch size of 512 ensured stable gradient updates while efficiently utilizing GPU memory, and mean absolute error (L1 loss) was selected as the objective function due to its robustness to outliers in anomaly detection tasks.

The model was trained for 20 epochs to achieve convergence without overfitting, while the EBAQ module operated under a 16-bit total budget with 1 to 4 bits allocated per latent dimension for fine-grained, entropy-aware control. All experiments were implemented in PyTorch with fixed random seeds on standard GPU hardware, reflecting a rigorous yet human-centered approach to building efficient and reliable anomaly detection systems.

2.5 Validation Scenario

The validation scenario was designed to reflect real-world clinical settings, where normal physiological patterns are common and anomalies are relatively rare. The autoencoder was trained exclusively on normal ECG signals so that it could learn a compact representation of healthy cardiac activity without exposure to abnormal patterns. During testing, the dataset included both normal and anomalous signals to evaluate the model’s ability to generalize beyond its training distribution.

The model was expected to reconstruct normal samples accurately while producing higher reconstruction errors for unfamiliar anomalous signals, which served as the detection criterion. This approach aligns with a human-centered perspective on clinical monitoring, where reliable recognition of typical patterns and sensitivity to rare but critical deviations are equally essential.

2.6 Evaluation Metrics

To rigorously assess the effectiveness of the proposed EBAQ, this study employed standard evaluation metrics, including Accuracy, Precision, Recall, and a reconstruction-error-based threshold. Accuracy offers an overall measure of correct classifications, but it may be misleading in imbalanced anomaly detection settings where normal samples dominate. Therefore, greater emphasis was placed on Precision and Recall, which provide more meaningful insight into detection quality. Precision reflects the proportion of correctly identified anomalies among all predicted anomalies, indicating the model’s ability to reduce false alarms, while Recall measures the proportion of true anomalies successfully detected, capturing sensitivity to rare but critical events. Together, these metrics ensure a balanced, reliable, and human-centered evaluation aligned with the practical and ethical demands of real-world monitoring systems.

Mathematically, let TP (true positives) denote correctly identified anomalies, TN (true negatives) denote correctly identified normal samples, FP (false positives) denote normal samples misclassified as anomalies, and FN (false negatives) denote anomalies missed by the model. Then:

Accuracy is defined in Equation 6.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (6)$$

Precision is given in Equation 7.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (7)$$

Recall is defined in Equation 8.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (8)$$

In addition, anomaly detection decisions were determined by comparing the reconstruction error of each sample against a decision threshold, as shown in Equation 9.

$$\text{Anomaly}(x) = \begin{cases} 1, & \text{if } \mathcal{L}(x, \hat{x}) > \tau \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

$\mathcal{L}(x, \hat{x})$ represents the reconstruction error (using L1 loss in this study) and τ is the threshold, empirically set to the mean reconstruction loss of the training data plus one standard deviation. This strategy ensures sensitivity to deviations while maintaining robustness against noise.

3. Results and Discussions

This chapter presents and interprets the experimental findings, providing a structured and human-centered analysis of the model's performance across multiple evaluation perspectives.

3.1 Training and Validation Loss

Figures 4 and 5 present the training and validation loss curves over 20 epochs for the proposed EBAQ and the standard autoencoder baseline, respectively, offering a clear view of convergence behavior and generalization performance. In Figure 4, both training and validation losses decrease smoothly and remain closely aligned throughout the learning process, indicating stable optimization and minimal overfitting despite the integration of entropy-guided quantization. This consistent alignment suggests that EBAQ preserves essential latent information while maintaining strong reconstruction fidelity on unseen data. In contrast, Figure 5 shows that although the baseline autoencoder achieves a steady reduction in training loss, its validation loss plateaus at a higher level and gradually diverges, signaling modest overfitting and weaker generalization. The comparison highlights how entropy-aware bit allocation promotes a more discriminative and efficient latent representation, improving robustness without sacrificing stability.




Figure 4. Training and Validation Loss for EBAQ




Figure 5. Training and Validation Loss for AE

To complement this analysis, Table 1 summarizes the anomaly detection performance of EBAQ and the standard autoencoder using reconstruction-error-based thresholds, reporting Accuracy, Precision, and Recall for each model. These metrics provide an objective basis for evaluating how effectively each approach distinguishes normal from anomalous ECG signals. By presenting both models side by side, the table offers a transparent comparison of their strengths and limitations in practical detection scenarios. Together with Figures 4 and 5, these results establish a comprehensive and human-centered evaluation framework that balances technical rigor with real-world relevance.

3.2 Reconstruction Signals

Figures 6 and 7 provide side-by-side visualizations of signal reconstructions for both normal and anomalous ECG inputs, generated by the EBAQ and the standard autoencoder (AE Standard), respectively. Each figure consists of two panels: the left panel displays a representative normal signal, and the right panel shows an anomalous signal. In both figures, the original input signal is overlaid with the reconstructed output, allowing for a direct visual comparison of waveform fidelity. At the same time, shaded areas between the two curves highlight the reconstruction error. These figures aim to provide a qualitative understanding of each model's reconstruction capability under varying physiological conditions, illustrating how each approach responds to typical versus atypical patterns in the data. By presenting both normal and anomalous examples under identical visualization schemes, Figures 5 and 6 support a human-centered evaluation of model performance and lay the groundwork for further interpretive analysis of latent behavior and detection reliability.

Figure 6 offers a compelling visual representation of the EBAQ model's ability to distinguish between normal and anomalous ECG signals through reconstruction. In the left panel, which depicts a normal signal, the reconstructed waveform closely follows the input signal with minimal deviation, and the error band remains narrow demonstrating EBAQ's capacity to retain critical signal features when the input aligns with the learned distribution. This alignment highlights the model's preservation of high-information dimensions, made possible by its adaptive bit allocation strategy. In contrast, the right panel shows the reconstruction of an anomalous ECG signal, where the divergence between input and output becomes visibly more pronounced, particularly in regions with atypical waveform structures. The broader shaded error regions in these segments reflect EBAQ's inability to accurately reconstruct unfamiliar patterns, a desirable characteristic in unsupervised anomaly detection, where deviation from normality is used as a proxy for identifying anomalies. The contrast between the two panels affirms that EBAQ not only compresses the

latent space efficiently but also enhances its expressiveness by selectively preserving dimensions rich in diagnostic information. This behavior illustrates a human-aligned model response: maintaining fidelity to expected patterns while amplifying reconstruction errors when confronted with the unexpected.

Figure 7 illustrates the reconstruction performance of the standard autoencoder when applied to both normal and anomalous ECG signals, offering valuable insight into the model's inherent strengths and limitations. In the left panel, which depicts a normal signal, the reconstruction closely follows the input waveform in general shape; yet, subtle deviations are visible, particularly in regions with sharper transitions, indicating that even familiar patterns may not be perfectly preserved without adaptive latent control. The right panel, representing an anomalous signal, reveals

more pronounced discrepancies between the input and output, with broad shaded regions denoting elevated reconstruction error across multiple segments. This behavior aligns with the expectation that an autoencoder trained exclusively on normal data will struggle to faithfully reproduce unfamiliar inputs; however, the reconstruction error lacks sharp contrast, suggesting that the model may not distinctly emphasize anomalous features. Without a mechanism to prioritize information-rich latent dimensions, the standard autoencoder distributes representational precision uniformly, potentially diluting its sensitivity to clinically relevant variations. While the model offers a baseline level of functionality, its reconstructions appear less discriminative, reinforcing the need for more intelligent quantization strategies such as those introduced by EBAQ to enhance both fidelity and anomaly separation in unsupervised settings.




Figure 6 Reconstruction Signals for EBAQ




Figure 7 Reconstruction Signals for Standard AE

3.3 Quantitative Evaluation

Table 2 summarizes the quantitative comparison between the proposed EBAQ and the standard autoencoder (AE Standard) using four evaluation metrics: threshold, accuracy, precision, and recall. The classification threshold for each model was defined as the mean reconstruction loss plus one standard deviation from the normal training data, ensuring consistent decision boundaries under identical testing conditions. EBAQ achieves a lower threshold (0.0302) than the baseline (0.0348), indicating greater sensitivity

to subtle reconstruction deviations. In terms of detection performance, EBAQ slightly outperforms the standard model with higher accuracy (0.9490 vs. 0.9450), precision (0.9942 vs. 0.9922), and recall (0.9143 vs. 0.9089). Although these improvements appear numerically modest, they are meaningful in anomaly detection contexts such as healthcare, where small gains can translate into fewer false alarms and more reliable identification of critical events.

Beyond metric differences, the results in Table 2 highlight the practical value of entropy-aware

quantization. By allocating bit precision according to the informational richness of each latent dimension, EBAQ preserves essential features while reducing redundancy, leading to a more discriminative and efficient representation. This adaptive strategy enhances generalization without modifying the underlying autoencoder architecture or requiring retraining. Importantly, the performance gains occur within the same computational framework, reinforcing EBAQ's suitability for resource-constrained environments such as edge or clinical monitoring systems. Overall, the consistent improvement across all reported metrics demonstrates that integrating information-theoretic principles into latent space design can produce measurable and human-centered benefits in real-world anomaly detection.

Table 2. Quantitative Evaluation

Method	Threshold	Accuracy	Precision	Recall
EBAQ	0.030237969	0.9490	0.9942	0.9143
AE Standard	0.034770142	0.9450	0.9922	0.9089

3.4 Latent Space Visualization

Figures 8 and 9 visualize the two-dimensional latent space representations of the EBAQ-enhanced autoencoder and the standard autoencoder, respectively, using Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE). These techniques project high-dimensional encoded ECG features into an interpretable space, enabling a clearer understanding of how each model organizes normal and anomalous signals. In Figure 8, the PCA view shows moderate overlap, which is expected given its linear projection, but the t-SNE visualization reveals well-defined and distinct clusters between normal and anomalous samples. This clearer separation indicates that EBAQ strengthens the discriminative structure of the latent space by preserving high-entropy, information-rich dimensions while compressing less relevant features. The resulting clusters provide an intuitive and human-readable confirmation that entropy-aware quantization supports meaningful differentiation between healthy and abnormal cardiac patterns.




Figure 8. Latent Space Visualization for EBAQ




Figure 9. Latent Space Visualization for Standard AE

In contrast, Figure 9 illustrates that the standard autoencoder produces more overlapping latent representations. Both the PCA and t-SNE projections display substantial mixing between normal and

anomalous samples, suggesting weaker class disentanglement and reduced anomaly separability. The intertwined distribution indicates that the baseline model compresses signals into similar latent regions

without selectively emphasizing informative features. This uniform treatment of latent dimensions limits the model's ability to amplify subtle but clinically important deviations. Compared with the clearer structure observed in Figure 8, these results underscore how entropy-guided precision allocation enhances interpretability, robustness, and alignment with human needs for transparent anomaly detection.

3.5 Key Finding

The central finding of this study is that integrating entropy-aware quantization through the EBAQ improves both the efficiency and interpretability of latent representations in autoencoder-based anomaly detection. By allocating bit precision according to the informational contribution of each latent dimension, EBAQ preserves critical signal features while reducing redundancy, leading to higher precision and recall without increasing architectural complexity.

Compared with the standard autoencoder, the EBAQ-enhanced model achieves more stable reconstruction, clearer latent space separation, and stronger generalization to anomalous patterns. These results demonstrate that embedding information-theoretic principles into deep learning frameworks offers a practical, efficient, and human-centered pathway for developing reliable anomaly detection systems in real-world settings.

3.6 Limitations and Future Work

Although the proposed EBAQ improves anomaly detection performance and latent space efficiency, several limitations remain. The current implementation employs non-differentiable bucket-based quantization, which prevents end-to-end optimization of bit allocation alongside encoder training. In addition, entropy estimation depends on sufficient statistical diversity within mini-batches, which may limit effectiveness in small or highly imbalanced datasets. This constraint could lead to suboptimal bit distribution under data scarcity or distribution shifts. Furthermore, the framework has been validated only on one-dimensional ECG signals, restricting its demonstrated applicability to other data types such as images, multi-channel biosignals, or structured tabular data.

Future research will focus on developing a differentiable version of EBAQ that supports joint optimization with the autoencoder through gradient-based methods. Techniques such as soft quantization or straight-through estimators may enable this integration while preserving computational efficiency. Expanding the framework to handle multimodal and high-dimensional datasets will also be essential to evaluate its robustness across domains. Such extensions could broaden its impact in areas including medical imaging, industrial monitoring, and cybersecurity.

Investigating dynamic bit allocation in streaming or online settings also represents a promising direction, particularly for edge devices that require real-time

adaptation under varying resource constraints. Lastly, future studies may benefit from incorporating explainability modules to better understand how entropy-aware representations contribute to anomaly detection, thereby bridging the gap between algorithmic performance and clinical or operational trust.

4. Conclusions

This study introduced and validated the EBAQ, a lightweight and principled framework that enhances anomaly detection in autoencoder-based models by allocating bit precision according to the entropy of each latent dimension. By aligning representational accuracy with informational richness, EBAQ overcomes the limitations of uniform quantization and delivers measurable improvements in efficiency and interpretability. In ECG anomaly detection experiments, EBAQ increased accuracy from 94.50% to 94.90%, precision from 99.22% to 99.42%, and recall from 90.89% to 91.43% compared with a standard autoencoder, gains that are clinically meaningful because they reduce false alarms and improve the detection of critical abnormalities. Latent space visualizations further confirmed that entropy-guided quantization produces more discriminative representations without requiring architectural changes or retraining, making the method suitable for resource-constrained environments such as wearable or edge devices. Overall, this work advances the integration of information-theoretic principles into neural network design and offers a practical, human-centered pathway toward efficient, reliable, and trustworthy anomaly detection systems, while future research will explore differentiable variants, multimodal extensions, and adaptive real-time deployment strategies.

Acknowledgements

The author declares no conflict of interest. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] M. Nawaz and J. Ahmed, "Cloud-based healthcare framework for real-time anomaly detection and classification of 1-D ECG signals," *PLOS ONE*, vol. 17, no. 12, p. e0279305, Dec. 2022, doi: [10.1371/journal.pone.0279305](https://doi.org/10.1371/journal.pone.0279305).
- [2] J. Qin et al., "A novel temporal generative adversarial network for electrocardiography anomaly detection," *Artificial Intelligence in Medicine*, vol. 136, p. 102489, Feb. 2023, doi: [10.1016/j.artmed.2023.102489](https://doi.org/10.1016/j.artmed.2023.102489).
- [3] H. Zhou and C. Kan, "Tensor-Based ECG Anomaly Detection toward Cardiac Monitoring in the Internet of Health Things," *Sensors*, vol. 21, no. 12, 2021, doi: [10.3390/s21124173](https://doi.org/10.3390/s21124173), <https://doi.org/10.3390/s21124173>.
- [4] A. Raza, K. P. Tran, L. Koehl, and S. Li, "AnoFed: Adaptive anomaly detection for digital health using transformer-based federated learning and support vector data description," *Engineering Applications of Artificial Intelligence*, vol. 121, p. 106051, May 2023, doi: [10.1016/j.engappai.2023.106051](https://doi.org/10.1016/j.engappai.2023.106051).
- [5] T. Dissanayake, T. Fernando, S. Denman, S. Sridharan, and C. Fookes, "DConv-LSTM-Net: A Novel Architecture for Single- and 12-Lead ECG Anomaly Detection," *IEEE Sensors Journal*, vol. 23, no. 19, pp. 22763–22776, Oct. 2023, doi: [10.1109/JSEN.2023.3300752](https://doi.org/10.1109/JSEN.2023.3300752).

- [6] M. Roy, S. Majumder, A. Halder, and U. Biswas, "ECG-NET: A deep LSTM autoencoder for detecting anomalous ECG," *Engineering Applications of Artificial Intelligence*, vol. 124, p. 106484, Sep. 2023, doi: [10.1016/j.engappai.2023.106484](https://doi.org/10.1016/j.engappai.2023.106484).
- [7] M. R. Azghadi et al., "Hardware Implementation of Deep Network Accelerators Towards Healthcare and Biomedical Applications," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 6, pp. 1138–1159, Dec. 2020, doi: [10.1109/TBCAS.2020.3036081](https://doi.org/10.1109/TBCAS.2020.3036081).
- [8] Y. Liu et al., "Deep Anomaly Detection for Time-Series Data in Industrial IoT: A Communication-Efficient On-Device Federated Learning Approach," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6348–6358, Apr. 2021, doi: [10.1109/JIOT.2020.3011726](https://doi.org/10.1109/JIOT.2020.3011726).
- [9] Q. Zhang, R. Han, G. Xin, C. H. Liu, G. Wang, and L. Y. Chen, "Lightweight and Accurate DNN-Based Anomaly Detection at Edge," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 11, pp. 2927–2942, Nov. 2022, doi: [10.1109/TPDS.2021.3137631](https://doi.org/10.1109/TPDS.2021.3137631).
- [10] Q. Yu, C. Li, Y. Zhu, and T. Kurita, "Convolutional autoencoder based on latent subspace projection for anomaly detection," *Methods*, vol. 214, pp. 48–59, Jun. 2023, doi: [10.1016/j.ymeth.2023.04.007](https://doi.org/10.1016/j.ymeth.2023.04.007).
- [11] A. Ghafourian, H. Shui, D. Upadhyay, R. Gupta, D. Filev, and I. Soltani, "Targeted Collapse Regularized Autoencoder for Anomaly Detection: Black Hole at the Center," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 6, pp. 10348–10358, Jun. 2025, doi: [10.1109/TNNLS.2024.3472456](https://doi.org/10.1109/TNNLS.2024.3472456).
- [12] K. A. Alaghbari, H.-S. Lim, M. H. M. Saad, and Y. S. Yong, "Deep Autoencoder-Based Integrated Model for Anomaly Detection and Efficient Feature Extraction in IoT Networks," *IoT*, vol. 4, no. 3, pp. 345–365, Aug. 2023, doi: [10.3390/iot4030016](https://doi.org/10.3390/iot4030016).
- [13] M. Ivanovska and V. Štruc, "Y-GAN: Learning dual data representations for anomaly detection in images," *Expert Systems with Applications*, vol. 248, p. 123410, Aug. 2024, doi: [10.1016/j.eswa.2024.123410](https://doi.org/10.1016/j.eswa.2024.123410).
- [14] A. Niaz, S. Ul Amin, S. Soomro, H. Zia, and K. Nam Choi, "Spatially Aware Fusion in 3D Convolutional Autoencoders for Video Anomaly Detection," *IEEE Access*, vol. 12, pp. 104770–104784, 2024, doi: [10.1109/access.2024.3435144](https://doi.org/10.1109/access.2024.3435144).
- [15] H. Torabi, S. L. Mirtaheri, and S. Greco, "Practical autoencoder based anomaly detection by using vector reconstruction error," *Cybersecurity*, vol. 6, no. 1, Jan. 2023, doi: [10.1186/s42400-022-00134-9](https://doi.org/10.1186/s42400-022-00134-9).
- [16] F. Alshameri and R. Xia, "An Evaluation of Variational Autoencoder in Credit Card Anomaly Detection," *Big Data Min. Anal.*, vol. 7, no. 3, pp. 718–729, Sep. 2024, doi: [10.26599/bdma.2023.9020035](https://doi.org/10.26599/bdma.2023.9020035).
- [17] M. H. Soleimani-Babakamali, R. Soleimani-Babakamali, R. Sarlo, M. F. Farghally, and I. Lourentzou, "On the effectiveness of dimensionality reduction for unsupervised structural health monitoring anomaly detection," *Mechanical Systems and Signal Processing*, vol. 187, p. 109910, Mar. 2023, doi: [10.1016/j.ymssp.2022.109910](https://doi.org/10.1016/j.ymssp.2022.109910).
- [18] J. Shin, Y. Kang, and Y.-S. Jeon, "Vector Quantization for Deep-Learning-Based CSI Feedback in Massive MI MO Systems," *IEEE Wireless Commun. Lett.*, vol. 13, no. 9, pp. 2382–2386, Sep. 2024, doi: [10.1109/lwc.2024.3415633](https://doi.org/10.1109/lwc.2024.3415633).
- [19] S. Kotyan, T. Ueda, and D. V. Vargas, "k* Distribution: Evaluating the Latent Space of Deep Neural Networks Using Local Neighborhood Analysis," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2024, doi: [10.1109/TNNLS.2024.3446509](https://doi.org/10.1109/TNNLS.2024.3446509).
- [20] Z. Ge, S. Ma, W. Gao, J. Pan, and C. Jia, "NLIC: Non-Uniform Quantization-Based Learned Image Compression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 10, pp. 9647–9663, Oct. 2024, doi: [10.1109/tcsvt.2024.3401872](https://doi.org/10.1109/tcsvt.2024.3401872).
- [21] L. Gu, J. Liu, X. Liu, W. Wan, and J. Sun, "Entropy-Optimized Deep Weighted Product Quantization for Image Retrieval," *IEEE Trans. on Image Process.*, vol. 33, pp. 1162–1174, 2024, doi: [10.1109/tip.2024.3359066](https://doi.org/10.1109/tip.2024.3359066).
- [22] Z. Chen et al., "Learning vector quantized representation for cancer subtypes identification," *Computer Methods and Programs in Biomedicine*, vol. 236, p. 107543, Jun. 2023, doi: [10.1016/j.cmpb.2023.107543](https://doi.org/10.1016/j.cmpb.2023.107543).
- [23] S. Lee, J.-B. Jeong, and E.-S. Ryu, "Entropy-Constrained Implicit Neural Representations for Deep Image Compression," *IEEE Signal Process. Lett.*, vol. 30, pp. 663–667, 2023, doi: [10.1109/lsp.2023.3279780](https://doi.org/10.1109/lsp.2023.3279780).
- [24] G. C. Marinò, F. Furia, D. Malchiodi, and M. Frasca, "Efficient and Compact Representations of Deep Neural Networks via Entropy Coding," *IEEE Access*, vol. 11, pp. 106103–106125, 2023, doi: [10.1109/access.2023.3317293](https://doi.org/10.1109/access.2023.3317293).
- [25] M. Cao et al., "End-to-End Optimized Image Compression With Deep Gaussian Process Regression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 35, no. 4, pp. 3770–3785, Apr. 2025, doi: [10.1109/tcsvt.2022.3209661](https://doi.org/10.1109/tcsvt.2022.3209661).
- [26] H. Li, J. Zhang, Z. Li, J. Liu, and Y. Wang, "Improvement of Min-Entropy Evaluation Based on Pruning and Quantized Deep Neural Network," *IEEE Trans. Inform. Forensic Secur.*, vol. 18, pp. 1410–1420, 2023, doi: [10.1109/tifs.2023.3240859](https://doi.org/10.1109/tifs.2023.3240859).
- [27] S. Choi, H. Jin, and E. Kim, "Task-Aware Dynamic Model Optimization for Multi-Task Learning," *IEEE Access*, vol. 11, pp. 137709–137717, 2023, doi: [10.1109/access.2023.3339793](https://doi.org/10.1109/access.2023.3339793).
- [28] B. S. Sharmila and R. Nagapadma, "Quantized autoencoder (QAE) intrusion detection system for anomaly detection in resource-constrained IoT devices using RT-IoT2022 dataset," *Cybersecurity*, vol. 6, no. 1, Sep. 2023, doi: [10.1186/s42400-023-00178-5](https://doi.org/10.1186/s42400-023-00178-5).
- [29] M. S. Alam, W. A. Misba, and J. Atulasimha, "Quantized non-volatile nanomagnetic domain wall synapse based autoencoder for efficient unsupervised network anomaly detection," *Neuromorph. Comput. Eng.*, vol. 4, no. 2, p. 024012, Jun. 2024, doi: [10.1088/2634-4386/ad49ce](https://doi.org/10.1088/2634-4386/ad49ce).