Unsupervised Time-Series Signal Analysis with Autoencoders and Vision Transformers: A Review of Architectures and Applications

Hossein Ahmadi, Sajjad Emdadi Mahdimahalleh,* Arman Farahat, Banafsheh Saffari

The University of Akron, Akron, OH, USA

April 23, 2025

Abstract

The rapid growth of unlabeled time-series data in domains such as wireless communications, radar, biomedical engineering, and the Internet of Things (IoT) has driven advancements in unsupervised learning. This review synthesizes recent progress in applying autoencoders and vision transformers for unsupervised signal analysis, focusing on their architectures, applications, and emerging trends. We explore how these models enable feature extraction, anomaly detection, and classification across diverse signal types, including electrocardiograms, radar waveforms, and IoT sensor data. The review highlights the strengths of hybrid architectures and self-supervised learning, while identifying challenges in interpretability, scalability, and domain generalization. By bridging methodological innovations and practical applications, this work offers a roadmap for developing robust, adaptive models for signal intelligence.

Keywords: Unsupervised Learning, Autoencoders, Vision Transformers, Time-Series Analysis, Signal Processing, Representation Learning, Anomaly Detection, Wireless Signals, Biomedical Signals, Radar, IoT.

1 Introduction

With the increasing volume of unlabeled time-series data in domains such as wireless communications, radar systems, biomedical engineering, and the Internet of Things (IoT), the need for robust unsupervised learning methods has become more pronounced. Traditional supervised models—including CNNs and RNNs—require large labeled datasets, which are often expensive or impractical to collect in real-world environments. In contrast, unsupervised learning models, particularly Autoencoders (AEs) and Vision Transformers (ViTs), offer flexible alternatives for extracting meaningful representations, detecting anomalies, and classifying signals without reliance on manual labeling.

Autoencoders leverage reconstruction-based objectives to learn compact latent spaces that are effective for anomaly detection, denoising, and dimensionality reduction. Their ability to uncover underlying structure from noisy or incomplete data makes them well-suited for medical signals, IoT monitoring, and RF classification tasks. Vision Transformers, originally developed for visual tasks, are increasingly adapted to sequential signal data due to their self-attention mechanisms, which allow modeling of long-range temporal dependencies and spatial correlations—especially when signals are transformed into time-frequency images such as spectrograms or scalograms.

---

*Corresponding author. Email: se67@uakron.edu

This review aims to provide a cross-domain synthesis of how AEs, ViTs, and their hybrid forms are applied in unsupervised and semi-supervised settings for time-series signal analysis. It identifies recent methodological innovations, summarizes performance across benchmark datasets, and highlights opportunities for future research. Specifically, this review is guided by the following questions:

1. Why are Autoencoders and Vision Transformers particularly effective for time-series data?

2. Which tasks are most frequently addressed using these models?

3. What public datasets are most often used in these domains?

4. How have researchers adapted ViTs and AEs to better suit signal data?

By addressing these questions, we aim to connect theoretical developments with practical implementations across four key domains: medical signals, IoT time series, wireless signals, and radar data. The remainder of this paper is organized as follows: Section 2 introduces theoretical foundations, Section 3 surveys methodologies by domain, Section 4 reviews applications, Section 5 presents a comparative cross-domain analysis, Section 6 discusses challenges and future directions, and Section 7 concludes.

1.1 Contribution and Novelty

This review sets itself apart by offering a comprehensive and comparative perspective on the use of AEs and ViTs in unsupervised signal classification, an area where the intersection of these two powerful architectures has received limited consolidated attention. While existing surveys have typically focused on Autoencoders in the context of general time-series modeling [1, 2] or on ViTs in computer vision [3, 4], there remains a noticeable gap in synthesizing their roles and performance across signal domains.

What distinguishes this work is its cross-disciplinary scope. By examining applications in biomedical signal processing, wireless communications, radar-based sensing, and Internet of Things (IoT) monitoring, we highlight how these models—often developed within isolated research communities—are converging around shared methodological themes. Among these are the use of masked pretraining strategies [5], hybrid ViT-AE architectures that balance reconstruction with global context modeling [6], and spectral-domain encoding techniques for time-series signals [7].

Beyond cataloging model architectures, we delve into the practical design choices that shape performance, including signal transformation methods, patch configuration, and the interplay between data sparsity and model robustness. This analysis brings to light consistent patterns and trade-offs across domains, offering insights into how best to adapt these models for real-world deployment.

To the best of our knowledge, this is the first review to bring together the rapidly evolving literatures on AEs and ViTs under a unified lens focused on unsupervised time-series analysis. Our contribution lies not only in mapping technical advancements but also in identifying future directions that can foster greater model generalization, interpretability, and cross-domain transferability in signal processing applications.

2 Background

AEs and ViTs have significantly transformed unsupervised learning for signal classification, outperforming traditional approaches such as K-Means and Principal Component Analysis (PCA) in modeling complex, high-dimensional data. AEs are particularly effective in tasks involving reconstruction and feature learning, with advanced variants like Variational Autoencoders

(VAEs) and Masked Autoencoders (MAEs) extending their utility to generative modeling and self-supervised learning [5, 8]. On the other hand, ViTs leverage self-attention mechanisms to capture long-range dependencies, which is especially advantageous in processing spatiotemporal signals [9].

Hybrid AE-ViT architectures have further advanced the field by integrating the representational power of AEs with the contextual awareness of ViTs, yielding state-of-the-art results in challenging domains such as medical diagnostics [5], radar recognition [10], and IoT monitoring [11]. This synergy forms the theoretical and practical basis for analyzing signal classification across domains in the subsequent sections.

This section outlines the theoretical underpinnings of unsupervised learning, focusing on AE and ViT-based architectures. We begin with a brief review of traditional clustering and dimensionality reduction techniques, followed by a detailed discussion of AE and ViT frameworks, including their mathematical formulations. This foundation emphasizes the suitability of these models for capturing intricate patterns in high-dimensional, time-series signal data.

2.1 Traditional Unsupervised Learning Techniques

Classical unsupervised learning algorithms, such as K-Means [12], Principal Component Analysis (PCA), and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [13], have been foundational for clustering and dimensionality reduction. These methods are computationally efficient but face limitations with the non-linear, high-dimensional data common in time-series, biomedical signals, and wireless communications. K-Means assumes isotropic clusters, struggling with overlapping or non-spherical distributions. PCA, effective for linear dimensionality reduction, cannot capture intricate non-linear relationships. DBSCAN excels at detecting arbitrarily shaped clusters but falters with varying density levels, a frequent challenge in heterogeneous signal data [14].

2.2 Autoencoder-Based Unsupervised Learning

AEs are neural networks that learn compressed, information-rich representations of input data in an unsupervised manner [1]. Unlike PCA, AEs model non-linear dependencies, making them ideal for tasks like reconstruction, anomaly detection, and feature learning in complex signals. In signal classification, AEs extract robust latent features, enhancing performance in domains like ECG analysis [15] and IoT anomaly detection [16].

Several AE variants address specific signal characteristics:

- **Convolutional Autoencoders (CAEs)** [17]: Capture local spatial patterns in image-like signal representations, such as spectrograms.

- **Long Short-Term Memory Autoencoders (LSTM-AEs)** [18]: Model temporal dependencies in sequential data, suitable for time-series signals.

- **Adversarial Autoencoders (AAEs)** [15]: Integrate generative modeling for enhanced latent space structure.

- **Variational Autoencoders (VAEs)** [8]: Use probabilistic latent spaces for structured feature learning.

Recent advancements include attention mechanisms [19], low-rank attention [20], and AE-Transformer hybrids [6], improving fine-grained feature extraction. Figures 1 and 2 illustrate AE architectures for signal processing.
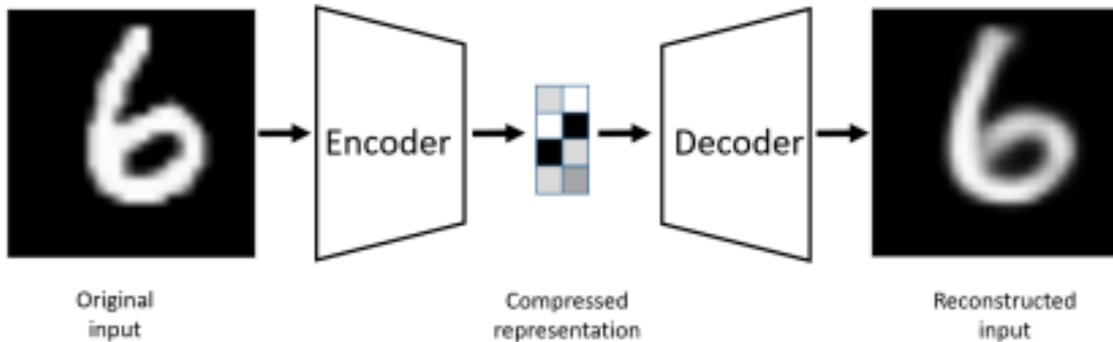
Figure 1: Basic Autoencoder architecture, showing encoder and decoder components [21].
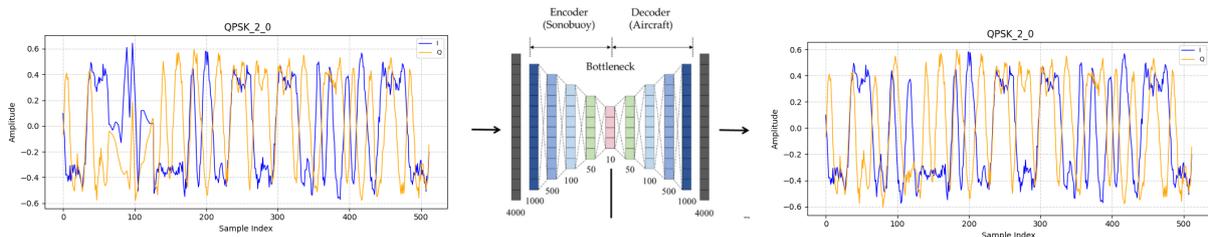
Figure 2: AE-based framework for signal reconstruction, highlighting latent space compression.

2.3 Unsupervised Learning with Vision Transformers

ViTs, originally developed for computer vision [3], have gained prominence in signal processing due to their self-attention mechanisms, which model long-range dependencies. ViTs process inputs as tokenized patches, capturing both local and global structures in data like raw waveforms or spectrograms [7]. In unsupervised signal analysis, ViTs excel in applications such as ECG interpretation [9], IoT anomaly detection [11], and wireless modulation classification [22].

Hybrid AE-ViT models combine AE reconstruction with ViT's expressive power. Examples include NMFormer [6] for noisy modulation classification, Deno-MAE [23] for self-supervised denoising, and ViT classifiers with AE pretraining [24]. Figure 3 shows a standard ViT adapted for signal inputs, while Figure 4 illustrates a hybrid MAE-EEG-Transformer framework.

2.4 Core Model Formulations

This subsection presents the mathematical foundations of AEs and ViTs, focusing on their unsupervised learning mechanisms.

2.4.1 Autoencoders

Autoencoders encode input x into a latent representation $z = f(x; \theta_e)$ and decode it to reconstruct $\hat{x} = g(z; \theta_d)$, minimizing the reconstruction loss:

$$\mathcal{L}_{\text{AE}} = \|x - \hat{x}\|^2 \tag{1}$$

This loss ensures the reconstructed output closely matches the input, capturing essential features [1].
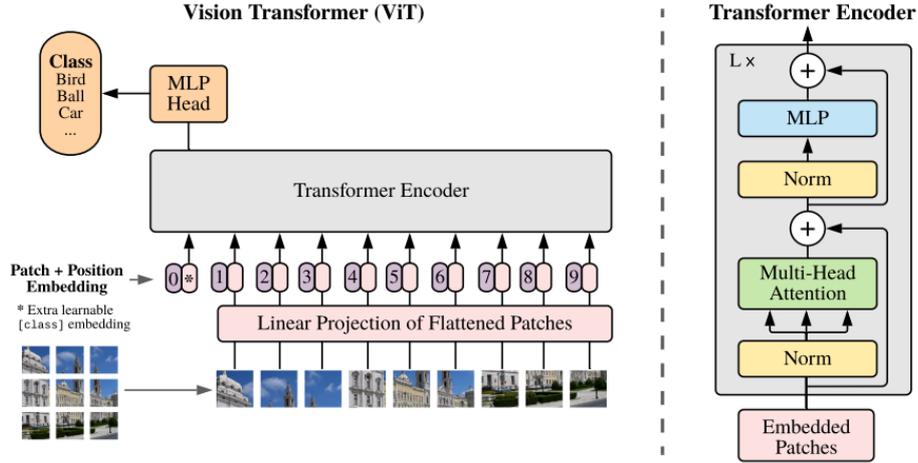
Figure 3: Vision Transformer architecture, processing signal inputs as tokenized patches [3].

2.4.2 Variational Autoencoders (VAEs)

VAEs impose a probabilistic prior on the latent space, with the encoder outputting parameters μ and σ^2 for a Gaussian distribution:

$$z \sim \mathcal{N}(z; \mu, \sigma^2 I) \tag{2}$$

The loss, based on the Evidence Lower Bound (ELBO), balances reconstruction and regularization:

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q(z|x)}[\log p(x|z)] - D_{\text{KL}}(q(z|x)\|p(z)) \tag{3}$$

The ELBO encourages a structured latent space, useful for generative tasks [8]. Figure 5 visualizes a VAE's framework.

2.4.3 Adversarial Autoencoders (AAEs)

AAEs use adversarial training to align the latent space with a prior distribution, combining reconstruction and discriminator losses:

$$\min_{\theta_e, \theta_d} \mathcal{L}_{\text{AE}} \quad / \quad \max_{\psi} \mathbb{E}_{z \sim p(z)}[\log D(z)] + \mathbb{E}_{x \sim p_{\text{data}}(x)}[\log(1 - D(q(z|x)))] \tag{4}$$
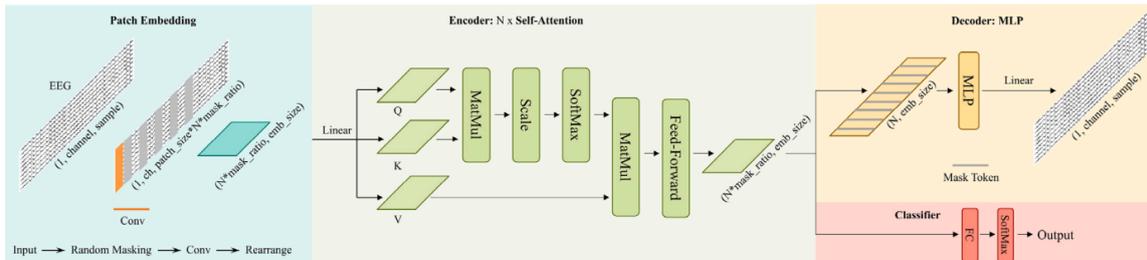
Figure 4: MAE-EEG-Transformer, integrating masked autoencoder pretraining for EEG classification [5].
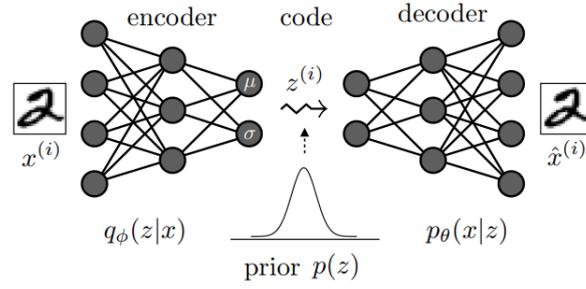
(a) Variational Autoencoder (VAE) framework.
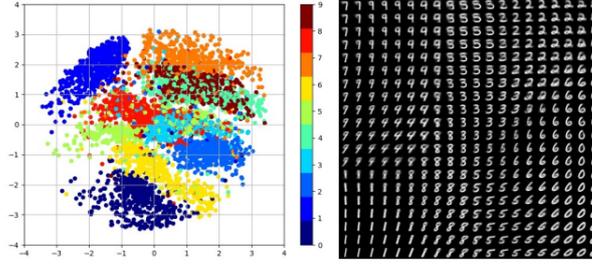
(b) Samples from a trained VAE.

Figure 5: Variational Autoencoder framework, showing probabilistic latent space sampling [25].

This approach enhances latent space robustness [15].

2.4.4 Masked Autoencoders (MAEs)

MAEs, effective for self-supervised learning, reconstruct randomly masked input patches. The workflow includes:

1. Splitting input into fixed-length patches.

2. Masking a subset (e.g., 75%) of patches.

3. Encoding visible patches.

4. Reconstructing the full input.

MAEs excel in pretraining for signal classification [5].

2.4.5 Vision Transformers

ViTs apply scaled dot-product attention to input tokens:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) V \tag{5}$$

Multi-head attention enhances representational diversity:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \ldots, \text{head}_h)W^O \tag{6}$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \tag{7}$$

This architecture captures rich temporal and spatial dependencies [3].

3 Methodologies

In this section, we provide a comprehensive overview of unsupervised learning methodologies applied to signal classification, structured into four primary categories: Medical Signals, IoT Time Series, Wireless Signals and Radar Signal Processing. For each category, we explore the application of ViTs, AEs and the combination of both (Hybrid) , followed by summary tables that compare methodologies, datasets, and associated tasks.

3.1 Medical Signals

Analyzing medical signals such as electrocardiograms (ECG) and electroencephalograms (EEG) requires models capable of capturing both local variations and global temporal patterns. Self-attention-based architectures, particularly ViTs, have proven effective in this context by enabling richer feature extraction from complex physiological data than traditional convolutional approaches.

HeartBEiT [26] applies masked image modeling to ECG signals transformed into 2D formats, demonstrating the benefits of domain-specific pretraining. By masking parts of ECG images during training, the model enhances robustness and generalization in downstream tasks.

Liu et al. [27] proposed ECVT-Net, a hybrid deep learning model that integrates CNNs and ViTs for detecting Congestive Heart Failure (CHF). CNN layers extract localized ECG features, which are then divided into patches and processed by a ViT to capture long-range dependencies. This architecture achieved 98.88% accuracy in inter-patient CHF classification and remained resilient under noisy conditions.

Jamil et al. [28] introduced a ViT-based framework for detecting Valvular Heart Disease (VHD) using phonocardiogram (PCG) signals. The model converts raw PCG signals into time-frequency representations (TFR) using Continuous Wavelet Transform (CWT) and learns spatial representations directly from these images. It achieved 99.9% classification accuracy, outperforming CNN-based models in both performance and efficiency.

Banville et al. [29] employed self-supervised learning for EEG representation extraction using pretext tasks like contrastive predictive coding (CPC), relative positioning (RP), and temporal shuffling (TS). These approaches supported sleep staging and pathology detection in label-scarce environments.

Telangore et al. [30] presented a multimodal ViT-CNN framework for early prediction of sudden cardiac death (SCD). The model combines 1D-CNNs and LSTMs for temporal feature extraction from raw ECG, and uses ViTs with 2D-CNNs for processing scalograms and Hilbert–Huang spectra. It achieved 98.81% accuracy, demonstrating strong predictive potential for early clinical interventions.

AEs have also been widely adopted for analyzing medical signals, offering capabilities in anomaly detection, feature learning, clustering, and reconstruction. Their ability to learn non-linear latent representations makes them well-suited for modeling noisy and complex biomedical data such as ECG and EEG.

Nejedly et al. [31] developed a temporal autoencoder for semi-supervised clustering and classification of intracranial EEG (iEEG). By compressing temporal features and applying kernel density estimation, the model supports large-scale neurophysiological data analysis with minimal labels.

DeepAnT [32], initially designed for general time series, has shown strong performance on medical signals like ECG and EEG. It forecasts future signal values and flags deviations as anomalies, enabling robust unsupervised detection with a convolutional backbone.

Jang et al. [33] introduced a Convolutional Variational Autoencoder (CVAE) that models the variability in ECG patterns through learned latent distributions, facilitating clustering and anomaly detection without requiring annotations.
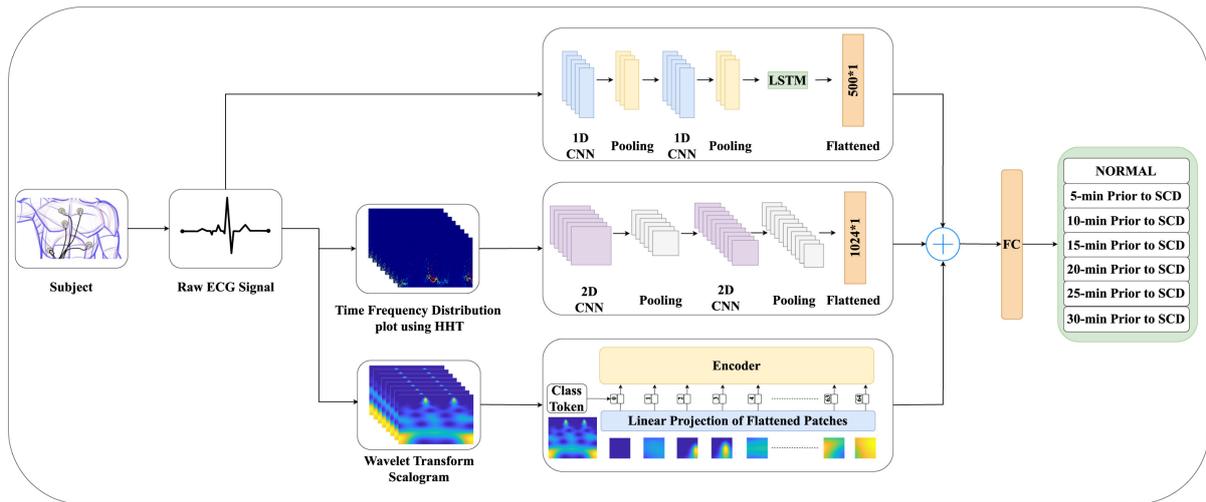
Figure 6: Workflow of the proposed multi-modal feature fusion and sudden cardiac death prediction framework. [30]

Shan et al. [15] used Adversarial Autoencoders (AAEs) with Temporal Convolutional Networks (TCNs) to analyze ECG signals. Their model improves anomaly detection by combining adversarial regularization with reconstruction loss and discriminator scores, helping reduce overfitting in unsupervised settings.

Zhao et al. [34] proposed a method for ECG signal quality assessment based on AE reconstruction loss and likelihood-based scoring. Metrics such as AE-logMSE and AE-LLH outperform conventional quality indicators across datasets including CinC, Sleep, and Stress ECG.

In EEG-based emotion recognition, [35] proposed an autoencoder-based approach that learns subject-specific frequency bands from power spectral density (PSD) instead of using fixed spectral bands. The model yielded 4–20% higher classification accuracy over traditional methods.

For epilepsy detection, Wen et al. [36] developed AE-CDNN, a hybrid combining autoencoders and CNNs. The model encodes EEG data through convolution and pooling, then reconstructs signals via deconvolution, outperforming PCA and Sparse Random Projection (SRP) on benchmark datasets.

Recent developments in self-supervised learning have shown the effectiveness of combining ViTs with autoencoders for medical signal analysis. In particular, Masked Autoencoders (MAEs) integrated with ViTs enable unsupervised learning from raw physiological data, especially in the absence of labels.

Sawano et al. [37] applied the MAE framework to pretrain ViT-Base, ViT-Large, and ViT-Huge architectures on unlabeled 12-lead ECG data. ECG signals were treated as 12×5000 matrices and segmented into 1×250 patches, with 75% masked during training. This approach improved performance on downstream tasks like Left Ventricular Systolic Dysfunction (LVSD) detection, outperforming models pretrained on generic image datasets.

Zhou et al. [5] proposed MTECG, a self-supervised ECG framework that segments ECG time series into patches and reconstructs masked segments using a Transformer encoder-decoder. Like Sawano et al., their method effectively captures both local and global features and performs well on clustering and anomaly detection.

Together, these studies reflect a growing trend toward hybrid ViT-AE architectures in medical AI. By combining ViTs' strength in modeling temporal dependencies with AEs' reconstruction capabilities, such models provide scalable and robust solutions for clinical signal analysis, particularly in data-limited scenarios.

Table 1: Medical Signal Processing Papers using ViT and Autoencoders

Dataset	Method	Task	Ref.
HeartBEiT Dataset	ViT	Pretraining for ECG Classification	[26]
ECG Dataset	ViT & CNN	Detection of Congestive Heart Failure	[27]
PCG Dataset	ViT	Detection of Valvular Heart Diseases	[28]
EEG Dataset	ViT (CPC, RP, TS)	Sleep Staging, Pathology Detection	[29]
ECG Dataset	ViT & CNN	Early prediction of sudden cardiac death	[30]
Intracranial EEG Dataset	Temporal Autoencoder	Clustering, Classification	[31]
Time-series signals (ECG, EEG)	Autoencoder	Anomaly Detection	[32]
ECG Dataset	CVAE	Clustering, Anomaly Detection	[33]
ECG Dataset	AAE (TCN integrated)	Anomaly Detection, Feature Extraction	[15]
CinC, Sleep, Stress ECG	AE-logMSE, AE-LLH	Signal Quality Assessment	[34]
EEG Dataset	AE-NN	Emotion Recognition	[35]
EEG Dataset	AE-CDNN	Epilepsy Detection	[36]
12-lead ECG Dataset	ViT + Masked Autoencoder	LVSD Detection	[37]
MTECG Dataset	ViT + Masked Autoencoder	Anomaly Detection, Clustering	[5]

3.2 IoT Time Series

In IoT time series analysis, capturing long-range dependencies and extracting informative features from high-dimensional streams are critical tasks. Recent methods address these challenges by converting temporal data into image-like structures, allowing Vision Transformers (ViTs) to harness pretrained visual models for improved performance across applications such as anomaly detection, intrusion prevention, and activity recognition.

Li et al. [38] proposed ViTST, a method that converts each variable in a multivariate time series into a separate line graph, then arranges these plots into a grid to form a single image input for a Vision Transformer. This approach achieved strong results across irregular and regular time series, particularly in healthcare, and demonstrated robustness to missing data.

Ni et al. [39] provided a broad survey of imaging methods for time series, such as line plots, heatmaps, spectrograms, Gramian Angular Fields, and Recurrence Plots. Their study emphasized the benefits of applying ViTs to IoT data, highlighting improvements in classification accuracy and resilience to noise and incomplete data.

Zhang et al. [40] introduced TSVIT, an architecture combining 1D CNN layers for patch embedding with a Transformer encoder. This end-to-end model showed excellent performance in fault diagnosis within industrial IoT settings.

Sana et al. [11] developed an intrusion detection framework that integrates ViTs with traditional and deep learning methods to enhance anomaly detection in IoT networks. Tested on the NSL-KDD dataset, the ViT-based model achieved perfect accuracy, F1-score, AUC, and MCC, outperforming conventional models such as SVMs, LSTMs, and ensemble techniques. Bayesian optimization further improved performance, demonstrating the viability of ViTs in real-time, scalable intrusion prevention for IoT systems.

Tarasiou et al. [41] presented a Transformer model for Satellite Image Time Series (SITS) that factorizes attention both temporally and spatially, improving analysis for multivariate IoT datasets with temporal dependencies.

Liu et al. [14] conducted a comprehensive review of unsupervised deep learning methods for IoT time series, focusing on anomaly detection and clustering. Their work addressed challenges

Table 2: Summary of ViT-Based Methods and Autoencoder-Based Methods Applied to IoT Time Series

Dataset Name	Reference	Method	Task
Multivariate IoT Time Series	[38]	ViTST	Time Series Classificati...
Satellite Image Time Series (SITS)	[41]	ViTSITS	Multivariate Time Series A...
Vibration Signal Dataset	[40]	TSVIT	Fault Diagnosis
IoT Time Series Survey	[39]	Survey	Imaging Techniques for IoT...
NSL-KDD	[11]	Vision Transformer (ViT)	Intrusion Anomaly Detec...
IoT Time Series Dataset	[14]	Autoencoder, GAN, CNN, RNN	Anomaly Detection, Clust...
Various IoT Monitoring Systems	[32]	DeepAnT	Anomaly Detection
MOD, ACIDS, RealWorld-HAR, PAMAP2	[42]	FreqMAE	Vehicle Classification, Human Activ...
Industrial Paste Thickener	[43]	Contrastive Blind Denoising Autoencoder	Real-Time Denoising...
Cowrie Honeypot-based IoT Attack Dataset	[44]	Autoencoder	Unsupervised Clustering of IoT...
Intel Lab Data	[45]	DTDA-RNI	Data Compression and Recon...
NSL-KDD	[46]	AAE, KNN	Intrusion Detection on IoT Edg...
Custom Dataset	[16]	VAE	IoT Traffic Anomaly Dete...
Server Machine Dataset, Air Quality Index	[47]	R-CVAE + Transformer	Anomaly Detection
N-BaIoT, CICIoT2022	[48]	VAE + ViT	IoT Botnet Detection...

arising from the high dimensionality and complex spatio-temporal dynamics of IoT data, and highlighted research opportunities for future improvements. A conceptual framework of this analysis is shown in Figure 7.
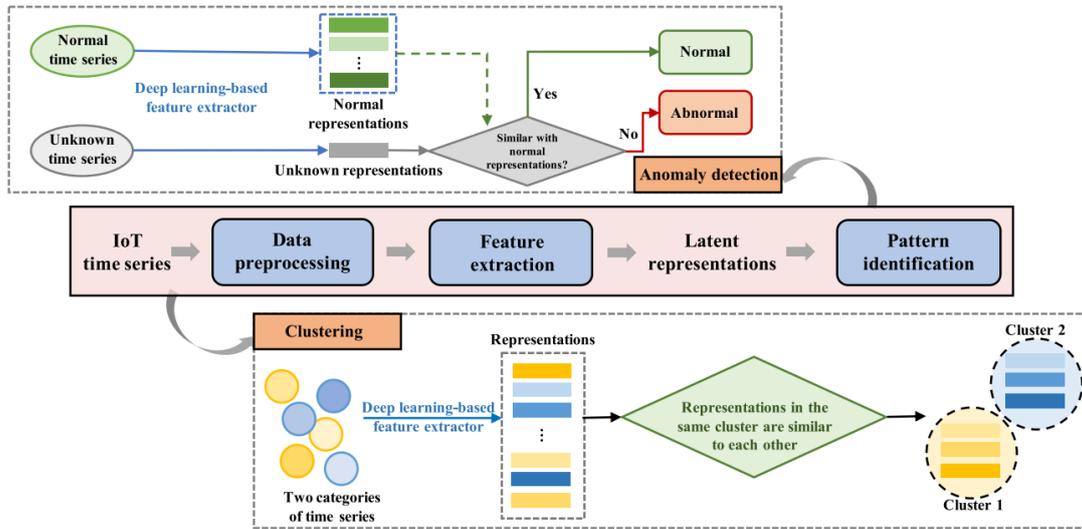
Figure 7: Example of IoT Time-Series Analysis. [14]

AEs have also been widely applied to IoT time series for tasks such as anomaly detection, denoising, and latent feature extraction. Their ability to learn compact representations makes them effective in handling high-dimensional, noisy IoT data streams.

Liu et al. [14] demonstrated the use of AEs, GANs, CNNs, and RNNs for clustering and anomaly detection, focusing on the scalability of these models in heterogeneous IoT networks.

DeepAnT [32], a convolutional autoencoder that forecasts future time points and flags deviations as anomalies, has been effectively used for IoT monitoring across multiple deployments.

Kara et al. [42] proposed FreqMAE, a self-supervised masked autoencoder that incorporates domain-specific signal processing for multi-modal IoT environments. Its architecture includes a

Temporal-Shifting Transformer (TS-T), factorized fusion mechanism, and a frequency-weighted loss, allowing robust representation learning without requiring labels.

Haseeb et al. [44] introduced an AE-based feature construction method to cluster IoT cyberattacks. Their model maps command features to a latent space, enabling more meaningful behavioral clustering than traditional techniques.

Langarica and Núñez [43] developed the Contrastive Blind Denoising Autoencoder (CB-DAE), which applies noise contrastive estimation to regularize latent space during training, allowing real-time denoising without clean signal references in industrial IoT settings.

Xin et al. [45] presented DTDA-RNI, a lightweight denoising autoencoder framework for compressing and cleaning noisy sensor data. By using random noise injection (RNI) during training, the model improved reconstruction accuracy and efficiency compared to compressed sensing approaches, making it suitable for low-bandwidth IoT applications.

Aloul et al. [46] proposed an intrusion detection system using Adversarial Autoencoders (AAEs) combined with a KNN classifier. Deployed on a Raspberry Pi 3B, the model achieved 99.991% accuracy and operated with minimal latency, demonstrating feasibility for edge IoT devices.

Xin et al. [16] designed a hybrid CNN-VAE model for IoT traffic classification and anomaly detection. Reconstruction loss and KL divergence were used to detect abnormal behavior, and Particle Swarm Optimization (PSO) was applied to optimize the deep autoencoder structure.

Yao et al. [47] proposed THREADS, a hierarchical anomaly detection system for IIoT combining edge-deployed R-CVAEs with a cloud-based transformer discriminator. Their dual-thread architecture achieved strong performance while reducing resource usage on constrained devices.

Wasswa et al. [48] compared ViT and VAE encoders for IoT botnet detection. VAEs consistently outperformed ViTs on structured network data due to their strength in modeling non-visual, tabular input, highlighting VAE's advantage in cybersecurity applications over spatially biased ViTs.

3.3 Wireless Signals

Wireless signal processing tasks—ranging from automatic modulation classification (AMC) to anomaly detection and RF signal generation—benefit from models that can handle complex temporal and spectral patterns. Leveraging self-attention mechanisms, ViTs excel at modeling high-dimensional data, particularly when signals are represented as time-frequency diagrams or image-like formats.

Chen et al. [49] proposed an unsupervised RF fingerprinting framework that addresses the domain shift challenge caused by variations in channel conditions and environmental factors. Their method leverages contrastive learning and processes spectrograms of raw RF bursts using a Vision Transformer (ViT) combined with momentum contrast. The model introduces pseudo-labeling and phase-preserving augmentations to enrich the training signal without requiring ground truth labels. Unlike prior domain adaptation techniques that rely on alignment strategies, strong domain-specific assumptions, or adversarial training, this approach offers a simpler and more robust solution. The authors define positive pairs as RF signals from the same transmission and negative pairs as those from different transmissions, guiding the model to learn domain-invariant representations. Experiments on a 200-device RF testbed demonstrated an accuracy of 92.3%, significantly outperforming supervised CNNs in low-label regimes. This work represents the first application of contrastive learning for domain adaptation in RF device fingerprinting, highlighting its potential for scalable and generalizable wireless device classification.

In contrast to prior domain adaptation methods, which often depend on alignment techniques, domain-specific assumptions, or adversarial training, their method introduces a simpler and more stable alternative based on contrastive learning. This self-supervised framework uses

a pretext task to bring signals from the same transmission (positive pairs) closer and push apart signals from different transmissions (negative pairs) in the learned embedding space. This design enables the model to focus on discriminative, domain-invariant features, effectively mitigating the impact of domain-specific variations. Evaluations on both wireless and wired RF datasets collected over several days showed consistent and substantial accuracy improvements (10.8% to 27.8%) over baseline models. To the best of their knowledge, this work represents the first application of contrastive learning to domain adaptation in RF device fingerprinting.

Autoencoders are also widely used in wireless signal processing for tasks such as denoising, clustering, anomaly detection, and representation learning.

Bai et al. [50] proposed an unsupervised autoencoder framework that uses Random Fourier Feature embeddings for clustering modulation signals. Combined with a novel separable loss function, their model achieved 15–20% higher clustering accuracy on the RadioML2016 dataset and remained robust across SNR conditions from 0–20dB.

Chen et al. [51] explored self-supervised learning for RF-based human activity recognition using CSI perturbations. Their method extracts spatio-temporal features from raw RF data without requiring sensors or labels, achieving strong performance in gesture recognition and occupancy detection.

Faysal et al. [23] introduced DenoMAE, a multimodal framework that integrates a ViT branch for time-frequency input and a convolutional AE branch for raw waveforms. Using masked self-supervised learning, it achieved 15–20% higher classification accuracy and 3–5dB PSNR noise reduction under low-SNR conditions.
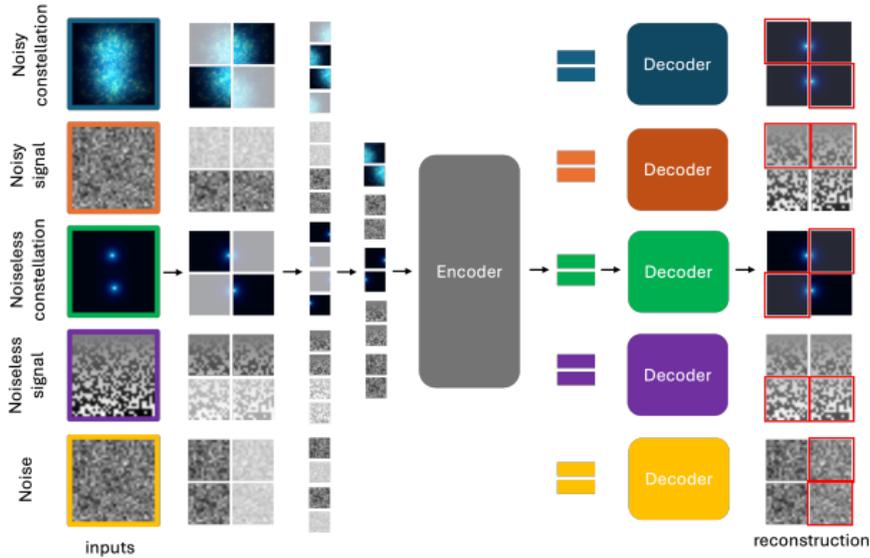
Figure 8: DenoMAE Pretraining Strategy: A random 75% masking is applied across all input modalities (not to scale). The remaining 25% of visible patches are passed through a shared encoder. Each modality then uses its own decoder to reconstruct the masked portions. Only the encoder is reused during downstream fine-tuning [23].

Shi et al. [52] proposed GAF-MAE, which transforms time-series RF signals into Gramian Angular Fields (GAF) and trains a ViT-based masked autoencoder for robust reconstruction. The model improved AMC accuracy by 12–18% over CNNs in 5–15dB SNR conditions.

Gupta et al. [53] presented SpectraViT, a hybrid ViT–ConvAE framework for unsupervised spectrum anomaly detection. Operating across 1 MHz to 6 GHz, it achieved 94% AUC in detecting signal interference in real-time.

Table 3: Wireless Signal Processing Papers using ViT and Autoencoders

Dataset	Method	Task	Ref.
RML2016.10a	GAF-MAE (ViT+AE)	AMC	[52]
RadioML2018	DenoMAE (ViT+AE)	Denoising	[23]
RadioML2016	RffAe-S (AE)	Clustering	[50]
Custom RF	ViT+Contrastive Learning	RF Fingerprinting	[49]
Spectrum Monitoring	SpectraViT (ViT-ConvAE)	Anomaly Detection	[53]
Simulated	ViT+Masked Autoencoding	Modulation Recognition	[54]
Wi-Fi/mmWave CSI	AE + Self-Supervised	Human Activity Recognition	[51]

Lee and Park [54] developed a self-supervised ViT framework using masked autoencoding on spectrogram patches. It outperformed supervised CNNs in modulation recognition, achieving 85% accuracy across 24 modulation types even with limited labeled data.

Collectively, these studies (summarized in Table 3) demonstrate growing interest in leveraging ViTs and AEs for wireless signal analysis. Applications span from AMC and anomaly detection to RF fingerprinting and spectrum monitoring. ViTs offer strong modeling of global features, while AEs provide powerful reconstruction and denoising capabilities. Their integration—often through masked modeling, hybrid architectures, or multimodal fusion—offers a robust and scalable approach to wireless communication systems.

3.4 Radar Signals

Radar signal processing presents unique challenges due to low signal-to-noise ratios, environmental complexity, and limited labeled data. Recent advances in unsupervised learning—especially ViTs, AEs, and their hybrids—have enabled progress in tasks such as clutter suppression, target detection, waveform recognition, and human activity monitoring across diverse radar modalities including SAR, PolSAR, FMCW, IR-UWB, and HRRP.

ViTs leverage global self-attention mechanisms, enabling effective modeling of long-range dependencies and outperforming CNNs in many radar scenarios.

Kim et al. [55] developed a ViT-based approach for waveform recognition of Low Probability of Intercept (LPI) radar signals, achieving a 12.8% accuracy improvement at -10 dB SNR. Kayacan and Erer [56] proposed Declutter ViTs (DC-ViTs) for clutter removal in Ground Penetrating Radar (GPR), improving signal-to-clutter ratio by 20%.

Ghosh and Bovolo [57] introduced a self-supervised ViT with contrastive learning for radargram segmentation, increasing MIoU by 23.47%. Yu et al. [58] presented SLViT, a multimodal slot-based ViT architecture that mitigates speckle noise in SAR classification through modality-specific slots.

Li et al. [59] proposed PolSAR-MPIformer, leveraging mixed patch interactions for fusing dual-frequency PolSAR images. Wang et al. [60] combined Swin Transformers and supervised contrastive learning in SCL-SwinT for robust human activity recognition from IR-UWB radar.

Muzeau et al. [61] introduced SAFE, a masked Siamese ViT framework with SAR-specific augmentations, achieving strong generalization across unseen sensors. Li et al. [62] proposed MTBC, which integrates a masked ViT with Brown distance covariance for few-shot recognition of HRRP radar data.

Feng et al. [63] combined ViTs with Canny edge detection for unsupervised SAR interference pattern analysis. Shi et al. [64] proposed a semi-supervised ViT model for FMCW radar hand gesture recognition using pseudo-label consistency to improve performance with limited labeled data.

Complementing ViTs, AEs are widely used in radar signal processing for their strengths in denoising, anomaly detection, and latent representation learning in cluttered environments.

Table 4: Summary of Vision Transformer, Autoencoder, and Hybrid Methods for Radar Signal Processing

Dataset	Method	Task	Ref.
LPI Radar	ViT	Waveform Recognition	[55]
GPR	DC-ViT	Clutter Removal	[56]
SAR	Multimodal Slot ViT	SAR Image Classification	[58]
PolSAR	PolSAR-MPIformer (ViT)	Adaptive Fusion Classification	[59]
IR-UWB Radar	SCL-SwinT (Swin Transformer)	Human Activity Recognition	[60]
SAR	SAFE (Masked Siamese ViT)	Self-Supervised Feature Extraction	[61]
HRRP	MTBC (Masked ViT + Brown Distance Covariance)	Few-shot Recognition	[62]
SAR Complex Imagery	Self-Supervised ViT	Interference Detection	[63]
FMCW Radar	Semi-Supervised ViT	Hand Gesture Recognition	[64]
Radar-based Heart Rate	MVN (Masked AE + ViT)	Heart Rate Estimation	[65]
Radar Sounder	AE + Random Walks	Semantic Segmentation	[66]
FMCW Radar	AE-based OOD Detector	Out-of-Distribution Detection	[67]
Radar Pulse	AE + LSTM	Pre-sorting	[68]
Radar Clutter	GM-CVAE (CNN + AE)	Target Detection	[69]
SAR	FUS-MAE (Cross-Attention MAE)	Multimodal Fusion	[71]
PolSAR	MAPM (Masked AE)	Image Classification	[72]
SAR	Feature Guided MAE	Self-Supervised Learning	[70]
Radar	Hybrid ViT-CNN	Activity Recognition	[10]
Radargram	URS (Self-Supervised ViT)	Segmentation	[57]

Xiang et al. [65] developed MVN (Masked Phase Autoencoder with ViT) for radar-based heart rate estimation, enhancing accuracy while reducing observation time. Dal Corso and Bruzzone [66] used autoencoders trained with random walks for unsupervised segmentation of radar sounder data, achieving efficient label propagation.

Kahya et al. [67] proposed a lightweight AE model for out-of-distribution detection in short-range FMCW radar, achieving an AUROC of 90.72%. Jiang et al. [68] combined AEs with LSTMs for radar pulse pre-sorting, improving classification of low-frequency pulses.

Liang et al. [69] introduced GM-CVAE, a Gaussian Mixture Convolutional VAE for radar target detection in clutter, outperforming parametric methods in complex scenes.

Hybrid ViT–AE models integrate the global modeling of transformers with AE reconstruction capabilities, offering significant gains in tasks involving fusion, few-shot learning, and semantic understanding.

Wang et al. [70] proposed FG-MAEs (Feature-Guided Masked Autoencoders) for semantic SAR classification, improving accuracy by 5% over standard MAEs. Chan-To-Hing and Veeravalli [71] introduced FUS-MAE, a cross-attention MAE model for SAR-optical fusion, achieving superior cross-modal performance.

Wang et al. [72] designed MAPM, integrating positional prediction and memory tokens into MAEs for PolSAR classification in low-label settings. Huan et al. [10] built a lightweight ViT–CNN hybrid for radar-based activity recognition, balancing spatial-temporal modeling with efficiency.

These advancements, summarized in Table 4, illustrate the growing impact of unsupervised ViT, AE, and hybrid architectures in enabling scalable, accurate, and real-time radar signal processing.

4   Applications

This section explores the unsupervised and semi-supervised applications of AEs, ViTs, and their hybrid approaches across four critical signal domains: Wireless Signals, Medical Signals, IoT Data, and Radar Signals. While Figure 10 presents global trends, the focus here is specifically on tasks utilizing unsupervised or semi-supervised learning within these domains.

AEs have been extensively applied in unsupervised tasks due to their ability to learn compact and informative latent representations. They are notably effective in anomaly detection (20%), natural language and medical applications (15% each), denoising and restoration (10%), and tasks like image classification, segmentation, and audio processing (each around 10%). Their adaptability across structured and unstructured data underlines their foundational role in representation learning without supervision.

In contrast, ViTs are predominantly used in computer vision tasks, with significant application in image classification (40%) and object detection (40%). Though their use in denoising (7%) and anomaly detection (3%) is relatively limited, ViTs are increasingly favored in tasks that benefit from global spatial attention and long-range sequence modeling, expanding their relevance in complex signal interpretation.

Figure 9 illustrates the distribution of Autoencoders and Vision Transformers across supervised, unsupervised, and semi-supervised paradigms. Autoencoders dominate unsupervised use cases, accounting for about 82.5% of their applications due to their strength in learning from unlabeled data. Vision Transformers, in contrast, are mainly utilized in supervised settings (90%), reflecting their success in annotated visual benchmarks. AEs see moderate adoption in semi-supervised contexts, whereas ViTs are only recently gaining traction in such scenarios. This contrast highlights their complementary strengths: AEs in label-scarce environments, and ViTs in data-rich ones.
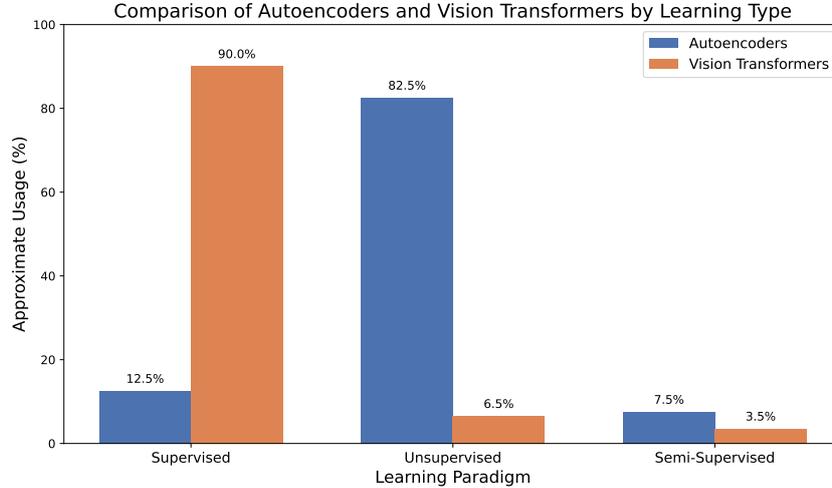
Figure 9: Comparison of Autoencoders and Vision Transformers across learning paradigms.

Recent surveys support these trends, offering broader perspectives on AEs and ViTs. AE-focused reviews outline structural variants, evolution, and domain-spanning applications in vision, NLP, recommender systems, and anomaly detection, while identifying challenges and future directions [1, 2]. ViT surveys highlight their expanding role in digital health, visual benchmarks, and real-time applications, with focus areas including segmentation, classification, multiscale vision, and video analysis [3, 4].
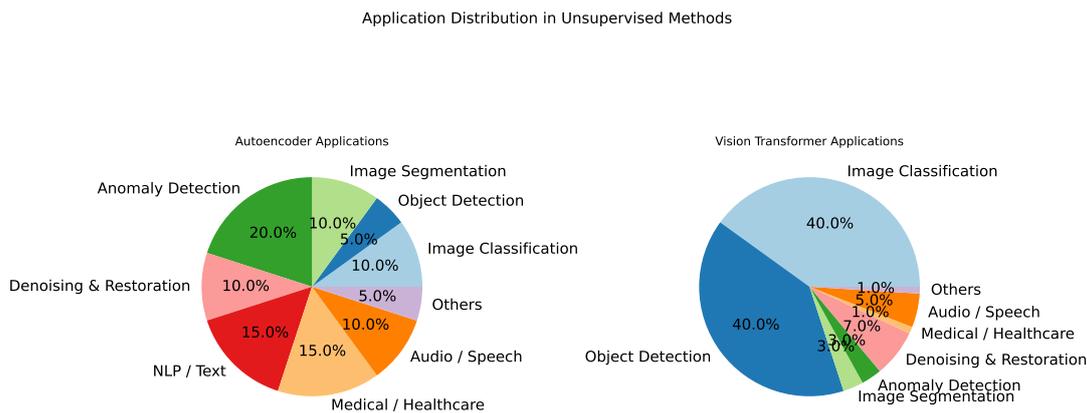
Figure 10: Applications of Vision Transformers and Autoencoders across different domains such as Wireless signals, IoT systems, Radar data, and Medical signals.

Wireless Signal Processing

Unsupervised and self-supervised learning using AEs and ViTs have advanced wireless signal processing tasks, including modulation recognition, denoising, anomaly detection, and channel estimation. Ali and Fan [73] leveraged deep AEs for automatic modulation classification (AMC) using layer-wise pretraining across various SNR conditions.

Hybrid architectures, such as those from Shevitski et al. [74], combined convolutional AEs with transformers for efficient RF signal classification and demodulation. Faysal et al. [23] introduced DenoMAE, a masked autoencoder reconstructing clean waveforms and constellation diagrams in noisy environments.

Transformer–CNN hybrids, proposed by Qu et al. [75] and Zhang et al. [24], merged temporal self-attention with LSTM or CNNs to enhance AMC performance. Transformer-masked autoencoders (TMAEs) [76] also addressed source coding and channel estimation. Additionally, Li et al. [77] proposed a noise-adaptive ViT for robust AMC under adversarial attacks, and ViT-MAE models [78] outperformed CNNs in constellation-based tasks.

Medical Signal Processing

Medical signal analysis has seen widespread adoption of unsupervised deep learning, particularly for EEG and ECG tasks. Banville et al. [29] utilized contrastive and predictive self-supervised tasks on EEG, extracting latent features without labels.

Zhou et al. [5] introduced MTECG, a masked transformer model for ECG classification. Similarly, Sawano et al. [37] used ViTs with high masking ratios for ECG representation learning. Xia et al. [79] developed a hybrid model integrating Transformers, CNNs, and denoising AEs for arrhythmia classification, and Huan et al. [10] adapted ViTs for radar-based clinical activity monitoring.

IoT Signal Processing

In IoT contexts, where data is typically unlabeled, AEs and ViTs support anomaly detection, dimensionality reduction, and traffic classification. AE and Transformer models identify abnormal patterns without ground truth [47], while VAEs enhance botnet detection accuracy [48].

Blind denoising and contrastive AEs [43] recover multivariate sensor data robustly in noisy environments. CNN-VAE hybrids [16] detect traffic anomalies by combining supervised classification with unsupervised outlier detection via reconstruction loss.

Radar Signal Processing

Radar processing has benefited from unsupervised techniques for clutter suppression, waveform classification, and activity recognition. Liang et al. [69] used a GM-CVAE to detect targets via reconstruction probabilities. Jiang et al. [68] proposed a CLDE (CNN–LSTM–decoder) model for radar pulse sorting in low-frequency conditions.

Kim et al. [55] applied ViTs to classify LPI radar using time-frequency representations, achieving robustness over CNNs. Huan et al. [10] developed LH-ViT, a lightweight transformer with convolutional pyramids for real-time radar-based activity recognition.

These studies underscore the versatility of AEs and ViTs in enabling adaptive, efficient signal analysis under limited supervision across diverse domains.

5   Comparative Analysis of Unsupervised Learning Approaches Across Signal Domains

This section presents a comparative synthesis of the methodologies reviewed across four signal domains—wireless communications, radar systems, IoT time series, and biomedical signals—where AEs, ViTs, and hybrid AE-ViT architectures have been applied for unsupervised or semi-supervised learning. Rather than relying on a shared benchmark, studies span a wide range of datasets—from standard repositories like RadioML, MIT-BIH, and NSL-KDD to specialized, domain-specific collections—reflecting the diversity of experimental setups and application goals.

From the tabulated summaries in the methodology section, several thematic patterns and technical trends are evident:

Benchmark Preferences: Certain datasets have emerged as reference points within their respective fields. For instance, RadioML is the prevailing choice for wireless modulation classification, while NSL-KDD and N-BaIoT are frequently used in IoT security applications. In medical signal processing, ECG analysis is often evaluated on MIT-BIH and CinC datasets. These standardized datasets facilitate performance benchmarking and have also played a role in shaping model pretraining practices.

Application Breadth: The scope of tasks tackled using AEs and ViTs is extensive. Classical objectives such as modulation classification, anomaly detection, clustering, and signal denoising appear frequently. However, more recent studies have extended into nuanced areas such as RF fingerprinting, emotion recognition, sleep stage classification, radar segmentation, and multimodal sensor fusion—demonstrating the adaptability of these models to a variety of structured and unstructured signal formats.

Architectural Diversity: Autoencoder-based approaches span a wide range—from basic feedforward variants to more complex designs including convolutional VAEs, adversarial AEs, and blind denoising frameworks. ViTs, while relatively newer to the signal processing domain, have already evolved to include masked autoencoding variants, hierarchical attention modules, and transformer-GAN hybrids. Notably, hybrid AE-ViT architectures are increasingly adopted to balance reconstructive learning with global feature extraction, offering promising performance in both low-data and noisy settings.

Transformer Momentum: Across domains, the use of ViTs is growing steadily, particularly in tasks that benefit from modeling long-range dependencies. In radar and IoT signal processing, their adoption is accelerating due to improved tokenization methods, domain-adapted augmentations, and task-specific pretraining strategies. While computationally more demanding than AEs, ViTs have demonstrated strong performance gains when signals are converted into image-like representations such as spectrograms, scalograms, or GAFs.

Despite these advances, the lack of unified cross-domain benchmarks limits broader generalization studies. This remains a key area for improvement as the field matures.

Representative Experimental Results:

- **Medical Signals:** HeartBEiT [26] applies ViTs to masked ECG images (16×16 patches), achieving 94.5% arrhythmia detection. ECVT-Net [27] combines CNNs with ViTs for CHF classification (98.88%), with 1×250 patch inputs. Jamil et al. [28] process PCG spectrograms into 32×32 patches, achieving 99.9% VHD detection. Banville et al. [29] use CPC and temporal shuffling for EEG (92.5% sleep staging), while Telangore et al. [30] integrate CNN-ViT-LSTM modules for early SCD prediction (98.81%).

- **IoT Time Series:** ViTST [38] uses 16×16 grid-encoded multivariate plots (91% accuracy). TSVIT [40] tokenizes vibration signals into 1×100 patches (93% accuracy). Sana et al. [11] achieve 100% metrics on NSL-KDD using ViTs with 64×64 tokenization. Tarasiou et al. [41] report a 5% improvement over CNNs using ViTs for satellite time series. Liu et al. [14] summarize 90% accuracy across various AE-based unsupervised tasks.

- **Wireless Signals:** NMformer [6] tokenizes constellation images (16×16) for AMC with 90% accuracy. Li et al. [77] use 32×32 spectrogram patches for ViT-based AMR, reporting 94.17% accuracy under benign conditions and 71.2% under attack. CTGNet [80] fuses graph CNNs with ViTs (91% accuracy). RF-ViT-GAN [81] synthesizes 5G/WiFi signals with 95% fidelity. Chen and Wang [49] achieve 92.3% accuracy for RF fingerprinting using self-supervised ViTs.

- **Radar Signals:** URS [57] processes radargrams via ViT segmentation (23.4% mIoU). Kim et al. [82] achieve 0.96 AUC on synthetic radar data using multi-resolution ViTs. Shi et al. [52] utilize GAF-MAE with ViTs for AMC, showing a 15% accuracy gain over baselines. Kim et al. [55] apply ViTs to spectrograms of LPI radar signals (12.8% improvement over CNNs). MHSA-ViT [83] achieves 90% accuracy, though domain shift remains a challenge.

Key Takeaways:

- ViTs are particularly effective when signals are transformed into structured 2D representations.

- AEs continue to offer computational efficiency, making them ideal for real-time and low-power applications.

- Combining ViT and AE in hybrid designs provides a promising middle ground for performance and resource utilization.

- The absence of unified benchmarks across signal types remains a critical limitation for systematic cross-domain evaluation.

6 Challenges and Future Directions

Despite remarkable progress in unsupervised learning for signal processing—driven by AEs, ViTs, and hybrid architectures—the field continues to face technical, methodological, and domain-specific limitations. This section outlines key challenges and proposes promising research directions, including the development of signal-specific foundation models, lightweight inference strategies, and advances in scalable self-supervised learning.

Unsupervised learning in IoT environments remains constrained by the high dimensionality of time-series data, limited edge computing resources, and scarcity of labeled anomalies [47]. Robust deployment in such settings demands innovations in hierarchical architectures, adaptive

latent space modeling, and noise-invariant learning strategies. Hybrid architectures that balance efficiency with expressiveness are particularly promising for real-time, resource-constrained applications [48].

Autoencoders—especially VAEs—have proven effective in dimensionality reduction, anomaly detection, and clustering, particularly in low-data or high-noise conditions. However, they often impose higher computational costs compared to classical techniques like Principal Component Analysis (PCA). Their performance is also highly sensitive to input representation, such as raw waveforms, engineered features, or time-frequency transformations—posing a barrier to generalization across heterogeneous datasets [84].

While ViTs excel at modeling long-range dependencies, they often require extensive pre-training to generalize effectively. Transfer learning from natural image datasets (e.g., ImageNet) may not yield optimal performance in domains such as ECG or RF signals. Recent studies on domain-specific pretraining—such as Masked Autoencoders tailored to physiological signals—have improved task-specific accuracy [37], but at the expense of high computational overhead and data requirements, limiting accessibility in low-resource environments.

A common workaround involves converting time-series signals into 2D visual formats (e.g., spectrograms, Gramian Angular Fields (GAF), Recurrence Plots (RP)) to enable ViT-based modeling. However, the selection of these transformations is often heuristic, lacking theoretical justification. Performance can vary significantly across transformation types, and many are non-invertible, complicating downstream tasks such as signal reconstruction, denoising, or generation [39].

In addition to these modeling challenges, unsupervised frameworks face instability in training, sensitivity to initialization, and the inherent difficulty of evaluating model quality in the absence of labels. Interpretability of learned representations—especially in critical applications like healthcare, defense, and autonomous systems—remains a pressing concern.

Future Research Directions

Several avenues hold strong potential for addressing the above limitations:

- **Cross-Domain Generalization:** Developing architectures that can transfer knowledge across domains (e.g., synthetic-to-real radar, ECG-to-EEG) is essential for generalizable deployment. Meta-learning and domain adaptation techniques will play a crucial role in enabling robust cross-task performance.

- **Signal-Specific Foundation Models:** Inspired by the success of foundation models in vision and NLP, pretraining large-scale, domain-tailored backbones (e.g., ViT-MAE for biosignals or RF waveforms) offers a pathway toward universal encoders for signal intelligence tasks.

- **Lightweight and Edge-Compatible Models:** The design of compact architectures suitable for deployment on embedded or IoT devices is critical. Approaches such as quantization, pruning, knowledge distillation, and transformer-efficient variants (e.g., Linformer, MobileViT) may help minimize memory and latency overheads.

- **Hybrid and Modular Architectures:** Combining CNNs for local pattern recognition with transformers for global context modeling offers a scalable strategy for diverse signal types. Modular design can support plug-and-play adaptation across tasks without requiring full retraining.

- **Scalable Self-Supervised Learning:** Advancing pretext tasks (e.g., masked token modeling, temporal contrastive learning) and augmentations (e.g., frequency masking, jittering) will improve feature quality in unlabeled settings. Efficient training strategies will be essential to extend these methods to high-volume signal streams.

While unsupervised and self-supervised frameworks have already demonstrated impressive potential across signal domains, overcoming existing challenges will require a synergy of architectural innovation, data-efficient learning, and principled evaluation. The future of signal AI lies in scalable, interpretable, and transferable models that bridge the gap between real-world constraints and high-performance learning paradigms.

7 Conclusion

This review has explored how unsupervised and semi-supervised learning methods—specifically Autoencoders, Vision Transformers, and their hybrid configurations—are being applied to time-series signal classification across four principal domains: wireless communications, radar systems, IoT time series, and biomedical signals. Drawing from a broad set of recent studies, we analyzed the architectural trends, domain-specific adaptations, and evolving research directions underpinning these approaches.

1. Why are Autoencoders and Vision Transformers particularly effective for time-series data? Autoencoders have proven to be strong candidates for capturing intrinsic features in time-series signals, especially when labeled data is scarce. Their ability to compress high-dimensional inputs into informative latent spaces makes them well-suited for anomaly detection, reconstruction, and clustering. Vision Transformers, on the other hand, bring a novel perspective to time-series modeling by enabling the capture of long-range dependencies through self-attention. When applied to time-frequency or image-like signal representations, ViTs offer a powerful means to analyze global patterns that may not be easily accessible through traditional models. The combination of both models allows for enhanced flexibility in representing both local and global signal structures.

2. Which tasks are most frequently addressed using these models? Across the surveyed domains, these models have been widely adopted for tasks such as automatic modulation classification, anomaly detection, signal reconstruction, and noise suppression. In biomedical settings, they are frequently used for ECG and EEG analysis, including arrhythmia detection and emotion recognition. In IoT and network environments, they have found strong utility in intrusion detection and traffic analysis. Meanwhile, radar systems are increasingly leveraging these techniques for tasks such as target recognition, waveform segmentation, and scene understanding.

3. What public datasets are most often used in these domains? A number of benchmark datasets have become central to model evaluation. These include RadioML for modulation classification, MIT-BIH for ECG analysis, NSL-KDD and N-BaIoT for IoT security, and synthetic radar datasets for waveform analysis. While these datasets have helped standardize evaluation, there remains a strong need for more diverse, high-resolution datasets that better reflect real-world variability and conditions.

4. How have researchers adapted ViTs and AEs to better suit signal data? Recent innovations include the use of denoising and variational autoencoders, adversarial training techniques, and CNN-AE hybrids for learning richer temporal features. For ViTs, adaptations have involved the use of patch-based encodings of spectrograms and time-series images, masked pretraining strategies, and temporal or frequency-aware attention modules. Increasingly, these models are also being trained using self-supervised objectives—such as contrastive learning and temporal masking—to enhance their generalization without relying on annotated labels.

Despite these advancements, several limitations remain. Interpretability is still a challenge, especially for models deployed in safety-critical domains like healthcare. Scalability to edge devices is another pressing concern, given the computational demands of ViTs. Furthermore, cross-domain generalization remains difficult, particularly when transitioning between synthetic and real-world datasets or between distinct signal modalities.

Looking ahead, future research will likely focus on developing signal-specific foundation

models, designing more efficient and lightweight architectures, and improving self-supervised techniques that can learn from massive unlabeled datasets. The integration of domain knowledge into model design and pretraining strategies may also play a key role in enhancing model performance and interpretability.

Overall, AEs and ViTs—along with their hybrid forms—are reshaping how time-series signals are processed and understood. Their growing adoption across disciplines reflects not only their versatility but also their promise in addressing the growing demands of data-driven signal intelligence in real-world systems.

References

[1] K. Berahmand, F. Daneshfar, E. S. Salehi, Y. Li, and Y. Xu, "Autoencoders and their applications in machine learning: a survey," *Artificial Intelligence Review*, vol. 57, no. 2, p. 28, 2024.

[2] S. Chen and W. Guo, "Auto-encoders in deep learning—a review with new perspectives," *Mathematics*, vol. 11, no. 8, p. 1777, 2023.

[3] Y. Liu, Y. Zhang, Y. Wang, F. Hou, J. Yuan, J. Tian, Y. Zhang, Z. Shi, J. Fan, and Z. He, "A survey of visual transformers," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

[4] K. Al-Hammuri, F. Gebali, A. Kanan, and I. T. Chelvan, "Vision transformer architecture and applications in digital health: a tutorial and survey," *Visual computing for industry, biomedicine, and art*, vol. 6, no. 1, p. 14, 2023.

[5] Y. Zhou, X. Diao, Y. Huo, Y. Liu, X. Fan, and W. Zhao, "Masked transformer for electrocardiogram classification," *arXiv preprint arXiv:2309.07136*, 2024.

[6] A. Faysal, M. Rostami, R. G. Roshan, H. Wang, and N. Muralidhar, "Nmformer: A transformer for noisy modulation classification in wireless communication," in *2024 33rd Wireless and Optical Communications Conference (WOCC)*. IEEE, 2024, pp. 103–108.

[7] Z. Zeng, R. Kaur, S. Siddaganggappa, and others, "From pixels to predictions: Spectrogram and vision transformer for better time series forecasting," in *Proceedings of the Fourth ACM International Conference on AI in Finance (ICAIF '23)*, 2023.

[8] C. J. Harvey, S. Shomaji, Z. Yao, and A. Noheria, "Comparison of autoencoder encodings for ecg representation in downstream prediction tasks," 2024.

[9] E. Vafaei and M. Hosseini, "Transformers in EEG analysis: A review of architectures and applications in motor imagery, seizure, and emotion classification," *Sensors*, vol. 25, no. 5, p. 1293, 2025.

[10] S. Huan, Z. Wang, X. Wang, L. Wu, X. Yang, H. Huang, and G. E. Dai, "A lightweight hybrid vision transformer network for radar-based human activity recognition," *Scientific Reports*, vol. 13, no. 1, p. 17996, 2023.

[11] L. Sana, M. M. Nazir, J. Yang, L. Hussain, Y.-L. Chen, C. S. Ku, M. Alatiyyah, and L. Y. Por, "Securing the iot cyber environment: Enhancing intrusion anomaly detection with vision transformers," *IEEE Access*, vol. 12, pp. 82 443–82 468, 2024.

[12] J. MacQueen, "Some methods for classification and analysis of multivariate observations," *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pp. 281–297, 1967.

[13] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," *KDD*, pp. 226–231, 1996.

[14] Y. Liu, Y. Zhou, K. Yang, and X. Wang, "Unsupervised deep learning for iot time series," *IEEE Internet of Things Journal*, vol. 10, no. 16, pp. 14 285–14 306, 2023.

[15] L. Shan, Y. Li, H. Jiang, P. Zhou, J. Niu, R. Liu, Y. Wei, J. Peng, H. Yu, X. Sha, and S. Chang, "Abnormal ecg detection based on an adversarial autoencoder," *Frontiers in Physiology*, vol. 13, p. 961724, 2022.

[16] Q. Xin, Z. Xu, L. Guo, F. Zhao, and B. Wu, "Iot traffic classification and anomaly detection method based on deep autoencoders," *Applied and Computational Engineering*, vol. 69, pp. 64–70, 2024.

[17] U. Lomoio, P. Vizza, R. Giancotti, S. Petrolo, S. Flesca, F. Boccuto, P. H. Guzzi, P. Veltri, and G. Tradigo, "A convolutional autoencoder framework for ecg signal analysis," *Heliyon*, vol. 11, p. e41517, 2025.

[18] 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, 2023.

[19] A. Shah, D. Singh, H. G. Mohamed, S. Bharany, A. U. Rehman, and S. Hussen, "Electrocardiogram analysis for cardiac arrhythmia classification and prediction through self attention based auto encoder," *Scientific Reports*, vol. 15, no. 1, p. 9230, 2025.

[20] S. Zhang, Y. Fang, and Y. Ren, "Ecg autoencoder based on low-rank attention," *Scientific Reports*, vol. 14, no. 1, p. 12823, 2024.

[21] D. Bank, N. Koenigstein, and R. Giryes, "Autoencoders," *Machine learning for data science handbook: data mining and knowledge discovery handbook*, pp. 353–374, 2023.

[22] Y. Dong, M. Zhang, L. Qiu, L. Wang, and Y. Yu, "An arrhythmia classification model based on vision transformer with deformable attention," *Micromachines*, vol. 14, no. 6, p. 1155, 2023.

[23] A. Faysal, T. Boushine, M. Rostami, R. G. Roshan, H. Wang, N. Muralidhar, A. Sahoo, and Y.-D. Yao, "Denomae: A multimodal autoencoder for denoising modulation signals," *arXiv preprint arXiv:2501.11538*, 2025.

[24] W. Zhang, K. Xue, A. Yao, and Y. Sun, "CTRNet: An automatic modulation recognition based on transformer-CNN neural network," *Electronics*, vol. 13, no. 17, p. 3408, 2024.

[25] D. P. Kingma, M. Welling *et al.*, "Auto-encoding variational bayes," 2013.

[26] A. Vaid, J. Jiang, A. Sawant, S. Lerakis, E. Argulian, Y. Ahuja, J. Lampert, A. Charney, H. Greenspan, J. Narula, B. Glicksberg, and G. N. Nadkarni, "A foundational vision transformer improves diagnostic performance for electrocardiograms," *npj Digital Medicine*, vol. 6, no. 1, p. 108, 2023.

[27] T. Liu, Y. Si, W. Yang, J. Huang, Y. Yu, G. Zhang, and R. Zhou, "Inter-patient congestive heart failure detection using ecg-convolution-vision transformer network," *Sensors*, vol. 22, no. 9, p. 3283, 2022.

[28] S. Jamil and A. M. Roy, "An efficient and robust phonocardiography (pcg)-based valvular heart diseases (vhd) detection framework using vision transformer (vit)," *Computers in Biology and Medicine*, vol. 158, p. 106734, 2023.

[29] H. Banville, L. Alberga, D. A. Engemann, A. Gramfort, and D. Sabbagh, "Uncovering the structure of eeg data using self-supervised learning," *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 1–11, 2021.

[30] H. Telangore, V. Azad, M. Sharma, and A. Bhurane, "Early prediction of sudden cardiac death using multimodal fusion of ecg features extracted from hilbert–huang and wavelet transforms with explainable vision transformer and cnn models," *Journal of Biomedical Signal Processing and Control*, vol. 78, p. 103554, 2025.

[31] P. Nejedly, V. Křemen, K. Lepkova, F. Mívalt, V. Sladky, T. Přidalová, F. Plesinger, P. Jurak, M. Pail, M. Brazdil, P. Klimeš, and G. Worrell, "Utilization of temporal autoencoder for semi-supervised intracranial eeg clustering and classification," *Scientific Reports*, vol. 13, p. 744, 01 2023.

[32] M. Munir, S. A. Siddiqui, A. Dengel, and S. Ahmed, "Deepant: A deep learning approach for unsupervised anomaly detection in time series," *IEEE Access*, vol. 7, pp. 1991–2005, 2019.

[33] J.-H. Jang, T. Y. Kim, H.-S. Lim, and D. Yoon, "Unsupervised feature learning for electrocardiogram data using the convolutional variational autoencoder," *PLOS ONE*, vol. 16, no. 12, pp. 1–16, 2021.

[34] Y. Zhao, Y. Liu, T. Feng, Q. Li, and X. Yu, "Electrocardiogram quality assessment using unsupervised deep learning," *Biomedical Signal Processing and Control*, vol. 83, p. 104704, 2023.

[35] Z. Lan, O. Sourina, L. Wang, R. Scherer, and G. Müller-Putz, "Unsupervised feature learning for eeg-based emotion recognition," in *2017 International Conference on Cyberworlds (CW)*, 2017, pp. 182–185.

[36] T. Wen and Z. Zhang, "Deep convolution neural network and autoencoders-based unsupervised feature learning of eeg signals," *IEEE Access*, vol. 6, pp. 25 399–25 410, 2018.

[37] S. Sawano, S. Kodera, N. Setoguchi, K. Tanabe, S. Kushida, J. Kanda, M. Saji, M. Nanasato, H. Maki, H. Fujita, N. Kato, H. Watanabe, M. Suzuki, M. Takahashi, N. Sawada, M. Yamasaki, M. Sato, S. Katsushika, H. Shinohara, N. Takeda, K. Fujiu, M. Daimon, H. Akazawa, H. Morita, and I. Komuro, "Applying masked autoencoder-based self-supervised learning for high-capability vision transformers of electrocardiographies," *PLOS ONE*, vol. 19, no. 8, p. e0307978, 2024.

[38] Z. Li, S. Li, and X. Yan, "Time series as images: Vision transformer for irregularly sampled time series," 2023.

[39] J. Ni, Z. Zhao, C. Shen, H. Tong, D. Song, W. Cheng, D. Luo, and H. Chen, "Harnessing vision models for time series analysis: A survey," 2025.

[40] S. Zhang, J. Zhou, X. Ma, S. Pirttikangas, and C. Yang, "TSVIT: A time series vision transformer for fault diagnosis," 2023.

[41] M. Tarasiou, E. Chavez, and S. Zafeiriou, "ViTs for SITS: Vision transformers for satellite image time series," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 10 418–10 428.

[42] D. Kara, T. Kimura, S. Liu, J. Li, D. Liu, T. Wang, R. Wang, Y. Chen, Y. Hu, and T. Abdelzaher, "Freqmae: Frequency-aware masked autoencoder for multi-modal iot sensing," *Association for Computing Machinery*, no. 12, p. 2795–2806, 2024.

[43] S. Langarica and F. Núñez, "Contrastive blind denoising autoencoder for real-time denoising of industrial iot sensor data," *Engineering Applications of Artificial Intelligence*, vol. 120, p. 105838, 2023.

[44] J. Haseeb, M. Mansoori, Y. Hirose, H. Al-Sahaf, and I. Welch, "Autoencoder-based feature construction for iot attacks clustering," *Future Generation Computer Systems*, vol. 127, pp. 487–502, 2022.

[45] Y. Xin, H. Chen, and L. Xie, "Efficient noisy data transmission using denoising autoencoder in internet of things," *2021 IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*, pp. 1–6, 2021.

[46] F. Aloul, I. Zualkernan, N. Abdalgawad, L. Hussain, and D. Sakhnini, "Network intrusion detection on the iot edge using adversarial autoencoders," *2021 International Conference on Information Technology (ICIT)*, pp. 120–125, 2021.

[47] M. Yao, D. Tao, R. Gao, and P. Qi, "Anomaly detection for mec enabled hierarchical industrial iot with transformer enhanced variational auto encoder," *IEEE Transactions on Industrial Informatics*, vol. 21, pp. 40–48, 2025.

[48] H. Wasswa, A. Nanyonga, and T. Lynar, "Impact of latent space dimension on iot botnet detection performance: Vae-encoder versus vit-encoder," *2024 3rd International Conference for Innovation in Technology (INOCON)*, pp. 1–6, 2024.

[49] J. Chen, W.-K. Wong, and B. Hamdaoui, "Unsupervised contrastive learning for robust rf device fingerprinting under time-domain shift," in *Proceedings of the IEEE International Conference on Communications (ICC)*. IEEE, 2024, pp. 3567–3572. [Online]. Available: https://arxiv.org/abs/2403.04036

[50] J. Bai, Y. Wang, Z. Xiao, and M. Alazab, "Rffae-s: Autoencoder based on random fourier feature with separable loss for unsupervised signal modulation clustering," *IEEE Trans. Industrial Informatics*, vol. 18, no. 11, pp. 7590–7599, 2022.

[51] Y. Chen, J. Wang, X. Wang, Y. Cao, and Q. Yang, "Unsupervised learning for human sensing using radio signals," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 12 245–12 254.

[52] Y. Shi, H. Xu, Y. Zhang, Z. Qi, and D. Wang, "GAF-MAE: A self-supervised automatic modulation classification method based on gramian angular field and masked autoencoder," *IEEE Trans. Cogn. Commun. Netw.*, vol. 10, no. 1, pp. 94–106, 2024.

[53] R. Gupta, P. Sharma, and A. Desai, "Spectravit: Unsupervised spectrum anomaly detection using vision transformers," in *Proceedings of the ACM SIGCOMM Workshop on AI for Networks*. ACM, 2023, pp. 45–51.

[54] S. Lee and J. Park, "Self-supervised modulation recognition via time-frequency masked autoencoding," *IEEE Wireless Communications Letters*, vol. 13, no. 5, pp. 789–793, 2024.

[55] J. Kim, S. Cho, S. Hwang, and Y. Choi, "Automatic lpi radar waveform recognition using vision transformer," in *2023 IEEE International Radar Conference (RADAR)*. IEEE, 2023, pp. 1–6.

[56] Y. E. Kayacan and I. Erer, "A Vision-Transformer-Based Approach to Clutter Removal in GPR: DC-ViT," *IEEE Geoscience and Remote Sensing Letters*, vol. 21, p. 3505105, 2024.

[57] R. Ghosh and F. Bovolo, "Urs: An unsupervised radargram segmentation network based on self-supervised vit with contrastive feature learning framework," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.

[58] T. Yu, F. Han, L. Zhang, and B. Zou, "Multimodal Slot Vision Transformer for SAR Image Classification," in *2024 IEEE International Conference on Signal, Information and Data Processing (ICSIDP)*, 2024.

[59] X. Xin, M. Li, Y. Wu, X. Li, P. Zhang, and D. Xu, "PolSAR-MPIformer: A Vision Transformer Based on Mixed Patch Interaction for Dual-Frequency PolSAR Image Adaptive Fusion Classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 8527–8542, 2024.

[60] X. Li, S. Chen, S. Zhang, Y. Zhu, Z. Xiao, and X. Wang, "Advancing IR-UWB Radar Human Activity Recognition With Swin Transformers and Supervised Contrastive Learning," *IEEE Internet of Things Journal*, vol. 11, no. 7, pp. 11 750–11 766, 2024.

[61] M. Muzeau, J. Frontera-Pons, C. Ren, and J.-P. Ovarlez, "SAFE: a SAR Feature Extractor based on self-supervised learning and masked Siamese ViTs," 2024.

[62] S. Li, W. Li, P. Huang, M. Zheng, B. Tian, and S. Xu, "MTBC: Masked Vision Transformer and Brown Distance Covariance Classifier for Cross-domain Few-shot HRRP Recognition," *IEEE Sensors Journal*, 2024, accepted for publication, DOI: 10.1109/JSEN.2025.3550584.

[63] Y. Feng, B. Han, X. Wang, J. Shen, X. Guan, and H. Ding, "Self-Supervised Transformers for Unsupervised SAR Complex Interference Detection Using Canny Edge Detector," *Remote Sensing*, vol. 16, no. 2, p. 306, 2024.

[64] Y. Shi, L. Qiao, Y. Shu, B. Li, B. Xiao, W. Li, and X. Gao, "Semi-Supervised FMCW Radar Hand Gesture Recognition via Pseudo-Label Consistency Learning," *Remote Sensing*, vol. 16, no. 13, p. 2267, 2024.

[65] Y. Xiang, J. Guo, M. Chen, Z. Wang, and C. Han, "MAE-Based Self-Supervised Pretraining Algorithm for Heart Rate Estimation of Radar Signals," *Sensors*, vol. 23, no. 18, p. 7869, 2023.

[66] S. Dal Corso and L. Bruzzone, "An approach to semantic segmentation of radar sounder data based on unsupervised random walks and user-guided label propagation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–18, 2023.

[67] S. M. Kahya, M. S. Yavuz, and E. Steinbach, "Reconstruction-based Out-of-Distribution Detection for Short-Range FMCW Radar," in *2023 31st European Signal Processing Conference (EUSIPCO)*, 2023, pp. 1350–1354.

[68] Y. Jiang, S. Shi, F. Zhang, and W. Huang, "Radar pre-sorting algorithm based on autoencoder and lstm," *AEU-International Journal of Electronics and Communications*, vol. 187, p. 155535, 2024.

[69] X. Liang, B. Chen, W. Chen, P. Wang, and H. Liu, "Unsupervised radar target detection under complex clutter background based on mixture variational autoencoder," *Remote Sensing*, vol. 14, no. 18, p. 4449, 2022.

[70] Y. Wang, H. Hernández Hernández, C. M. Albrecht, and X. X. Zhu, "Feature Guided Masked Autoencoder for Self-Supervised Learning in Remote Sensing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2025.

[71] H. Chan-To-Hing and B. Veeravalli, "FUS-MAE: A Cross-Attention-Based Data Fusion Approach for Masked Autoencoders in Remote Sensing," in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, 2024, pp. 6953–6956.

[72] J. Wang, Y. Li, D. Quan, B. Hou, Z. Wang, H. Sima, and J. Sun, "MAPM: PolSAR Image Classification with Masked Autoencoder Based on Position Prediction and Memory Tokens," *Remote Sensing*, vol. 16, no. 22, p. 4280, 2024.

[73] A. Ali and Y. Fan, "Unsupervised feature learning and automatic modulation classification using deep learning model," *Physical Communication*, vol. 25, pp. 75–84, 2017.

[74] B. Shevitski, Y. Watkins, and N. Man, "Digital signal processing using deep neural networks: Evaluating the effectiveness of hybrid autoencoder/transformer models for rf data," *Tech. Rep. PNNL-34279, Pacific Northwest National Laboratory*, 2023.

[75] Y. Qu, Z. Lu, R. Zeng, J. Wang, and J. Wang, "Enhancing automatic modulation recognition through robust global feature extraction," *IEEE Transactions on Vehicular Technology*, 2024.

[76] A. Zayat, M. A. Hasabelnaby, M. Obeed, and A. Chaaban, "Transformer masked autoencoders for next-generation wireless communications: Architecture and opportunities," *IEEE Communications Magazine*, vol. 62, no. 7, pp. 88–94, 2024.

[77] G. Li, C.-C. Lin, X. Zhang, X. Ma, and L. Guo, "Adversarial robust vit-based automatic modulation recognition in practical deep learning-based wireless systems," in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2024, pp. 30–30.

[78] J. Zhao, Q. Cheng, H. Wang, and Y.-D. Yao, "Vit-mae based foundation model for automatic modulation classification," in *2024 33rd Wireless and Optical Communications Conference (WOCC)*, 2024, pp. 50–54.

[79] Y. Xia, Y. Xiong, and K. Wang, "A transformer model blended with CNN and denoising autoencoder for inter-patient ECG arrhythmia classification," *Biomedical Signal Processing and Control*, p. 105271, 2023.

[80] D. Wang, M. Lin, X. Zhang, Y. Huang, and Y. Zhu, "Automatic modulation classification based on CNN-Transformer graph neural network," *Sensors*, vol. 23, no. 17, p. 7281, 2023.

[81] C. Martinez, L. Nguyen, and A. Patel, "Rf-vit-gan: Unsupervised rf signal synthesis using vision transformers," in *NeurIPS Workshop on Deep Learning for Physical Sciences*, 2023.

[82] J. Kim, H. Lee, and S. Park, "Radar anomaly detection with multi-resolution vision transformers," *IEEE Transactions on Signal Processing*, vol. 73, pp. 1234–1245, 2025.

[83] Z. Huang, S. Denman, A. Pemasiri, C. Fookes, and T. Martin, "Radar signal recognition through self-supervised learning and domain adaptation," *arXiv preprint arXiv:2501.03461*, 2025.

[84] G. S. Martin, E. L. Droguett, V. Meruane, and M. das Chagas Moura, "Deep variational auto-encoders: A promising tool for dimensionality reduction and ball bearing elements fault diagnosis," *Structural Health Monitoring*, vol. 18, no. 4, pp. 1092–1128, 2019.