



IMPROVING AUTONOMOUS VEHICLE SAFETY WITH DF-GRU AND VAE-POWERED ANOMALY DETECTION

Keerthana B ^{*1}, and. Jayanthi B ²

¹Department of Computer Applications, Kongu Arts and Science College (Autonomous), Erode – 638107, Tamil Nadu, India

²Department of Computer Science, Kongu Arts and Science College (Autonomous), Erode – 638107, Tamil Nadu, India

¹<http://orcid.org/0009-0002-0161-9561>, ²<http://orcid.org/0000-0001-6153-5760>

E-mail: *keerthikec@gmail.com, sjaihere@gmail.com

ARTICLE INFO

Article History

Received: February 4, 2026

Reviewed: April 11, 2026

Accepted: May 11, 2026

Published: June 30, 2026

Keywords:

Autonomous Vehicles,
Convolutional Dual Flow Gated
Recurrent Unit,
Variational Autoencoder.
Multivariate Time Series Data.
Anomaly Detection.

ABSTRACT

As autonomous vehicles (AVs) have become more widespread, there is a mightier necessity for upgraded security actions to assure credible and secure operation in challenging environments. Guaranteeing AV safety needs anomaly detection, which helps find irregularities in sensor data, system performance, and vehicle behavior. In an effort to strengthen AV safety, this work introduces an anomaly detection approach based on a Convolutional Dual Flow Gated Recurrent Unit (DF-GRU) and Variational Autoencoder (VAE) model. The first convolutional layers emphasize collecting spatial attributes and highlighting semantic elements, while the DF-GRU is conscientious for extracting the temporal context from multivariate time series data. Additionally, the presented construction integrates a VAE to detect anomalies. The VAE generates a threshold setting mode to develop anomaly recognition performance. Anomalies are discovered by paralleling the log-likelihood score and the reconstruction loss of the output, with a pre-established threshold value. Estimate the proposed method's competence using a range of metrics, such as accuracy, precision, recall, F1 score, and a loss. The experiment results show that the recommended system delivers better than the striving procedures. Ultimately, the model's output proves to have the capacity to successfully spot and avoid cyberattacks on AV.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

Autonomous vehicles have been made possible by recent developments in automation and communication. This system seeks to lessen collisions, cut down on energy use, lessen noise, clear obstructions, and improve accessible mobility. Because autonomous vehicles can perceive their surroundings, they can operate and carry out essential duties without the need for human interaction. An autonomous vehicle (AV) uses a fully automated driving system to manage scenarios that often require human assistance. AVs also provide advantages like better land use, increased mobility, and improved communication. By ensuring effective and immediate data transfer among the vehicles, automated cars are revolutionizing the automobile sector. While there are numerous advantages to automated and connected cars, there may also be new privacy, security, and safety concerns. Furthermore, AVs rely exclusively on sensor data and data from supplementary vehicles. However, errors, malfunctions, and hacks can cause irregularities in the sensor data, which can result in accidents and fatalities.

Therefore, techniques for identifying abnormalities and figuring out their origins need to be developed before AVs are extensively deployed [1]. The detection of anomalies is a well-known yet significant subject, and many deep learning-based algorithms have been created that generally perform better in terms of detection accuracy than conventional techniques. Serious circumstances, odd occurrences, or failures in a variety of fields, like anomalous data or network cyber-attacks, can be indicated by anomalies [2]. Accurate finding of anomalies in multivariate time series data is vital for many applications, which has led to substantial interest in this area. For multivariate time series, an amount of unsupervised outlier detection methods have been created, as a result of the challenge of acquiring well-labeled data. However, developing an efficient system is difficult since it must encode the inter-correlations between various time series pairs and capture the temporal dependencies within each time series [3] Due to high labor costs and the requirement for significant domain expertise, anomaly labels are often absent from the majority of time series.

As a result, unsupervised methods are crucial for practical industrial applications. However, existing reconstruction-based anomaly detection methods struggle to capture intricate temporal dependencies and underlying correlations in time series [4]. Autonomous vehicle navigation, traffic flow inspection, obstacle appreciation, and passenger safety all depend on time series sensor data. Anomalies brought on by sensor failures, environmental factors, or unforeseen traffic accidents may have an impact on the data's dependability. Anomalies brought on by environmental factors, unanticipated traffic accidents, or sensor failures may affect the data's dependability. To guarantee safety, it is therefore essential to detect abnormalities in both spatial-temporal and semantic aspects. The reconstruction-based By using deep neural networks' reconstruction capabilities, anomaly detection finds anomalous patterns. This variation endeavors to detect and focus irregularities by comparing the rebuilt data with the original input after generative deep models have been used to recreate the input data.

Although there are many reconstruction-based networks that are appropriate for the AD task [5], the Variational Autoencoder and its variants are the main focus of this work. Although the aforementioned deep learning techniques have demonstrated efficacy, a notable limitation is their dependence on labeled normal or abnormal data for training, which categorizes them as either supervised or semi-supervised models. The difficulty of identifying irregularities in multivariate time series sensor data gathered by self-driving cars is examined in this work. Convolutional DF-GRU and VAE model is a new low-complexity unsupervised technique. The suggested method combines the learning powers of dual flow GRU networks for temporal dependencies from multivariate time series data and convolutional neural networks (CNN) for spatial and semantic features with the generative capabilities of variational auto-encoders. The suggested combination of deep learning methods aims to address the difficulties that come with AD in pre-processed multivariate time series data, when traditional methods might not be sufficient.

An overview of the main contributions this study made is provided below:

- The suggested technique applies lower and upper bound normalization during preprocessing, which grows the model's capability to improve the quality of input data. This stage guarantees consistent and dependable results.
- Feature extraction procedure utilizing the Convolutional DF-GRU Model. The Dual Flow GRU layer summarizes the pertinent contextual information for time series data by utilizing the geographical, semantic, and temporal links present in the input data. Dual-flow data management is useful for improving the perception of multivariate time series data and for ensuring the safety of autonomous cars.
- Using a VAE model, anomaly detection in latent space is accomplished by calculating the reconstruction loss and log-likelihood score, which is used as the anomaly score.

The project's composition is as follows: Section 2 describes the Literature Review, Provide examples of the suggested approach to AV anomaly detection in Section 3. Section 4 provides an explanation of the findings. Finally, the recommended work's conclusion is given in Section 5.

II. LITERATURE REVIEW

The goal of this research is to present an extensive analysis of the body of current literature, emphasizing important discoveries, research approaches, and knowledge gaps. This survey lays the groundwork for comprehending anomaly detection in autonomous vehicles and identifying possible research directions by combining these efforts. In terms of predictive vehicle maintenance, the work [6] was a major breakthrough. This technique has the feasibility to significantly increase the consistency and proficiency of transportation services by integrating graph outlier identification with the inimitable features of vehicle sensor data. Better results can be obtained by representing the sensor network as a graph and applying self-supervised learning to understand the temporal correlations among the features. This framework isolates faults in the sensor data. In [7], a novel unsupervised framework for highway anomaly detection was hosted: Structural Attention-Based Recurrent VAE (SABerVAE). This technique increases anomaly identification by using the environmental structure. It features a lane-vehicle attention module that gauges the connotation of permitted lanes for trajectories and a vehicle self-attention module that comprehends how cars transmit with one another on the road. The architecture hires a recurrent encoder-decoder strategy and a stochastic Koopman operator-propagated latent space to forecast the vehicles' adjacent states based on the outputs from the attention modules. This model proficiently detects anomalies in both typical and distinct situations and is qualified end-to-end to decrease prediction loss supplementary with typical vehicle actions. A method for detecting anomalies in data having a temporal dimension that is unsupervised is presented [8]. In order to learn the temporal sequence continuation proxy problem, study [9] modified the VAE-GAN architecture. Rather than replicating the input, the variational decoder decodes to a future sequence forecast.

This method incorporated an extra decoder that, once the optimal latent-variable has been identified, outputs the pointwise confidence of the prediction in order to distinguish between stochastic and structural uncertainty (which the model may rebuild by fitting to observed data). Without any instances, this model can be used for zero-shot anomaly detection, which separates anomalies from stochastic variance that cannot be modeled. This study [10] introduced a DSA-CNN, a method based on CNN and the self-attention mechanism for identifying anomalous behavior of CAV velocity and acceleration sensors. The model's sensitivity to minor anomalies is improved by a dual-channel self-attention structure, which makes it possible to incorporate spatiotemporal feature learning into the deep learning training procedure. An inventive incremental training technique is part of a new design for a Variational Autoencoder (VAE) targeted at Multivariate Sequential Modeling that was introduced [11]. A progress-based latent space is a feature of this method that works well for anomaly identification in real time as well as offline. This method uses unlabeled data from a robot performing tasks that include both normal and anomalous behaviors, in contrast to many existing methods that are trained in a semi-supervised fashion and depend on large assemblages of normal observations. Additionally, [12] proposed VSAD, a scheme for discovering abnormalities in multivariate time series (MTS) that incorporates spatial-temporal graph networks with a VAE. This method syndicates spatial-temporal graph networks and an autonomous threshold selection technique to arrest convoluted spatial and temporal relationships in MTS data.

Its VAE framework simplifies unsupervised learning of data distribution patterns. A traffic flow prediction idea called the Dual Spatial Convolution Gated Recurrent Unit (DSC-GRU) was initially announced in [13]. It covers both local and global spatiotemporal communication by merging the DSC unit with the Gated Recurrent Unit (GRU). Traditional forecasting models that only emphasis on nearest-neighbor spatial variables sporadically oversee or just overlay global spatial characteristics when accounting for spatial link. The DSC-GRU method hires a superior reliant graph that abuses the parallel coefficient among nodes to signify global geographical links, whilst the conventional static graphs in the DSC unit internments adjacent spatial dependencies. The DSC unit evaluates the effects of both local and global spatial correlations using a modified gated mechanism. Conversely, [14] offered a VAE prophecy model that leverages LSTM for the encoder and decoder units and Planar Flow. By substituting the inner assembly of the VAE with planar flow, this prototypical increases its capability to learn intricate time series data. Planar flow corrects the VAE to better handle the nonlinearity in time series data, permitting more accurate projections.

In [15], an anomaly detection technique based on the reconstruction probability of a Variational Autoencoder (VAE) was presented. By embracing inconsistency and containing the probabilistic characteristics of VAE, this framework delivers the reconstruction probability a more detached and moral anomaly score than reconstruction errors from autoencoders or PCA-based systems. It deals a more trustworthy evaluation of anomalies using a probabilistic metric. It offers a more truthful valuation of anomalies using a probabilistic metric. Additionally, the generative nature of VAE agrees for data reconstruction, which enables the examination of the essential causes of anomalies. In [16], the unsupervised anomaly detection model called GGU-VAE was established by uniting a GRU-based architecture with a Gaussian Mixture VAE for multidimensional time-series data. This model employs a deep latent embedding based on GRU to seizure temporal connections. Against traditional VAE models, which rely on a single Gaussian distribution as the previous model, GGU-VAE uses a model of Gaussian Mixture, which more precisely captures the space of latent exploiting several Gaussian distributions, to growth its aptitude to detect anomalies.

II.1 RESEARCH GAP AND NOVELTY

Current deep learning frameworks for AV safety mostly concentrate on temporal or geographical data. However, their capacity to fully capture the dynamics of intricate driving scenarios is limited by the absence of integrated spatiotemporal modeling. Conventional anomaly detection techniques in AV systems are either deterministic or depend on preset criteria, making it difficult for them to generalize to failure modes that are uncommon or invisible. Real-time inference on resource-constrained AV systems is difficult because to the computationally demanding nature of numerous top-of-the-line designs (such as Transformers and 3D CNNs). A large percentage of current approaches rely on a single sensor modality, such as cameras or LiDAR, ignoring the benefits of multi-modal data integration for more dependable safety precautions. Models that have been trained on particular datasets frequently don't adjust to changes in driving conditions, such bad weather or moving traffic. In order to overcome the shortcomings of AV safety systems, this study presents a unique architecture that combines the Convolutional DF-GRU and VAE models. The robustness, effectiveness, and safety of autonomous vehicle systems are enhanced by this method's resolution of research gaps, particularly in dynamic and uncertain situations.

III. MATERIALS AND METHODS

This study's methodology is structured around a few key components. The main aim of the model's early training and testing phase is to acquire an important dataset regarding autonomous driving. The conclusions are next normalized applying lower and upper boundaries by scaling the multivariate time series data to a predetermined range. The methodology is depends on the provided convolutional DF-GRU and VAE model, a deep learning framework aimed for anomaly recognition. Applying CNNs' strengths in feature extraction, DF-GRU's temporal modeling competencies, and VAE's generative learning features is the intent of this convolutional DF-GRU and VAE conjunction for multivariate time series data analysis. For detecting unusual patterns, this architecture understands the temporal dynamics and the latent representations from typical data. By striving to prevent flips and errors, this kind of technology eventually seeks to increase the safety of autonomous vehicles. Furthermore, it is pretty active at identifying anomalies, which helps make interacting with autonomous cars on the road safer.




Figure 1: Architecture Diagram of Proposed Model.
Source: Authors, (2026).

Figure 1 exhibits the convolutional DF-GRU's complete architecture in combination with a variational autoencoder. The following outlines its four primary stages.

- Data Collection
- Data Preprocessing with Normalization
- Extracting Features with Convolutional DF-GRU
- Anomaly Detection with VAE

III.1 DATA COLLECTION

To identify anomalous actions that deviate from the norm, data collecting for anomaly detection in autonomous cars entails compiling large and varied datasets. Data is gathered from a variety of sensors that record the vehicle's environment in real time, including obstacles, speeds, and distances. Recording the internal conditions of the car, including steering angle, brake pressure, acceleration, and speed. This aids in comprehending the operational dynamics of the vehicle. logging events like as collisions, abrupt stops, or system failures in order to spot possible irregularities. Six public databases are used to evaluate how well DF-GRU and VAE detect anomalies. The statistical information about these sediments is shown in Table 1. Determine how well the model works on univariate time series while being built for multivariate data. Therefore, NAB and UCR were chosen as the two databases[17], [18]. The time series features selected from this data cover the designated time period and can be used to assess the consistency of the model and inform multivariate statistics.

Table 1: Datasets.

Type	Dataset	Test	Dimension	Anomalies
Univariate	NAB	4022	1	0.93
	UCR	5800	1	1.87
Multivariate	MBA	100000	2	0.13
	MSL	73724	54	10.73
	SMAP	427618	26	13.13
	WADI	172802	124	5.98

Source: Authors, (2026).

III.1.1 NAB (Numenta Anomaly Benchmark)

NAB offers an extensive set of real-world time-series data with classified abnormalities, which enables the development of anomaly detection methods. It is would be suitable for testing anomaly detection algorithms in autonomous vehicles since it contains a several kind of domains, such as server metrics, network traffic, and environmental sensors.

III.1.2 UCR (Hexagonml)

A variety of time-series datasets are accessible for used in anomaly detection, clustering, and classification applications via the UCR Time Series Classification/Clustering Repository. It is supportive for estimating anomaly recognition algorithms in autonomous car systems since it holds datasets with several aspects and levels of difficulty.

III.1.3 MBA (Multi-Source Background Autoencoders)

The MBA dataset was developed exclusively for applications that use data from various sensors or sources for multi-source anomaly detection. Robust anomaly detection algorithms for self-driving automobiles may be settled to the labeled data it offers for training and assessment.

III.1.4 MSL (Mars Science Laboratory)

The Mars Science Laboratory mission's telematics data, which includes assessments from spacecraft and environmental sensors, is included in the MSL dataset. It deals an opportunity to test anomaly detection algorithms in inflexible and insulated environments, simulating the challenges faced by autonomous vehicles in unexplored areas.

III.1.5 SMAP (Soil Moisture Active Passive)

The SMAP dataset is composed of soil moisture readings collected by satellite sensors. It produces valuable environmental data that can be used to identify anomalies harming vehicle performance, such as variations in weather patterns or topographical conditions, even if it is not directly associated with autonomous vehicles.

III.1.6 WADI

The WADI dataset covers authentic sensor data from a water distribution network, such as flow rates, pressure readings, and valve statuses. It provides a distinctive environment for testing anomaly detection algorithms important to autonomous vehicles exploiting infrastructure networks, such as noticing leaks in the water distribution system or other problems.

III.2 DATA PRE-PROCESSING

To strengthen the robustness of the suggested model, it's vital to change the data into time-series windows for training and testing after normalizing it. A data processing procedure labeled lower bound and upper bound scaling, commonly referred to as normalizing, adapts the values in a time series dataset such that they lie inside a particular range, typically between 0 and a higher value in computer vision applications. This strategy certifies that every data point lies inside a certain range, which is essential for this study. By doing this, we can develop a deep learning model that runs smoothly [19].

$$x_n = \frac{x - x_{bl}}{x_{bu} - x_{bl}} \quad (1)$$

Eq (1) Determines, that in the time series x proves the Current value, x_{bl} specifies the lower bound, x_{bu} indicates the upper bound, x_n represents the normalized value. This strategy adjusts every single data point equally accordance to its proximity to the upper and bottom bounds of the time series. Post normalization, a score of 0 will be applied to the lowest value and a score of 1 to the highest.

III.3 EXTRACTING FEATURES WITH CONVOLUTIONAL DF-GRU

Locating the most appropriate features in a dataset in order to boost the model performance and limit overfitting is known as feature selection. By leveraging pre-processed feature vectors in the input layer and a use of deep learning technique for feature extraction and anomaly detection, the Convolutional DF-GRU Model optimizes the safety of autonomous vehicles. These vectors capture essential features that are important for tracing irregularities in multivariate time series sensor information. Based on the DF-GRU structure, the offered model assimilates temporal patterns (time-wise focus) with spatial patterns (sensor-wise focus). This study is capable of excellently survey and interpret the many characteristics of time series thanks to these preloaded vectors that retain information from sensor data. Sequential data points' temporal dependencies are the main focus of time series operations. Without doing any calculations, the input layer serves as the entry point for inputs into the neural network.

$$T_i = (t_1, t_2, \dots, t_{n-1}, t_n) \quad (2)$$

In Eq. (2), the feature vector for the a^{th} instance of time series data is denoted as X_a , while the attribute vector for the a^{th} characteristic in the time series data is symbolized as x_a . L is an estimation of the input data's length or dimensionality. The temporal dependencies within a time sequence are meant to be captured by the Dual Flow Gated Recurrent Unit layer. It may gradually understand the context by analyzing data in both forward and backward flows. The output of this layer is an accumulation of feature vectors that depict the temporal characteristics of the incoming data. As explained in Figure 2, the input data for this model consists of characteristics that depict time series scenarios. After passing through Dual Flow GRU layers, the input x exits the convolutional layer. Each of the two parts that make up the DF-GRU layer simultaneously evaluates the features in both forward and backward flows. The output generated by the two GRU layers is denoted by Y .




Figure 2: Structure of DF-GRU Model

Source: Authors, (2026).

These sequential features are processed using the GRU algorithm [8] to produce an output vector with specified dimensions. There are four primary components involved in its operation. The initial element includes the reset gate. The GRU computes the output based on the results of the operations that were previously explained. Contextual information about each piece in the incoming data is efficiently captured by the DF-GRU layer. It is standard procedure in Convolutional Neural Networks (CNNs) to apply the Convolutional Layer on the spatial representation of input data. The convolution layer takes in its input and generates intermediate semantic information. In Eq. (3), $CL_{i:j}$ represents the cluster of features related to the i^{th} and j^{th} items in the context of time series data analysis. The computation within the convolution layer is performed as follows:

$$CL_{i:j} = CL_i \oplus CL_{i+1} \oplus \dots \oplus CL_j \quad (3)$$

The model then combines the convolution results and highlights the key features using maximum pooling. These pooled results are used to create the convolution layer's final output. Here, number of convolution results is denoted as ' rn ', while number of convolution kernels denoted as ' k_i '. Following equations represents the computations.

$$k_i = [k_1, k_2, \dots, k_{rn-1}, k_{rn}] \quad (4)$$

$$\hat{k}_i = \max(k_i) \quad (5)$$

$$\hat{k} = [\hat{k}_1, \hat{k}_2, \hat{k}_3] \quad (6)$$

The convolutional layer's output assists as the input for the fully connected layer in this investigation. Then, use the SoftMax tool to optimize each segment's outcomes. The most efficient portion of the input is the one that recognizes spatial-temporal elements. The computation is expressed in the following formula. Where w_{tp} indicates the weighting parameters of the mesh and bi_p denotes the bias term.

$$p_i = w_m \hat{k} + bi_p \quad (7)$$

$$ok_i = \frac{\exp(p_i)}{\sum_i \exp(p_i)} \quad (8)$$

By using weights that the model assigns to key attributes, the convolutional layer finds important features. The SoftMax activation and the fully linked layer it supply for the output then properly classify these traits. By assisting in collision avoidance and strengthening anomaly identification, this approach illustrates the possibility of elevating the safety of autonomous vehicles.

III.4 VAE BASED ANOMALY DETECTION

Variational Autoencoders are a type of generative model. The inputs from the Convolutional DF-GRU model are encoded into a latent space during the VAE's training process. This is accomplished by integrating two primary aspect encoders and decoders. The encoder uses a probabilistic distribution (mean and variance) to project high-dimensional (multivariate time series) input data into a latent space. For every input sample x , the encoder learns parameters $\mu(z)$ (mean) and σ^2 (variance). The decoder recreates the true time series data after decoding these input samples using the sampled latent points. When the VAE model is used to identify anomalies in multivariate time series data, anomalies can be found using the reconstruction loss and the approximate log-likelihood score (via the ELBO). With this configuration, the VAE may produce precise reconstructions and regularize the latent space to more closely resemble the typical data transmission.

Reconstruction should closely resemble the original sequence if the input is standard. Fig. 3 illustrates that the design of a VAE comprises two fundamental concepts: encoder and decoder. The reconstruction will be subpar if it contains abnormalities. The encoder creates the constraints of a probability distribution by transforming the input data into a latent space. The encoder converts a multivariate time series input $x = \{x_s\}_{s=1}^S$ where $x_s \in \mathbb{R}^i$ (with S time steps and i features at each time step) x into a latent variable z , which is defined by a variance vector (σ) and a mean vector (μ). **Decoder** uses samples of latent variables extracted from the encoder's distribution to reconstruct the input data. The decoder utilized the latent representation z to reconstruct every single time step of the sequence for multivariate time series data. The decoder i_θ (with parameters θ) takes a latent variable z and maps it to a reconstruction $\hat{x} = \{\hat{x}_s\}_{s=1}^S$ which is usually produced by feeding z .




Figure 3: VAE Architecture with Anomaly Detection.
Source: Authors, (2026).

III.4.1 Reconstruction Loss

The Mean Squared Error (MSE) between x , the original input, and $\hat{x}$, the reconstructed output is commonly used to compute the reconstruction loss for VAE. The following formula is used to evaluate the reconstruction loss using Mean Squared Error:

$$L_{cr} = \frac{1}{T_{s=1}} \|x_s - \hat{x}_s\|^2 \quad (9)$$

Where:

x_s = Time series data from the original input at each time step s .

$\hat{x}_s$ = Time series data from the original input at each time step s .

S = Number of the sequence's time steps.

E Establish a threshold (t_{cr}) based on the reconstruction loss distribution for normal data. If the reconstruction loss values are higher than this threshold, the sample may be anomalous.

$$\text{Anomaly if: Reconstruction } r\text{Loss} > t_{cr} \quad (10)$$

KL Divergence:

The Kullback-Leibler (KL) divergence term reflects the disparity of the acquired latent distribution from a previous distribution. For every data instance, the KL divergence term L_{LK} between the learned distribution $y(z|x)$ and the prior $p(z)$.

$$L_{LK} = Di_{KL}(q(z|x)||p(z)) = \frac{1}{2} \sum_{j=1}^{di_z} 1 + \log \sigma_j^2 - \mu_j^2 - \sigma_j^2 \quad (11)$$

Where di_z - latent space dimensionality (or number of latent features).

III.4.2 Log-Likelihood Score (ELBO)

As a rough log-likelihood score, Evidence Lower Bound (ELBO) measures how well an input fits the learnt distribution of "normal" data. The ELBO is used to find samples that deviate from this distribution, which may indicate an abnormality, in an anomaly detection setting.

$$\log p(x) \approx ELBO(x) = T_{q(z|x)}[\log p(x|z)] - Di_{KL}(q(z|x)||p(z)) \quad (12)$$

A lower likelihood is indicated by a lower ELBO score, which raises the possibility that the sample is abnormal.

Anomaly Score:

Utilize both the reconstruction loss and ELBO to generate a weighted composite score. A threshold is then established on the anomaly score once it has been calculated using the weighting variables α and β .

$$SA = \alpha \times L_{cr} - \beta \times ELBO \quad (13)$$

Reconstruction loss and ELBO criteria [12] provide more precise segregation between normal and abnormal data, lowering false positives and improving detection accuracy. By rendering the model more adaptable to missing delicate patterns, the use of both measures promotes anomaly detection accuracy in complicated datasets.

IV EXPERIMENTAL RESULTS

This section explores the outcomes of an innovative approach called the Convolutional DF-GRU and VAE Model, which was put forth upon an analysis of various anomaly detection methods for AVs that have been detailed in the literature. Both prevalent and infrequent data were separated in the test dataset. Separate the dataset into three sections—training, validation, and test sets—before you start training the model. Training is done using the validation and training sets to lower the reconstruction loss. Fig. 4 shows the result of the loss functions during the training period. The loss values gradually drop during the training and validation phases over 50 epochs, eventually arriving at a minimum value for both.




Figure 4: Proposed model loss curves.
Source: Authors, (2026).

IV.1 PERFORMANCE METRICS FOR EVALUATION

The measurements are used to assess the suggested plan during the testing stage. In order to assess a test's efficacy, the F1 score (F1)—also known as the true positive rate—is computed by taking the harmonic mean of precision and recall (R). The evaluation metrics are precision (P) and F1. These metrics, which are derived from both identified anomalies and ground-truth labels, assess the model's anomaly detection skills.

IV.2 COMPARISON WITH EXISTING WORK

This section presents the results of the proposed context testing, validation, and training processes and compares them to those of other deep learning models. Four current models and the suggested work are contrasted for this assessment. LSTM-AE [20], CNN-BiLSTM VAE [8], Modified-Convolutional Neural Network (M-CNN) [21], and GRU-AE, an existing models are completely unsupervised anomaly detection method, are examples of deep auto encoder-based models.




Figure 5: Comparison of Proposed Model with existing Models.
Source: Authors, (2026).

Figure 5 displays the outcome of the suggested model along with the evaluation metrics of current deep learning techniques. The values demonstrate how well the recommended work performs in comparison to a number of current models. The comparisons show how much better the suggested effort is in terms of F1 score, memory, accuracy, and precision. In addition to achieving the best accuracy of 97.8%, the Convolutional DF-GRU and VAE model also led in F1 (95.6%), recall (94.7%), and precision (94.8%). The time series anomaly score generated by the suggested design is shown in Fig 6. The red dotted line stood for the ideal threshold. It is clear that while the scores for anomalous points typically above this barrier, the scores for normal points typically fall below it. The anomaly score curves closely resemble the shape of the time series due to the reconstruction variations that result in the anomaly score, especially since the reconstructed segments are relatively smooth [22].




Figure 6: The Anomaly Score for the Time Series.
Source: Authors, (2026).

In order to concurrently manage the spatial, semantic, and temporal connections, a Strong Anomaly Detection approach was enhanced for pre-processed multivariate time series data by combining VAE architecture with CNN and DF-GRU networks. This method improves latent representation learning, which results in precise anomaly identification. By processing data in both forward and backward flows, a DF-GRU can simultaneously comprehend the connections between historical and prospective data points. This is particularly useful for self-driving automobiles since understanding the series of occurrences demands contextual information from the past and future. Because DF-GRUs presents a balance among calculating economy and dependable temporal modeling, they are important element of autonomous vehicle safety system design. To undertake an all-inclusive and intricate approach to anomaly identification, VAEs may employ two techniques: reconstruction loss and log-likelihood score.

Reconstruction loss shines in isolating explicit faults in the input-output mapping, while log-likelihood finds delicate indiscretions that digress from the probabilistic structure of the data. Together, these metrics guarantee a more truthful, flexible, and trustworthy anomaly detection system. The final result of this work has expressively enhanced when compared to the procedures that were previously used. The judgments of the experiment exhibit that the recommended policy is very effective at identifying abnormalities in multivariate time series data and suggestively growths sensitivity under all known conditions. This indicates that the submitted methodology performs superior than existing methods in the literature in identifying anomalies that could be hazardous in systems seen in the real world.

V CONCLUSION

This methodology offers a thorough and productive way to improve self-sustaining vehicle defense. Convolutional DF-GRU and VAE incorporation provides a responsible method for identifying anomalies in autonomous vehicle safety. Convolutional DF-GRU excellently captured the intricate spatiotemporal and semantic links seen in multivariate time series data, certifying a complete understanding of sensor dynamics. Robust and trustworthy anomaly identification is made possible by VAE's probabilistic framework for distinguishing intricate variations and indeterminacies. Together, these improvements enhanced the toughness and compliance of autonomous systems, leading to safer and more reliable work. This unified methodology displays prominent improvements over current DL methodologies. Using 5-fold cross-validation, the extraordinary presentation of the suggested model was considered. The model outperformed previous methods in terms of average F1-Score, sensitivity (recall), detection accuracy, and precision.

These results demonstration how well the proposed method can detect anomalies, enhancing the safety and consistency of autonomous vehicle systems. Discuss how the system can be improved to navigate data from numerous sensor types and rigid functioning scenarios that arise in the real world, such as bad climate or sensor failure, in future studies. The model could be superior with weather-alert sensors that are adapted for definite situations (e.g., radar for rain or infrared sensors for fog). By combining these sensor types and the information they compromise into the anomaly detection algorithm, autonomous cars will be better able to handle a range of environmental situations. The VAE can imitate weather tempted sensor noise and learn to discriminate between actual abnormalities and environmental distortions because it is probabilistic. This endless adaptation will be necessary to certify that autonomous vehicles are not just safe and reliable, but also proficient of directing the dynamic and changeable nature of human driven environments.

VI. AUTHOR'S CONTRIBUTION

Conceptualization: Ms. Keerthana B.

Methodology: Jayanthi B.

Investigation: Ms. Keerthana B. and Jayanthi B.

Formal Analysis: Ms. Keerthana B.

Validation: Ms. Keerthana B.

Writing – Original Draft: Ms. Keerthana B. and Jayanthi B.

Writing – Review and Editing: Ms. Keerthana B.

Resources: Jayanthi B.

Supervision: Jayanthi B.

Project Administration: Jayanthi B.

Approval of the Final Text: Keerthana B, and. Jayanthi B.

VII. ACKNOWLEDGMENTS

The author would like to appreciate the effort of the editors and reviewers.

VIII. REFERENCES

- [1] S. Rajendar and V. K. Kaliappan, "Sensor data-based anomaly detection in autonomous vehicles using modified convolutional neural network," *Intell. Autom. Soft Comput.*, vol. 32, no. 2, pp. 859–875, 2022.
- [2] C. Zhang, S. Li, H. Zhang, and Y. Chen, "VELC: A new variational autoencoder-based model for time series anomaly detection," *arXiv preprint arXiv:1907.01702*, 2019.
- [3] U. Yokkampon, A. Mowshowitz, S. Chumkamon, and E. Hayashi, "Robust unsupervised anomaly detection with variational autoencoder in multivariate time series data," *IEEE Access*, vol. 10, pp. 57835–57849, 2022.
- [4] S. He, M. Du, X. Jiang, W. Zhang, and C. Wang, "VAEAT: Variational autoencoder with adversarial training for multivariate time series anomaly detection," *Inf. Sci.*, vol. 676, Art. no. 120852, 2024.
- [5] H. H. Nguyen, C. N. Nguyen, X. T. Dao, Q. T. Duong, D. P. T. Kim, and M.-T. Pham, "Variational autoencoder for anomaly detection: A comparative study," *arXiv preprint arXiv:2408.13561*, 2024.
- [6] H. Hojjati, M. Sadeghi, and N. Armanfard, "Multivariate time-series anomaly detection with temporal self-supervision and graphs: Application to vehicle failure prediction," in *Lect. Notes Comput. Sci.*, Cham, Switzerland: Springer, 2023, pp. 242–259.
- [7] N. Chakraborty, A. Hasan, S. Liu, T. Ji, W. Liang, D. L. McPherson, *et al.*, "Structural attention-based recurrent variational autoencoder for highway vehicle anomaly detection," *arXiv preprint arXiv:2301.036*, 2023.
- [8] N. Sboui, M. Hadded, H. Ghazzai, M. Elhadef, and G. Setti, "Anomaly detection in autonomous vehicle lidar sensor data using variational autoencoders," *TechRxiv*, 2024.
- [9] Z. Xing, O. Mehmood, and W. A. P. Smith, "Unsupervised anomaly detection with a temporal continuation, confidence-aware VAE-GAN," *SSRN Electron. J.*, 2024.
- [10] Z. Zhang, *et al.*, "Time series anomaly detection in vehicle sensors using self-attention mechanisms," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 11, pp. 15964–15976, 2024.
- [11] D. Azzalini, L. Bonali, and F. Amigoni, "A minimally supervised approach based on variational autoencoders for anomaly detection in autonomous robots," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2985–2992, 2021.

- [12] K. Zhang, Y. Jiang, L. Seversky, C. Xu, D. Liu, and H. Song, "Federated variational learning for anomaly detection in multivariate time series," in *Proc. IEEE Int. Performance, Computing, and Communications Conf. (IPCCC)*, 2021.
- [13] Q. E. Zhang, "Gated recurrent unit embedded with dual spatial convolution for long-term traffic flow prediction," *ISPRS Int. J. Geo-Inf.*, vol. 12, no. 9, 2023.
- [14] X. B. Jin, *et al.*, "PFVAE: A planar flow-based variational auto-encoder prediction model for time series data," *Mathematics*, vol. 10, no. 4, Art. no. 610, 2022.
- [15] S. I. An, "Variational autoencoder-based anomaly detection using reconstruction probability," in *Special Lecture on Industrial Engineering*, 2015.
- [16] Y. Guo, *et al.*, "Multidimensional time series anomaly detection: A GRU-based Gaussian mixture variational autoencoder approach," *Proc. Mach. Learn. Res.*, vol. 95, pp. 97–112, 2001.
- [17] Numenta, "Numenta anomaly benchmark (NAB) dataset," GitHub Repository, 2015.
- [18] H. A. Dau, E. Keogh, K. Kamgar, C.-C. M. Yeh, Y. Zhu, S. Gharghabi, *et al.*, "The UCR time series classification archive," University of California, Riverside, 2019.
- [19] W. A. AlZoubi, G. B. Desale, S. Bakyarani, C. U. K., D. Nimma, K. Swetha, *et al.*, "Attention-based deep learning approach for pedestrian detection in self-driving cars," *Int. J. Adv. Comput. Sci. Appl.*, vol. 15, no. 8, 2024.
- [20] M. Keerthana and B. Jayanthi, "Enhancing autonomous vehicle safety through LSTM-AE based anomaly detection," *Library Progress (International)*, vol. 44, no. 3, 2024.
- [21] A. Garg, W. Zhang, J. Samaran, R. Savitha, and C.-S. Foo, "An evaluation of anomaly detection and diagnosis in multivariate time series," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2508–2517, 2022.
- [22] Z. Niu, K. Yu, and X. Wu, "LSTM-based VAE-GAN for time-series anomaly detection," *Sensors*, vol. 20, no. 13, Art. no. 3738, 2020.