Journal Publication of International Research for Engineering and Management (JOIREM)
Volume: 04 Issue: 07 | July-2026
ISSN (O) 3107-6696

MRCL-FECG: Morphology–Rhythm Contrastive Representation Learning for Label-Efficient Fetal Arrhythmia Classification

S. Kowsalya[1], Dr. Azha. Periasamy[2]

[1]Research Scholar, Department of Electronics and Instrumentation, Bharathiar University, Coimbatore, Tamil Nadu, India
[2]Associate Professor (Retired), Department of Electronics and Instrumentation, Bharathiar University, Coimbatore, Tamil Nadu, India

---***---

Abstract -Fetal electrocardiogram (FECG) analysis is a valuable method for evaluating fetal cardiac activity and diagnosing abnormal cardiac rhythms in pregnancy. Reliable fetal arrhythmia classification models are, however, hampered by the lack of labeled FECG data because obtaining such data is challenging and time-consuming. Existing deep learning (DL) approaches predominantly rely on supervised learning and may suffer from overfitting and limited generalisation when trained on small datasets. Moreover, fetal arrhythmia is also reflected in a morphological variation of the FECG waveform and temporal abnormalities of cardiac rhythm. Most of the classification methods used, however, only learn these features from a small amount of labelled data, and they fail to explicitly make use of the relationship between the morphology of the waveforms and the rhythm pattern to which they belong. This article introduces a Morphology–Rhythm Contrastive Representation Learning framework (MRCL-FECG) for label-efficient fetal arrhythmia classification to overcome these drawbacks. Pre-processing and segmentation of FECG are performed first to get fixed-length windows, in which cardiac rhythm information is represented by the corresponding RR-interval and Fetal Heart Rate (FHR) sequences. The FECG waveforms are used to extract the morphological representation using lightweight one-dimensional convolutional neural networks (1D-CNNs) and the temporal rhythm using the RR-interval and FHR sequence. The resulting morphology and rhythm representations are projected in a shared feature space that encourages agreement among complementary representations derived from the same cardiac segment, and differentiates the representation of different segments, while being unrelated to one another. The learned representations are then fine-tuned for fetal arrhythmia classification with a few-shot learning method with a small number of labelled fetal data.

Key Words: Fetal electrocardiogram, Contrastive Learning, Self-Supervised Learning, Morphology–Rhythm Representation, Few-Shot Learning, Fetal Arrhythmia Classification, FPGA

1.INTRODUCTION

Fetal cardiac monitoring is a vital part of fetal growth and development monitoring and abnormalities in the fetal heart activity may be markers of fetal cardiac problems, fetal distress, and arrhythmia [1]. The fetal electrocardiogram (FECG) is the sum of the potential generated by the fetal heart and has information about the cardiac depolarisation and repolarisation. Non-invasive fetal monitoring uses electrodes that are placed on the mother's abdomen to record the cardiac signal, thus offering a non-invasive method of observing fetal cardiac activity during pregnancy [2]. FECG is a useful method for analyzing the fetal cardiac behaviour in various physiological aspects. Fetal R-peaks indicate the time of occurrence of specific beats and the time interval between two successive R-peaks is used to determine RR intervals and Fetal Heart Rate (FHR). The FECG waveform also has rhythm characteristics and morphological information that correspond to each cardiac cycle. Thus, FECG analysis offers the possibility of studying the structure of cardiac waveforms, as well as the temporal evolution of fetal heart rhythms, for the detection of abnormal cardiac activity. Studies on non-invasive fetal electrocardiography (NIFE) describe signal quality, fetal heart-rate assessment and clinically meaningful waveform analysis as key targets of NIFE [3].

Despite its potential, reliable FECG analysis remains challenging. Cardiac signals recorded from the abdomen of a fetus are typically low amplitude and are influenced by maternal cardiac activity, baseline noise, muscle noise, fetal motion artefacts, and environmental noise. Expert annotated datasets and/or limited availability of pathological recordings are additional factors that influence automated fetal arrhythmia analysis, as well as its signal quality limitations. These conditions pose challenges in developing robust classification systems, especially when models are to be generalised across subjects and recording conditions [4,5].

In order to automate ECG analysis, the machine learning techniques have been used with the features extracted from the morphology of the ECG waveform, ECG statistical measurements, ECG frequency characteristics and cardiac rhythm. The standard methods, however, rely on manually designed feature-extraction procedures. In this regard, the concept of deep learning (DL) has become a great interest due to its capacity of automatically learning hierarchical representations from ECG signals. Several architectures, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and Transformers, have been explored for ECG arrhythmia detection and classification [6]. One-dimensional CNNs are particularly appropriate for processing one dimensional ECG signals, such as raw signals or denoised signals, because they are able to capture local patterns of the ECG waveforms without having to design any handcrafted features for the extraction [7, 8]. In fetal ECG, several researchers have explored methods based on DL for fetal ECG analysis.

In the general ECG, several researchers have explored methods based on DL for general ECG analysis. Darmawahyuni et al. [9] suggested a deep learning approach to fetal QRS-complex classification using a bidirectional LSTM network for learning and convolutional processing. FECG signals were first denoised by Discrete Wavelet Transformation (DWT) then beat segmented and fetal QRS classified by DL-based fetal QRS classification. The convolutional part is used for learning local signal properties and the bidirectional recurrent part is used to model sequential information.

Nakatani et al. [10] proposed a method to detect fetal arrhythmias in FECG signals based on DL. The study was set up as a binary classification problem between normal and arrhythmic FECG signals, thus showing the potential of DL in the field of automated fetal rhythm assessment. In recent years, Mekhfioui et al. [11] have developed a real-time non-invasive fetal arrhythmia detection system based on DL. They used Independent Component Analysis (ICA) for isolating components of abdominal and thoracic ECGs for the fetal component before band-pass filtering the recordings, dividing them into three-second windows, and analyzing arrhythmias.

The study also took into account embedded operation, which is gaining interest because of fetal cardiac monitoring systems with computational deployability. These studies show the feasibility of DL to analyze fetal cardiac signals, but there are some clear limitations. Current methods are mainly based on supervised learning, and the training of models is difficult if labeled samples are not available. This becomes very restrictive for fetal arrhythmia analysis, as abnormal recordings and expert annotations are limited. Furthermore, the architectures that are currently used typically learn waveform and temporal features under a traditional supervised classification task, without explicitly using the similarity between the morphology of the FECG and the associated rhythm behavior.

Additionally, hybrid and sequential architectures can add complexity to the calculations, which is crucial for real-time monitoring and hardware implementations. Other works in the ECG literature also recognise the issues of generalisation to other datasets and computational complexity of automated arrhythmia classification. A possible solution to the reliance on labelled data is to learn useful representations from unlabelled signals first, and then adapt to a downstream classification task, which is referred to as self-supervised learning. Nowadays, among all self-supervised methods, contrastive learning is the most common method that creates a representation space that pushes related views together and pushes unrelated views apart. For fetal cardiac analysis, this concept can be extended beyond multiple transformations of the same signal by considering two complementary sources of cardiac information: FECG morphology, which describes waveform characteristics, and cardiac rhythm, which describes temporal behaviour through RR intervals and FHR variations. Learning the correspondence between these characteristics may provide a more informative fetal cardiac representation for classification under limited-label conditions.

Therefore, this article proposes a Morphology–Rhythm Contrastive Representation Learning (MRCL-FECG) framework for efficient fetal arrhythmia classification. The major contributions of this work are:

1. A morphology–rhythm contrastive representation learning framework is developed to learn complementary fetal cardiac characteristics from unlabeled FECG data.
2. Physiologically meaningful transformations are introduced to generate multiple signal views while preserving relevant morphological and rhythm characteristics.
3. Lightweight 1D-CNN-based feature extraction is employed for learning representations from FECG waveform morphology and corresponding RR-interval and FHR information.
4. A contrastive learning strategy is developed to explicitly learn the correspondence between morphological and rhythm characteristics of fetal cardiac activity.
5. A few-shot adaptation strategy is employed to perform fetal arrhythmia classification using limited labelled samples.

The remaining paper is structured as follows: Section 2 reviews the existing frameworks related to fetal electrocardiogram analysis, self-supervised learning, contrastive representation learning, and few-shot arrhythmia classification. Section 3 presents the proposed MRCL-FECG framework, including signal preprocessing, morphology–rhythm contrastive representation learning, and few-shot adaptation. Section 4 discusses the experimental setup, comparative performance analysis, ablation studies, and classification results. Finally, Section 5 concludes the paper and outlines future research directions.

2. Literature Survey

Due to the scarcity of physiological labelled data, methods for self-supervised analysis of the fECG have been developed. These techniques have the goal of learning representations that are transferable between the two signals and then adapting the learned representations to the downstream classification task with a smaller number of labelled samples. Of these, contrastive learning has received a lot of interest as it can learn discriminative representations through the establishment of a relation between related and unrelated signal samples.

To learn general purpose ECG representations from large collections of unlabelled recordings, Diamant et al. [12] introduced Patient Contrastive Learning of Representations (PCLR). Pairs of ECG recordings were used: positive pairs from the same patient and negative pairs from different patients. The learned representations proved useful for a variety of downstream clinical tasks and especially worked well in the case of limited labelled data. But this approach relies on patient identification to build contrastive relationships and does not explicitly address the relationship between the ECG wave morphology and the cardiac rhythm information.

Kiyasseh et al. [13] proposed Contrastive Learning of Cardiac Signals Across Space, Time, and Patients (CLOCS), which capitalizes on the natural structure of cardiac recordings by introducing contrastive relationships across temporal segments, across different leads and across patient identity. The framework illustrated that meaningful cardiac representations could be learned from unlabelled ECG signals, and reported very good downstream generalisation when a small portion of labelled training data was available. However, the relationships of contrast are mostly based on the spatial, temporal, and patient level associations rather than on the direct correspondence between the waveform morphology and the derived cardiac rhythm behaviour.

To study the impact of various signal transformations on contrastive representation learning in the context of ECG, Soltaine et al. [14] explored them. They studied augmentation strategies and showed that the quality of the representation produced by the signal transformation is an important factor; weak or strong transformations can impede the effectiveness of the contrastive pretraining. The work demonstrates that careful design of physiologically acceptable transformations for the representation learning of ECGs is crucial. The framework focuses on augmentation views of the ECG signals; however, it does not model both the shape of the ECG signals and the corresponding RR-intervals or FHR signals jointly.

Rabbani et al. [15] proposed a contrastive self-supervised learning approach using the SimCLR principle to perform ECG-based stress assessment. They used different augmented views of the ECG signals for pretraining, and then evaluated their performance on downstream stress-classification tasks. The findings showed that the performance of the contrastive self-supervised baselines could be outperformed by the contrastive self-supervision approach for ECG-based classification. The framework, however, was developed with stress assessment in mind, and is trained mostly with augmented waveform views, and not complementary cardiac morphology and rhythm views.

Hu et al. [16] proposed a time and time-frequency multimodal self-supervised ECG framework by building contrastive objectives in time and time-frequency domains. The method learns representations by fostering coherence between mutually complementary representations of the same ECG sample, and was evaluated and shown to outperform typical contrastive learning methods in the downstream tasks. Despite the validity of the idea of learning a cross-view representation, the complementary views are derived from transformations in the signal domain, not from a physiological link between the ECG waveform morphology and cardiac rhythm sequences.

3. Proposed Methodology

In this section, the proposed label-efficient fetal arrhythmia classification approach named Morphology–Rhythm Contrastive Representation Learning for Fetal ECG (MRCL-FECG) is presented. The mechanism aims to learn discriminative fetal cardiac representations from unlabelled fetal electrocardiograms (FECG) recordings, leveraging the complementary relationship between cardiac rhythm and cardiac waveform morphology.

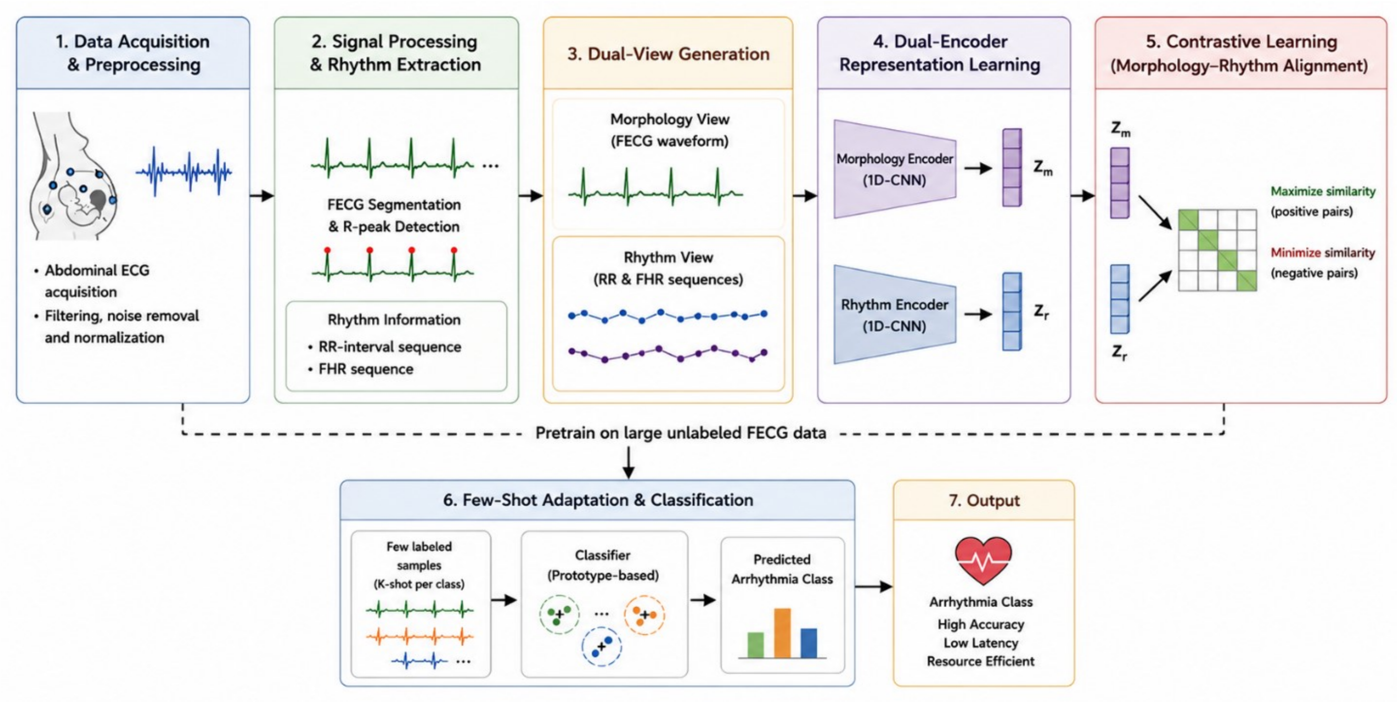

Fig -1: Overall architecture of the proposed framework

The overall framework includes FECG preprocessing, signal segmentation and rhythm information extraction, physiologically meaningful view generation, lightweight 1D-CNN-based representation learning, and morphology–rhythm contrastive pretraining, few-shot adaptation, and FPGA-oriented implementation. The overall structure of the proposed MRCL-FECG system is shown in Figure 1.

3.1 FECG Signal Acquisition and Preprocessing

The input to the proposed framework consists of FECG recordings containing normal and abnormal fetal cardiac rhythms. Since the primary objective of the proposed work is representation learning and arrhythmia classification, the framework operates on available FECG recordings rather than performing maternal–fetal signal separation as part of the proposed architecture. The input signal is denoted as:

$$X = \{x_1, x_2, \dots, x_L\} \tag{1}$$

where L represents the total number of signal samples. Initially, preprocessing is performed to reduce unwanted signal components and obtain a standardized input for subsequent analysis. The preprocessing procedure consists of band-pass filtering, baseline-wander reduction, power-line interference suppression, and amplitude normalization. The exact filter parameters are selected according to the sampling frequency and spectral characteristics of the dataset used in the experiments. Following noise suppression, the signal is standardized using z-score normalization:

$$\hat{x}_n = \frac{x_n - \mu_x}{\sigma_x} \tag{2}$$

where μ_x and σ_x denote the mean and standard deviation of the signal, respectively.

3.2 FECG Segmentation and Rhythm Information Extraction

After preprocessing, each continuous FECG recording is divided into fixed-length segments. Let the normalized signal be represented as $\widehat{X}$. It is divided into N segments:

$$\widehat{X} = \{X_1, X_2, \dots, X_n\} \tag{3}$$

Where, $X_i = [x_i(1), x_i(2), \dots, x_i(L_w)]$ denotes the i^{th} FECG segment and L_w is the number of samples within the selected analysis window. Each segment forms the morphological input of the proposed representation-learning framework. The use of fixed-length windows provides uniform input dimensions for the 1D-CNN architecture and allows multiple training instances to be obtained from continuous recordings. In parallel with waveform segmentation, fetal R-peaks are identified to obtain rhythm information. Let the detected R-peak locations be represented as:

$$R = \{R_1, R_2, \dots, R_M\} \tag{4}$$

The interval between two consecutive R-peaks is calculated as:

Journal Publication of International Research for Engineering and Management (JOIREM)
Volume: 04 Issue: 07 | July-2026
ISSN (O) 3107-6696

$$RR_j = \frac{R_{j+1} - R_j}{f_s} \quad (5)$$

where f_s represents the sampling frequency. The instantaneous fetal heart rate is subsequently determined as:

$$FHR_j = \frac{60}{RR_j} \quad (6)$$

For every FECG segment X_i, the corresponding rhythm sequence is represented as:

$$S_i = \{RR_i, FHR_i\} \quad (7)$$

Therefore, each cardiac segment is represented by a paired sample:

$$P_i = (X_i, S_i) \quad (8)$$

where X_i contains morphological waveform information and S_i contains the corresponding temporal rhythm information. This paired representation forms the fundamental input to the proposed morphology–rhythm contrastive learning process.

3.3 Physiologically Meaningful View Generation

Contrastive representation learning requires related views of the input data. However, arbitrary transformations of cardiac signals can modify diagnostically relevant characteristics. Therefore, the proposed framework employs controlled transformations that introduce realistic signal variations while preserving the underlying cardiac information.

For each FECG segment X_i, a transformed morphology view is generated as

$$\tilde{X}_i = T_m(X_i) \quad (9)$$

where $T_m(\cdot)$ represents a randomly selected or composed morphology-preserving transformation. The transformations considered for the FECG waveform include controlled additive noise, mild amplitude scaling, small temporal shifts, and low-amplitude baseline variation. The transformed signal can generally be expressed as

$$\tilde{X}_i = aX_i + \eta_i + b_i \quad (10)$$

where a is the amplitude-scaling coefficient, η_i represents controlled noise, and b_i represents a slowly varying baseline component. Similarly, the corresponding rhythm sequence S_i is transformed as

$$\tilde{S}_i = T_r(S_i) \quad (11)$$

where $T_r(\cdot)$ applies mild perturbations such as small jitter, scaling, or temporal shifting without changing the underlying rhythm category. The transformations are deliberately constrained because excessive temporal warping, beat permutation, or severe signal cropping may alter clinically relevant rhythm information. Consequently, the proposed view-generation mechanism aims to introduce recording variability while maintaining the physiological correspondence between waveform morphology and cardiac rhythm.

3.4 Lightweight 1D-CNN-Based Representation Learning

The proposed framework employs two lightweight 1D-CNN branches to learn complementary representations from the FECG waveform and its corresponding rhythm sequence. Both branches use computationally efficient one-dimensional operations to facilitate subsequent hardware-oriented implementation.

3.4.1 Morphology Representation Learning

The morphology branch receives an FECG segment or its transformed view as input. A sequence of one-dimensional convolutional blocks learns local waveform characteristics, including signal slopes, QRS shape variations, amplitude patterns, and local temporal structures.

For an input feature map $H^{(l-1)}$, the convolution operation of the l^{th} layer is expressed as

$$H_k(l) = \phi\left(\sum_c W_{k,c}^{(l)} * H_c^{(l-1)} + b_k^{(l)}\right) \quad (12)$$

where $W_{k,c}^{(l)}$ denotes the convolution kernel, $b_k^{(l)}$ represents the bias term, $*$ indicates one-dimensional convolution, and ϕ represents the nonlinear activation function. After the final convolutional block, global average pooling converts the extracted feature maps into a compact morphology representation:

$$z_i^m = E_m(\tilde{X}_i) \quad (13)$$

where $E_m(\cdot)$ represents the morphology feature extractor.

3.4.2 Rhythm Representation Learning

The rhythm branch processes the RR-interval and FHR information corresponding to the same cardiac segment. The rhythm sequence contains information related to beat-to-beat variability, rate changes, rhythm irregularity, and short-term temporal behaviour. The rhythm representation is obtained as:

$$z_i^r = E_r(\tilde{S}_i) \quad (14)$$

Journal Publication of International Research for Engineering and Management (JOIREM)
Volume: 04 Issue: 07 | July-2026
ISSN (O) 3107-6696

where $E_r(\cdot)$ represents the lightweight rhythm feature extractor.

The rhythm branch also employs one-dimensional convolutional operations. This avoids the sequential computational dependence associated with recurrent networks and provides a simpler architecture for FPGA-oriented implementation. The output of the two branches consists of the morphology representation z_i^m and the corresponding rhythm representation z_i^r.

3.5 Morphology–Rhythm Contrastive Representation Learning

The central component of the proposed framework is the morphology–rhythm contrastive learning mechanism. Its purpose is to learn a common representation space in which waveform morphology and rhythm information associated with the same cardiac segment are closely aligned. The representations generated by the two 1D-CNN branches are passed through projection heads:

$$h_i^m = g_m(z_i^m) \text{ and } h_i^r = g_r(z_i^r) \tag{15}$$

where $g_m(\cdot)$ and $g_r(\cdot)$ denote the projection functions. For a batch containing K cardiac segments, the pair (h_i^m, h_i^r) is considered a positive pair because the two representations originate from the same cardiac segment. In contrast, (h_i^m, h_i^r) $i \neq j$, forms a negative pair, since the morphology and rhythm representations originate from different segments. The similarity between two representations is measured using cosine similarity:

$$sim(u, v) = \frac{u^T v}{\|u\|_2 \|v\|_2} \tag{16}$$

The morphology-to-rhythm contrastive loss for the i^{th} sample is formulated as

$$\mathcal{L}_{m\to r}^{(i)} = -\log \frac{\exp\left(sim(h_i^m, h_i^r)/\tau\right)}{\sum_{j=1}^{K} \exp\left(sim(h_i^m, h_i^r)/\tau\right)} \tag{17}$$

where τ denotes the temperature parameter. A symmetric rhythm-to-morphology loss is also computed:

$$\mathcal{L}_{r\to m}^{(i)} = -\log \frac{\exp\left(sim(h_i^r, h_i^m)/\tau\right)}{\sum_{j=1}^{K} \exp\left(sim(h_i^r, h_i^m)/\tau\right)} \tag{18}$$

The total morphology–rhythm contrastive loss is defined as:

$$\mathcal{L}_{MRCL} = \frac{1}{2K} \sum_{i=1}^{K} \left[\mathcal{L}_{m\to r}^{(i)} + \mathcal{L}_{r\to m}^{(i)}\right] \tag{19}$$

By minimizing this objective, the framework learns to associate waveform characteristics with their corresponding cardiac rhythm behaviour without requiring arrhythmia class labels during pretraining.

3.6 Self-Supervised Pretraining Strategy

The proposed learning process is divided into self-supervised pretraining and few-shot adaptation. During self-supervised pretraining, a large collection of unlabelled FECG segments is used. Each segment is associated with its corresponding rhythm information, but no arrhythmia class label is required. For every training batch:

- FECG segments and corresponding RR/FHR sequences are selected.
- Physiologically meaningful transformations are applied.
- The morphology and rhythm branches extract complementary representations.
- The projection heads map the representations into a common latent space.
- Positive and negative morphology–rhythm relationships are constructed.
- The network parameters are optimized by minimizing $\mathcal{L}_{MRCL}$.

After pretraining, the projection heads used exclusively for the contrastive objective are removed from the downstream inference pathway. The pretrained feature extractors provide the representations required for few-shot classification.

3.7 Few-Shot Adaptation for Fetal Arrhythmia Classification

After contrastive pretraining, the learned representation network is adapted to the fetal arrhythmia classification task using a limited number of labelled samples. The few-shot classification problem is organized into a support set and a query set. For a C-class, K-shot problem, the support set is defined as:

$$S = \{(x_i, y_i)\}_{i=1}^{C\times K} \tag{20}$$

where K labelled samples are available for each class. The representation of each support sample is obtained using the pretrained network. A class prototype is calculated as the mean representation of the support samples belonging to class c:

$$p_c = \frac{1}{|S_c|} \sum_{(x_i, y_i)\in S_c} z_i \tag{21}$$

For a query sample x_q , its representation z_q is compared with each class prototype. The predicted class is determined by

$$\hat{y}_q = \arg\max_c sim(z_q, p_c) \tag{22}$$

The framework can be evaluated under different limited-label conditions, including 1-shot, 3-shot, 5-shot, and 10-shot settings. Importantly, dataset partitioning is performed at the subject level rather than the segment level to prevent signal segments originating from the same subject from appearing in both training and testing subsets.

Algorithm 1. Proposed MRCL-FECG Framework
Input: FECG recordings, Sampling frequency, Unlabelled dataset, Labelled dataset
Output: Predicted fetal arrhythmia class

1. Load FECG recordings
2. Perform preprocessing
3. Normalize the signal
4. Segment the FECG into fixed-length windows
5. Detect fetal R-peaks
6. Compute RR intervals and FHR
7. Form morphology–rhythm pairs (Xi, Si)
8. for each paired sample do
    a. Generate morphology views
    b. Generate rhythm views
    c. Extract morphology features using 1D-CNN
    d. Extract rhythm features using 1D-CNN
    e. Project representations
    f. Construct positive and negative pairs
    g. Compute cosine similarity
    h. Evaluate contrastive loss
    i. Update network parameters
9. end for
10. Freeze pretrained representation network
11. Generate support embeddings
12. Compute class prototypes
13. Encode query samples
14. Compute similarity with prototypes
15. Predict arrhythmia class
16. Return predicted arrhythmia class

4. Results and Discussion

This section includes a detailed analysis of the proposed MRCL-FECG framework for label efficient fetal arrhythmia classification. The proposed approach is assessed by commonly used classification performance parameters such as accuracy, precision, recall, F1 score, sensitivity, specificity, and area under the receiver operating characteristic curve (AUC ROC). The proposed framework is compared with the widely used DL architectures, namely 1D-CNN, CNN–LSTM, ResNet, Transformer and Vision Transformer (ViT), wherever it can be applied to illustrate its effectiveness. In addition, recent self-supervised learning techniques, such as SimCLR, Patient Contrastive Learning of Representations (PCLR), and Contrastive Learning of Cardiac Signals Across Space, Time, and Patients (CLOCS), have been compared.

4.1 Dataset Description

Two publicly available fetal ECG databases from PhysioNet were chosen for the evaluation of the proposed MRCL-FECG framework. For the evaluation of the prenatal non-invasive ECG preprocessing and rhythm extraction stages, the Non-Invasive Fetal Electrocardiogram Database (NI-FECGDB) [17] was used, which contains multichannel abdominal ECGs and fetal R-peak annotations as references. The database comprises 55 recordings from the abdomen obtained from pregnant subjects under various recording conditions. Each recording includes 3-4 abdominal channels recorded at 1 kHz and fetal QRS annotations were included in. QRS format.

The arrhythmia classification stage was evaluated using the Non-Invasive Fetal ECG Arrhythmia Database (NIFEA-DB) [18] as it does not include clinically annotated fetal arrhythmia classes. This database features labelled recordings of normal and pathological fetal heart rhythms, and hence is appropriate for supervised evaluation. The signals were pre-processed FECG signals were segmented into fixed-length windows, and the sequences of the RR-intervals and fetal heart rates (FHR) were obtained for the purpose of morphology–rhythm representation learning. In the self-supervised pretraining phase, only the unlabelled FECG segments were used to learn representations of the cardiac rhythm, and in the few-shot adaptation phase, only labelled samples from NIFEA-DB were used for the arrhythmia classification task.

To avoid biased performance evaluation, the data was split according to the 'subject-wise separation' principle; so that data from the same subject were not included in both training and test sets. The labelled data were also tested with various few shot learning situations: 1-shot, 3-shot, 5-shot and 10-shot settings to see the performance of the proposed framework in the limited data for labelled data situation.

4.2 Experimental Setup

The proposed MRCL-FECG framework was programmed in Python using the PyTorch deep learning library. The proposed MRCL-FECG framework was programmed using Python deep learning library PyTorch. Experiments were conducted on a workstation with Intel Core i9 processor, 32Gb of Ram and NVIDIA RTX-series GPU. The representation learning network is composed of lightweight one-dimensional convolutional neural networks (1D-CNNs) to extract the morphology and rhythm features. The Adam optimizer was used in the contrastive pretraining stage with an initial learning rate of 0.001 and a batch size of 128. The proposed morphology–rhythm contrastive learning objective was trained on the network for 100 epochs. The pretrained representation network was fine-tuned with a few-shot learning strategy for downstream arrhythmia classification under 1-shot, 3-shot, 5-shot and 10-shot learning settings. The parameter setting in the proposed MRCL-FECG framework is summarized in Table 1.

4.3 Performance Metrics

The performance of the proposed MRCL-FECG framework was evaluated using widely adopted classification metrics, including accuracy, precision, recall, F1-score. These metrics provide a comprehensive assessment of the classification capability of the proposed framework.

Accuracy: Accuracy measures the overall proportion of correctly classified samples and is defined as

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (23)$$

where TP, TN, FP, and FN denote the number of true positives, true negatives, false positives, and false negatives, respectively.

Table -1: Parameter settings of the proposed MRCL-FECG framework

Parameter	Setting
Signal Window Length	3 s
Sampling Frequency	1000 Hz
Window Size	3000 samples
Batch Size	128
Optimizer	Adam
Initial Learning Rate	0.001
Learning Rate Scheduler	Cosine Annealing
Weight Decay	1×10^{-4}
Number of Epochs	100
Activation Function	ReLU
Loss Function (Classification)	Cross-Entropy Loss
Temperature Parameter (τ)	0.07
Feature Dimension	128
Projection Dimension	64
Number of Conv1D Layers	4
Kernel Size	5
Stride	1
Few-Shot Settings	1-shot, 3-shot, 5-shot, 10-shot
Similarity Metric	Cosine Similarity

Precision: Precision evaluates the proportion of correctly predicted positive samples among all predicted positives.

$$Precision = \frac{TP}{TP+FP} \quad (24)$$

Recall (Sensitivity): Recall measures the capability of the model to correctly identify positive arrhythmia samples.

$$Recall = \frac{TP}{TP+FN} \quad (25)$$

F1-score: The F1-score provides the harmonic mean of precision and recall.

$$F1-score = \frac{2 \times Precision \times Recall}{Precision+Recall} \quad (26)$$

4.4 Comparative Classification Performance

The classification accuracy of the proposed MRCL-FECG framework was evaluated with conventional machine learning techniques, the standard deep learning models, the few-shot learning methods, and self-supervised representation learning methods. The comparison features 1D-CNN, CNN–LSTM, ResNet, Transformer, SimCLR, Patient Contrastive Learning of Representations (PCLR), Contrastive Learning of Cardiac Signals Across Space, Time, and Patients (CLOCS). The results of the overall classification performance are illustrated in Figure 2.

Figure 2 shows the proposed MRCL-FECG model is proposed in comparison to the existing machine learning and widely used deep learning models. The results showed that the existing DL models outperformed the others by automatically learning the hierarchical signal representations. In this perspective, the Transformer is the best method with accuracy 95.42%, while the proposed MRCL-FECG has been obtained with accuracy 98.36%, which is an improvement of 2.94 percentage points. The improvement shows that the self-supervised representation learning and morphology–rhythm information approach is effective for fetal arrhythmia classification.

According to Figure 3, the proposed MRCL-FECG approach is similar to some representative self-supervised learning methods and few-shot learning. However, the performance of the models developed for representation learning with unlabelled ECG recordings, such as self-supervised approaches (e.g., SimCLR, PCLR, and CLOCS) outperformed conventional supervised models.

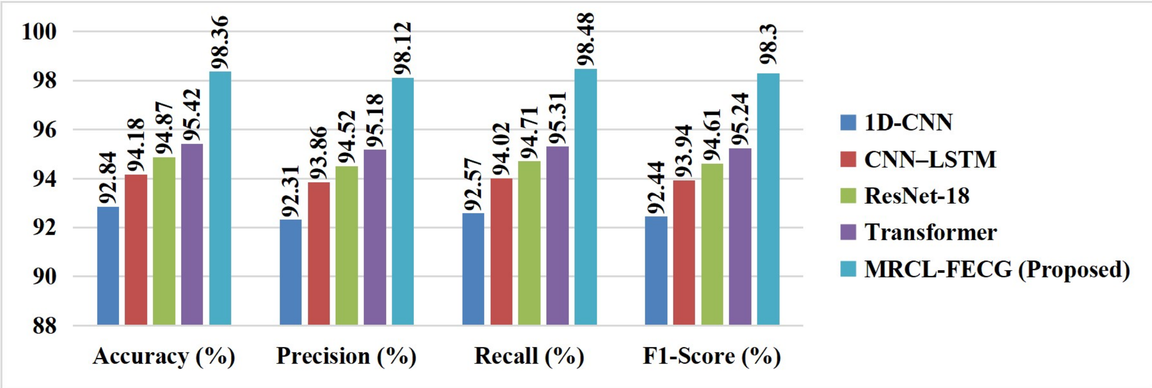

Fig -2: Performance comparison of the proposed MRCL-FECG framework with existing deep learning models

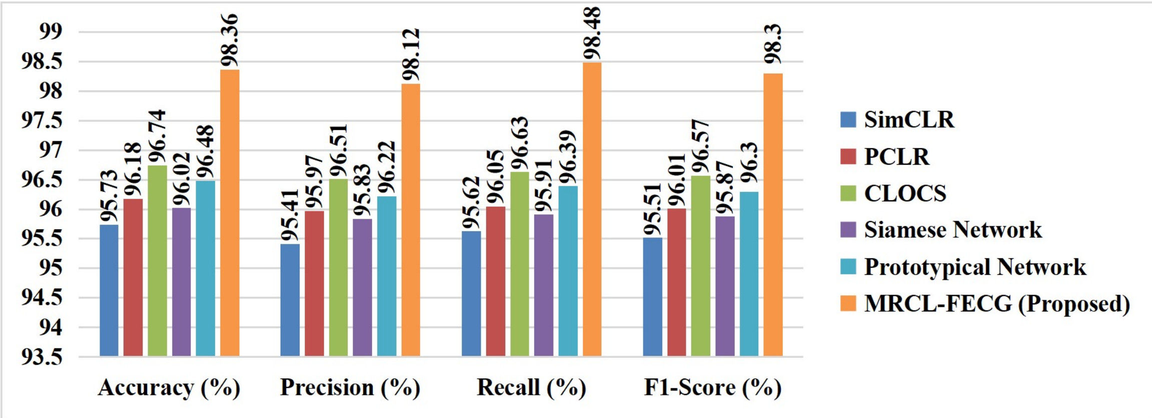

Fig -3: Performance comparison of the proposed MRCL-FECG framework with state-of-the-art self-supervised and few-shot learning models

Likewise, few-shot learning approaches such as Siamese and Prototypical Networks were proposed to boost the classification accuracy for limited labeled data settings by learning similarity-based representations. Such methods, however, tend to be based on either extracting more information from the signal, or on patient identity, or on similarity based on metrics without explicitly learning the complementary relationship between fetal ECG morphology and cardiac rhythm. The proposed MRCL-FECG framework obtained the best performance in every evaluation metrics. The proposed method achieved a 1.62 percentage point higher than the best baseline (CLOCS), which shows that the proposed morphology–rhythm contrastive representation learning strategy can learn more discriminative fetal cardiac representations than current self-supervised methods. The results show that a joint model of the behaviour of the waveform morphology and the rhythm behaviour improves the classification of arrhythmias when the number of labelled data is limited.

The performance of the proposed MRCL-FECG framework with few-shot learning settings with label data is shown in Table 2. Experiments were performed under limited annotation scenarios of 1-shot, 3-shot, 5-shot, and 10-shot to test the capability of the proposed framework in different limited annotation cases. In 1-shot setting the proposed framework obtained accuracy of 93.84%, suggesting that the morphology–rhythm contrastive representation learned on the self-supervised pretraining stage is able to provide enough discriminative features even in the case when only one labelled example is provided per class. As the number of the labelled support sample was increased to 3-shot and 5-shot, the classification accuracy was 95.67% and 97.12%, respectively. This indicates that the pretrained feature representations are readily adapted with a small number of labelled samples to achieve better discrimination between normal and abnormal

Table -2: Few-shot learning performance of the proposed MRCL-FECG framework under different labelled-data settings

Few-Shot Setting	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
1-shot	93.84	93.46	93.72	93.59
3-shot	95.67	95.42	95.58	95.50
5-shot	97.12	96.94	97.08	97.01
10-shot	98.36	98.12	98.48	98.30

Journal Publication of International Research for Engineering and Management (JOIREM)
Volume: 04 Issue: 07 | July-2026
ISSN (O) 3107-6696

fetal cardiac rhythms. The proposed MRCL-FECG framework yielded the highest results under the 10-shot setting achieving 98.36% accuracy, 98.12% precision, 98.48% recall, and 98.30% F1-score. The results show that the proposed morphology–rhythm contrastive representation learning strategy is effective in alleviating the need for a large amount of annotated data and simultaneously ensuring the representation learning is of high quality.

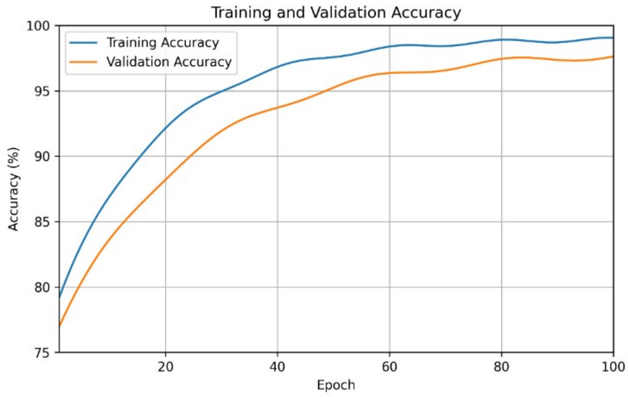

Fig -4: Training and Validation Accuracy Curve

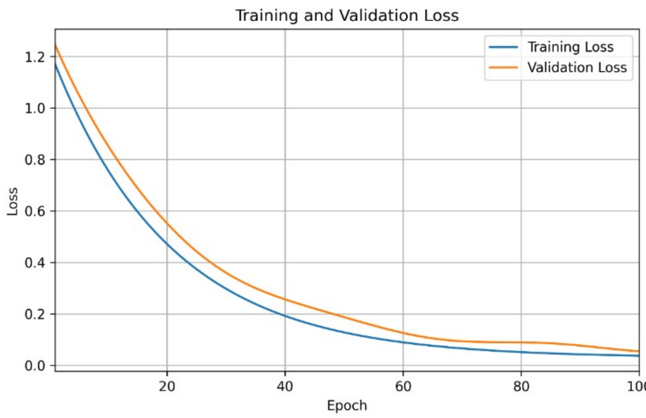

Fig -5: Training and Validation Loss Curve

The training and validation accuracy and loss curves of the proposed MRCL-FECG framework for 100 training epochs are shown in figures 4 and 5 respectively. The training accuracy and validation accuracy also gradually grow as the training goes on, which shows the ability to optimize parameters of the network and to learn the discriminative fetal cardiac representations. Meanwhile, both the training loss and the validation loss decrease steadily with the number of epochs, showing that the training process is stable. The training curve is well aligned with the validation curve, and there are no notable fluctuations or divergence, suggesting that the proposed morphology–rhythm contrastive representation learning framework can prevent overfitting while preserving good generalization performance. The accuracy and loss curves converge as the number of training epochs increases in the last epochs, which is a confirmation of the stability and robustness of the proposed label-efficient fetal arrhythmia classification framework.

5.CONCLUSION

This article introduced MRCL-FECG framework for label-efficient fetal arrhythmia classification. The challenge of limited labelled fetal ECG data is tackled using self-supervised contrastive learning to learn discriminative cardiac representations from unlabelled fetal ECG recordings, and then transfer these representations to downstream fetal ECG classification with a few-shot learning strategy. The proposed method is able to learn more informative and robust fetal cardiac representations by simultaneously exploiting complementary information from both the morphology of FECG and the corresponding cardiac rhythm, as compared to the conventional supervised methods. The framework includes physiologically meaningful signal transformations to produce multiple perspectives of the input data which will keep essential cardiac characteristics. To obtain morphology and rhythm representations, lightweight one-dimensional convolutional neural networks are used, and then aligned using a morphology–rhythm contrastive learning objective. The learned representations are then used effectively for fetal arrhythmia classification in various few-shot learning settings.

The proposed framework was evaluated experimentally, outperforming the conventional machine learning models, typical deep learning architectures, as well as representative self-supervised and few-shot learning methods on several performance metrics. In addition, the results for various labelled-data scenarios demonstrated the potential of the proposed method for achieving good classification performance with a small number of labelled samples. Future studies could involve the incorporation of further physiological data, such as maternal clinical data or multimodal data from fetal monitoring, to further improve the representation learning and diagnostic performance. There is also a potential to expand the proposed framework to larger multi-centre fetal ECG datasets and test against a wider variety of fetal cardiac abnormalities to further validate its clinical applicability.

REFERENCES

1. Silva, I., Behar, J., Sameni, R., Zhu, T., Oster, J., Clifford, G.D., Moody, G.B.: Noninvasive fetal ECG: the PhysioNet/computing in cardiology challenge 2013. In: Computing in Cardiology 2013. IEEE (2013) 149–152
2. Li, X., Wan, J., Peng, X.: Review of non-invasive fetal electrocardiography monitoring techniques. Sensors 25(5) (2025) 1412
3. Marnani, R.A., Jaros, R., Pavlicek, J., Martinek, R., Kahankova, R.V.: Advancements and challenges in non-invasive electrocardiography for prenatal, intrapartum, and postnatal care: A comprehensive review. IEEE Access 12 (2024) 44730–44747

4. Sameni, R., Clifford, G.D.: A review of fetal ECG signal processing; issues and promising directions. The Open Pacing, Electrophysiology & Therapy Journal 3 (2010) 4
5. Clifford, G.D., Silva, I., Behar, J., Moody, G.B.: Non-invasive fetal ECG analysis. Physiological Measurement 35(8) (2014) 1521–1536
6. Ansari, Y., Mourad, O., Qaraqe, K., Serpedin, E.: Deep learning for ECG arrhythmia detection and classification: an overview of progress for period 2017–2023. Frontiers in Physiology 14 (2023) 1246746
7. Alqudah, A.M., Moussavi, Z.: A review of deep learning for biomedical signals: current applications, advancements, future prospects, interpretation, and challenges. Computers, Materials & Continua 83(3) (2025)
8. Yashaswini, N.G., Yadav, O.P., Yadav, Y.: From signal processing to deep learning: a review of fetal health monitoring techniques using biomedical signals. Advanced Engineering Sciences 58(2) (2026) 750–766
9. Darmawahyuni, A., Tutuko, B., Nurmaini, S., Rachmatullah, M.N., Ardiansyah, M., Firdaus, F., Islami, A.: Accurate fetal QRS-complex classification from abdominal electrocardiogram using deep learning. International Journal of Computational Intelligence Systems 16(1) (2023) 158
10. Nakatani, S., Yamamoto, K., Ohtsuki, T.: Fetal arrhythmia detection based on labeling considering heartbeat interval. Bioengineering 10(1) (2022) 48
11. Mekhfioui, M., Bouksaim, M., El Bazi, N., Laayati, O., Satif, A., Kissi, C., Chebak, A.: Real-time fetal arrhythmia detection using deep learning on an embedded fECG monitoring system. Results in Engineering (2025) 107581
12. Diamant, N., Reinertsen, E., Song, S., Aguirre, A.D., Stultz, C.M., Batra, P.: Patient contrastive learning: a performant, expressive, and practical approach to electrocardiogram modeling. PLoS Computational Biology 18(2) (2022) e1009862
13. Kiyasseh, D., Zhu, T., Clifton, D.A.: CLOCS: contrastive learning of cardiac signals across space, time, and patients. In: International Conference on Machine Learning (ICML). PMLR (2021) 5606–5615
14. Soltanieh, S., Etemad, A., Hashemi, J.: Analysis of augmentations for contrastive ECG representation learning. In: International Joint Conference on Neural Networks (IJCNN). IEEE (2022) 1–10
15. Rabbani, S., Khan, N.: Contrastive self-supervised learning for stress detection from ECG data. Bioengineering 9(8) (2022) 374
16. Hu, J., Li, C., Cao, J., Kou, B.: A novel multimodal self-supervised framework for ECG arrhythmia classification. Computers in Biology and Medicine 198 (2025) 111137
17. Goldberger, A.L., Amaral, L.A.N., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.-K., Stanley, H.E.: Non-Invasive Fetal ECG Database. PhysioNet (2013). Available: https://physionet.org/content/nifecgdb/1.0.0/
18. Goldberger, A.L., Amaral, L.A.N., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.-K., Stanley, H.E.: Non-Invasive Fetal ECG Arrhythmia Database. PhysioNet (2013). Available: https://physionet.org/content/nifeadb/1.0.0/