Detection of cardiac arrhythmias using RNA-MLP with PSO

Augusto P. Zatt[1], **Juliano C. Machado**[2], **Alexandre Balbinot**[1]

[1]Department of Electrical Engineering – Federal University of Rio Grande do Sul (UFRGS)
Porto Alegre – RS – Brazil

[2]Federal Institute of Education, Science and Technology of Southern Rio Grande do Sul (IFSUL)
Charqueadas – RS – Brazil

{augusto.zatt,alexandre.balbinot}@ufrgs.br, julianomachado@ifsul.edu.br

Abstract. *This work proposes and evaluates a cardiac arrhythmia classifier based on a Multilayer Perceptron (MLP) Artificial Neural Network (ANN) optimized by the Particle Swarm Optimization (PSO) algorithm. The model was trained and validated using a 10-fold cross-validation scheme, achieving remarkable performance metrics, with an average F1-score of 99.17%, precision of 99.2%, and sensitivity of 99.16%. The analysis of the learning curves, which show the training and validation loss converging in parallel, refutes the overfitting hypothesis. Although PSO is computationally more expensive than traditional gradient-based optimizers, this study demonstrates that this cost is a strategic investment. The meta-heuristic approach promotes robust convergence and the ability to escape local minima, resulting in a superior solution with a high generalization capability. The results confirm the effectiveness and robustness of the MLP-PSO combination for high-performance classification.*

1. Introduction

According to the World Health Organization [WHO 2021], cardiovascular diseases (CVDs) remain the leading cause of global mortality, responsible for approximately 17.9 million deaths in 2019. Cardiac arrhythmias have clinical relevance as they pose risks ranging from benign disturbances to potentially fatal events [HRS 2017]. The electrocardiogram (ECG) is the most widely used diagnostic tool for detecting these anomalies, but its manual interpretation is subject to inter-observer variability [Acharya et al. 2007]. The correct classification of arrhythmias is essential, as each category requires distinct therapeutic approaches [Acharya et al. 2007]. To reduce subjectivity and support specialists, Computer-Aided Diagnosis (CAD) systems have been widely used [Martis et al. 2014]. Traditionally, these systems involve signal pre-processing, manual feature extraction, and the application of classifiers such as Artificial Neural Networks (ANNs), Support Vector Machines, and Decision Trees [Duda et al. 2006]. Although these approaches provide good results, performance can be limited by the effort of manual parameterization and the difficulty in efficiently adjusting hyperparameters, especially in large biomedical datasets. In this scenario, classic architectures like the Multilayer Perceptron (MLP) still represent competitive alternatives, particularly when associated with nature-inspired optimization techniques. Particle Swarm Optimization (PSO), an algorithm based on the collective behavior of swarms, has stood out for its implementation simplicity, low dependence on tuning parameters, and effectiveness in optimizing nonlinear problems [Duda et al. 2006]. Its use in conjunction with MLP can overcome limitations of traditional gradient-based training, promoting more robust convergence and reducing the dependence on manual adjustments.

This work proposes the development of an arrhythmia classifier based on an MLP optimized by PSO, applied to ECG signals. The objective is to explore the balance between performance and computational cost, in order to make its application viable in real-world scenarios that require efficient and fast processing, such as in intelligent healthcare systems and remote monitoring [Mansour et al. 2021].

2. Methodology

To train the model, the MIT-BIH Arrhythmia Database [Moody and Mark 1990] was used, which contains ECG signals from 47 patients, 25 of whom have cardiac arrhythmias and 22 have only normal signals. For this study, the focus was on the MLII (Modified Lead II) derivation, which is commonly used in exams and provides good signal quality. To eliminate noise, a three-stage digital filter in cascade was applied to the signal, consisting of a high-pass filter to attenuate signals with frequencies below 0.5 Hz that can be caused by patient movement, and a notch filter to eliminate 60 Hz frequencies caused by the electrical grid. After this, the signals were segmented into 260-sample windows with the R peak centered and then subjected to Z-Score normalization, in order to standardize the signals in terms of amplitude and displacement, where the mean of each signal is transformed to 0 and the standard deviation to 1. This way, data consistency increases and allows the model to learn features from signals in a uniform way.

As expected from an ECG signal, most peaks belong to a signal without anomalies. Thus, synthetic data generation techniques were used so that the number of signals representing arrhythmias was sufficient to train the model adequately and avoid a class bias. To do this, each original signal was used to generate enough synthetic data so that all signals with anomalies reached the same number of normal signals existing in the database. The original segments were "denormalized"by multiplying by the standard deviation and adding the mean. Then, slight variations of up to 10% in the mean and standard deviation were applied, creating new examples without altering the essential morphology of the signal. The model used is an artificial neural network (ANN) of the multilayer perceptron (MLP) type, and the hyperparameters were optimized by the particle swarm optimization (PSO) algorithm. Among the optimized hyperparameters are the number of neurons in the first hidden layer with a search range from 16 to 128 neurons; number of neurons in the second hidden layer with a search range from 16 to 128 neurons; learning rate with a search range from 0.0001 to 0.01; L2 regularization rate with a search range from 0.00001 to 0.01, which helps to reduce the risk of overfitting by penalizing large weights.

The adopted MLP structure contains an input layer with 260 neurons (corresponding to the number of attributes in the database). Two hidden layers, the first with 108 neurons and the second with 109 neurons (optimized via PSO). In addition, the output layer has 5 neurons with softmax activation. The structure also has ReLU (Rectified Linear Unit) activation functions in the hidden layers, ideal for avoiding the vanishing gradient problem and accelerating convergence. Softmax in the output layer is used for multiclass classification, as it transforms the values into normalized probabilities. The training was supervised, using the categorical crossentropy loss function, suitable for classification with more than two classes. Optimization was also performed with the Adam algorithm, which is a robust technique based on adaptive gradient moments. The training was carried out in 20 epochs using the ten-fold cross-validation method, with 9 parts for training and 1 for validation. For each iteration, the values of sensitivity, precision, and F1-Score were calculated.

3. Results and Discussion

The experiment obtained expressive results for the metrics shown in Table 1. In addition to the average metrics calculated and presented in the table, it was also possible to extract the graph of the loss curves for each fold from the training, as shown in Figure 1.

Tabela 1. Results obtained from the experiment. Source: prepared by the author.

Average F1-Score (%)	Average Sensitivity (%)	Average Precision (%)
99,17	99,16	99,22

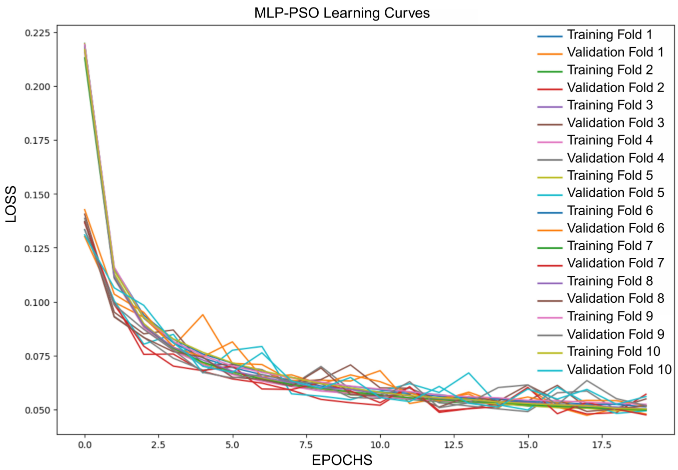

Figura 1. Loss curves for training and validation, plotted as a function of the number of epochs. Source: prepared by the author.

A legitimate concern in research, given such high performance results, is the possibility of overfitting. However, an analysis of the graph shown in Figure 1 provides evidence that contradicts this hypothesis. The most critical observation is the parallel and convergent behavior of the training and validation losses in all folds. Throughout the training, both loss curves decrease sharply in the initial epochs and then stabilize at a plateau with consistently low values (around 0.05 to 0.06). In an overfitting scenario, the training loss would continue to decrease while the validation loss would begin to diverge and increase, indicating that the model was losing its generalization ability. The absence of such divergence in the graph is the main evidence against overfitting. The joint convergence and stable behavior of the loss curves demonstrate that the optimization process was effective in finding a solution that minimizes error on both training and validation data. The stability of the curves in the final epochs suggests that the model found an ideal balance between fitting and generalization, validating the training process and the robustness of the solution found. The use of the PSO algorithm to adjust the weights and biases of the

MLP neural network constitutes a significant methodological choice with implications for computational cost. It is important to recognize that PSO, as a population-based metaheuristic, is inherently more costly than gradient-based optimizers. This cost is a function of three main factors: swarm size, the complexity of the loss function, and the number of iterations required for convergence [Chen et al. 2023]. However, the higher computational cost should not be seen as a disadvantage, but rather as the price of a substantial methodological benefit. While gradient-based optimizers are fast, they are notably susceptible to local minima, where the optimization process can stagnate in a suboptimal solution and fail to reach the true global minimum. In contrast, PSO, inspired by the social behavior of swarms, explores the search space more broadly and globally. PSO's lower sensitivity to its own control parameters, compared to traditional optimizers, also contributes to the robustness of the final solution. In essence, the computational cost of PSO is a strategic investment that is justified by the quality and robustness of the solution obtained.

4. Conclusion

In summary, the discussion of the results and the methodology employed validates the high performance of the RNA-MLP model optimized by PSO. The performance metrics of over 99% are a testament to the model's effectiveness in solving the classification problem. The analysis of the learning curves refutes the overfitting hypothesis, as both training and validation losses converged in parallel and stabilized, indicating that the model learned the essential characteristics of the dataset without memorizing noise. Although the use of PSO entails a higher computational cost compared to traditional optimizers, this investment is fully justified. PSO offers a more robust solution, with a higher probability of finding the global optimum and lower sensitivity to local minima. For future work, it is essential to observe the detriment of using synthetic data and the possibility that they are responsible for the high numbers in the model's performance.

Referências

Acharya, U. R., Krishnan, S. M., Spaan, J. A., and Suri, J. S. (2007). *Advances in cardiac signal processing*. Springer.

Chen, J., Zhang, R., and Liu, Y. (2023). An adam-enhanced particle swarm optimizer for latent factor analysis.

Duda, R. O., Hart, P. E., et al. (2006). *Pattern classification*. John Wiley & Sons.

HRS (2017). Heart diseases and disorders. Technical report, Heart Rhythm Society.

Mansour, R., Amraoui, A., Nouaouri, I., García Díaz, V., Gupta, D., and Kumar, S. (2021). Artificial intelligence and internet of things enabled disease diagnosis model for smart healthcare systems. *IEEE Access*, PP:1–1.

Martis, R. J., Acharya, U. R., and Adeli, H. (2014). Current methods in electrocardiogram characterization. *Computers in Biology and Medicine*, 48:133–149.

Moody, G. and Mark, R. (1990). The mit-bih arrhythmia database on cd-rom and software for use with it. In *[1990] Proceedings Computers in Cardiology*, pages 185–188.

WHO (2021). Cardiovascular diseases (cvds). Technical report, World Health Organizaton.