

Brief Report

Not peer-reviewed version

A Proof-of-Concept Multimodal Framework for Transthyretin Amyloidosis Classification from Clinical Narratives and ECG Images

Sofia Pagano , Cristina Perri , Marika Votta , [Ugo Lomoio](#) , [Pierangelo Veltri](#) , [Pietro Hiram Guzzi](#) *

Posted Date: 25 May 2026

doi: 10.20944/preprints202605.1601.v1

Keywords: transthyretin amyloidosis; rare diseases; multimodal deep learning; ECG; BERT; ResNet50; proof of concept; clinical decision support



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Brief Report

A Proof-of-Concept Multimodal Framework for Transthyretin Amyloidosis Classification from Clinical Narratives and ECG Images

Sofia Pagano¹, Cristina Perri¹, Marika Votta¹, Ugo Lomoio² and Pierangelo Veltri², and Pietro Hiram Guzzi^{1,*}

¹ Department of Surgical and Medical Sciences, Magna Graecia University of Catanzaro, 88100 Catanzaro, Italy

² Department of Computer Science, Modelling and Electronics (DIMES), University of Calabria, 87036 Rende, Italy

* Correspondence: m.milano@unicz.it

Abstract

Transthyretin amyloidosis (ATTR) is a rare, progressive, and frequently under-recognized systemic disease whose early clinical manifestations may overlap with common cardiac and neurological conditions. This work presents a proof-of-concept multimodal deep-learning framework for the binary classification of ATTR and non-ATTR cases by jointly analysing structured clinical narratives and electrocardiogram (ECG) images. A curated dataset of 100 cases was assembled, including 60 literature-derived cases and 40 synthetic cases generated through a multimodal Generative Adversarial Network (GAN) to mitigate data scarcity and class imbalance. The proposed architecture combines a frozen ResNet-50 visual encoder with a frozen Italian BERT textual encoder; the resulting modality-specific embeddings are projected into a shared latent space and fused through a multilayer perceptron classifier. In a held-out test subset, the final multimodal configuration achieved an overall accuracy of 73%, an AUC-ROC of 0.78, and an ATTR recall of 0.89, indicating promising sensitivity for a preliminary screening scenario. These results should be interpreted as evidence of methodological feasibility rather than clinical readiness, given the limited cohort size, the inclusion of synthetic data, and the absence of external multicentre validation. Overall, the study supports the feasibility of multimodal AI for rare-disease triage and provides a foundation for larger, fully validated clinical investigations.

Keywords: transthyretin amyloidosis; rare diseases; multimodal deep learning; ECG; BERT; ResNet50; proof of concept; clinical decision support

1. Introduction

Artificial intelligence (AI) and machine learning (ML) are increasingly being explored as tools to support clinical decision-making, particularly when heterogeneous biomedical information must be interpreted jointly [1–3]. Recent progress in combining large language models and deep visual encoders has enabled multimodal architectures capable of processing textual and imaging data within a single analytical pipeline [4,5]. Such approaches are especially relevant in complex diagnostic settings, where clinical narratives, instrumental findings, and subtle visual patterns may collectively contribute to disease recognition [6,7].

This study focuses on transthyretin amyloidosis (ATTR), a rare systemic disorder characterized by extracellular deposition of amyloid fibrils derived from misfolded transthyretin [8,9]. ATTR may involve the heart, peripheral nervous system, and other organs, and its early manifestations are often nonspecific. As a result, diagnosis is frequently delayed, limiting the timely use of disease-modifying therapies [10,11]. For this reason, computational tools able to support preliminary triage and guide clinical suspicion may be valuable, provided that their performance is rigorously validated before any real-world deployment.

We propose a proof-of-concept multimodal framework for automated ATTR classification based on clinical text and ECG images. The framework uses a hybrid dataset composed of textual clinical descriptions and electrocardiogram (ECG) images, with additional synthetic cases generated through Generative Adversarial Networks (GANs) to address the data scarcity typical of rare diseases [12]. Architecturally, the model combines a CNN-based visual encoder, ResNet50 [13], and a transformer-based textual encoder, Italian BERT [14], through a neural fusion classifier designed to integrate complementary visual and linguistic representations. Figure 1 summarizes the analytical workflow.

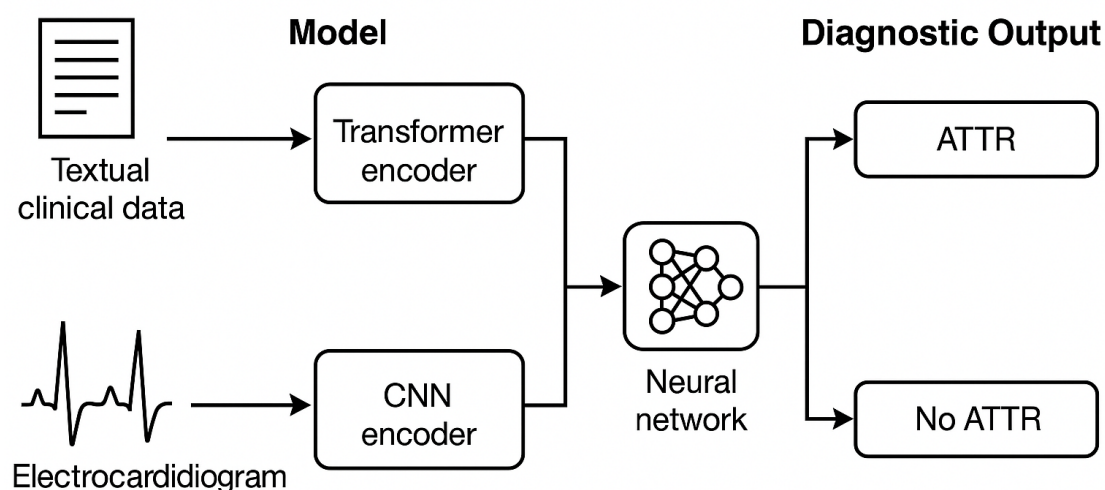


Figure 1. Overview of the proof-of-concept multimodal deep-learning framework for transthyretin amyloidosis (ATTR) classification. Textual clinical descriptions are encoded using a transformer-based language model, while electrocardiogram (ECG) images are processed through a convolutional neural network. The resulting representations are fused by a neural classifier that estimates the probability of ATTR versus NO_ATTR presentation.

The contribution of this work is twofold. Methodologically, it evaluates whether pretrained textual and visual encoders can be combined to generate a unified representation for rare-disease classification. Clinically, it investigates whether such a multimodal strategy can support early suspicion of ATTR in a controlled experimental setting. Importantly, the study is intentionally framed as a proof of concept: the objective is to test feasibility, identify methodological constraints, and define the requirements for subsequent external validation, rather than to propose a clinically deployable diagnostic system.

2. Methods

2.1. Operational Context and Dataset

For model development and preliminary validation, we constructed an original multimodal dataset consisting of 100 clinical cases designed to integrate structured textual descriptions and ECG tracings. The dataset included 60 real cases extracted from specialized scientific literature [15–58] and 40 synthetic cases generated using GAN-based augmentation. The synthetic cases were used to improve class balance and increase the variability available for this proof-of-concept experiment.

ECG images were standardized by converting the tracings into digital images with a resolution of 224×224 pixels, matching the input requirements of the visual encoder. Clinical information was organized into structured fields, including case identifier, age, sex, symptoms, and ECG description. Each

record preserved a one-to-one correspondence between the tabular description and the corresponding ECG image, thereby enabling paired multimodal learning.

case_id	type_case	age	sex	symptoms	ecg_description	diagnosis	ecg_image	
0	NA63_C01	NO_ATTR	27	Male	acute, stabbing chest pain; shortness of breath, fa...	ECG: sinus rhythm with supraventricular extrasyst...	Non-specific ST-T changes.	image84
1	A38_C01	ATTR	86	Male	shortness of breath at rest; chest pain; palpitations;...	ECG: atrial fibrillation with rapid ventricular respo...	Atrial fibrillation with rapid ventricular response.	image54
2	A70_C01	ATTR	90	Male	shortness of breath on exertion; palpitations; parox...	ECG: atrial fibrillation with frequent extrasystoles.	Atrial fibrillation with frequent ventricular extrasystoles.	image71
3	NA18_C01	NO_ATTR	60	Female	chest pain on exertion or stress relieved by rest.	ECG: T wave inversion in leads I, II, V2-V4; Wellens...	Wellens' syndrome type A with ST elevation in lead aVR.	image46
4	NA17_C01	NO_ATTR	45	Male	retrosternal chest pain for 2 days.	ECG: biphasic T waves in V2-V3 (Wellens type 1).	Wellens' syndrome with 75% risk of proximal LAD stenosis.	image45

Figure 2. Preview of the multimodal dataset structure. Each case links a structured textual record to the corresponding ECG image.

2.1.1. Data Augmentation

To mitigate the data scarcity that commonly affects rare-disease studies, we augmented the dataset using a multimodal Generative Adversarial Network. The generator received a latent noise vector ($z = 100$) and a semantic condition as input, enabling the generation of paired artificial clinical descriptions and ECG tracings. The GAN was trained for 500 epochs with a batch size of 16 using the Adam optimizer, a learning rate of 2×10^{-4} , $\beta_1 = 0.5$, and $\beta_2 = 0.999$. A fixed random seed of 42 was used to support reproducibility.

After generation, records were shuffled at the case level to preserve alignment between textual and visual modalities. The resulting dataset provided a balanced experimental setting for comparing unimodal and multimodal classifiers. Because synthetic data may not fully capture the biological and clinical variability of real patients, all results are interpreted as feasibility evidence and require confirmation on larger real-world cohorts.

As a preliminary exploratory step, sample images were inspected and relationship networks were constructed using NetworkX to connect patients by sex and diagnostic class. These graphs, shown in Figures 3 and 4, were used only for descriptive visualization and were not included as model features.




Figure 3. Exploratory relationship network connecting patients by sex and ATTR status.




Figure 4. Exploratory patient network stratified by sex and diagnostic class (ATTR and NO_ATTR).

2.2. Pre-Processing

Textual and imaging data underwent harmonized pre-processing to ensure consistency across modalities. Clinical narratives were normalized by removing non-alphabetic characters, standardizing spacing, and lowercasing. Tokenization was followed by Italian stop-word removal and lemmatization. Key attributes, including age, sex, symptoms, and ECG description, were integrated into a single structured text string, while diagnostic labels were binarized as ATTR or NO_ATTR.

	case_id	label	text	image_path
0	NA63_C01	0	Male patient, 27 years old. Symptoms: chest pain t...	ECG_images/image84.png
1	A38_C01	1	Male patient, 86 years old. Symptoms: shortness of...	ECG_images/image54.png
2	A70_C01	1	Male patient, 90 years old. Symptoms: shortness of...	ECG_images/image71.png
3	NA18_C01	0	Female patient, 60 years old. Symptoms: chest pai...	ECG_images/image46.png
4	NA17_C01	0	Male patient, 45 years old. Symptoms: chest pain t...	ECG_images/image45.png

Total available rows: 100

Figure 5. Structure of the consolidated multimodal dataset used as model input. Heterogeneous features were mapped into a standardized format containing the case identifier (case_id), binary diagnostic label (label), structured clinical text (text), and ECG image path (image_path).

Text sequences were processed using the Italian BERT tokenizer with fixed-length padding and truncation. ECG images were converted to RGB and resized to 224×224 pixels. During training, light image augmentation, including rotation, horizontal flipping, and controlled brightness/contrast variation, was applied to increase visual variability. Validation and test images were only resized and normalized using ImageNet statistics to ensure compatibility with the pretrained ResNet50 encoder.

2.3. Multimodal Architecture

The model adopted a multimodal design integrating two pretrained encoders: a visual branch and a textual branch. Both encoders were frozen during training to reduce computational cost, limit overfitting, and focus learning on cross-modal fusion. The visual encoder was a ResNet50 network pretrained on ImageNet1K_V2, with the classification head replaced by an identity layer to output 2,048-dimensional feature vectors describing high-level ECG morphology. The textual encoder was Italian BERT, which produced a 768-dimensional embedding derived from the [CLS] token.

Feature vectors from the two branches were independently projected to 512-dimensional embeddings using linear layers with ReLU activation and dropout ($p = 0.2$). The resulting embeddings were concatenated and passed to a multilayer perceptron composed of a Linear–ReLU–Dropout–Linear sequence, generating logits for binary classification. Softmax transformation was applied during inference to obtain normalized class probabilities.

Data handling was managed using a custom `ECGMultimodalDataset` class responsible for loading ECG images, tokenizing textual inputs, and returning the corresponding tensors and labels. A dedicated collation function ensured homogeneous batching for efficient parallel computation. Training and evaluation relied on PyTorch `DataLoader` objects with batch size 8; shuffling was enabled during training and disabled during validation and testing.

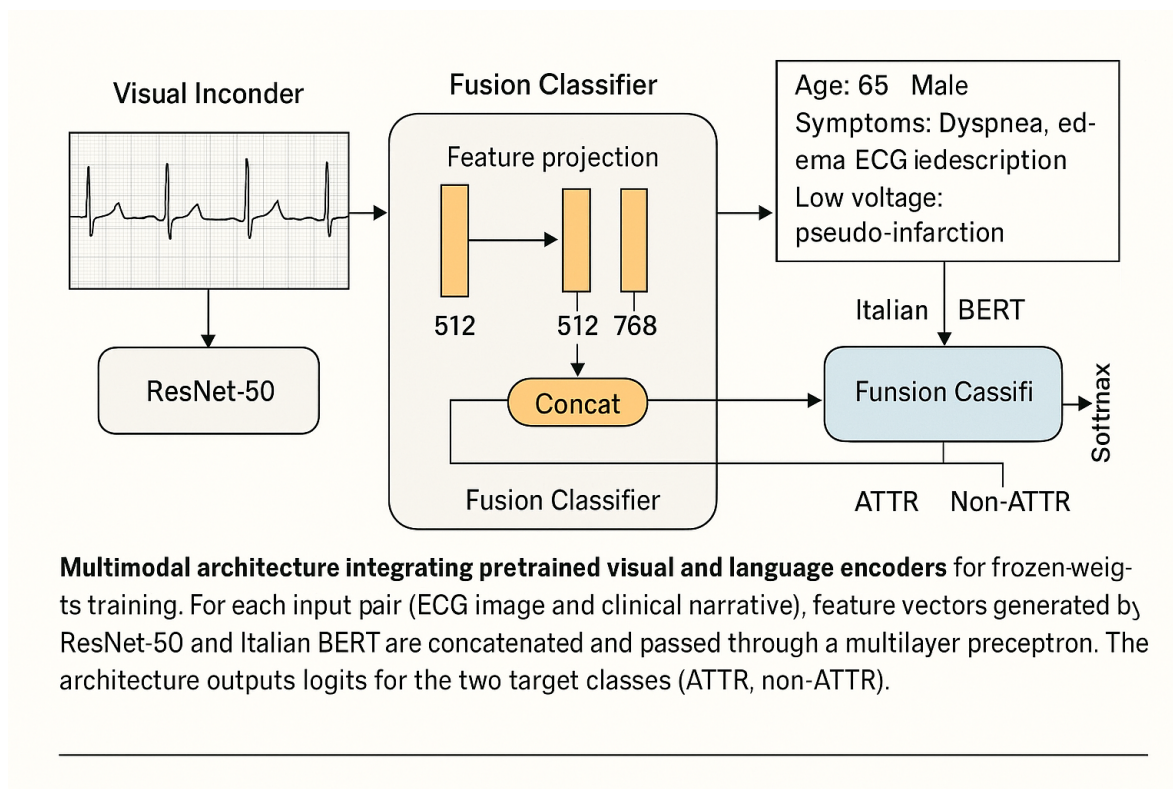


Figure 6. Overview of the multimodal architecture for automated ATTR classification. A frozen ResNet50 encoder extracts visual features from ECG images, while a frozen Italian BERT model extracts semantic representations from clinical narratives. Each modality is projected into a 512-dimensional latent space, concatenated, and processed by a multilayer perceptron for binary classification.

2.3.1. Detailed Architecture

The processing pipeline comprises three stages: extraction, projection, and classification. During feature extraction, ResNet50 produces 2,048-dimensional ECG representations and Italian BERT produces 768-dimensional textual embeddings. During projection, both representations are mapped into a common 512-dimensional latent space using Linear-ReLU-Dropout blocks. This step balances the contribution of each modality before fusion.

The concatenation of the two projected embeddings produces a 1,024-dimensional multimodal representation, which is then passed to the final classification module. A sequence of dense layers maps the fused vector to a two-dimensional logit output. Softmax calibration provides the predicted class (ATTR or NO_ATTR) and an associated confidence score, supporting transparent interpretation of the model output within this experimental setting.

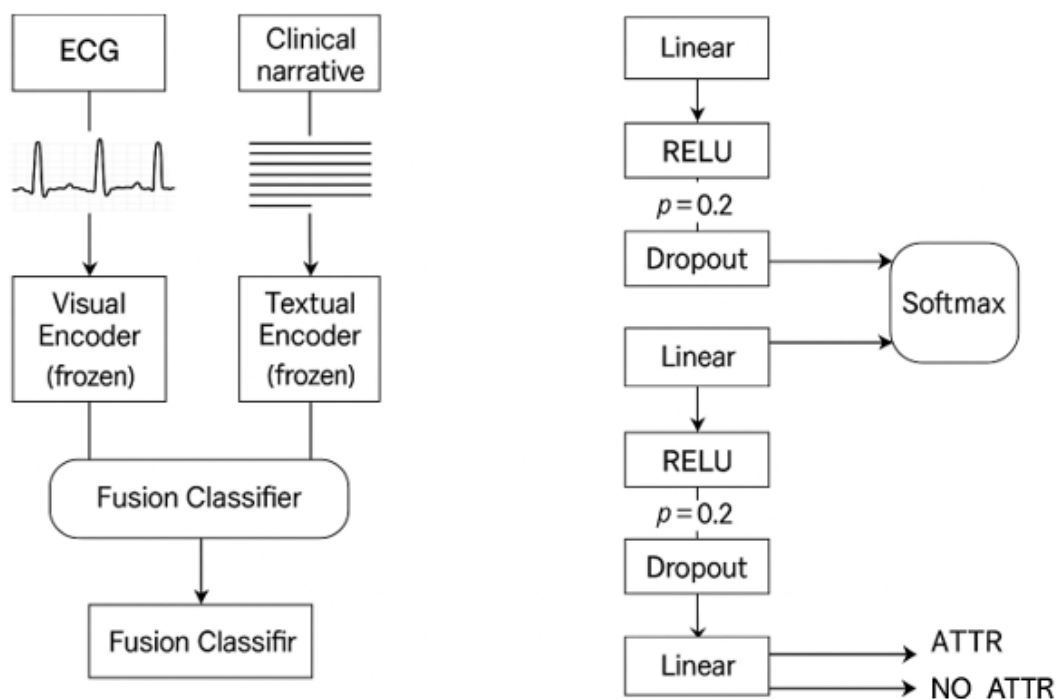


Figure 7. Implementation details of the multimodal architecture. (a) ECG images and structured clinical narratives are processed by two frozen encoders. The visual branch uses ResNet50 pretrained on ImageNet1K_V2; the textual branch uses Italian BERT. (b) Modality-specific feature vectors are independently projected into a shared 512-dimensional latent space. (c) The concatenated 1,024-dimensional representation is passed through a multilayer perceptron that outputs class logits and Softmax probabilities for ATTR and NO_ATTR.

2.4. Experimental Protocol

After pre-processing and consistency checks, incomplete records and records with missing image references were removed. The cleaned dataset was partitioned into a training-validation subset (85%) and a held-out test subset (15%) using a stratified split to preserve the class distribution between ATTR and NO_ATTR cases. All experiments were performed in Python on Google Colab with a fixed random seed of 42.

2.5. Cross-Validation and Training

Model development followed a five-fold stratified cross-validation design on the training-validation subset. For each fold, a fresh instance of the multimodal classifier was initialized with frozen encoders and randomly initialized fusion layers. Training was conducted for six epochs using the AdamW optimizer with a learning rate of 2×10^{-3} and Cross-Entropy Loss. Gradient clipping with a threshold of 1.0 was applied to improve numerical stability, and a linear learning-rate scheduler without warm-up was used.

For each fold, the model achieving the highest validation F1-score was checkpointed. After cross-validation, the best-performing model was reloaded and evaluated on the held-out test subset to estimate generalization within the constructed proof-of-concept dataset.

3. Results

3.1. Unimodal Comparative Analysis and Baseline

The initial modelling phase evaluated baseline classifiers trained on the 60 real cases only. This step was used to quantify the limitations imposed by small sample size and class imbalance before introducing GAN-based augmentation. Three representative algorithms were tested: Random Forest, Multi-Layer Perceptron, and Logistic Regression.

Random Forest (RF). Random Forest was selected because ensemble decision trees can model nonlinear relationships and are relatively robust to noisy features. In this setting, however, the limited number of cases favoured overfitting and reduced generalization to unseen samples.

```

=== REPORT RANDOM FOREST MULTIMODALE ===
              precision    recall  f1-score   support

   NO_ATTR         1.00      0.33      0.50         3
     ATTR         0.75      1.00      0.86         6

 accuracy                   0.78         9
 macro avg         0.88      0.67      0.68         9
 weighted avg         0.83      0.78      0.74         9

 Balanced Accuracy: 0.6666666666666666
 ROC AUC: 0.9444444444444444

```

Figure 8. Classification report for the Random Forest baseline trained on real cases only.

Multi-Layer Perceptron (MLP). The MLP was used to assess whether a neural classifier could capture nonlinear relationships among clinical and ECG-derived features. Despite its flexibility, convergence was unstable because the number of available samples was insufficient to reliably optimize the hidden-layer weights.

```

Report finale sul TEST SET:
              precision    recall  f1-score   support

   NO_ATTR         0.50      0.67      0.57         3
     ATTR         0.80      0.67      0.73         6

 accuracy                   0.67         9
 macro avg         0.65      0.67      0.65         9
 weighted avg         0.70      0.67      0.68         9

 Accuracy: 0.6667

```

Figure 9. Classification report for the MLP baseline trained on real cases only.

Logistic Regression (LR). Logistic Regression was included as an interpretable linear baseline. Although useful for comparison, it showed limited capacity to separate classes when symptom descriptions and ECG patterns overlapped between ATTR and NO_ATTR cases.

```

=== REPORT MULTIMODAL LOGISTIC REGRESSION ===
              precision    recall  f1-score   support

   NO_ATTR         1.00      0.67      0.80         3
     ATTR         0.86      1.00      0.92         6

 accuracy                   0.89         9
 macro avg         0.93      0.83      0.86         9
 weighted avg         0.90      0.89      0.88         9

 Balanced Accuracy: 0.8333333333333333
 ROC AUC: 1.0

```

Figure 10. Classification report for the Logistic Regression baseline trained on real cases only.

Overall, the baseline models achieved suboptimal F1-score and recall values, indicating that the real-only dataset was insufficient for a reliable screening-oriented classifier. This finding motivated the use of GAN-based augmentation as an experimental strategy for class balancing and proof-of-concept evaluation.

3.2. Impact of Data Augmentation on Classifiers

Expanding the dataset to 100 instances (60 real and 40 synthetic) improved the stability of the tested classifiers and reduced the imbalance between diagnostic classes. The augmented dataset allowed a more controlled comparison of algorithms, while remaining limited in scale and therefore unsuitable for definitive clinical conclusions.

Random Forest with GAN. The augmented dataset reduced overfitting relative to the real-only baseline. The additional variability introduced by synthetic examples acted as a form of regularization and improved the ability of the classifier to identify more stable patterns.

```

=== REPORT RANDOM FOREST MULTIMODALE ===
              precision    recall  f1-score   support

   NO_ATTR         0.75         0.50         0.60         6
    ATTR         0.73         0.89         0.80         9

 accuracy                   0.73         15
 macro avg         0.74         0.69         0.70         15
 weighted avg         0.74         0.73         0.72         15

 Balanced Accuracy: 0.6944444444444444
 ROC AUC: 0.9259259259259259

```

Figure 11. Classification report for the Random Forest classifier after GAN-based augmentation.

Logistic Regression with GAN. Logistic Regression benefited from improved class balance, producing more stable coefficients and better linear separation. Nevertheless, the model remained limited in its ability to capture complex nonlinear interactions between textual and imaging information.

```

=== REPORT MULTIMODAL LOGISTIC REGRESSION ===
              precision    recall  f1-score   support

   NO_ATTR         1.00         0.50         0.67         6
    ATTR         0.75         1.00         0.86         9

 accuracy                   0.80         15
 macro avg         0.88         0.75         0.76         15
 weighted avg         0.85         0.80         0.78         15

 Balanced Accuracy: 0.75
 ROC AUC: 0.8333333333333334

```

Figure 12. Classification report for the Logistic Regression classifier after GAN-based augmentation.

Multi-Layer Perceptron with GAN. The MLP produced the strongest overall performance among the tested classifiers. The larger and more balanced dataset allowed the neural classifier to optimize its weights more effectively, improving the balance between precision and recall.

The main improvement associated with augmentation was observed in the recall of the ATTR class. While models trained only on real data tended to favour the majority class, the synthetic contribution increased sensitivity to ATTR cases. This characteristic is relevant for screening-oriented applications, where false negatives may carry greater clinical risk than false positives. Based on its ability to fuse

features derived from BERT and ResNet50, the MLP with GAN-based augmentation was selected as the final proof-of-concept configuration.

3.3. Performance Evaluation and Metrics

The final multimodal framework was evaluated on a held-out test subset to estimate generalization within the constructed dataset. Performance was quantified using standard binary classification metrics, including accuracy, precision, recall, F1-score, and AUC-ROC.

The model achieved an overall accuracy of 73% and an AUC-ROC of 0.78. These values suggest encouraging discrimination between ATTR and NO_ATTR cases in this preliminary setting, but they should not be interpreted as evidence of clinical validity without external validation.

Consistent with a screening-oriented objective, the model achieved a recall of 0.89 for the ATTR class, corresponding to 8 correctly identified ATTR cases out of 9. Precision was 0.73 for the ATTR class and 0.75 for the NO_ATTR class, indicating a reasonable balance between sensitivity and false-positive management within the small test subset.

```
Report finale sul TEST SET:
      precision    recall  f1-score   support

   NO_ATTR      0.75      0.50      0.60         6
    ATTR       0.73      0.89      0.80         9

 accuracy              0.73         15
 macro avg           0.74      0.69      0.70         15
 weighted avg       0.74      0.73      0.72         15
```

Accuracy: 0.7333

Figure 13. Classification report for the final multimodal proof-of-concept classifier.

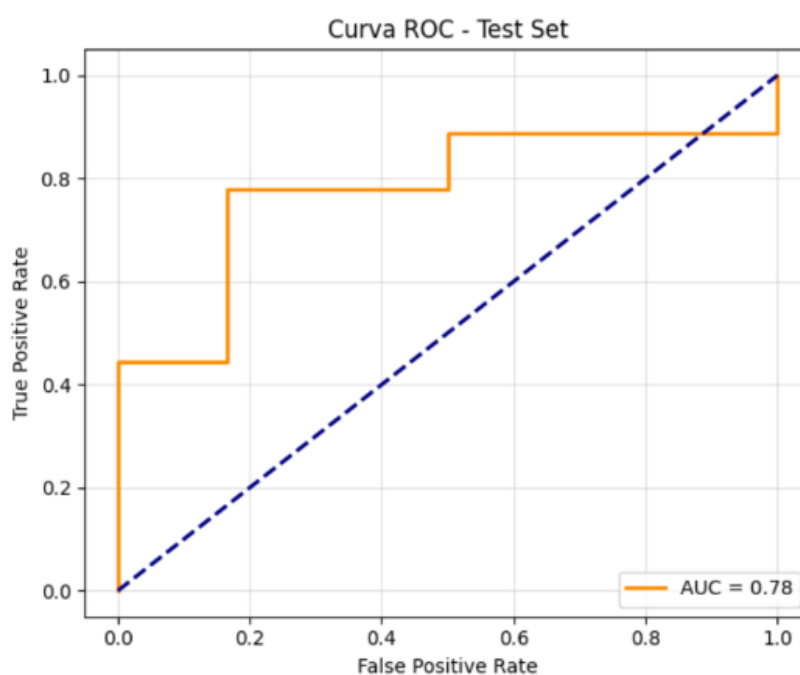


Figure 14. Receiver Operating Characteristic curve for the final multimodal classifier.

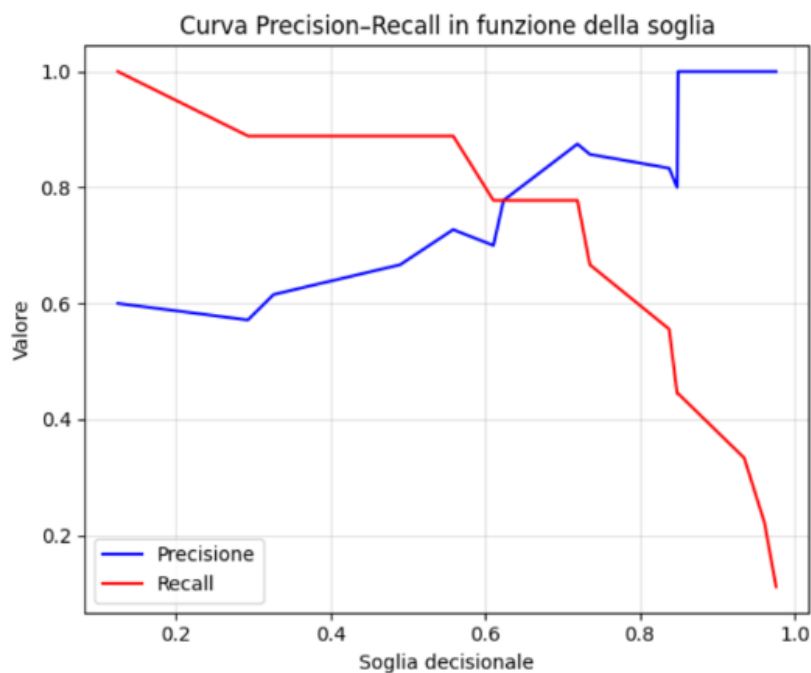


Figure 15. Precision–Recall curve for the final multimodal classifier.

The normalized confusion matrix showed lower recall for the NO_ATTR class (0.50), with three false positives among six negative cases. This pattern indicates a cautious decision profile that favours ATTR detection over specificity. Such behaviour may be acceptable in an early triage scenario, provided that positive predictions are followed by confirmatory diagnostic tests. However, the small number of test cases makes this interpretation preliminary.

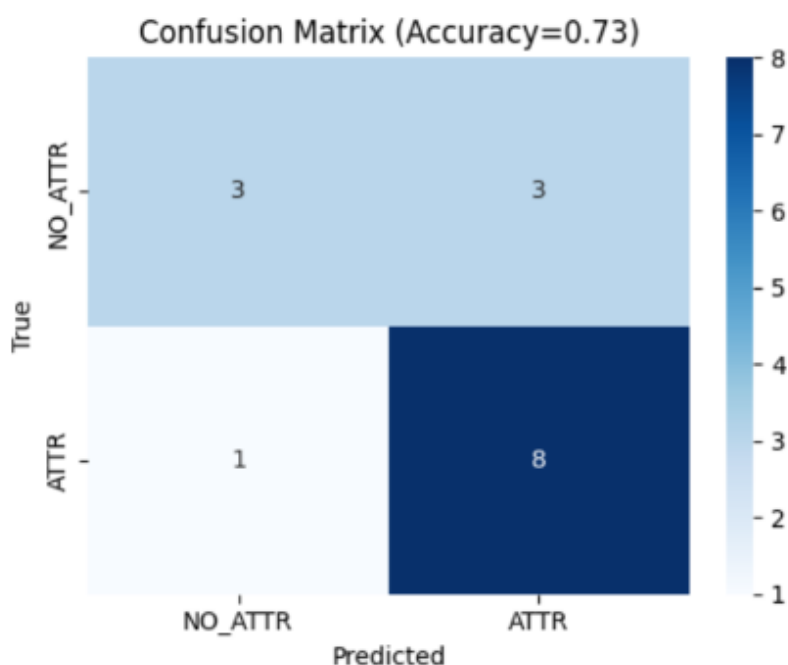


Figure 16. Normalized confusion matrix for the final multimodal classifier.

3.4. Inference Pipeline

To emulate a possible single-case workflow, we implemented an inference module that receives an ECG image and the corresponding clinical narrative [59]. The module encodes both modalities, applies multimodal fusion, and outputs the predicted class (ATTR or NO_ATTR) together with a

Softmax confidence score. In the illustrative case analysed here, the model returned a confidence score of 0.88. This interface supports qualitative inspection of model behaviour, but it should be considered an experimental prototype rather than a clinically validated tool.

4. Discussion

This proof-of-concept study suggests that integrating clinical narratives and ECG images through a multimodal deep-learning pipeline can improve the preliminary classification of ATTR compared with unimodal or real-only baselines. The final configuration combined frozen pretrained encoders with a lightweight fusion classifier, allowing the model to exploit complementary textual and visual cues while limiting the number of trainable parameters. This design is appropriate for small-data contexts, where full fine-tuning of large encoders would increase the risk of overfitting.

The most relevant result was the high recall observed for the ATTR class. In a screening-oriented scenario, sensitivity is particularly important because missed cases may prolong diagnostic delay. The model's tendency to produce false positives should therefore be interpreted in relation to its intended role: it is not a replacement for specialist assessment, but a possible triage component that could prompt confirmatory testing when clinical suspicion is high.

The Precision–Recall analysis further supports this interpretation, showing that decision-threshold selection can favour sensitivity while maintaining acceptable precision in a preliminary setting. Nevertheless, these results are based on a very small test subset and must be interpreted cautiously. In particular, the step-like shape of the Precision–Recall curve reflects the limited number of test examples and underscores the instability of performance estimates in small cohorts.

4.1. Limitations and Future Work

Several limitations define the scope of this study. First, the dataset is small and includes only 60 real cases, making it insufficient for robust clinical validation. Second, although GAN-based augmentation improves experimental balance, synthetic samples may introduce artefacts or fail to reproduce the full heterogeneity of real-world ATTR presentations. Third, the test subset contained only 15 cases, so the reported metrics are sensitive to individual classification errors. Fourth, the framework has not yet been evaluated on external multicentre cohorts, prospective data, or data acquired under routine clinical conditions.

Future work should therefore focus on expanding the dataset with fully real, multicentre cases; separating synthetic data from final clinical testing; evaluating calibration and uncertainty; and integrating additional modalities such as echocardiographic findings, laboratory markers, genetic information, and longitudinal follow-up. Explainability analyses should also be incorporated to identify which textual and ECG patterns drive model predictions, thereby increasing clinical interpretability and supporting expert review.

5. Conclusion

This work presents a proof-of-concept multimodal framework for ATTR classification based on the fusion of clinical narratives and ECG images. By combining frozen ResNet50 and Italian BERT encoders with an MLP fusion classifier, the model achieved promising preliminary sensitivity for ATTR detection in a small held-out test subset. The findings support the methodological feasibility of multimodal AI for rare-disease triage, but they do not establish clinical readiness. Larger real-world cohorts, external validation, and rigorous evaluation of robustness, calibration, and interpretability are necessary before the framework can be considered for clinical decision-support use. Within these limits, the study provides a structured starting point for the development of multimodal diagnostic support systems for ATTR and other rare diseases characterized by heterogeneous clinical and instrumental evidence.

References

1. Rajpurkar, P.; Irvin, J.; Ball, R.L. Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PLoS Medicine* **2018**, *15*, e1002686.
2. Esteva, A.; Robicquet, A.; Ramsundar, B. A guide to deep learning in healthcare. *Nature Medicine* **2019**, *25*, 24–29.
3. Lomoio, U.; Veltri, P.; Guzzi, P.H.; Liò, P. Design and use of a Denoising Convolutional Autoencoder for reconstructing electrocardiogram signals at super resolution. *Artificial Intelligence in Medicine* **2025**, *160*, 103058.
4. Brown, T.; Mann, B.; Ryder, N. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems* **2020**, *33*, 1877–1901.
5. Vaswani, A.; Shazeer, N.; Parmar, N. Attention is all you need. *Advances in Neural Information Processing Systems* **2017**, *30*.
6. Huang, Y.; Dong, Z.; Wang, J. Fusion of Visual and Language Model Features for Medical Diagnosis. *IEEE Transactions on Medical Imaging* **2023**, *42*, 233–244.
7. Hosseinzadeh, M.M.; Cannataro, M.; Guzzi, P.H.; Dondi, R. Temporal networks in biology and medicine: a survey on models, algorithms, and tools. *Network Modeling Analysis in Health Informatics and Bioinformatics* **2022**, *12*, 10.
8. Maurer, M.; Elliott, P.; Merlini, G. Diagnosis, Management, and Clinical Trials in Transthyretin Amyloidosis Cardiomyopathy. *JACC: Heart Failure* **2020**, *8*, 197–202.
9. Ruberg, F.; Berk, J. Transthyretin (TTR) Cardiac Amyloidosis. *Circulation* **2012**, *126*, 1286–1300.
10. Maurer, M.; Rapezzi, C.; Planté-Bordeneuve, V. How to diagnose cardiac amyloidosis: a practical approach. *European Heart Journal* **2022**, *43*, 1539–1549.
11. Gillmore, J.; Damy, T.; Fontana, M. Diagnosis of transthyretin amyloidosis in clinical practice. *European Journal of Heart Failure* **2022**, *24*, 759–771.
12. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M. Generative Adversarial Nets. *Advances in Neural Information Processing Systems* **2014**, *27*.
13. He, K.; Zhang, X.; Ren, S. Deep Residual Learning for Image Recognition. *IEEE Conference on Computer Vision and Pattern Recognition* **2016**, pp. 770–778.
14. Devlin, J.; Chang, M.; Lee, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805* **2018**.
15. Tanaka, K.; et al. Wild-type transthyretin amyloidosis: A case report. *World Journal of Clinical Cases* **2019**, *7*, 742–748.
16. Kittleston, M.M.; et al. Diagnosis and management of transthyretin amyloid cardiomyopathy. *Frontiers in Cardiovascular Medicine* **2023**.
17. Gagliardi, C.; et al. Transthyretin cardiac amyloidosis in the elderly. *European Heart Journal* **2020**.
18. Niu, X.; et al. Transthyretin cardiac amyloidosis: a review. *Journal of Geriatric Cardiology* **2023**, *20*, 150–162.
19. Staff, C.E. Cardiac Amyloidosis: A Case Report. *Cureus Journal of Medical Science* **2024**.
20. Staff, C.E. A Case of Transthyretin Cardiac Amyloidosis. *Cureus Journal of Medical Science* **2024**.
21. Staff, C.E. Wild-Type Transthyretin Cardiac Amyloidosis. *Cureus Journal of Medical Science* **2024**.
22. Garcia-Pavia, P.; et al. Diagnosis and treatment of transthyretin cardiac amyloidosis. *Journal of Clinical Medicine* **2024**.
23. Zhao, L.; et al. Heart failure as the initial presentation of transthyretin cardiac amyloidosis. *Medicine* **2019**.
24. Martini, N.; et al. Atrial fibrillation and transthyretin cardiac amyloidosis. *Annals of Noninvasive Electrocardiology* **2022**.
25. Staff, C.E. Transthyretin Cardiac Amyloidosis: A Case Report on Diagnostic Challenges. *Cureus Journal of Medical Science* **2024**.
26. Fujimoto, K.; et al. Case with transthyretin amyloid cardiomyopathy complicated with rapidly progressive aortic stenosis. *Journal of Cardiology Cases* **2021**.
27. AbCHF, J. Clinical characteristics of transthyretin cardiac amyloidosis. *ABC Heart Failure* **2022**.
28. Wang, Z.; et al. Diagnostic value of echocardiography in cardiac amyloidosis. *Cardiovascular Innovations and Applications* **2020**.
29. Reports, C.C. ATTR-CA and renal involvement: a case study. *Clinical Case Reports* **2024**.
30. Journal, E.H. Multimodality imaging in cardiac amyloidosis. *European Heart Journal - Case Reports* **2024**.
31. Board, M.E. Systemic amyloidosis and cardiac involvement. *Medicine (Baltimore)* **2021**.
32. Reports, O.C. Rare presentations of ATTR-CA. *Oxford Medical Case Reports* **2024**.

33. Journal, I.H. Prevalence of ATTR-CA in South Asian populations. *Indian Heart Journal* **2025**.
34. Supplements, E.H.J. Advances in TTR stabilization therapy. *European Heart Journal Supplements* **2021**.
35. Cunha, C.; et al. Transthyretin Cardiac Amyloidosis: A Review of Current and Emerging Treatment Strategies. *Journal of Cardiovascular Development and Disease* **2023**.
36. Khouri, M.G.; et al. The role of multimodality imaging in the diagnosis of cardiac amyloidosis. *NIHMS/Global Cardiology Science and Practice* **2023**.
37. Philippakis, A.; Falk, R.H. Cardiac amyloidosis mimicking hypertrophic cardiomyopathy with obstruction. *European Heart Journal* **2021**.
38. of Clinical Medical Case Reports, I.J. Diagnostic challenges in ATTR-CA: A clinical prediction study. *IJCMCR* **2024**.
39. Yousaf, A.; et al. Fabry Disease: A critical differential diagnosis in restrictive cardiomyopathy. *European Heart Journal - Case Reports* **2024**.
40. Main, A.; et al. Hypertrophic Cardiomyopathy: Clinical characteristics and differentiation from amyloidosis. *Cardiology Case Reports* **2024**.
41. of the American College of Cardiology, J. ATTR-CM: Contemporary diagnosis and management. *JACC: Case Reports* **2022**.
42. Group, P.R. Pathological findings in transthyretin cardiac amyloidosis. *Pathology International* **2023**.
43. Xie, L.; et al. Chest pain in a patient with transthyretin cardiac amyloidosis: A case report. *Clinical Case Reports* **2024**.
44. Science, G.C.; Practice. Epidemiology and outcomes of cardiac amyloidosis. *GCSP* **2018**.
45. Staff, C.E. Hypertrophic Cardiomyopathy: A Clinical Case Study. *Cureus Journal of Medical Science* **2023**.
46. of Cardiology, J. Aortic Stenosis mimicking restrictive patterns. *Journal of Cardiology Cases* **2021**.
47. Board, H.F. Case reports in restrictive cardiomyopathy. *ESC Heart Failure* **2022**.
48. Group, C.R. Hypertensive heart disease vs. Amyloidosis. *Journal of Clinical Medicine* **2021**.
49. Reports, M.C. Sarcoidosis with cardiac involvement. *Journal of Medical Case Reports* **2024**.
50. Journal, E.H. Clinical features of Fabry disease. *EHJ Case Reports* **2024**.
51. Reports, O.M. Endomyocardial fibrosis: a rare differential. *Oxford Medical Case Reports* **2024**.
52. Board, J.R. Hemochromatosis and the heart. *Journal of Integrative Cardiology* **2024**.
53. Practice, C. Athletic heart syndrome vs. Pathology. *Cardiology in Practice* **2021**.
54. Investigations, C. Myocarditis and restrictive patterns. *Journal of Clinical Investigation* **2021**.
55. Zhang, J.; et al. Machine Learning for Cardiac Disease Diagnosis. *arXiv preprint arXiv:2506.06315* **2025**.
56. Health, B.; Informatics, C. AI and Rare Disease Diagnostics in Clinical Practice. *BMJ HCI* **2024**.
57. Board, C.R. Restricted Cardiomyopathy: Emerging Trends. *Heart and Vessels* **2023**.
58. Experimental.; Medicine, T. Myocardial involvement in systemic diseases. *ETM* **2020**.
59. Sato, T.; et al. Severe cibenzoline toxicity in hypertrophic obstructive cardiomyopathy: A clinical case study. *Journal of Cardiology Cases* **2021**.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.