

# Early Explanatory Prediction of Cardiovascular Risk

Ane G. Domingo-Aldama

*University of the Basque Country, Bilbao, Biscay, Spain*

## Abstract

Early diagnosis (ED) of cardiovascular diseases (CVDs), particularly atrial fibrillation, is critical for improving patient outcomes and optimizing healthcare resource utilization. This research leverages artificial intelligence and natural language processing techniques to enhance ED by analyzing structured and unstructured Electronic Health Records (EHR) and electrocardiograms. The study focuses on three main objectives: (1) developing models for CVD prediction, (2) applying explainability techniques to improve model transparency, and (3) refining clinical guidelines using insights derived from predictive analytics.

The methodology involves multimodal machine learning approaches to integrate diverse data sources, including tabular EHR data, free-text discharge reports, and electrocardiograms. Key challenges include managing data heterogeneity and ensuring clinical interpretability of model predictions. Preliminary experiments demonstrate the potential of large tabular models and large language models in outperforming traditional clinical scoring systems. Future work will explore multimodal architectures and automated annotation techniques to further improve predictive accuracy and explainability.

This interdisciplinary research, conducted in collaboration with clinical experts, aims to bridge the gap between artificial intelligence-driven diagnostics and practical medical applications, ultimately contributing to more precise and actionable CVD risk assessment.

## Keywords

Early Diagnosis, Natural Language Processing, Machine Learning, Explainability, Cardiovascular Diseases, Atrial Fibrillation

## 1. Introduction and Rationale

Early diagnosis (ED) is gaining prominence in medical research due to advancements in biomedical informatics and the digitization of patient medical records. ED aims to detect diseases or conditions at an early stage, often before symptoms appear, enhancing treatment outcomes and patient prognosis. It supports medical professionals by providing alerts based on predictive indicators, helping to mitigate disease complications and improve treatment effectiveness while optimizing healthcare resource utilization. Currently, medical guidelines are responsible for providing recommendations for diagnosing, managing, and treating diseases, including risk assessment. However, they are often generalized, lacking specificity for local populations and clear conclusions on the superiority of one strategy over another [1].

The integration of Artificial Intelligence (AI) in biomedicine has further accelerated ED [2], improving diagnostic accuracy [3] and advancing predictive analytics and personalized medicine. Techniques such as Natural Language Processing (NLP) enable the analysis of patient medical history to identify key indicators, further enhancing early detection capabilities.

The research project outlined here focuses on ED of cardiovascular disease (CVD), crucial for proactive intervention, treatment adjustment, and preventive measures. As CVDs become a leading cause of morbidity and mortality globally [4, 5], initiatives like the EU's drive to reduce non-communicable disease burden underscore the significance of ED in healthcare. Utilizing AI and NLP techniques expedites diagnostic processes and aids informed decision-making, with a proposed emphasis on explaining prediction rationale to enhance transparency and empower personalized healthcare decisions. Using advanced NLP techniques, this project will explore the ED of CVD paying special attention to the explainability by means of the secondary use of information contained in Electronic Health Reports

---

*Doctoral Symposium on Natural Language Processing, 25 September 2025, Zaragoza, Spain.*

✉ [ane.garciad@ehu.eus](mailto:ane.garciad@ehu.eus) (A. G. Domingo-Aldama)

id 0009-0000-8202-1099 (A. G. Domingo-Aldama)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

(EHRs) in Spanish and Electrocardiograms (ECGs). Further, the insights derived from the resulting Risk Prediction explanations will be used to enhance existing medical guidelines, with the aim of uncovering new knowledge.

The origin of this research line was proposed by cardiology researchers from BioBizkaia in collaboration with the HiTZ Basque Center for Language Technology. These clinicians and researchers identified the need of deepening research on clinical AI and provided the necessary data for their creation. This is a practical, interdisciplinary research where Health and Artificial Intelligence converge towards Explainable Artificial Intelligence (XAI). The entire research is driven by real clinical needs in mind, ensuring that these tools have the potential for future application in medical practice. To achieve this, ongoing medical guidance is to be provided by the collaborators from BioBizkaia.

## 2. Background and Related Work

ED of health risks has become a critical focus in modern medical research. Cardiological disease prediction [6, 7, 8] has emerged as a critical area of research in recent years, owing to its potential to enhance ED and management. The success of ED models heavily depends on the diversity, volume, and granularity of the data available [9, 10]. Modern AI systems for personalized, predictive, participatory, and preventive medicine increasingly rely on EHRs [11]. However, the growing availability of diverse and large-scale data sources has resulted in increasingly complex AI model architectures. This complexity exacerbates the black-box nature of these models, raising concerns about interpretability and transparency, particularly in the medical domain where decision-making must be transparent and trustworthy.

Considering this, the focus of this research lies on three major topics: Generating AI models capable of predicting CVDs, XAI and enhancing Clinical Guidelines.

Recent advancements in CVD prediction have explored a variety of approaches, leveraging structured data, free-text clinical narratives, and imaging-based features. Traditional Machine Learning (ML) techniques, such as Support Vector Machines (SVMs), K-Nearest Neighbors (KNN), and ensemble models, have been widely applied to structured EHR data and patient vectors, achieving high predictive performance [12, 13]. Large Language Models (LLMs) have also been employed to process textual discharge reports, demonstrating strong capabilities in risk stratification and outperforming conventional classifiers [14, 15]. Additionally, studies incorporating ECG data have shown that integrating these signals with either ML or deep learning (DL) models can enhance predictive accuracy [12, 13]. Recent reviews [16, 17, 18] highlight the growing interest in multimodal learning, where different data modalities—structured vectors, textual reports, and ECGs—are combined to improve disease detection and prognosis, emphasizing the importance of explainability and interpretability in these models.

Among different CVDs, Atrial Fibrillation (AF) is one of the most prevalent. The first steps of the thesis project primarily focuses on AF, with plans for future expansion to include other CVDs. Research in this domain typically focuses on two primary scenarios: the prediction of new-onset AF and the recurrence of AF following therapeutic interventions. AI-driven models have demonstrated remarkable success in predicting incident and recurrent AF, often outperforming conventional methods [19]. These studies leverage diverse datasets, including clinical data, cardiac imaging data, and electrophysiological data [20].

Regarding the second objective of this thesis XAI [21] consists of integrating low-level attributes of DL with higher-level schemas inherent to human argumentation capabilities. XAI aims to bridge the gap between high-performance DL models and the human need for interpretability by explaining model decisions in a manner aligned with human reasoning processes [21]. In medical applications, understanding the mechanisms triggering models' predictions is crucial for fostering trust, especially in life-critical scenarios [22, 23].

Addressing the third objective of this study, current CVD guidelines cover the diagnosis and management of CVD in adults and provide recommendations for delivering optimal care and treatment for individuals with CVD, including the assessment and management of complication risks. However,

these guidelines are designed to be general in nature. Consequently, they may appear overly broad and ambiguous due to their generalized approach. Additionally, they are subscribed to clinical studies where the cohorts are limited in size. Therefore, important gaps exist in the evidence on effectiveness of implementation interventions, especially regarding broad clinical outcomes.

Regarding AF, guidelines offer clear indications for anticoagulation, but the choice between rate control and rhythm control depends on various factors. The 2020 European guidelines recommend rhythm control over rate control in symptomatic AF cases, especially for patients with paroxysmal forms and no risk factors for arrhythmia recurrence post-surgery [24]. Comparatively, catheter ablation has shown superiority over drug therapy in promoting sinus rhythm recovery [25]. Consequently, ED systems that guide the treatment selection have a huge importance.

### 3. Description of the Research

As outlined in section 1, the primary objective of this research is to advance AI-driven ED in CVDs by utilizing complex datasets from EHRs and ECGs. With a particular focus on AF, the most common arrhythmia worldwide, the study aims to refine predictive modeling techniques while emphasizing explainability. The project is set to a duration of 3-4 years to achieve the following objectives:

- **Research Goal 1:** Development of AI models for the ED and prediction of CVDs.
- **Research Goal 2:** Application of XAI techniques to enhance transparency and interpretability in CVD prediction systems.
- **Research Goal 3:** Comparison and refinement of clinical guidelines using insights generated by our predictive models and XAI algorithms.

Building on these objectives, the current thesis and experiments aim to address three key research questions:

**Research Question 1: Structured EHR data and free-text discharge reports capture patient clinical information in different formats. How can Language Models (LMs) and ML methods leverage these resources to improve the prediction of CVDs?**

*Hypothesis:* Free-text discharge reports, authored by clinicians, provide rich, narrative descriptions of a patient's clinical history, while structured EHR data is produced through a coding process in which documentation services convert information from clinical notes into coded entries using standardized codes. Both formats present challenges: discharge reports are subject to the complexities of natural language, whereas structured data can suffer from inaccuracies and missing values. Consequently, this research aims to evaluate whether these data sources, when combined with appropriate model architectures and data preprocessing techniques, can be effectively utilized to train predictive models for CVDs. Structured tabular data will serve as input for traditional ML and Large Tabular Models (LTMs), while discriminative and generative LMs will be employed to process free-text discharge reports. Beyond developing models capable of ED of CVDs, this research seeks to compare different architectures to determine their relative effectiveness.

**Research Question 2: EHRs, both structured and unstructured, and ECGs contain complementary clinical information that healthcare professionals use to assess patient history, guide diagnosis, and inform treatment decisions. How can Multimodal Machine Learning Models (MMLMs) leverage these data sources to enhance the prediction and reasoning of CVD risks, particularly AF recurrence?**

*Hypothesis:* Combining complementary information of structured EHRs and ECGs has shown to be an effective approach when predicting the risk of CVD using traditional ML techniques and hopefully this holds true for multimodal approaches when discharge reports are combined with ECGs and other structured data. The proposed research will be conducted using a two-stage, triple-approach method: Firstly, CVD risk prediction through structured EHR data processed by ML methods and, concurrently, discharge reports processed by LLMs; secondly, discharge reports and ECGs processed jointly by MMLMs and, conversely, structured EHR data enriched with ECG features processed by ML methods,

and, finally, the combination of both structured and free-text EHR data and the combination of all data sources with MMLM models will be studied. This strategy follows two purposes: on the one hand, it provides a way to compare LLM and MMLM with ML methods; on the other hand, it allows to measure the impact of the addition of different information in the risk prediction of CVD.

**Research Question 3: Is it possible to enhance existing clinical guidelines for the diagnosis and management of cardiovascular disease (CVD)?** Current guidelines encompass the diagnosis and management of CVD in adults and provide recommendations for delivering optimal care and treatment for individuals with CVD, including the assessment and management of complication risks. However, these guidelines typically do not account for geographic variations or detailed phenotypic characteristics. Furthermore, they often do not address patient-specific factors, including certain comorbidities, which fall outside their scope because they are based on limited size cohort studies.

*Hypothesis:* Secondary use of clinical information within the Basque region allows to uncover local population phenotyping and characteristics not captured in the existing guidelines and might also help better adjust recommendations and assessments included in the guidelines by the secondary use of vast amounts of health data by means of NLP and ML techniques.

## 4. Methodology and Experiments

In the current section we explain the methodology and experiments planned to answer the previous research questions and to achieve the main objectives of the study (see Figure 1). The experiments and methodologies proposed are based on the available resources kindly provided by the Basque Public Healthcare System (Osakidetza), more concretely clinical information from patients treated at Basurto Hospital in the Basque Country. The resources include:

- *Discharge reports in Spanish:* A pool of  $1.2 \times 10^6$  discharge reports dating from 2015 to 2020, corresponding to 305,358 unique patients (both with and without atrial fibrillation).
- *Codified Tabular Data from the Osakidetza Business Intelligence (OBI) System:* Structured clinical data for 9,191 patients who experienced an AF debut. The information is encoded by healthcare professionals using standardized coding systems and stored in the OBI platform.
- *Electrocardiograms (ECGs):* Patient ECGs stored in MUSE XML format; currently not available for analysis.

The research methodology follows a structured approach that integrates data preparation (package 1. in Figure 1), predictive modelling (package 2. in Figure 1), explainability (package 3 in Figure 1), and medical guideline enhancement (package 4 in Figure 1). While each component has a specific objective, they are not strictly sequential; multiple tasks can be conducted in parallel to optimize research efficiency and facilitate iterative improvements.

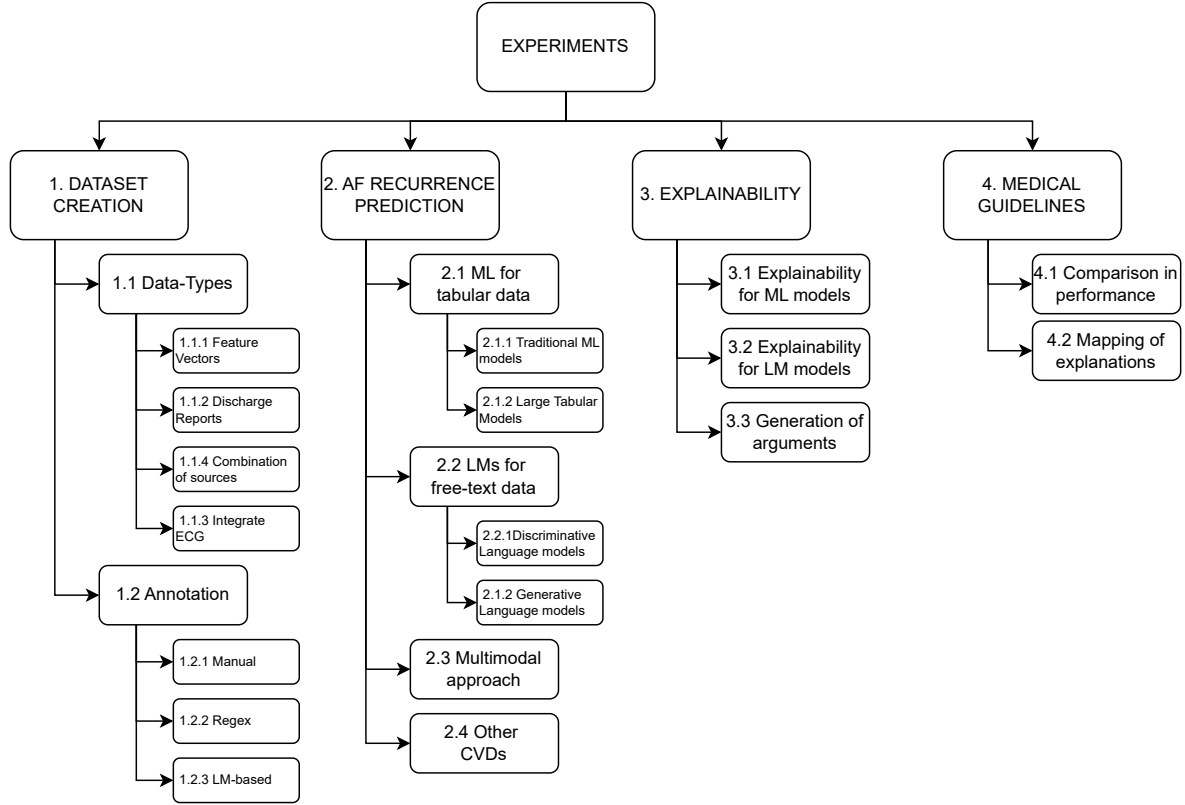
Moreover, each package is aligned with one of the research questions outlined in section 3. The first two packages correspond to the first research question, while the third and fourth packages are linked to the second and third research questions, respectively.

In addition, the models and techniques proposed in this work are intentionally lightweight and computationally efficient, making them suitable for deployment in real-world clinical settings where access to high-end GPUs may be limited.

### 4.1. Dataset Creation

The first package focuses on constructing high-quality datasets to train, validate, and test predictive models. Distinct versions of Clinical Histories (CHs) are generated leveraging both unstructured (free-text) and structured (tabular data) and ECGs. This dual approach ensures comprehensive data representation, facilitating both traditional ML techniques and advanced language modelling for disease prediction and analysis.

The dataset preparation is conducted in two distinct modalities: tabular (vectorized) CHs and natural language-based CHs. The first version (1.1.1) generates tabular representations of patients' health status



**Figure 1:** Schema of the experiments planned for the current thesis, consisting on 4 main packages and their sub-tasks.

by combining structured and unstructured EHR data into a structured format. The second version (1.1.2) focuses on creating clinical histories (CHs) in natural language, utilizing discharge reports. Another version (1.1.3) adopts a multimodal approach, combining both data sources to overcome the limitations of each individually. Finally, the integration of ECG information to enhance these representations (1.1.4) will be explored, contingent on the availability of ECG data in this phase of the research.

Moreover, this study aims to address the limitations of ED systems that rely solely on structured EHR data or manually annotated datasets. Structured EHR data, while valuable, is often incomplete or prone to errors due to inconsistencies in data entry and missing information [26, 27]. On the other hand, manual annotation is time-consuming, labor-intensive, and not easily scalable. To overcome these challenges, the research includes experiments on the automatic generation of patient cohorts and disease detection in free-text clinical histories. These experiments will be conducted using both rule-based algorithms (1.2.2) and generative LMs (1.2.3), enabling a more efficient and accurate identification of relevant clinical patterns while reducing the dependence on manual intervention.

## 4.2. AF Recurrence Prediction Experimentation

The second package involves training and evaluating predictive models for AF recurrence based on the datasets created in the previous step. Different modelling approaches are applied depending on the data format:

- Vector-based ML models: CHs represented as structured vectors are used to train traditional ML algorithms (2.1.1) and Large Tabular Models (LTM) (2.1.2).
- NLP-based models: Prediction models are developed using LMs trained on CHs (2.2) derived from discharge reports. In this stage, due to the nature of the task and the emphasis on learning from the training set, primarily discriminative LMs of different pretrainings and architectures will be



employed—such as the EriBERTa Spanish clinical LM [28]. However, the fine-tuning of generative LLMs, such as Gemma3 or Aloe-Beta, using techniques like QLoRA, is also being considered.

- Multimodal approaches: Prediction models based on MMLM for multimodal inputs (2.3).
- Generalization to other cardiovascular diseases (CVDs): The methodology is extended to explore its applicability to other CVD conditions (2.4).

### 4.3. Explainability of Model Predictions

To improve the interpretability of AF recurrence predictions, explainability mechanisms are incorporated into the modelling framework. Two experiments are conducted: the application of explainability algorithms to ML models (3.1) and the application of explainability algorithms to LM models (3.2). To address these objectives, we plan to employ model-agnostic explainability techniques, including local surrogate methods such as LIME [29], contrastive approaches like Shapley values, and example-based strategies such as counterfactual explanations.

These algorithms typically provide numerical attributions for each feature or token, consequently, the generation of argumentation based on these local and global explanations will be explored (3.3).

### 4.4. Enhancing Medical Guidelines

The outcomes of the prediction and explainability experiments are compared against current clinical guidelines. Presently, widely used clinical scoring systems, such as CHADS2-VASc, HATCH, and APPLE, are employed to predict new-onset AF and postoperative AF recurrence [30, 31, 32].

The predictive performance of these scores can be assessed by defining appropriate classification thresholds and comparing their predictions against the annotated ground truth labels in the test set. Consequently, our primary objective (4.1) is to develop AI-based models that outperform these established clinical scoring systems. A secondary objective (4.2) is to identify key predictive features and extract decision boundaries from our models, with the goal of enhancing the accuracy of current clinical tools. While the initial focus is on AF, the methodology is designed to be extendable to other cardiovascular diseases (CVDs).

### 4.5. Current stage of the study

Currently, a manuscript detailing experiments on tabular data generation (2.1) using rule-based annotations (1.2.2) and AF recurrence prediction based on structured and unstructured EHR data (1.1.1) has been submitted to a journal. These experiments compare traditional ML algorithms (2.1.1) to TabPFN [33] (2.1.2), a LTM specialized in small tabular datasets. The LTM approach surpasses the performance of current clinical scores across multiple subsets when evaluated on a manually annotated test set (4.1).

Consequently, the study now focuses on generating automatic annotations using LLMs and encoder-based architectures (1.2.3). This shift was motivated by the identification of noise introduced through silver rule-based labelling in the tabular data experiments.

In parallel, experiments utilizing discharge reports for AF recurrence prediction are in the preliminary stages, with both generative LLMs and discriminative LMs being explored (2.2).

Explainability experiments are currently on hold (3), with only preliminary attempts made using SHAP values for tabular models and LIME [29] and Integrated Gradients [34] for LM models, yielding no conclusive results thus far.

## 5. Discussion and challenges

ED of CVDs is crucial for improving patient outcomes, reducing complications, and optimizing treatment strategies. AI-driven predictive models have the potential to enhance ED by analyzing complex and heterogeneous data sources. However, developing accurate models for CVD prediction remains a significant challenge due to the multifactorial nature of these diseases. Many CVDs, including AF

recurrence, have underlying causes that are not fully understood, making it difficult to establish clear predictive patterns. Additionally, factors such as genetic predisposition, lifestyle influences, and comorbid conditions contribute to disease progression in ways that are not always captured by available clinical data. Despite these challenges, AI offers a promising avenue for uncovering hidden patterns, identifying potential risk factors, and refining current clinical guidelines to improve early diagnosis and patient care.

All data types used in this project present challenges related to relevance, missing information, outliers, and noise. This issue is particularly evident in discharge reports, where excessive and unstructured content creates long, complex contexts that hinder the extraction of key clinical information. However, similar challenges exist in tabular data, which may contain missing values and outliers, and in ECGs, which might be unavailable or of inconsistent quality. Therefore, identifying appropriate architectures or preprocessing techniques is crucial to ensure that model decisions are based on meaningful and reliable information. In prior experiments, LLMs have demonstrated an ability to manage these limitations in tabular data. However, for discharge reports, although LLMs can process longer contexts than discriminative models, preliminary experiments indicate that the high level of noise in medical reports continues to hinder their effective extraction of relevant information.

Another key discussion point is the application of explainability algorithms that are not only relevant but also interpretable for clinicians. While it is crucial to use methods that accurately represent the model's decision boundaries, commonly used algorithms such as LIME [29] and Integrated Gradients [34] often produce numerical representations of tokens or features that lack clinical interpretability. Therefore, developing strategies to translate these outputs into meaningful and actionable explanations remains a critical challenge.

Moreover, integrating text, tabular data, and ECG information into a unified multimodal model presents several challenges. Each data modality has distinct characteristics: textual data from discharge reports is unstructured and often contains noise and redundancy, tabular data is structured but may include missing values and outliers, and ECG signals are time-series data with high dimensionality. Effectively combining these heterogeneous sources requires architectures capable of learning meaningful cross-modal representations while preserving the unique features of each modality. Additionally, aligning temporal and contextual relationships between these data types is complex, as ECGs and structured EHRs provide quantitative physiological measurements, whereas clinical text contains qualitative descriptions and reasoning.

In conclusion, addressing these challenges and developing models with genuine predictive capability, comparable to or surpassing existing clinical guidelines, are the core objectives of this thesis.

## Declaration of Generative AI

Declaration on Generative AI. During the preparation of this work, the author(s) used ChatGPT in order to: Grammar and spelling check, Paraphrase, translate and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

## References

- [1] V. C. Pereira, S. N. Silva, V. K. Carvalho, F. Zanghelini, J. O. Barreto, Strategies for the implementation of clinical practice guidelines in public health: an overview of systematic reviews, *Health research policy and systems* 20 (2022) 13.
- [2] Y. Kumar, A. Koul, R. Singla, M. F. Ijaz, Artificial intelligence in disease diagnosis: a systematic literature review, synthesizing framework and future research agenda, *Journal of Ambient Intelligence and Humanized Computing* 14 (2023) 8459–8486. doi:<https://doi.org/10.1007/s12652-021-03612-z>.

- [3] Z. Kanjee, B. Crowe, A. Rodman, Accuracy of a generative artificial intelligence model in a complex diagnostic challenge, *JAMA* 330 (2023) 78–80. URL: <https://doi.org/10.1001/jama.2023.8288>. doi:10.1001/jama.2023.8288.
- [4] D. S. Celermajer, C. K. Chow, E. Marijon, N. M. Anstey, K. S. Woo, Cardiovascular disease in the developing world: prevalences, patterns, and the potential of early disease detection, *Journal of the American College of Cardiology* 60 (2012) 1207–1216.
- [5] G. Lippi, F. Sanchis-Gomar, G. Cervellin, Global epidemiology of atrial fibrillation: an increasing epidemic and public health challenge, *International journal of stroke* 16 (2021) 217–221.
- [6] M. A. Naser, A. A. Majeed, M. Alsabah, T. R. Al-Shaikhli, K. M. Kaky, A review of machine learning's role in cardiovascular disease prediction: Recent advances and future challenges, *Algorithms* 17 (2024) 78.
- [7] C. M. Bhatt, P. Patel, T. Ghetia, P. L. Mazzeo, Effective heart disease prediction using machine learning techniques, *Algorithms* 16 (2023) 88.
- [8] J. Soni, U. Ansari, D. Sharma, S. Soni, et al., Predictive data mining for medical diagnosis: An overview of heart disease prediction, *International Journal of Computer Applications* 17 (2011) 43–48.
- [9] K. Ng, S. R. Steinhubl, C. DeFilippi, S. Dey, W. F. Stewart, Early detection of heart failure using electronic health records: practical implications for time before diagnosis, data diversity, data quantity, and data density, *Circulation: Cardiovascular Quality and Outcomes* 9 (2016) 649–658.
- [10] A. A. Alzu'bi, V. J. Watzlaf, P. Sheridan, Electronic health record (ehr) abstraction, *Perspectives in health information management* 18 (2021).
- [11] B. Ristevski, M. Chen, Big data analytics in medicine and healthcare, *Journal of integrative bioinformatics* 15 (2018) 20170030.
- [12] S. Sharma, M. Parmar, Heart diseases prediction using deep learning neural network model, *International Journal of Innovative Technology and Exploring Engineering (IJITEE)* 9 (2020) 2244–2248.
- [13] F. Ali, S. El-Sappagh, S. R. Islam, D. Kwak, A. Ali, M. Imran, K.-S. Kwak, A smart healthcare monitoring system for heart disease prediction based on ensemble deep learning and feature fusion, *Information Fusion* 63 (2020) 208–222.
- [14] C. Han, D. W. Kim, S. Kim, S. C. You, S. Bae, D. Yoon, Large-language-model-based 10-year risk prediction of cardiovascular disease: insight from the uk biobank data, *medRxiv* (2023) 2023–05.
- [15] J. James, H. Brooks, I. J. Kullo, Leveraging large language models for cardiovascular mortality prediction from ct chest reports, *Journal of the American College of Cardiology* 83 (2024) 2441–2441.
- [16] N. E. Almansouri, M. Awe, S. Rajavelu, K. Jahnavi, R. Shastry, A. Hasan, H. Hasan, M. Lakkimsetti, R. K. AlAbbasi, B. C. Gutiérrez, et al., Early diagnosis of cardiovascular diseases in the era of artificial intelligence: An in-depth review, *Cureus* 16 (2024).
- [17] A. B. Teshale, H. L. Htun, M. Vered, A. J. Owen, R. Freak-Poli, A systematic review of artificial intelligence models for time-to-event outcome applied in cardiovascular disease risk prediction, *Journal of medical systems* 48 (2024) 68.
- [18] A. Di Costanzo, C. A. M. Spaccarotella, G. Esposito, C. Indolfi, An artificial intelligence analysis of electrocardiograms for the clinical diagnosis of cardiovascular diseases: a narrative review, *Journal of Clinical Medicine* 13 (2024) 1033.
- [19] K. C. Siontis, X. Yao, J. P. Pirruccello, A. A. Philippakis, P. A. Noseworthy, How will machine learning inform the clinical care of atrial fibrillation?, *Circulation research* 127 (2020) 155–169.
- [20] A. S. Tseng, P. A. Noseworthy, Prediction of atrial fibrillation using machine learning: a review, *Frontiers in Physiology* 12 (2021) 752317.
- [21] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, *Artificial intelligence* 267 (2019) 1–38.
- [22] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, H. Müller, Causability and explainability of artificial intelligence in medicine, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 9 (2019) e1312.
- [23] R. L. Pierce, W. Van Biesen, D. Van Cauwenberge, J. Decruyenaere, S. Sterckx, Explainability in



- medicine in an era of ai-based clinical decision support systems, *Frontiers in genetics* 13 (2022) 903600.
- [24] G. Hindricks, T. Potpara, N. Dagres, E. Arbelo, J. J. Bax, C. Blomström-Lundqvist, G. Boriani, M. Castella, G.-A. Dan, P. E. Dilaveris, et al., 2020 esc guidelines for the diagnosis and management of atrial fibrillation developed in collaboration with the european association for cardio-thoracic surgery (eacts) the task force for the diagnosis and management of atrial fibrillation of the european society of cardiology (esc) developed with the special contribution of the european heart rhythm association (ehra) of the esc, *European heart journal* 42 (2021) 373–498.
  - [25] J. G. Andrade, M. W. Deyell, P. Khairy, J. Champagne, P. Leong-Sit, P. Novak, L. Sterns, J.-F. Roux, J. Sapp, R. Bennett, et al., Atrial fibrillation progression after cryoablation vs. radiofrequency ablation: the circa-dose trial, *European heart journal* 45 (2024) 510–518.
  - [26] A. Garcia Olea, J. Ormaetxe Merodio, A. Atutxa Salazar, I. Diez Gonzalez, I. Fernandez De La Prieta, M. Maeztu Rada, E. Amuriza De Luis, K. Ugedo Alzaga, U. Idiazabal Rodriguez, I. Pereiro Lili, et al., The role of congestive heart failure at atrial fibrillation onset in the data entry errors of electronic health records, in: *EUROPEAN JOURNAL OF HEART FAILURE*, volume 23, WILEY 111 RIVER ST, HOBOKEN 07030-5774, NJ USA, 2021, pp. 303–304.
  - [27] T. Botsis, G. Hartvigsen, F. Chen, C. Weng, Secondary use of ehr: data quality issues and informatics opportunities, *Summit on translational bioinformatics* 2010 (2010) 1.
  - [28] I. de la Iglesia, A. Atutxa, K. Gojenola, A. Barrena, Eriberta: A bilingual pre-trained language model for clinical natural language processing, 2023. URL: <https://arxiv.org/abs/2306.07373>. arXiv:2306.07373.
  - [29] M. T. Ribeiro, S. Singh, C. Guestrin, ” why should i trust you?” explaining the predictions of any classifier, in: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.
  - [30] O. Ramos, The use of hatch score for the prediction of post operative atrial fibrillation (poaf) after coronary artery bypass graft: a meta-analysis, *European Heart Journal* 44 (2023) ehac779–019.
  - [31] F. Vitali, M. Serenelli, J. Airaksinen, R. Pavasini, A. Tomaszuk-Kazberuk, E. Mlodawska, S. Jaakkola, C. Balla, L. Falsetti, N. Tarquinio, et al., Cha2ds2-vasc score predicts atrial fibrillation recurrence after cardioversion: Systematic review and individual patient pooled meta-analysis, *Clinical cardiology* 42 (2019) 358–364.
  - [32] W. Huang, H. Sun, Y. Luo, S. Xiong, Y. Tang, Y. Long, Z. Zhang, H. Liu, Better performance of the apple score for the prediction of very early atrial fibrillation recurrence post-ablation, *Hellenic Journal of Cardiology* (2024).
  - [33] N. Hollmann, S. Müller, L. Purucker, A. Krishnakumar, M. Körfer, S. B. Hoo, R. T. Schirrmeister, F. Hutter, Accurate predictions on small data with a tabular foundation model, *Nature* 637 (2025) 319–326.
  - [34] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks, in: *International conference on machine learning*, PMLR, 2017, pp. 3319–3328.