

Modeling Clinical Workflow for SYNTAX Scoring from Coronary Angiography Videos

Suzhong Fu^{1,2}, Jingqi Dong^{1,2}, Xuan Ding^{1,2}, Rui Sun^{1,2}, Yiming Yang^{1,2},
Shuguang Cui^{2,1}, and Zhen Li^{2,1*}

¹ FNii-Shenzhen, The Chinese University of Hong Kong, Shenzhen, China

² School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China

Abstract. The SYNTAX score is a clinically established tool for assessing anatomical lesion complexity in coronary artery disease and guiding subsequent treatment. However, automated SYNTAX scoring is commonly formulated as a direct regression problem from coronary angiography videos to patient-level scores. In this work, we reformulate SYNTAX scoring as a vessel segment identity-preserving anatomical reasoning problem and propose a hierarchical modeling framework that explicitly aligns learning with the clinical workflow. Our approach maintains vessel segment identity across frames and views, estimates stenosis severity at the segment level, and aggregates evidence hierarchically according to coronary anatomy. Simultaneously, to address the scarcity of domain-specific data, we integrate and complete multiple public coronary angiography datasets, constructing a large-scale resource featuring completed vessel segmentation and derived structural annotations. Experiments demonstrate that vessel segment-level stenosis embedding enhances explanatory power and reduces prediction variability compared to baseline models, with the R^2 score improving by 0.201 and dev STD decreasing by 18.4%. These results highlight the necessity of structure-aligned modeling for reliable and stable automated SYNTAX scoring from multi-view coronary angiography videos. The GitHub link is [SYNTAX_score_777](#).

Keywords: Coronary Angiography · SYNTAX Score Prediction · Multi-view Video Analysis.

1 Introduction

Coronary artery disease (CAD)[1,2] remains a leading cause of morbidity and mortality worldwide, and invasive coronary angiography is the clinical gold standard for visualizing coronary lesions and guiding subsequent treatment. The SYNTAX scoring[3,4] outlined in clinical guidelines is inherently structured: cardiologists reason over anatomically defined vessel segments across multiple views and integrate segment-level findings hierarchically into left/right and patient-level scores.

* Corresponding author: Zhen Li, email: lizhen@cuhk.edu.cn




Fig. 1: Conceptual comparison between naïve multi-view aggregation and identity-preserving hierarchical modeling for automated SYNTAX prediction. Left: Direct aggregation of multi-view features leads to anatomical identity collapse and structural ambiguity. Middle: Clinical reasoning operates on anatomically defined vessel segments and hierarchical integration. Right: The proposed framework preserves segment identity across frames and views, enabling structure-aware aggregation for patient-level SYNTAX prediction.

Despite recent advances in deep learning for coronary angiography analysis, most existing approaches address isolated components of this reasoning process, such as vessel segmentation[5,6,7,8,10,9], stenosis detection[11,12,13,14,18], or dominance classification[15,16]. A limited number of studies attempt direct end-to-end SYNTAX score prediction from angiographic videos[17]. However, these methods typically formulate the task as a black-box regression from high-dimensional spatiotemporal inputs to a low-dimensional clinical score, collapsing structured anatomical reasoning into flat feature representations. As a result, vessel identity correspondence across views, segment-level lesion attribution, and hierarchical aggregation are not explicitly modeled, limiting interpretability, robustness, and generalization across heterogeneous clinical data.

In this work, we argue that automated SYNTAX prediction should be formulated as a structured anatomical reasoning problem rather than a direct regression task. We propose a clinically inspired, vessel segment identity-preserving framework that explicitly models anatomical identity propagation across frames, segment-level lesion estimation, and hierarchical aggregation across vessels and views. Central to our approach is a segment-aware decision field that embeds

discrete segment-level stenosis assessments into a continuous, topology-aware representation, enabling spatially coherent reasoning while preserving anatomical correspondence throughout the modeling pipeline.

Simultaneously, to address the scarcity of domain-specific data, we integrate and complete multiple public coronary angiography datasets, constructing a large-scale resource featuring completed vessel segmentation and derived structural annotations. Experiments demonstrate that explicitly modeling anatomical structure and hierarchical aggregation leads to consistently improved patient-level SYNTAX score prediction performance and reduced prediction variability compared with strong baselines, with particularly pronounced benefits in anatomically complex coronary systems. Our contributions are summarized as follows:

1. Introduce a vessel segment identity-preserving modeling framework that aligns learning objectives with the degree of stenosis in each coronary vessel segment;
2. Construct a coronary artery representation network capable of jointly learning from vessel segment stenosis Attention Heatmaps and multi-view videos to predict SYNTAX scores;
3. Integrate and refine commercially available high-quality coronary X-ray data, utilizing progressive generation to complete segmentation labels for all 1.2M images;
4. Experiments validate that our framework yields more stable and interpretable SYNTAX predictions across heterogeneous multi-view data.

2 Method

Problem Formulation under Hierarchical Structure. A coronary angiography examination consists of multiple videos acquired from heterogeneous viewing angles, each containing hundreds of frames. Formally, an examination is represented as $\mathcal{V} = \{V_1, \dots, V_N\}$, where each video $V_i = \{I_{i,1}, \dots, I_{i,T_i}\}$ captures spatiotemporal vessel dynamics. Unlike conventional video regression problems, SYNTAX scoring is inherently hierarchical. The final clinical targets (y_{left} , y_{right} , y_{full}) are determined by localized lesion decisions over anatomically defined vessel segments, which are further organized within left and right coronary systems. Therefore, the learning objective is not merely to approximate a mapping $F : \mathbb{R}^{N \times T_i \times H \times W} \rightarrow \mathbb{R}^3$, but to construct intermediate representations that preserve segment-level identity and hierarchical structure before aggregation. This perspective motivates explicit enforcement of anatomical identity and structured feature integration as learning constraints, rather than direct end-to-end regression. Because raw angiography videos include non-informative phases caused by absent or dissipated contrast agent and cardiac motion, the model is trained on high-quality clips rather than uniformly sampling all frames. We annotated high-quality video segments from all frame-level data of 120 patients and trained a binary ResNet-34 classifier using more than 800 videos to distinguish informative frames from other frames. The complete supplementary videos are kept




Fig. 2: Schematic diagram of the proposed syntax score prediction framework for a single case. Key frames are selected from the patient’s videos to ensure the accuracy of feature information sources. High-quality frames from each perspective video serve as sources for generating narrow attention heatmaps. The stenosis attention heatmap is combined with video key frames as joint inputs to the network, comprising a video-based visual encoder and a sequential feature decoder that aggregates multi-video features for syntax prediction.

untrimmed for transparency, whereas training and heatmap construction use the selected high-quality clips.

Identity-Preserving Anatomical Modeling. To prevent segment anatomical identity collapse across frames and views, we represent vessel segments as topology-defined structural entities rather than appearance-dependent pixel regions. In coronary angiography, pixel-level representations are highly sensitive to projection angle, vessel foreshortening, and contrast dynamics, causing the same anatomical segment to exhibit drastically different appearances across views. Given a vessel segmentation mask M , we extract its centerline and represent the vascular tree as a connectivity graph $G = (\mathcal{N}, \mathcal{E})$, where nodes correspond to endpoints or branching points and edges correspond to centerline branches. By decomposing the graph at bifurcation nodes, we obtain a set of linear vessel segments $\{S_k\}_{k=1}^K$, each corresponding to an anatomically defined coronary segment. This topology-based decomposition transforms the vascular tree into structurally stable segment entities whose identities remain consistent across frames and views, despite variations in motion, contrast injection, and imaging projection. By anchoring segment identity to vascular topology rather than pixel appearance, the model retains anatomically meaningful correspondence throughout subsequent temporal propagation and multi-view aggregation. The segment

extraction follows the skeleton-graph splitting strategy in VesSAM [5]: bifurcation and endpoint nodes are first detected on the vessel skeleton, and connected skeleton branches are then used to define segment candidates. These segments are mainly used to produce stable, non-overlapping video crops for downstream stenosis estimation. To reduce error accumulation in this multi-stage pipeline, segmentation masks are iteratively refined, segment tracking is performed at the frame level, and highly overlapping vessel crops are removed before stenosis estimation.

Segment-Aware Decision Field. To preserve both the magnitude and anatomical support of segment-level clinical findings, we introduce a *segment-aware decision field*, which embeds discrete stenosis assessments into a continuous spatial representation while maintaining explicit segment correspondence. For each tracked vessel segment S_k , we estimate a stenosis severity score s_k from its corresponding vessel crop using the segment-level stenosis estimator, and map it to a non-negative modulation coefficient α_k . The spatial contribution of segment k at location $p \in \mathbb{R}^2$ is modeled as

$$H_k(p) = \alpha_k \exp \left(-\frac{1}{2} (p - \mu_k)^\top \Sigma_k^{-1} (p - \mu_k) \right),$$

where μ_k denotes the spatial center of segment S_k , and Σ_k encodes its spatial extent and orientation. In practice, (μ_k, Σ_k) are derived from the geometry of the tracked segment (e.g., centerline statistics) and are fixed given the segmentation and tracking results. The final decision field is obtained by superposing contributions from all segments, $H(p) = \sum_{k=1}^K H_k(p)$, resulting in a continuous spatial field that encodes the distribution of segment-level stenosis evidence over the vascular tree. We do not claim the Gaussian form is optimal; it serves as a continuous carrier of discrete clinical decisions.

This construction embeds discrete segment-level clinical decisions into a continuous, topology-aware representation that preserves both anatomical locality and segment identity. Unlike conventional attention or saliency maps, which are typically learned from appearance features and vary across views, the proposed decision field is explicitly anchored to tracked anatomical segments and remains consistent across frames and views. As a result, the decision field enforces identity-preserving modulation during downstream feature learning and hierarchical aggregation, rather than serving as a view-specific importance weighting. This superposition also makes the representation less sensitive to occasional pseudo-bifurcation points caused by 2D vessel overlap: even if a projected vessel branch is split into multiple crops, the local stenosis evidence is still mapped back to the same image support and contributes to the final decision field.

Hierarchical Multi-Video Aggregation. Case-level SYNTAX prediction requires integrating evidence across multiple angiographic videos while preserving vessel-level correspondence. In multi-view examinations, different videos often capture distinct subsets of coronary segments under heterogeneous projections. Directly pooling video-level embeddings implicitly assumes that such embeddings are anatomically aligned, an assumption that is violated when segment identity is not explicitly enforced.

Table 1: Summary of datasets used in this study. Segmentation availability is reported for original annotations, optimized and after completion by our pipeline.

Attribute	Syntax-score[18]	Cardio-Syntax[17]	ARCADE[6]	DCA1[7]	XCAD[8]
Dim	2D	3D	2D	2D	2D
Cases Imgs	237 3459	3018 1.2M	* 1500	* 130	* 126
Multi-view	✓	✓	×	×	×
Dominance	✓	1025	×	×	×
SYNTAX	✓	1844	×	×	×
Segmentation	Completed	Completed	Optimized	Original	Original
Stenosis	✓	×	✓	×	×

We therefore adopt a hierarchical aggregation strategy aligned with coronary anatomy, rather than performing direct, view-agnostic pooling. Each video is first encoded into a video-level representation using a feature extractor that is explicitly conditioned on the segment-aware decision field. The decision field serves as an identity-preserving modulation, emphasizing anatomically localized stenosis evidence while suppressing appearance-driven ambiguity.

Video-level representations are then grouped according to their associated coronary system (left or right), reflecting the anatomical organization underlying SYNTAX scoring, before subsequent aggregation. This hierarchy-aware aggregation prevents feature mixing between anatomically distinct coronary systems, ensuring that evidence from the left and right coronary arteries is not inadvertently entangled. By enforcing vessel-level correspondence prior to patient-level integration, the model preserves the anatomical structure required for reliable SYNTAX score estimation and mitigates instability caused by heterogeneous multi-view inputs.

3 Experiments

Datasets. Data barriers remain a significant obstacle to progress in this field, as most current research relies on proprietary datasets. No single dataset simultaneously contains large-scale multi-view videos, complete vessel segmentation, stenosis attributes, and patient-level SYNTAX scores. This fragmented supervision setting introduces structural inconsistency and increases the risk of anatomical identity collapse during learning. To enable structure-aware learning under this fragmented supervision regime, we utilized high-quality annotations from DCA1 and XCAD to complete missing segmentation labels in the CardiaSyntax and SYNTAX-Score datasets through progressive segmentation refinement and structural annotation transfer. Additionally, segmentation labels for the ARCADE dataset were optimized. The high-quality clip annotation set used for frame selection contains 120 patients and more than 800 videos, and will be released together with the code and trained weights. Dataset characteristics and segmentation availability are summarized in Table 1.

Baselines. We reproduce the prior state-of-the-art method (CarS)[17] under identical 5-fold patient-level splits. To isolate structural contributions, we evaluate: (i) CarS + SAH, which augments the baseline with segment-aware hierarchical guidance, (ii) Ours w/o SAH, which removes hierarchical guidance from our framework, and (iii) Ours, the complete identity-preserving model. We focus on CarS because it is the directly reproducible temporal coronary angiography benchmark for SYNTAX regression. CathAI-like systems target 2D angiographic stenosis detection and severity estimation rather than full-pipeline video-based SYNTAX scoring, and their code, weights, and datasets are not publicly available for direct adaptation.

Evaluation Performance is assessed from three complementary perspectives: (1) explanatory power measured by R^2 , (2) prediction bias measured by mean and median error, and (3) prediction dispersion measured by the standard deviation of prediction errors (STD). Unless otherwise specified, values reported in the main tables correspond to mean $\pm$ standard deviation across folds. These regression metrics are used to evaluate prediction reliability and to allow direct comparison with CardioSyntax [17]. We acknowledge that clinically risk-stratified metrics based on SYNTAX score thresholds are also important for downstream decision support and should be included in future external clinical validation.

4 Results and analysis

Table 2 summarizes patient-level SYNTAX score prediction performance on the test dataset. Compared with the baseline, structure-aware models consistently improve explanatory power while simultaneously reducing prediction variability. These gains directly reflect the benefit of explicitly preserving anatomical identity and hierarchical structure during multi-view aggregation.

The complete model achieves the highest R^2 for both left (0.469) and full (0.529) SYNTAX scores. Notably, performance gains are more pronounced for the left coronary system, which is characterized by more complex branching topology and a larger number of anatomically distinct segments. This trend is consistent with our formulation in Section 2, indicating that identity-preserving and topology-aware modeling is particularly advantageous under higher anatomical complexity. Structure-aware variants exhibit increased mean bias for the full SYNTAX score, indicating sensitivity to global score scaling.

Given that SYNTAX scoring is itself derived from discrete, segment-level clinical decisions, improved variance explanation with near-zero median bias reflects enhanced structural consistency rather than random prediction drift. As shown in Tables 2 and 3, STD consistently decreases for structure-aware models, with particularly pronounced reductions for the left and full SYNTAX scores. This reduction supports the claim that explicitly preserving anatomical identity mitigates variability introduced by heterogeneous multi-view aggregation.

Challenge SYNTAX cases experiments. We further evaluate model performance on non-zero SYNTAX cases, which represent a more challenging clinical setting. As shown in Table 4, overall performance decreases across all meth-

Table 2: Overall SYNTAX score prediction performance on the full test dataset. Mean $\pm$ standard deviation over 5-fold patient-level cross-validation are reported.

Model	$R^2 \uparrow$	Bias Mean	Bias Median	STD $\downarrow$
CarS	0.328 ± 0.035	-0.41 ± 1.56	0.01 ± 0.02	8.36 ± 0.22
CarS w SAH	0.420 ± 0.023	-1.05 ± 0.32	0.04 ± 0.08	7.83 ± 0.16
Ours w/o SAH	0.467 ± 0.047	-2.11 ± 1.19	0.06 ± 0.12	7.17 ± 0.49
Ours	0.529 ± 0.027	-1.96 ± 0.62	0.02 ± 0.03	6.82 ± 0.09
Δ vs. CarS	+0.201	-1.55	+0.01	-1.54(-18.4%)

Table 3: Left and right coronary SYNTAX prediction performance on the full dataset. Mean values over 5-fold patient-level cross-validation are reported.

Left coronary (LCA)				Right coronary (RCA)		
Model	$R^2 \uparrow$	Bias Mean	STD $\downarrow$	$R^2 \uparrow$	Bias Mean	STD $\downarrow$
CarS	0.164	0.36	7.71	0.253	-0.77	2.45
CarS w SAH	0.310	-0.38	7.14	0.239	-0.67	2.47
Ours w/o SAH	0.422	-1.40	6.29	0.232	-0.71	2.50
Ours	0.470	-1.15	6.14	0.254	-0.81	2.44
Δ vs. CarS	+0.306	-1.51	-1.57	+0.001	-0.04	-0.01

ods, reflecting the increased difficulty of this subset. Nevertheless, structure-aware models consistently improve both explanatory power and prediction stability compared with the baseline. In particular, the complete identity-preserving model achieves higher R^2 while simultaneously reducing deviation standard deviation for both left and overall SYNTAX scores, indicating more stable patient-level predictions under harder conditions. This suggests that explicitly modeling anatomical identity and hierarchical aggregation mitigates performance degradation when trivial cases are removed.

Effect of Removing Multi-view Aggregation. We further analyze a degenerate setting in which multi-view aggregation is removed and SYNTAX prediction is performed using single-view representations. As shown in Table 5, this setting leads to severely degraded performance, with consistently negative R^2 values for left and full SYNTAX scores. Notably, incorporating heatmap-based guidance without multi-view aggregation does not improve performance and often further destabilizes prediction. This observation indicates that reliable SYNTAX prediction fundamentally requires explicit aggregation of segment-level evidence across multiple angiographic views.

5 Conclusion

In this work, we proposed a vessel segment identity-preserving hierarchical framework for automated SYNTAX score prediction from coronary angiography videos, explicitly aligning learning with the structured clinical workflow underlying SYNTAX scoring. By maintaining vessel segment identity across frames and views

Table 4: SYNTAX prediction performance on non-zero cases. Mean values over 5-fold patient-level cross-validation are reported.

Method	Left SYNTAX		Overall SYNTAX	
	$R^2 \uparrow$	STD $\downarrow$	$R^2 \uparrow$	STD $\downarrow$
CarS	-0.338 $\pm$ 0.160	10.914 $\pm$ 0.512	-0.186 $\pm$ 0.063	11.786 $\pm$ 0.418
CarS w SAH	-0.096 $\pm$ 0.031	10.249 $\pm$ 0.169	-0.015 $\pm$ 0.042	10.983 $\pm$ 0.275
Ours w/o SAH	0.101 $\pm$ 0.107	8.456$\pm$0.959	0.089 $\pm$ 0.055	9.155 $\pm$ 0.917
Ours	0.160$\pm$0.033	8.500 $\pm$ 0.429	0.179$\pm$0.052	8.978$\pm$0.312
Δ vs. CarS	+0.498	-2.41	+0.365	-2.81

Table 5: Effect of removing multi-view aggregation on SYNTAX prediction. Mean R^2 values across 5-fold patient-level cross-validation are reported.

Setting	Left	Right	Full
Mean pooling (single-view)	-4.34 $\pm$ 5.97	0.07 $\pm$ 0.15	-4.47 $\pm$ 2.99
Mean pooling + SAH	-5.13 $\pm$ 5.61	-0.12 $\pm$ 0.24	-5.34 $\pm$ 7.68

and aggregating evidence according to the stenosis severity of coronary anatomy, our method avoids anatomical identity collapse inherent in direct end-to-end regression. We integrated and completed multiple public coronary angiography datasets to alleviate the current data constraints in this field. Experiments across multiple datasets demonstrate consistently improved explanatory power and reduced prediction variability compared with strong baselines, with particularly pronounced gains for left coronary systems and non-zero SYNTAX cases. These results highlight the importance of structure-aware and topology-informed modeling for reliable angiographic video analysis. Our current study is still limited by the absence of an independent external clinical cohort, the lack of risk-class-wise clinical evaluation, the relatively narrow comparison set, and the use of a Gaussian decision field as a practical continuous representation rather than an optimized spatial kernel. Future work will further validate the framework on external cohorts and explore more clinically grounded spatial representations. Our work provides a step toward more clinically faithful and robust automated assessment of coronary disease complexity.

Acknowledgments. This work was supported in part by Guangdong S&T Programme with Grant No. 2024B0101030002, the Basic Research Project No. HZQB-KCZY-2021067 of Hetao Shenzhen-HK ST Cooperation Zone, the Shenzhen Outstanding Talents Training Fund 202002, the NSFC with Grant No. 62293482, by NSFC with Grant No. 62573371, by the Guangdong Province Radio Science Data Center with grant No. 2025B1212070001, by the Shenzhen General Program No. JCYJ202205301436 00001, by the Guangdong Research Project No. 2017ZT07X152 and No. 2019CX01X104, by the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001), by the NSFC 61931024 & 12326610, by the Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No. SYSPG20241211173853027), by National Key Research and Development Program of China 2025YFF0515300 and 2025YFF0515304, the Open Project Program (Grant No. QHSFCS- 2606) of Key Lab-

oratory of Tibetan Information Processing, Ministry of Education, by the Shenzhen-Hong Kong Joint Funding No. SGDXX20211123112401002, and by Tencent & Huawei Open Fund 2024E0009 & 202301030019.

Declaration of Competing Interest The authors have no competing interests to declare that are relevant to the content of this article.

References

1. McCullough, Peter A.: Coronary artery disease. *Clinical Journal of the American Society of Nephrology* 2.3 (2007): 611-616.
2. Shahjehan, Rai Dilawar, Sanjeev Sharma, and Beenish S. Bhutta.: Coronary artery disease. StatPearls [Internet]. StatPearls Publishing, 2024.
3. Serruys, Patrick W., et al.: Assessment of the SYNTAX score in the Syntax study. *EuroIntervention* 5.1 (2009): 50-56.
4. Sianos, Georgios, et al.: The SYNTAX Score: an angiographic tool grading the complexity of coronary artery disease. *EuroIntervention* 1.2 (2005): 219-227.
5. Fu, Suzhong, et al.: VesSAM: Efficient Multi-Prompting for Segmenting Complex Vessel. 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). IEEE, 2025.
6. Popov, Maxim, et al.: Dataset for automatic region-based coronary artery disease diagnostics using X-ray angiography images. *Scientific data* 11.1 (2024): 20.
7. Cervantes-Sanchez, Fernando, et al.: Automatic segmentation of coronary arteries in X-ray angiograms using multiscale analysis and artificial neural networks. *Applied Sciences* 9.24 (2019): 5507.
8. Ma, Yuxin, et al.: Self-supervised vessel segmentation via adversarial learning. proceedings of the IEEE/CVF international conference on computer vision. 2021.
9. He, Yanglong, et al.: Conditional Virtual Imaging for Few-Shot Vascular Image Segmentation. *IEEE Transactions on Medical Imaging* (2025).
10. Gao, Yunlong, et al.: Spatio-Temporal correspondence attention network for vessel segmentation in X-ray coronary angiography. *Biomedical Signal Processing and Control* 99 (2025): 106792.
11. Avram, Robert, et al.: CathAI: fully automated coronary angiography interpretation and stenosis estimation. *npj Digital Medicine* 6.1 (2023): 142.
12. Labrecque Langlais, Élodie, et al.: Evaluation of stenoses using AI video models applied to coronary angiography. *NPJ digital medicine* 7.1 (2024): 138.
13. Li, Xinyu, et al.: STQD-Det: Spatio-temporal quantum diffusion model for real-time coronary stenosis detection in X-ray angiography. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46.12 (2024): 9908-9920.
14. Pang, Kun, et al.: Stenosis-DetNet: Sequence consistency-based stenosis detection for X-ray coronary angiography. *Computerized Medical Imaging and Graphics* 89 (2021): 101900.
15. Kruzhilov, Ivan, et al.: CoronaryDominance: Angiogram dataset for coronary dominance classification. *Scientific Data* 12.1 (2025): 341.
16. Akyürek, Hasan Ali.: Real-Time Coronary Artery Dominance Classification from Angiographic Images Using Advanced Deep Video Architectures. *Diagnostics* 15.10 (2025): 1186.
17. Ponomarchuk, Alexander, et al.: CardioSyntax: End-to-End SYNTAX Score Prediction-Dataset, Benchmark and Method. 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). IEEE, 2025.

18. Mahmoudi, Seyed Sajjad, et al.: X-ray coronary angiogram images and SYNTAX score to develop machine-learning algorithms for CHD diagnosis. *Scientific Data* 12.1 (2025): 471.