algorithms

MDPI

Article

A Data-Centric Algorithmic Pipeline for Enhancing Cardiac MRI Segmentation Using ViTUNeT and Quality-Aware Filtering

Salvador de Haro [1,*], Jesús Cámara [2], Pilar González-Férez [1], José Manuel García [1] and Gregorio Bernabé [1]

[1] Department of Computer Engineering, University of Murcia, 30100 Murcia, Spain; pilargf@um.es (P.G.-F.); jmgarcia@um.es (J.M.G.); gbernabe@um.es (G.B.)
[2] Department of Informatics, University of Valladolid, 47011 Valladolid, Spain; jesus.camara@uva.es
* Correspondence: salvadorde.haroo@um.es

Abstract

The performance of deep-learning-based segmentation models is strongly dependent on the quality of the input data, which is frequently heterogeneous or degraded in real-world medical imaging scenarios. This work presents a data-centric algorithmic pipeline designed to improve cardiac MRI segmentation accuracy through systematic image enhancement and automatic slice-quality filtering. The proposed method is formalized as deterministic algorithm that combines image processing and supervised learning components. The approach integrates a contrast- and structure-preserving enhancement stage, based on bilateral filtering and adaptive histogram equalization, with a quality-aware selection algorithm. Slice quality is assessed using anatomical attributes extracted via YOLOv11s-based localization and a supervised classification model trained to identify diagnostically reliable images. When applied to transformer-based segmentation architectures such as ViTUNeT, the pipeline yields consistent improvements across all evaluation metrics without increasing model complexity or training cost. These findings emphasize the importance of algorithmic data curation as an effective strategy for enhancing robustness and stability in deep-learning segmentation pipelines and demonstrate the broader applicability of the proposed approach to computer-vision tasks involving heterogeneous or low-quality image datasets.

Keywords: data-centric AI; medical image segmentation; cardiac MRI; ViTUNeT; slice quality assessment; YOLOv11; image enhancement

1. Introduction

Left ventricular non-compaction (LVNC) is a structurally distinct cardiomyopathy characterized by excessive myocardial trabeculation, deep intertrabecular recesses communicating with the ventricular cavity, and a thin compacted myocardial layer [1]. Clinically, LVNC exhibits a heterogeneous spectrum ranging from asymptomatic phenotypes to severe manifestations such as heart failure, ventricular arrhythmias, and thromboembolic events [2,3]. Cardiac magnetic resonance imaging (MRI), owing to its high spatial resolution and superior soft-tissue contrast, has become the reference modality for morphological assessment of LVNC and related cardiomyopathies.

Quantitative analysis of myocardial trabeculation plays a central role in LVNC evaluation. In this context, the trabeculated volume percentage (TV%) is one of the most objective and widely adopted metrics [4–6]. Accurate TV% estimation requires precise segmentation of multiple left ventricular structures, including the endocardial cavity, compacted

Check for updates

Academic Editor: Lei Fan

Received: 26 January 2026
Revised: 2 March 2026
Accepted: 2 March 2026
Published: 6 March 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

myocardium, and trabeculated region. However, manual segmentation remains common in clinical practice despite being time-consuming, operator-dependent, and poorly scalable.

These limitations have driven the development of automated segmentation pipelines. Convolutional neural networks, particularly U-Net-based architectures [7], enabled end-to-end segmentation and led to deep learning systems for trabecular quantification in cardiac MRI [8]. More recently, hybrid models combining convolutional encoders with global attention mechanisms based on Vision Transformers have been proposed to better capture the complex morphology of the left ventricle. Among them, ViTUNeT integrates convolutional feature extraction with long-range dependency modeling and has shown improved generalization across heterogeneous datasets [9]. Nevertheless, direct comparisons with purely convolutional architectures such as DL-LVTQ [10] indicate that increasing model complexity alone does not consistently translate into substantial performance gains, particularly in the trabeculated myocardium.

A closer analysis reveals that image quality constitutes a major bottleneck. Clinical cardiac MRI datasets frequently contain slices with low contrast, noise artifacts, or poorly defined anatomical boundaries, especially in basal and apical regions. Such low-quality images introduce ambiguity in both ground-truth annotations and the learning process, ultimately limiting segmentation accuracy and robustness. This observation aligns with broader evidence in medical image analysis, where data quality and preprocessing critically influence deep learning outcomes. In particular, edge-preserving denoising and contrast enhancement techniques, such as bilateral filtering and histogram-based methods, have been shown to improve downstream model performance while preserving anatomical structures [11,12].

In parallel, object detection frameworks from the YOLO family have demonstrated strong performance in medical imaging tasks requiring robust anatomical localization. Originally developed for real-time detection in natural images [13], YOLO-based models have been successfully applied to brain tumor detection in MRI [14], fracture localization in X-ray imaging [15], and multimodal classification of neurodegenerative diseases [16]. These results highlight the suitability of YOLO models as efficient tools for extracting anatomically meaningful features from medical images.

Motivated by these observations, this work adopts a data-centric algorithmic perspective and investigates whether systematic improvement of input data quality can unlock latent performance in state-of-the-art segmentation models. We propose a quality-aware pipeline that combines contrast- and structure-preserving visual enhancement, based on bilateral filtering [17] and contrast-limited adaptive histogram equalization (CLAHE) [18,19], with an automatic slice-quality selection algorithm driven by anatomically grounded features extracted using YOLOv11. Rather than increasing architectural complexity, the proposed approach focuses on algorithmic data curation and preprocessing, and its impact is evaluated on a ViTUNeT-based cardiac MRI segmentation pipeline, demonstrating improved robustness and segmentation performance in heterogeneous image datasets.

2. Materials and Methods

This work is framed as a data-centric, quality-aware preprocessing algorithm that operates before segmentation. The proposed pipeline has two complementary components: (i) a lightweight visual enhancement stage that increases anatomical contrast while preserving structure, and (ii) an automatic slice-quality estimator that uses anatomically grounded attributes extracted by a YOLOv11 detector to reject unreliable slices prior to segmentation. The overall goal is to improve robustness and stability of downstream models (here, ViTUNeT) without increasing segmentation-network complexity.

2.1. Dataset and Study Cohort

We use the latest version of an in-house short-axis cardiac MRI dataset curated by our research group, comprising 3459 diastolic-phase slices from 450 anonymized patients provided by collaborating hospitals. Each subject contributes approximately seven slices. A recent extension incorporated additional low-trabeculation cases, increasing morphological diversity and reducing bias in subsequent analyses (Figure 1).

The cohort is organized into four diagnostic groups:

- **P**: 293 patients diagnosed with hypertrophic cardiomyopathy (HCM) according to the Maron criteria [20].
- **H**: 78 patients presenting with LVNC, defined according to Petersen et al. [21].
- **X**: 69 patients exhibiting mixed or unclassifiable cardiomyopathy phenotypes.
- **T**: 10 patients affected by titin-related cardiomyopathy [22].

All images correspond to the diastolic phase of the cardiac cycle, which constitutes the clinical reference phase for LVNC assessment. At end-diastole, myocardial trabeculation is maximally expressed and diagnostic criteria such as trabeculated volume percentage are defined. Restricting the analysis to a single phase reduces variability due to cardiac motion and phase-dependent deformation, enabling a controlled evaluation of the proposed quality-aware preprocessing algorithm. Although this design limits direct applicability to other cardiac phases, the pipeline itself is phase-agnostic and can be extended through phase-specific retraining.

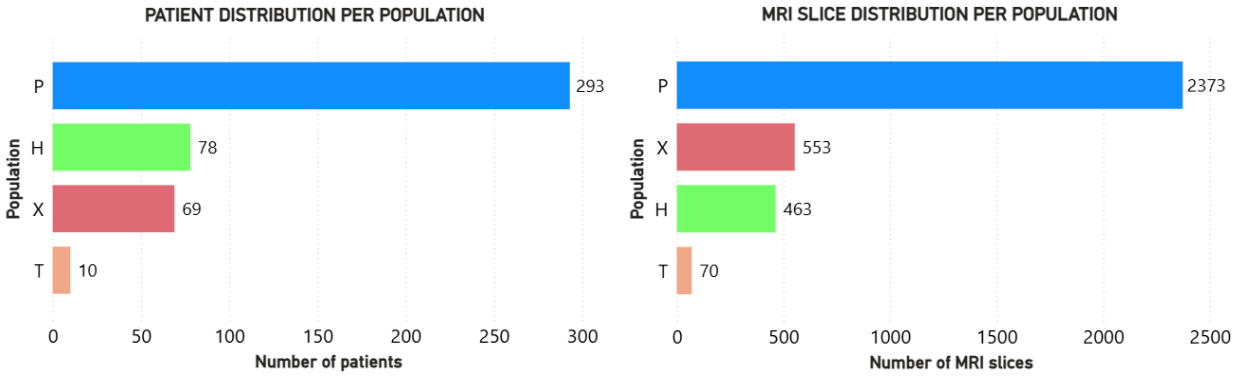


Figure 1. Distribution in the current dataset version. (**a**) Number of patients per diagnostic group. (**b**) Number of MRI slices per diagnostic group.

2.2. Visual Enhancement Preprocessing

Cardiac MRI slices frequently exhibit low contrast and acquisition-related artifacts that obscure myocardial boundaries, particularly in basal and apical regions. To address these limitations, we apply a three-stage visual enhancement pipeline designed to improve trabecular visibility while preserving anatomical edge integrity.

1. Intensity rescaling. Each slice is linearly mapped to an 8-bit intensity range $[0, 255]$ to standardize dynamic range across acquisitions and provide a consistent representation for subsequent processing. In our pipeline, MRI slices are already handled at 8-bit resolution; therefore, this operation does not reduce bit depth but homogenizes intensity ranges prior to OpenCV-based bilateral filtering and CLAHE. Subsequent z-score normalization ensures that all segmentation models operate on floating-point inputs.
2. Edge-preserving denoising. A bilateral filter [17] is applied to suppress impulsive noise and background artifacts while preserving the myocardium–blood pool boundary. The

bilateral filter is configured with a kernel diameter $d = 7$, $\sigma_{\text{color}} = 75$, and $\sigma_{\text{space}} = 75$, providing effective noise attenuation while preserving anatomical boundaries.

3. Local contrast normalization. Contrast-limited adaptive histogram equalization (CLAHE) [18,19] enhances local contrast and improves the visibility of subtle intensity variations, particularly within trabeculated regions. CLAHE is applied with a clip limit of 1.0 and a tile grid size of $(2, 2)$ to avoid over-enhancement artifacts while improving local structural contrast.

The sequential order of these operations is intentional. Intensity rescaling first homogenizes the input intensity distribution. Bilateral filtering is then applied to reduce noise while preserving anatomical edges, preventing noise amplification during contrast enhancement. Finally, CLAHE is applied once noise has been attenuated, ensuring that contrast enhancement primarily highlights anatomically meaningful structures rather than acquisition artifacts.

As illustrated in Figure 2, the enhanced images present improved intensity balance and clearer delineation of both compacted myocardium and trabeculated regions, reducing boundary ambiguity and facilitating downstream segmentation tasks.

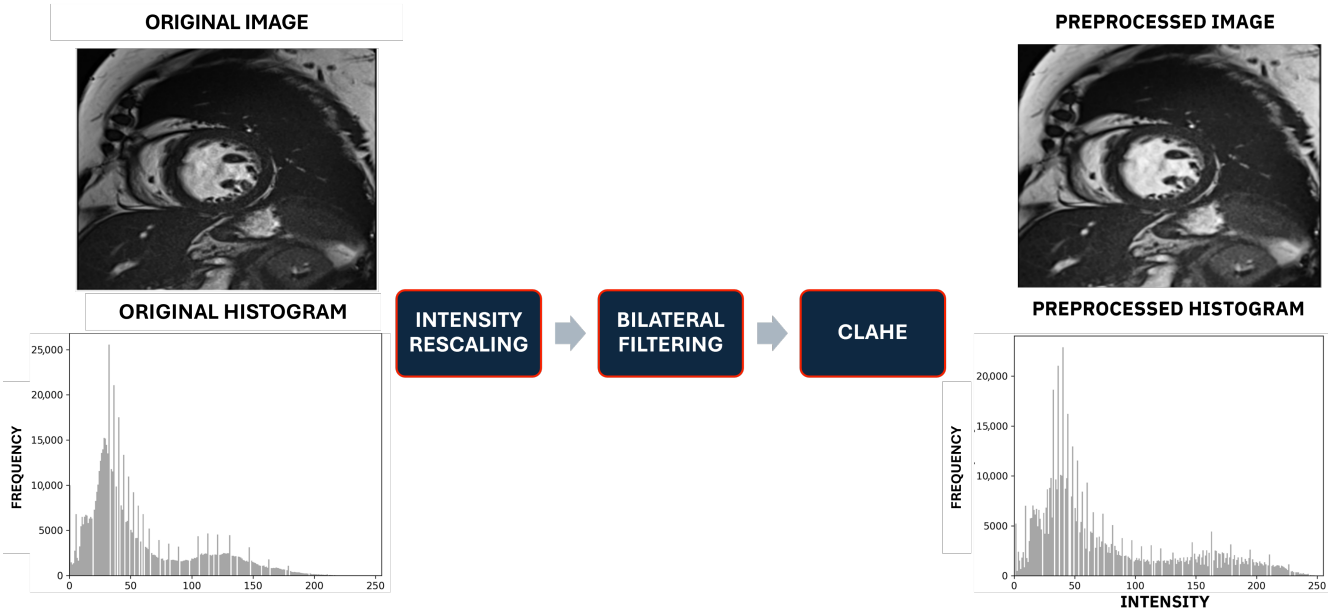


Figure 2. Overview of the visual enhancement pipeline. From left to right: original cardiac MRI slice and intensity histogram, followed by intensity rescaling, bilateral filtering, and CLAHE. The final enhanced image shows improved contrast and anatomical delineation, as reflected by the redistributed intensity histogram.

To complement this qualitative assessment, quantitative image quality metrics were computed across the full dataset (3459 slices). Shannon entropy [23], Response Surface Methodology (RSM) Contrast, and mean Sobel gradient [19] magnitude were measured before and after preprocessing. The aggregated results are summarized in Table 1.

The consistent increase in entropy reflects improved intensity dispersion, the rise in RSM contrast indicates enhanced global contrast, and the higher Sobel magnitude confirms stronger edge definition. Together, these quantitative results support the effectiveness of the proposed enhancement pipeline without evidence of excessive artificial amplification.

Table 1. Average image quality metrics computed across all 3459 slices before and after visual enhancement.

Metric	Before	After	Δ
Shannon Entropy	5.41	5.78	+0.37
RSM Contrast	39.13	49.63	+10.50
Sobel Gradient Magnitude	17.56	19.27	+1.71

The enhancement pipeline was implemented in Python (version 3.10, Python Software Foundation, Wilmington, DE, USA) using `OpenCV` (version 4.8, OpenCV Foundation, Mountain View, CA, USA) and `scikit-image` (version 0.21, scikit-image developers, USA). Parameter values were determined through an interactive tuning procedure on a representative subset of 20 diastolic-phase slices from different patients. A dedicated Python interface with adjustable sliders was developed to explore parameter configurations. Two experienced cardiologists evaluated myocardial boundary clarity and trabecular definition while varying the parameters. The final configuration corresponds to the most consistently selected settings and was fixed for all experiments to ensure reproducibility.

All slices are resampled to a fixed spatial resolution of 800×800 pixels prior to preprocessing and learning. Intensity images are resized using bicubic interpolation (order $= 3$). Segmentation masks are resized using label-preserving interpolation (nearest-neighbor, order $= 0$) when required, preventing the creation of spurious intermediate labels. The choice of 800×800 was made as a practical compromise between preserving fine trabecular detail and computational cost, and to provide a sufficiently dense spatial representation for the transformer-based components of ViTUNeT (i.e., a higher number of spatial tokens/patches), which improves the ability to model small anatomical structures. Although resampling can alter physical scale if pixel spacing metadata are ignored, our evaluation is based on overlap metrics (Dice) computed consistently in the same resampled space for both predictions and ground-truth masks; thus, comparisons across pipeline variants remain valid.

2.3. Supervised Slice-Quality Labels from Segmentation Behavior

Even after the enhancement stage, some slices remain unsuitable for reliable segmentation (e.g., due to truncated anatomy, signal dropouts, or severe motion artifacts). Since such failures should ideally be detected and discarded before applying the segmenter at inference time, we formulated a supervised machine learning approach aimed at filtering defective slices prior to segmentation. The first step, therefore, consisted of generating supervised slice-quality labels using a quantitative proxy derived from segmentation performance.

To ensure methodological rigor, the segmentation models used to construct these labels were trained under a strict patient-level separation protocol. First, an 80/20 train–test split was performed independently within each diagnostic group, guaranteeing that all slices from a given patient were assigned exclusively to one of the two subsets. Subsequently, within the training set, a 5-fold stratified cross-validation scheme was applied, also at the patient level. Stratification was based on the clinically established trabeculated volume (TV%) threshold of 27.4% [6], which discriminates between healthy subjects and LVNC phenotypes, thereby ensuring balanced pathological representation across folds.

For each fold, an independent segmentation model was trained from scratch. Final predictions were obtained through an ensemble strategy that averages, at inference time, the outputs of the five fold-specific models (Figure 3). This scheme reduces variance and improves robustness, providing more stable estimates of segmentation performance.

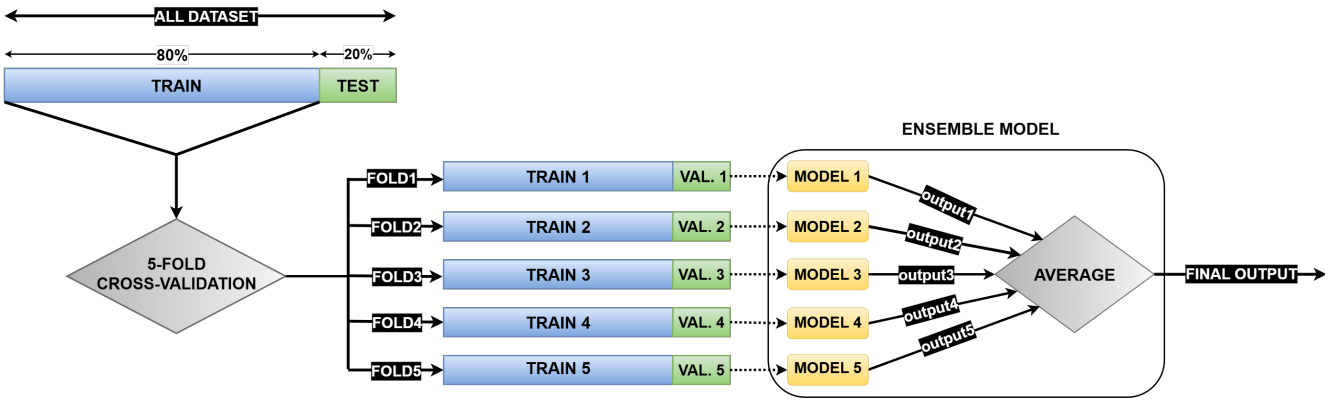


Figure 3. Scheme of the ensemble method.

Let D_{DL} and D_{ViT} denote the ordered sequences (from lowest to highest) containing the trabecular Dice coefficient values obtained for each image in the dataset using the DL-LVTQ and ViTUNeT segmentation models, respectively, following the ensemble approach described above. Slices were considered low quality if their performance fell below the first quartile of the distribution in either model; that is, images belonging to the lowest 25% according to D_{DL} or D_{ViT}:

$$\mathcal{L} = \{x : D_{DL}(x) < Q_1(D_{DL})\} \cup \{x : D_{ViT}(x) < Q_1(D_{ViT})\}. \quad (1)$$

For the remaining slices, we need to define a new criterion to determine medium and high quality slices. To do this, we performed an inflection-point analysis of the cumulative Dice distribution in the trabecular region.

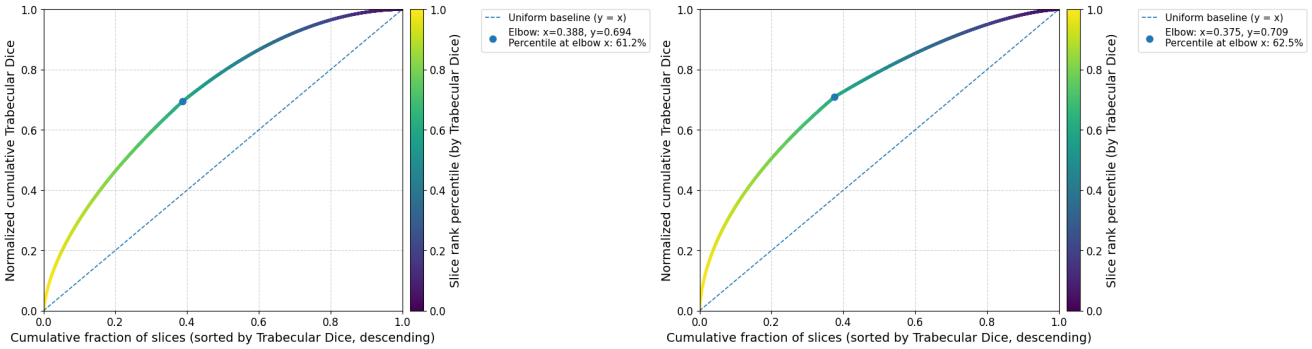
Specifically, the images that were not labeled as low quality were sorted in descending order according to their trabecular Dice score D_T. Let N denote the number of such slices and d_i the sorted Dice values ($d_1 \geq d_2 \geq \cdots \geq d_N$). We then computed the cumulative Dice curve

$$C(k) = \sum_{i=1}^{k} d_i, \quad k = 1, \ldots, N,$$

and normalized both axes as

$$x_k = \frac{k}{N}, \qquad y_k = \frac{C(k)}{C(N)}.$$

The resulting curve (Figure 4) represents the cumulative contribution of the top-performing slices relative to a uniform distribution.



(**a**) (**b**)

Figure 4. Inflection Point Analysis of the Trabecular Dice Cumulative Curve. (**a**) ViTUNeT. (**b**) DL_LVTQ.

To identify the structural transition point of this distribution, we applied a maximum-distance-to-diagonal criterion, computing $k^* = \arg\max_k(y_k - x_k)$, which corresponds to the classical knee-point detection in concave cumulative curves. This inflection point marks the transition from a steep accumulation region to a slower growth regime.

For ViTUNeT, Figure 4a shows that the elbow (inflection) point is located at $x \approx 0.388$, whereas for DL-LVTQ, Figure 4b indicates an elbow at $x \approx 0.375$.

In addition to the x-coordinate itself, we analysed the percentile position of the slice that defines the elbow. For ViTUNeT, the elbow corresponds to the 61.2th percentile of the ordered quality distribution, while for DL-LVTQ it corresponds to the 62.5th percentile.

The mean percentile value is approximately 0.618, which coincides remarkably with the golden ratio, $\phi = 0.618$. The golden ratio possesses well-established mathematical properties and naturally arises as the limiting ratio of consecutive terms in the Fibonacci sequence. Its appearance in natural growth processes and self-organizing systems has been extensively documented in mathematical and physical literature.

Given this empirical convergence and its theoretical significance, we adopted $\phi = 0.618$ as a stable and model-consistent threshold to distinguish between medium- and high-quality slices.

Building upon this principle, we further refined the stratification of the remaining slices by introducing a second cutoff defined through a golden-ratio quantile level $\phi = 0.618$ applied to the upper portion of the distribution, namely the interval between the first quartile and the maximum value. Since this interval contains 75% of the probability mass, the corresponding global quantile is

$$p_\phi = 0.25 + \phi \cdot 0.75 = 0.7135. \tag{2}$$

Accordingly, we define

$$Q_\phi^{DL} = Q_{0.7135}(D_{DL}), \qquad Q_\phi^{ViT} = Q_{0.7135}(D_{ViT}), \tag{3}$$

where $Q_p(\cdot)$ denotes the p-th empirical quantile and the maximum value is implicitly included as the upper bound of the distribution.

We then define two intermediate-quality subsets, one per model:

$$\mathcal{F} = \{x : Q_1(D_{DL}) \le D_{DL}(x) < Q_\phi^{DL}\}, \qquad \mathcal{G} = \{x : Q_1(D_{ViT}) \le D_{ViT}(x) < Q_\phi^{ViT}\}. \tag{4}$$

Slices labeled as *medium* quality are defined as the intersection of these subsets:

$$\mathcal{M} = \mathcal{F} \cap \mathcal{G}. \tag{5}$$

The intersection criterion was adopted to resolve disagreement between models in a principled manner. Specifically, if a slice does not belong to the intersection of the medium-quality sets defined by DL-LVTQ and ViTUNeT, it means that at least one of the models assigns it a trabecular Dice value consistent with medium-to-high quality. In such cases, we deliberately promote the slice to the high quality category rather than keeping it in the medium group.

Slices exceeding the upper threshold in both models, i.e.,

$$\mathcal{H} = \{x : D_{DL}(x) \ge Q_\phi^{DL}\} \cup \{x : D_{ViT}(x) \ge Q_\phi^{ViT}\}, \tag{6}$$

are labeled as high quality. This dual-model criterion reduces sensitivity to model-specific artifacts and promotes robust quality stratification.

Finally, we introduce an anatomical consistency safeguard using the Dice coefficient of the compact external layer, D_{EL}. We define the threshold

$$T_{EL} = \frac{Q_1(D_{DL}) + Q_1(D_{ViT})}{2}, \tag{7}$$

and downgrade any slice initially labeled as *medium* or *high* if $D_{EL} < T_{EL}$. Figure 5 summarizes the resulting labeling for supervised learning.

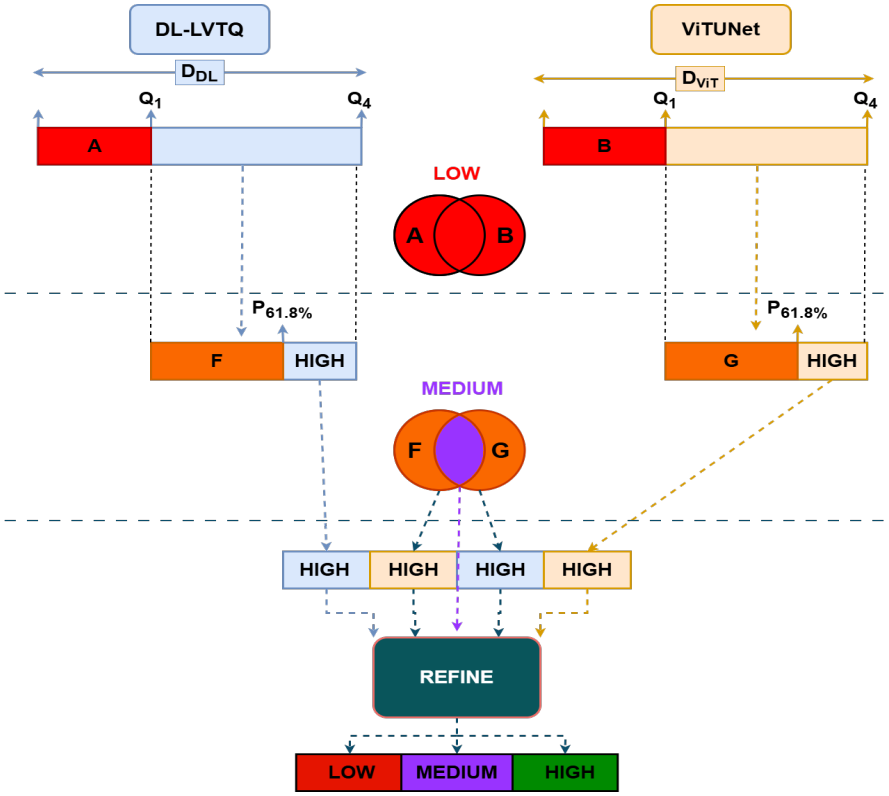


Figure 5. Supervised slice-quality labeling procedure based on the joint analysis of Dice coefficients obtained from DL-LVTQ and ViTUNeT.

2.4. YOLOv11-Based Anatomical Attribute Extraction

Segmentation quality metrics such as the trabecular Dice coefficient (D_T) provide an objective measure of performance, but they depend on ground-truth annotations and are, therefore, unavailable prior to inference. To enable quality-aware filtering before segmentation, we instead rely on anatomical cues that are directly observable in the image and can act as proxies for structural visibility and completeness.

To extract such cues, we train a YOLOv11s detector [24] to localize two anatomical entities in each slice: the full left ventricle (LV) and its internal cavity (IC). The detector is not used for classification or diagnosis; its sole purpose is to provide a compact set of geometrical and confidence-based attributes that characterize anatomical clarity.

The YOLOv11s model corresponds to the standard implementation provided by the Ultralytics framework, initialized from the official pretrained weights. No architectural modifications were introduced; the default backbone, neck, and detection head configurations were preserved.

2.4.1. Training Data and Annotations

YOLOv11s is trained exclusively on slices labeled as high quality according to the procedure described in Section 2.3. All images are first enhanced using the preprocessing pipeline in Section 2.2 and resized to 800 × 800 pixels.

Bounding-box annotations are generated automatically from the available segmentation masks using `OpenCV`. For each slice, the largest external contour defines the LV bounding box, while the largest enclosed region corresponds to the IC (Figure 6). The resulting coordinates are converted to the standard YOLO format `[class x y width height]`.

The dataset is divided at the patient level in an 80/20% ratio for train and test to avoid data leakage, using a fixed random seed to ensure reproducibility.

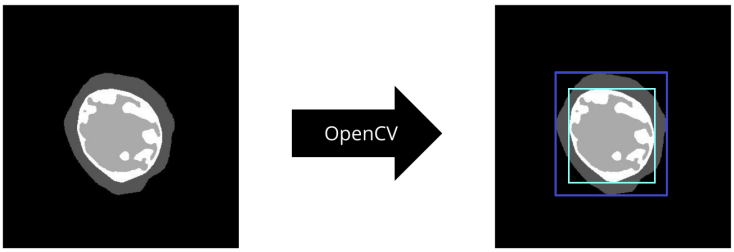


Figure 6. Bounding box annotations generated from segmentation masks using `OpenCV`. Dark blue: left ventricle. Light blue: internal cavity.

2.4.2. Model Training Qualitative Assessment

The YOLOv11s model is trained for 75 epochs using the `ultralytics` framework with a batch size of 4 and the RAdam optimizer.

Training is initialized from pretrained weights in detection mode, with input size 800×800 and mixed precision enabled (AMP). Deterministic training is enforced with fixed seed. The validation split is used during training, and training is configured with patience = 100. The optimization hyperparameters are: initial learning rate $\mathrm{lr}_0 = 0.01$, final learning-rate factor $\mathrm{lrf} = 0.01$, momentum $= 0.937$, weight_decay $= 5 \times 10^{-4}$, and warmup_epochs $= 3.0$ (warmup_momentum $= 0.8$, warmup bias LR $= 0.1$). The Intersection over Union (IoU) threshold for Non-Maximum Suppression (NMS) during evaluation is set to 0.7, with a maximum of 300 detections per image. Data augmentation follows the Ultralytics configuration: HSV jitter (h $= 0.015$, s $= 0.7$, v $= 0.4$), translation $= 0.1$, scale $= 0.5$, horizontal_flip_probability $= 0.5$, mosaic augmentation enabled (mosaic $= 1.0$) with mosaic disabled in the last 10 epochs, and random erasing_probability $= 0.4$. No vertical flips are applied, and mixup/copy-paste are disabled.

The training dynamics of the YOLOv11s model over the 75 epochs are shown in Figure 7.

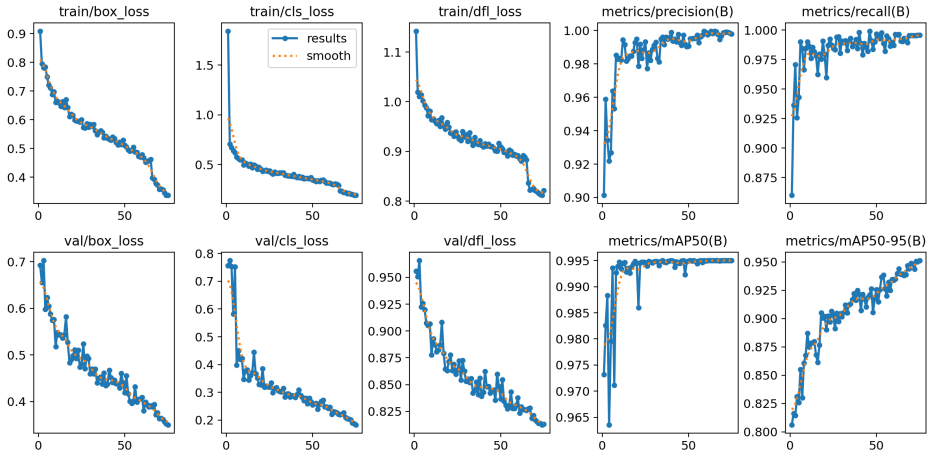


Figure 7. YOLOv11s training dynamics over 75 epochs, including loss components and detection metrics on training and validation sets.

Visual inspection of validation slices confirms that the trained model consistently localizes both LV and IC structures across a wide range of anatomical appearances and acquisition conditions (Figure 8). These results support the use of YOLOv11s as a reliable anatomical attribute extractor rather than as a diagnostic model.

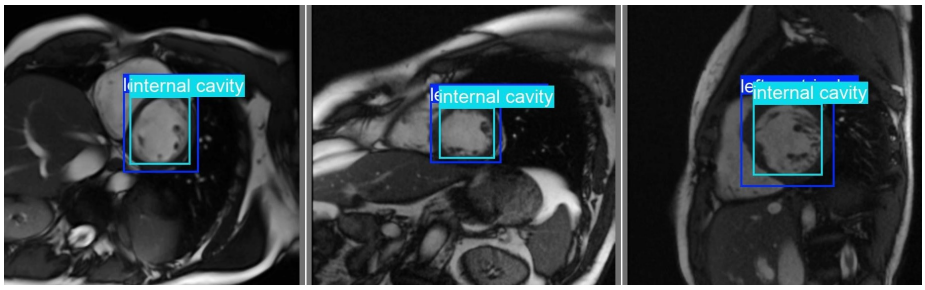


Figure 8. Representative YOLOv11s detections on enhanced validation slices. Dark blue boxes indicate the left ventricle; light blue boxes indicate the internal cavity.

2.4.3. Attribute Extraction

Once trained, the detector is applied to all MRI slices, independently of their assigned quality label. For each slice, four scalar attributes are extracted:

- Detection confidence for the left ventricle;
- Detection confidence for the internal cavity;
- Bounding-box area of the left ventricle;
- Bounding-box area of the internal cavity.

These attributes provide a compact and interpretable description of anatomical visibility and spatial extent. In the subsequent stage of the pipeline, they are used as input features for a supervised classifier that estimates slice quality prior to segmentation.

Although YOLOv11s is trained exclusively on high-quality slices to learn anatomically consistent localization patterns, it is intentionally applied to slices of all quality levels. When evaluated on medium- and low-quality slices, the detector typically produces lower confidence scores, unstable bounding boxes, or in extreme cases, missed detections. Rather than representing a failure of the pipeline, such degradation reflects reduced anatomical visibility and, therefore, constitutes informative signal for the slice-quality classifier. In this design, the quality-sensitive behavior of YOLOv11 becomes part of the discrimination mechanism itself, allowing unreliable slices to be identified without requiring additional adaptation.

2.5. Attribute–Quality Relationship

To verify that the extracted attributes are informative for anticipating segmentation quality, we quantify their relationship with the trabecular Dice coefficient D_T and include `slice_number` as a proxy for axial position within the volume.

Since the number of slices varies across patients, `slice_number` is defined as the normalized axial position of each slice within the patient-specific stack. Concretely, if k denotes the index of a slice in its ordered volume and N the total number of slices for that patient, we compute

$$\texttt{slice_number} = \frac{k}{N}.$$

This normalization ensures comparability across patients and prevents bias introduced by variable stack lengths.

Figure 9 reports the correlation matrix between D_T, YOLO confidences/areas for LV and IC, and `slice_number`. We observe moderate positive correlations between D_T and detection confidence (LV: $r = 0.57$, IC: $r = 0.58$), and weaker positive correlations

with bounding-box areas (LV: $r = 0.30$, IC: $r = 0.35$), supporting their use as predictors of slice quality. The negative association between box areas and slice number (approximately $r \approx -0.27$) is consistent with anatomy, as apical slices typically contain smaller ventricular structures.

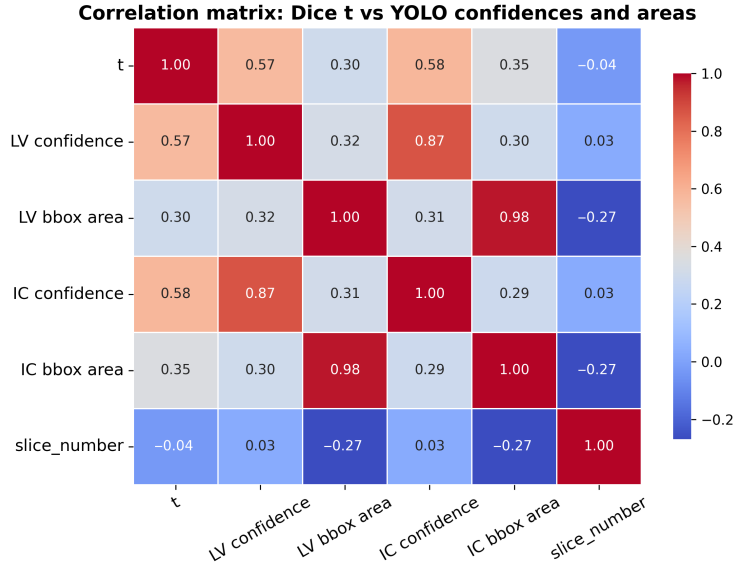


Figure 9. Correlation matrix between trabecular Dice (D_T), YOLO-derived anatomical attributes (confidence and bounding-box area for LV and IC), and slice number.

It is important to note that global image-quality descriptors such as entropy or Signal-to-Noise Ratio (SNR) primarily quantify photometric dispersion and overall contrast, but do not encode anatomical completeness or structural coherence. In contrast, the YOLO-derived attributes reflect geometrical extent and detection confidence of clinically meaningful structures. The observed correlations with D_T, therefore, suggest that the proposed features capture structurally relevant information beyond global intensity statistics.

2.6. Slice-Quality Classification and Model Training

Slice quality is estimated through supervised learning on tabular anatomical attributes extracted from YOLOv11 detections. Each MRI slice is represented by five predictors: detection confidence and bounding-box area for the left ventricle (LV) and its internal cavity (IC), together with the `slice_number` encoding the axial position within the volume. The target variable is the three-class quality label (low, medium, high) derived from segmentation behavior as described in Section 2.3.

The resulting dataset consists of 3459 samples with no missing values. The distribution of quality classes is shown in Figure 10, indicating a balanced class composition suitable for supervised learning.

To ensure fair comparison and reproducibility across learning paradigms, all candidate classifiers are trained and evaluated using a unified experimental protocol. The dataset is first split into stratified training (80%) and test (20%) subsets to preserve class proportions. Hyperparameter optimization is then performed exclusively on the training set using exhaustive grid search [25] combined with 5-fold stratified cross-validation. The macro-averaged F1-score is used as the primary selection criterion, as it balances performance across all classes irrespective of their frequency.

Prior to model training, input features were standardized using a zero-mean, unit-variance transformation. Since margin-based and neural models are sensitive to feature scale, this normalization ensures comparable feature magnitudes and stable optimization.

To prevent information leakage, the standardization parameters (mean and standard deviation) were computed exclusively on the training data within each cross-validation fold and subsequently applied to the corresponding validation subset. The same protocol was followed for the final training–test evaluation.

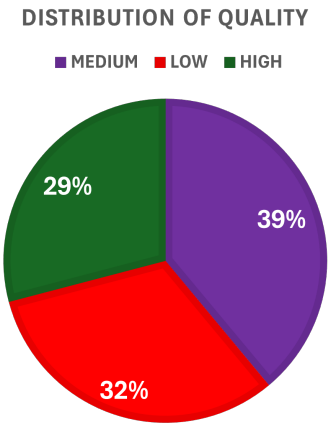


Figure 10. Class distribution for supervised slice-quality classification.

The complete training and selection procedure is summarized in Algorithm 1. Once the optimal configuration is identified for each model, the classifier is retrained on the full training set and evaluated on the held-out test set using F1-score, precision, and recall. The best-performing model is selected as the final slice-quality classifier and integrated into the preprocessing pipeline.

Algorithm 1 Training and Selection of the Slice-Quality Classifier

Require: Labeled dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$, candidate models $\{\mathcal{M}_k\}$, hyperparameter grids $\{\Theta_k\}$
Ensure: Trained slice-quality classifier $\mathcal{C}_{\text{ML}}$
Split $\mathcal{D}$ into stratified training set $\mathcal{D}_{\text{train}}$ (80%) and test set $\mathcal{D}_{\text{test}}$ (20%)
for each model $\mathcal{M}_k \in \{\text{RF, HGB, SVC, MLP}\}$ **do**
  Initialize best score $s_k \leftarrow 0$
  for each hyperparameter configuration $\theta \in \Theta_k$ **do**
    Perform 5-fold stratified cross-validation on $\mathcal{D}_{\text{train}}$
    Compute macro-F1 score f_θ
    if $f_\theta > s_k$ **then**
      $s_k \leftarrow f_\theta$
      $\theta_k^* \leftarrow \theta$
    end if
  end for
  Train $\mathcal{M}_k$ on $\mathcal{D}_{\text{train}}$ using θ_k^*
  Evaluate $\mathcal{M}_k$ on $\mathcal{D}_{\text{test}}$
end for
Select $\mathcal{M}_{k^*}$ with the highest test macro-F1 score
Train final classifier $\mathcal{C}_{\text{ML}} \leftarrow \mathcal{M}_{k^*}$ on $\mathcal{D}_{\text{train}}$
return $\mathcal{C}_{\text{ML}}$

The candidate learning algorithms evaluated in this study are:

- Random Forest (RF): an ensemble of decision trees trained via bagging.
- Histogram Gradient Boosting (HGB): a boosting-based method optimized for tabular data using histogram-based splits.
- Support Vector Classifier (SVC): a margin-based classifier maximizing inter-class separation.

- Multilayer Perceptron (MLP): a fully connected neural network modeling nonlinear interactions among features.

2.7. Quality-Aware Preprocessing Algorithm

Based on the components described above, we formalize the proposed method as a quality-aware preprocessing algorithm that operates at the slice level prior to neural-network-based segmentation. Algorithm 2 integrates deterministic image enhancement with two learned components: a YOLOv11-based anatomical attribute extractor and a supervised machine learning classifier for slice-quality estimation. Given a raw cardiac MRI slice, the algorithm outputs either a preprocessed slice suitable for segmentation or a rejection decision.

Algorithm 2 Quality-Aware Preprocessing Algorithm for Cardiac MRI Slices

Require: Raw cardiac MRI slice i, trained YOLOv11 detector $\mathcal{Y}_{\text{YOLO}}$, trained slice-quality classifier $\mathcal{C}_{\text{ML}}$
Ensure: Preprocessed slice $\tilde{i}$ or rejection decision
Input: Raw slice i
Output: Enhanced slice $\tilde{I}$ or `reject`
Step 1: Visual enhancement (deterministic preprocessing)
$I_1 \leftarrow \text{intensity rescaling}(i)$
$I_2 \leftarrow \text{bilateral filtering}(i_1)$
$\tilde{i} \leftarrow \text{CLAHE}(i_2)$
Step 2: Anatomical attribute extraction (YOLOv11)
$(b_{LV}, c_{LV}), (b_{IC}, c_{IC}) \leftarrow \mathcal{Y}_{\text{YOLO}}(\tilde{i})$
$a_{LV} \leftarrow \text{area}(b_{LV})$
$a_{IC} \leftarrow \text{area}(b_{IC})$
Step 3: Slice-quality estimation (supervised ML)
$\mathbf{x} \leftarrow [c_{LV}, c_{IC}, a_{LV}, a_{IC}, \texttt{slice_number_i}]$
$q \leftarrow \mathcal{C}_{\text{ML}}(\mathbf{x})$
Step 4: Quality-aware decision
if $q = \texttt{low}$ **then**
  return `reject`
else
  return $\tilde{i}$
end if

The computational overhead of the proposed preprocessing pipeline is negligible compared to the segmentation inference time.

2.8. Reproducibility and Code Availability

All experiments (preprocessing, YOLO training, feature extraction, classifier training, and evaluation) are implemented in reproducible Jupyter notebooks. The code and supporting resources are publicly available at https://github.com/salvadeharo10/LVNC_MRI_classification.git (accessed on 15 December 2025). The MRI dataset is subject to hospital data-sharing restrictions; therefore, it cannot be fully released, but we provide the complete preprocessing and training pipeline, along with detailed instructions to reproduce experiments on compatible datasets.

3. Results

This section presents the experimental results obtained with the proposed quality-aware preprocessing algorithm. The analysis focuses on two aspects: (i) the performance of supervised classifiers for slice-quality estimation based on anatomical attributes, and

(ii) the quantitative impact of quality-driven slice filtering on downstream segmentation performance.

For clarity, we define the following dataset variants:

- DS^0: original dataset containing all MRI slices.
- DS^0_{VIA}: dataset after applying the visual enhancement pipeline.
- DS^1_{VIA}: refined dataset obtained after removing slices predicted as low quality by Algorithm 2.

The visually enhanced dataset DS^0_{VIA} shows consistent but moderate improvements over DS^0. The remainder of this section focuses on the results obtained after quality-aware filtering, which leads to DS^1_{VIA}.

3.1. Slice-Quality Classification Performance

Four supervised classifiers are trained to predict slice quality using YOLOv11-derived anatomical attributes. Each model uses five input features: detection confidence and bounding-box area for the left ventricle and its internal cavity, together with the slice position within the cardiac volume.

Table 2 reports macro-averaged precision, recall, and F1-score on the held-out test set.

Table 2. Slice-quality classification performance on the test set using YOLO-derived anatomical attributes. Metrics are macro-averaged across the three quality classes.

Model	Precision	Recall	F1-Score
Random Forest (RF)	0.8482	0.8534	0.8504
Histogram Gradient Boosting (HGB)	0.8441	0.8459	0.8448
Support Vector Classifier (SVC)	0.8809	0.8860	0.8826
Multilayer Perceptron (MLP)	**0.9061**	**0.9113**	**0.9080**

The Multilayer Perceptron (MLP) achieves the highest performance across all metrics. Based on these results, the MLP classifier is selected and applied to the full dataset, yielding the filtered subset DS^1_{VIA}.

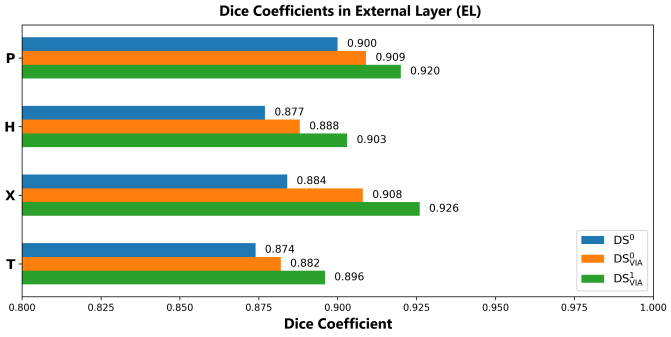
3.2. Impact of Quality-Aware Filtering on Segmentation Performance

Segmentation performance is evaluated using ViTUNeT on DS^0 and DS^1_{VIA}. The analysis focuses on the compact external layer (EL) and the trabecular region (T), which represent the most challenging anatomical structures due to diffuse boundaries and low contrast.

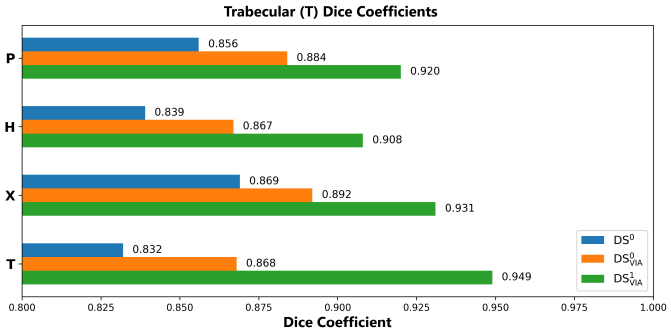
Figure 11 illustrates the effect of quality-aware slice filtering on segmentation accuracy. Improvements in the compact external layer are consistent across diagnostic groups but remain moderate, reflecting the relatively well-defined morphology of this structure. In contrast, larger gains are observed in the trabecular region, where anatomical variability and boundary ambiguity are more pronounced.

Overall, the filtering stage removes 1106 slices (approximately 32% of the dataset), distributed across diagnostic groups as follows: 709 (P), 193 (X), 172 (H), and 32 (T).

To formally assess statistical significance, paired comparisons were conducted between DS^0 and DS^1_{VIA} on the subset of retained slices ($n = 2353$), ensuring matched evaluation of identical anatomical slices. The unit of analysis was the slice. Since multiple slices originate from the same patient, observations are clustered and, therefore, not strictly independent. The reported p-values should be interpreted within this slice-level framework. Results are summarized in Table 3. Both anatomical regions exhibited statistically significant improvements with moderate effect sizes.



(**a**)



(**b**)

Figure 11. Dice coefficients across diagnostic groups before and after quality-aware filtering. (**a**) Compact external layer (EL); (**b**) Trabecular region (T) .

Table 3. Paired statistical comparison between DS^0 and DS^1_{VIA} on retained slices ($n = 2353$). ΔMean denotes the average Dice improvement. RBC: rank-biserial correlation.

Region	Mean DS^0	Mean DS^1_{VIA}	ΔMean	95% CI	p-Value	RBC
EL	0.9178	0.9259	+0.00811	[0.00739, 0.00882]	$<10^{-130}$	0.595
T	0.9166	0.9252	+0.00859	[0.00762, 0.00954]	$<10^{-70}$	0.440

Although the absolute Dice increase is numerically modest (approximately +0.008), the narrow confidence intervals indicate low variability in the gain, suggesting systematic rather than sporadic improvement across slices. The rank-biserial correlations (0.595 for EL and 0.440 for T) correspond to moderate effect sizes, confirming that a substantial proportion of slices benefit from the filtering strategy. Notably, these gains are achieved without modifying the segmentation architecture or increasing computational complexity, supporting the effectiveness of the proposed data-centric refinement approach within the evaluated experimental setting. Future analyses may incorporate hierarchical or mixed-effects modeling to explicitly account for intra-patient correlation.

3.3. Computational Environment and Runtime Analysis

All experiments were conducted in a high-performance computing environment equipped with an NVIDIA RTX 4090 GPU (24 GB), dual AMD EPYC processors (64 total threads), and 377 GB RAM. The software stack consisted of a 64-bit Linux distribution, PyTorch, and the Ultralytics YOLO framework executed within a Singularity container to ensure full reproducibility.

To quantify computational overhead, we evaluated the complete pipeline (visual enhancement preprocessing + YOLO inference + slice-quality classification + ViTUNeT

segmentation) on the full dataset comprising 3459 images. Two representative execution modes were analysed: CPU-only (batch size = 1) and GPU execution with batch size = 8.

The average runtime per image for the complete pipeline under CPU and GPU execution modes is summarized in Table 4.

Table 4. Average runtime per image for the complete pipeline.

Execution Mode	Mean Time (s/Image)	Throughput (img/s)
CPU (bs = 1)	0.1408	6.99
GPU (bs = 8)	0.0166	53.87

A detailed breakdown of the GPU configuration (batch size = 8) is provided in Table 5. Visual enhancement preprocessing accounts for the largest proportion of execution time, followed by ViTUNeT segmentation inference, while YOLO detection and slice-quality classification introduce comparatively minor overhead.

Table 5. Per-image runtime breakdown under GPU execution (batch size = 8).

Pipeline Component	Mean Time (ms/Image)
Visual enhancement preprocessing (CPU)	13–20
YOLO inference	$\sim$3
Slice-quality classification	$<$1
ViTUNeT segmentation inference	$\sim$10–15

Overall, GPU acceleration provides an approximately 8.5$\times$ speedup relative to CPU execution. Under GPU mode, the computational bottleneck shifts from neural inference to preprocessing, in agreement with Amdahl's Law. Peak GPU memory usage remained below 0.7 GB, leaving substantial capacity for larger batch sizes or concurrent workloads.

4. Discussion

The results demonstrate that slice-quality estimation based on anatomically grounded attributes constitutes an effective strategy for improving segmentation robustness in cardiac MRI. While visual enhancement alone provides limited gains, the quality-aware filtering stage produces substantial improvements, particularly in anatomically complex regions such as the trabecular myocardium.

The superior performance of the MLP classifier indicates that nonlinear interactions between anatomical confidence, spatial extent, and slice position are key for distinguishing reliable from unreliable slices. This supports the use of lightweight tabular learning models for quality assessment, avoiding the need for additional deep segmentation networks.

The largest segmentation improvements are observed in the trabecular region, where anatomical ambiguity and inter-slice variability are highest. This suggests that segmentation errors are driven primarily by input quality rather than model capacity. By selectively filtering structurally incomplete or noisy slices, the pipeline reduces error propagation and stabilizes downstream predictions.

An important methodological aspect is robustness to dataset imbalance across diagnostic groups. Although the cohort is dominated by hypertrophic cardiomyopathy cases, performance was analysed independently for each population rather than through pooled metrics. The consistent improvements observed across all groups indicate that the proposed filtering strategy is not driven by the dominant class and supports the statistical stability of the findings.

The filtered dataset (DS^1_{VIA}) should be interpreted as a quality-aware triage mechanism rather than a permanent exclusion of information. In practical deployment, slices identified as low quality would be flagged for manual review, excluded from automated quantification, or reacquired when feasible. The framework is, therefore, designed to prevent unreliable inputs from degrading automated segmentation while preserving clinical oversight.

We acknowledge that slice removal may raise concerns regarding artificial inflation of segmentation metrics. To mitigate this risk, improvements were additionally evaluated through paired comparisons on identical retained slices (Section 3.2), ensuring that reported gains reflect increased segmentation stability rather than post-hoc exclusion of difficult cases.

Importantly, these gains are achieved without modifying the segmentation architecture or increasing computational complexity. The approach, therefore, shifts the emphasis from model-centric optimization to data-centric algorithmic curation, which is particularly relevant for heterogeneous clinical datasets.

Although the present study focuses on cardiac MRI and the ViTUNeT architecture, the proposed preprocessing and quality-aware filtering framework is formulated independently of architecture-specific modifications. However, broader empirical validation across additional datasets and segmentation backbones is required to fully establish generalizability beyond the evaluated experimental setting.

From a deployment perspective, the computational overhead remains compatible with routine clinical workflows. Under GPU execution (batch size = 8), the complete process—including enhancement, YOLO inference, slice-quality classification, and segmentation—achieves approximately 54 slices per second. As a typical short-axis cardiac MRI study contains 10–20 slices, full quality assessment and segmentation can be completed within a few seconds, supporting real-time quality control without meaningful workflow disruption.

5. Conclusions

This work demonstrates that data quality constitutes a primary limiting factor in deep learning-based segmentation of cardiac MRI, even when using advanced architectures such as ViTUNeT. The results show that suboptimal performance does not necessarily originate from insufficient model capacity, but rather from heterogeneous and unreliable input data that hinder effective learning and inference.

To address this limitation, we proposed a quality-aware, data-centric preprocessing algorithm that operates prior to segmentation. The method combines a lightweight visual enhancement stage with an automatic slice-quality filtering mechanism. Anatomically meaningful attributes are extracted using a YOLOv11-based detector and subsequently exploited by a supervised machine learning classifier to identify and discard unreliable slices before segmentation. This design decouples data curation from the segmentation model itself and allows quality control to be treated as an independent algorithmic component.

Experimental results confirm that the proposed approach leads to statistically significant and consistent improvements in segmentation accuracy, particularly in anatomically complex regions such as the trabecular myocardium. Importantly, these gains are achieved without modifying the segmentation architecture or increasing its computational complexity, highlighting the effectiveness of algorithmic data curation over model-centric optimization strategies.

Although the present validation focuses on short-axis cardiac MRI and LVNC-related segmentation tasks using the ViTUNeT architecture, the proposed preprocessing and quality-aware filtering framework is designed independently of architecture-specific modifications. However, empirical validation has so far been limited to a single in-house dataset

and one segmentation backbone. The anatomical attributes and quality labels employed in this study are specific to left ventricular structures, and direct transfer to other anatomies or imaging modalities would require retraining the detection and classification components. Therefore, broader evaluation across additional datasets and alternative segmentation architectures is necessary to fully establish generalizability beyond the current experimental setting. Exploring such extensions constitutes an important direction for future research.

Beyond the specific application to cardiac MRI, formalizing quality assessment as an explicit algorithmic step reinforces the importance of data-centric methodologies for building robust, scalable, and clinically viable deep learning systems.

Author Contributions: Conceptualization, S.d.H. and G.B.; methodology, S.d.H.; software, S.d.H.; validation, S.d.H., J.C., G.B. and P.G.-F.; formal analysis, S.d.H.; investigation, S.d.H. and G.B.; resources, G.B. and J.M.G.; data curation, S.d.H.; writing—original draft preparation, S.d.H.; writing—review and editing, S.d.H., P.G.-F., J.C., J.M.G. and G.B.; visualization, S.d.H.; supervision, P.G.-F., J.M.G. and G.B.; project administration, J.M.G. and G.B.; funding acquisition, J.M.G. and G.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research has been partially funded by Grant TED2021-129221B-I00, which is funded by MCIN/AEI/10.13039/501100011033, and by the European Union NextGenerationEU/PRTR.

Institutional Review Board Statement: This study was conducted using anonymized retrospective cardiac MRI data. Ethical review and approval were waived in accordance with institutional guidelines and applicable regulations, as no identifiable patient information was used.

Informed Consent Statement: Patient consent was waived due to the retrospective nature of the study and the use of fully anonymized data.

Data Availability Statement: The preprocessing pipeline, quality-labeling strategy, and model configuration details are fully described in the manuscript. Source code for the visual enhancement and inference stages can be made available upon reasonable request to the corresponding author. The clinical MRI dataset was collected under institutional approval and contains anonymized patient data subject to ethical and data-protection regulations; therefore, it cannot be publicly released.

Acknowledgments: The authors thank the collaborating clinical centers for providing access to the cardiac MRI data used in this study. The authors acknowledge the use of ChatGPT (GPT-4, OpenAI, San Francisco, CA, USA) for language refinement during manuscript preparation. All scientific content, analysis, and conclusions remain the responsibility of the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

CLAHE	Contrast-Limited Adaptive Histogram Equalization
CNN	Convolutional Neural Network
DL	Deep Learning
EL	External Compacted Myocardial Layer
HCM	Hypertrophic Cardiomyopathy
HGB	Histogram Gradient Boosting
IC	Inner Cavity
IoU	Intersection over Union
LVNC	Left Ventricular Noncompaction
MLP	Multilayer Perceptron
MRI	Magnetic Resonance Imaging
NMS	Non-Maximum Suppression
RF	Random Forest
SNR	Signal-to-Noise Ratio

SVC	Support Vector Classifier
TV	Trabecular Volume
TZ	Trabeculated Zone

References

1. van Waning, J.I.; Caliskan, K.; Hoedemaekers, Y.M.; Van Spaendonck-Zwarts, K.Y.; Baas, A.F.; Boekholdt, S.M.; Van Melle, J.P.; Teske, A.J.; Asselbergs, F.W.; Backx, P.C.M.; et al. Genetic, Clinical and Imaging Characteristics in Noncompaction Cardiomyopathy. *J. Am. Coll. Cardiol.* **2018**, *71*, 711–722. [CrossRef] [PubMed]
2. Andreini, D.; Pontone, G.; Bogaert, J.; Roghi, A.; Barison, A.; Schwitter, J.; Mushtaq, S.; Vovas, G.; Sormani, P.; Aquaro, G.D.; et al. Long-Term Prognostic Value of Cardiac Magnetic Resonance in Left Ventricular Noncompaction. *J. Am. Coll. Cardiol.* **2016**, *68*, 2166–2181. [CrossRef] [PubMed]
3. Aung, N.; Doimo, S.; Ricci, F.; Sanghvi, M.M.; Pedrosa, C.; Woodbridge, S.P.; Al-Balah, A.; Zemrak, F.; Khanji, M.Y.; Munroe, P.B.; et al. Prognostic Significance of Left Ventricular Noncompaction. Systematic Review and Meta-Analysis of Observational Studies. *J. Am. Coll. Cardiol.* **2020**, *75*, 2295–2304 . [CrossRef]
4. Captur, G.; Muthurangu, V.; Cook, C.; Flett, A.S.; Wilson, R.; Barison, A.; Sado, D.M.; Anderson, S.; McKenna, W.J.; Mohun, T.J.; et al. Quantification of Left Ventricular Trabeculae Using Fractal Analysis. *J. Cardiovasc. Magn. Reson.* **2013**, *15*, 36. [CrossRef] [PubMed]
5. Choi, Y.; Choe, Y.H.; Kim, S.M.; Lee, S.-C.; Chang, S.-A. Quantification of left ventricular trabeculae using cardiac magnetic resonance imaging for the diagnosis of left ventricular non-compaction. *J. Cardiovasc. Magn. Reson.* **2015**, *17*, P301. [CrossRef]
6. Bernabé, G.; González-Carrillo, J.; Cuenca, J.; Rodríguez-Sánchez, D.; Saura-Espín, D.; Gimeno-Blanes, J.R. Performance of a new software tool for automatic quantification of left ventricular trabeculations. *Rev. Esp. Cardiol.* **2017**, *70*, 405–407. [CrossRef] [PubMed]
7. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*; Springer : Munich, Germany, 2015; pp. 234–241. [CrossRef]
8. Rodríguez-de-Vera, J.M.; Bernabé, G.; García, J.M.; Saura, D.; González-Carrillo, J. Left Ventricular Non-Compaction Cardiomyopathy Automatic Diagnosis Using a Deep Learning Approach. *Comput. Methods Programs Biomed.* **2022**, *214*, 106548. [CrossRef] [PubMed]
9. de Haro, S.; Bernabé, G.; González-Férez, P.; Cámara, J. ViTUNeT-Based Segmentation for LVNC Analysis. *J. Integr. Bioinform.* 2025 , *advance online publication*.
10. Bernabé, G.; González-Férez, P.; García, J.M.; Casas, G.; González-Carrillo, J. Expanding the Deep-Learning Model to Diagnose LVNC: Limitations and Trade-Offs. *Comput. Methods Biomech. Biomed. Eng. Imaging Vis.* **2024**, *12*, 1. [CrossRef]
11. Elhoseny, M.; Shankar, K. Optimal Bilateral Filter and CNN-Based Denoising of Medical Images. *Measurement* **2019**, *143*, 125–135. [CrossRef]
12. Awarayi, N.S.; Twum, F.; Ben Hayfron-Acquah, J.; Owusu-Agyemang, K. Bilateral Filtering-Based Image Enhancement for Alzheimer's Disease Classification. *PLoS ONE* **2024**, *19*, e0302358. [CrossRef] [PubMed]
13. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788. [CrossRef]
14. Almufareh, M.F.; Imran, M.; Khan, A.; Humayun, M.; Asim, M. Automated Brain Tumor Segmentation and Classification in MRI Using YOLO-Based Deep Learning. *IEEE Access* **2024**, *12*, 16189–16207. [CrossRef]
15. Wei, W.; Huang, Y.; Zheng, J.; Rao, Y.; Wei, Y.; Tan, X.; OuYang, H. YOLOv11-Based Multi-Task Learning for Enhanced Bone Fracture Detection and Classification in X-ray Images. *J. Radiat. Res. Appl. Sci.* **2025**, *18*, 101309. [CrossRef]
16. Hechkel, W.; Helali, A. Early Detection and Classification of Alzheimer's Disease through Data Fusion of MRI and DTI Images Using the YOLOv11 Neural Network. *Front. Neurosci.* **2025**, *19*, 1554015 . [CrossRef] [PubMed]
17. Tomasi, C.; Manduchi, R. Bilateral Filtering for Gray and Color Images. In Proceedings of the Sixth International Conference on Computer Vision (ICCV), Bombay, India, 4–7 January 1998; pp. 839–846. [CrossRef]
18. Zuiderveld, K. Contrast Limited Adaptive Histogram Equalization. In *Graphics Gems IV*; Heckbert, P.S., Ed.; Academic Press: San Diego, CA, USA, 1994; pp. 474–485.
19. Ulutas, G.; Ustubioglu, B. Underwater Image Enhancement Using Contrast Limited Adaptive Histogram Equalization and Layered Difference Representation. *Multimed. Tools Appl.* **2021**, *80*, 15067–15091. [CrossRef]
20. Maron, B.; Desai, M.; Nishimura, R.; Spirito, P.; Rakowski, H.; Towbin, J.A.; Rowin, E.J.; Maron, M.S.; Sherrid, M.V. Diagnosis and Evaluation of Hypertrophic Cardiomyopathy: JACC State-of-the-Art Review. *J. Am. Coll. Cardiol.* **2022**, *79*, 372–389. [CrossRef] [PubMed]

21. Petersen, S.E.; Selvanayagam, J.B.; Wiesmann, F.; Robson, M.D.; Francis, J.M.; Anderson, R.H.; Watkins, H.; Neubauer, S. Left Ventricular Non-Compaction: Insights from Cardiovascular Magnetic Resonance Imaging. *J. Am. Coll. Cardiol.* **2005**, *46*, 101–105. [CrossRef] [PubMed]
22. Tayal, U.; Newsome, S.; Buchan, R.; Whiffin, N.; Halliday, B.; Lota, A.; Roberts, A.; Baksi, A.J.; Voges, I.; Midwinter, W.; et al. Phenotype and Clinical Outcomes of Titin Cardiomyopathy. *J. Am. Coll. Cardiol.* **2017**, *70*, 2264–2274. [CrossRef] [PubMed]
23. Ismail, S.M.; Said, L.A.; Radwan, A.G.; Madian, A.H.; Abu-ElYazeed, M.F. A Novel Image Encryption System Merging Fractional-Order Edge Detection and Generalized Chaotic Maps. *Signal Process.* **2020**, *167*, 107280. [CrossRef]
24. Ultralytics. YOLOv11 Model Documentation. Available online: https://docs.ultralytics.com/es/models/yolo11/ (accessed on 15 December 2025).
25. Belete, D.M.; Huchaiah, M.D. Grid Search in Hyperparameter Optimization of Machine Learning Models for Prediction of HIV/AIDS Test Results. *Int. J. Comput. Appl.* **2022**, *44*, 875–886. [CrossRef]

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.