



Segmentation and quantification of atherosclerotic plaques in optical coherence tomography



Viacheslav V. Danilov ^{a,b,*}, Vladislav V. Laptev ^{c,d}, Kirill Yu. Klyshnikov ^d, Ivan S. Bessonov ^e, Nikita V. Litvinyuk ^f, Evgeny A. Ovcharenko ^d, Nikita A. Kochergin ^d

^a Pompeu Fabra University, Barcelona, Spain

^b Symfa, Miami, United States

^c Siberian State Medical University, Tomsk, Russia

^d Research Institute for Complex Issues of Cardiovascular Diseases, Kemerovo, Russia

^e Tyumen Cardiology Research Center, Tyumen, Russia

^f Krasnoyarsk Regional Clinical Hospital, Krasnoyarsk, Russia

ARTICLE INFO

Keywords:

Atherosclerosis
Optical coherence tomography
Machine learning
Deep learning
Plaque segmentation
Cardiovascular diagnostics
Automated imaging
Vulnerable plaque

ABSTRACT

Introduction: Cardiovascular disease remains a leading cause of mortality worldwide, with atherosclerosis at its core. Accurate identification of vulnerable plaques is critical for preventing acute cardiovascular events. While optical coherence tomography (OCT) provides high-resolution imaging of plaque features, manual analysis is labor-intensive and operator-dependent. This study addresses the need for automated, accurate segmentation of atherosclerotic plaques in OCT pullbacks using advanced machine learning (ML) techniques.

Methods: A comprehensive multi-center dataset of OCT pullbacks, encompassing 103 patients and annotated for key plaque morphological features (lumen, fibrous cap, lipid core, and vasa vasorum), was used to tune, train and evaluate nine ML models, including U-Net, U-Net++, DeepLabV3, DeepLabV3+, FPN, LinkNet, PSPNet, PAN, and MA-Net. To address dataset imbalances and optimize performance for each plaque feature, we introduced a hybrid segmentation strategy: single-class models were deployed for highly prevalent features (e.g., lumen) and underrepresented classes (e.g., vasa vasorum), while a multi-class model targeted morphologically complex features with overlapping boundaries (e.g., fibrous cap and lipid core). Hyperparameter tuning was performed using Bayesian optimization, and segmentation accuracy was assessed with the Dice Similarity Coefficient (DSC) and other metrics.

Results: The models achieved high segmentation accuracy for lumen (DSC: 0.987), fibrous cap (DSC: 0.736), and lipid core (DSC: 0.751), demonstrating the potential of leveraging ML techniques to enhance OCT's diagnostic capabilities. While lumen segmentation showed exceptional precision, the moderate accuracy for fibrous cap and lipid core highlights challenges with complex morphologies. Satisfactory results for vasa vasorum (DSC: 0.610) suggest areas for further refinement. By integrating these models into a weighted ensemble, taking into account class prevalence and model confidence, the combined system achieved a weighted DSC of 0.882 across all plaque features, a significant improvement over individual models. These findings confirm the hybrid strategy's ability to balance computational efficiency with accuracy through rigorous optimization, tailored model selection, and ensemble integration.

Conclusion: This study presents a robust ML-driven framework for automated OCT segmentation that uses a hybrid approach and weighted ensemble learning to address class imbalance and feature complexity, significantly improving the accuracy and efficiency of atherosclerotic plaque analysis. The findings suggest potential clinical implications, including improved detection of high-risk plaques and enhanced decision-making in cardiovascular care. However, further prospective validation is required before clinical adoption. Future research should focus on expanding datasets, integrating multimodal imaging, and refining models for real-time clinical use, paving the way for transformative advancements in cardiovascular diagnostics.

* Corresponding author. Pompeu Fabra University, Barcelona, Spain.

E-mail address: viacheslav.v.danilov@gmail.com (V.V. Danilov).

1. Introduction

Atherosclerosis is a leading cause of cardiovascular morbidity and mortality worldwide, underlying ischemic heart diseases such as myocardial infarction, stroke, and peripheral arterial disease. It is characterized by the accumulation of lipid-laden plaques in arterial walls, which narrow the vessel lumen, impair blood flow, and increase the risk of thrombotic events [1]. In 2021, ischemic heart disease, largely driven by atherosclerosis, accounted for over 10 million deaths globally, remaining the foremost cause of mortality across both developed and developing regions [2]. Despite progress in medical therapy and revascularization, the incidence of atherosclerotic cardiovascular diseases continues to rise, particularly in low- and middle-income countries, underscoring the need for improved diagnostic and therapeutic strategies [3]. The clinical danger of atherosclerosis stems from both progressive luminal obstruction and plaque instability [4]. While some plaques remain asymptomatic, others are prone to rupture, triggering acute coronary syndromes and stroke [5]. Plaque vulnerability, encompassing inflammation, endothelial dysfunction, and lipid content, contributes to their varied behavior [6]. In addition to established features such as thin fibrous cap and large lipid core, recent evidence highlights the role of neovascularization (specifically, the proliferation of vasa vasorum) in promoting intraplaque inflammation and hemorrhage, both of which are implicated in plaque destabilization and rupture [6,7]. Identifying rupture-prone plaques, therefore, increasingly involves the assessment of both structural and microvascular features.

Intravascular imaging techniques such as intravascular ultrasound (IVUS), optical coherence tomography (OCT), and near-infrared spectroscopy (NIRS) have become essential for diagnosing and managing atherosclerosis [5,8]. These methods offer advantages over traditional techniques like coronary angiography and coronary computed tomography angiography, which, while valuable for assessing luminal stenosis, have limitations in detecting plaque composition and vulnerability [9]. IVUS is useful for evaluating vessel wall morphology and remodeling, while OCT, with its superior resolution, is effective in visualizing microstructures like fibrous caps and thrombus formation, playing a pivotal role in percutaneous coronary interventions and plaque vulnerability assessment [10]. NIRS complements these by identifying lipid-core plaques, associated with increased cardiovascular risk [11].

Among these, intravascular OCT stands out for its high resolution, offering tenfold better spatial resolution compared to IVUS, making it ideal for visualizing thin fibrous caps, a key feature of plaque instability and acute coronary syndrome risk [12]. Importantly, intravascular OCT is an invasive imaging modality, as it requires catheterization to acquire high-resolution images of the coronary arteries. Recent studies highlight OCT's ability to predict cardiovascular events by measuring fibrous cap thickness and surface area, outperforming other modalities [13]. Moreover, ML-driven image analysis is enhancing OCT's clinical utility, automating the identification of vulnerable plaque features, and reducing operator dependence [14].

Deep learning has emerged as a powerful tool for automating the analysis of intravascular OCT data, enabling segmentation and classification of plaque types such as calcified and lipid-rich areas. For instance, Athanasiou et al. [15] introduced a fully automated methodology that significantly reduced manual intervention in plaque detection. Similarly, A-line-based classification by Zhu et al. [16] employed a deep neural network to classify plaques with high accuracy, offering rapid analysis of large datasets. Shi et al. [17] proposed a multi-task learning model that achieved high precision in detecting thin-cap fibroatheromas, a key indicator of plaque vulnerability.

Recent OCT studies have started to explore multi-tissue segmentation, yet each is constrained in one of three dimensions: dataset scale, tissue breadth, or code availability. Lee et al. [18] reported automated micro-vessel delineation in 30 patients (median DSC 0.73) but omitted lumen and plaque tissues. Chu et al. [19] treated plaque characterization as a classification problem, providing no pixel-wise masks. One main

difference between our study and such prior work is the emphasis on full-resolution tissue segmentation rather than classification alone, enabling more precise and clinically interpretable analysis. More recently, Wittmann et al. [20] proposed a simulation-based approach for 3D OCT angiography vessel segmentation, focusing on cerebral vasculature without addressing coronary plaque components. Wang et al. [21] introduced an attention-based network for segmenting hyperreflective foci in retinal OCT images, targeting diabetic macular edema. However, both studies did not consider comprehensive coronary plaque segmentation or the inclusion of vasa vasorum. Our work is the first to combine (i) an open, 103-patient multi-center dataset, (ii) four clinically relevant classes, including the rarely annotated vasa vasorum, and (iii) public release of models and code.

While prior studies have advanced OCT-based plaque analysis, critical gaps remain. Existing frameworks often prioritize classification over detailed segmentation [16], rely on homogeneous datasets [22], or lack computational efficiency for clinical deployment [23]. Our work addresses these limitations by introducing a hybrid segmentation framework optimized for multi-center OCT data. This approach combines task-specific models to balance class imbalance, leverages rigorous hyperparameter optimization for efficiency, and integrates explainability methods to enhance clinical trust. By focusing on both prevalent (e.g., lumen) and rare (e.g., vasa vasorum) plaque features, our methodology advances automated, generalizable plaque quantification.

2. Materials and methods

2.1. Data acquisition

To address the segmentation and quantification task, we collected a multi-center, multi-scanner dataset that includes OCT pullbacks of atherosclerotic plaques in blood vessels. Each OCT pullback represents a 3D volume with dimensions $N \times N \times M$, where N represents the height and width of the image in pixels, and M denotes the number of slices, capturing the depth of the OCT pullback. This comprehensive dataset spans various institutions, imaging devices, and patient demographics, making it ideal for training models capable of accurate and reliable segmentation and quantification of atherosclerotic plaques across diverse clinical settings. Below are the main characteristics of the dataset:

- **Institutions:** Data were collected from two leading cardiovascular research centers: Research Institute for Complex Issues of Cardiovascular Diseases (Kemerovo, Russia): 73 OCT pullbacks, and Tyumen Cardiology Research Center (Tyumen, Russia): 30 OCT pullbacks. Each pullback represents a single intravascular OCT acquisition and consists of cross-sectional slices.
- **Scanners:** The dataset includes OCT pullbacks acquired using two generations of OCT systems from the same manufacturer lineage. Specifically, the pullbacks were obtained using two systems: the ILUMIEN Intravascular Imaging System and the OPTIS Next Intravascular Imaging System, both marketed by Abbott. The use of both systems introduces technical heterogeneity in imaging protocols, supporting the model's ability to generalize across different device generations.
- **Data collection period:** OCT pullbacks were collected over 1 year, from June 30, 2022 to July 11, 2023, capturing temporal variability in imaging practices and patient characteristics.
- **Gender distribution:** The dataset comprises pullbacks from 77 male and 26 female patients, providing essential diversity for analyzing sex-specific variations in atherosclerotic plaque morphology.
- **Image sizes:** The height and width of the OCT slices vary between 704×704 pixels and 1024×1024 pixels, reflecting diverse anatomical features and imaging conditions.

- **Image depth:** The number of slices per OCT pullback ranges from 215 to 270, corresponding to the variable imaging depths employed during diagnostic procedures.

The study adhered to the principles outlined in the Declaration of Helsinki and received approval from the Local Ethics Committee of the Research Institute for Complex Issues of Cardiovascular Diseases (Kemerovo, Russia) under protocol code 2022/06 on June 30, 2022. All participants provided written informed consent prior to inclusion in the study.

Patients aged 18 years or older with stable coronary artery disease undergoing cardiac catheterization were considered for inclusion. All lesions exhibiting an angiographic diameter stenosis of 50–90%, as determined by site visual estimation, were functionally assessed using fractional flow reserve (FFR). Intermediate, non-flow-limiting lesions ($FFR > 0.80$) underwent further evaluation via OCT. Flow-limiting lesions ($FFR \leq 0.80$) were treated with percutaneous coronary intervention using drug-eluting stents. The main exclusion criteria were as follows:

- Acute coronary syndrome.
- Prior coronary artery bypass grafting.
- Previously stented target lesions.
- Heavily calcified or angulated lesions.

The study cohort consisted of 103 patients with stable coronary artery disease, with a median age of 69 years (range: 43–83). Of these, 77 (74.7%) were male and 26 (25.3%) were female, reflecting the gender distribution within the cohort. A significant portion of the cohort presented with comorbidities, including 22 (21.4%) patients with diabetes and a history of myocardial infarction.

Clinical presentations varied, with the majority of patients (94, 91.3%) experiencing angina pectoris classified as functional classes 1–3, while a smaller group (9, 8.7%) exhibited silent ischemia. These diverse patient characteristics, summarized in **Table 1**, enhance the dataset's clinical relevance and provide a comprehensive basis for exploring the variability in atherosclerotic plaque features.

2.2. Data annotation

After data collection, two cardiologists meticulously annotated 103 OCT pullbacks, represented by 25,698 slices, ultimately identifying four plaque morphological features (PMFs) associated with distinct patterns of atherosclerotic plaque development: lumen, fibrous cap, lipid core, and vasa vasorum. Each identified PMF was annotated using binary masks as illustrated in **Fig. 1**. The OCT annotations were conducted using the web-based computer vision platform, *Supervisely* [24].

Following the initial annotation, the labeled slices were reviewed and double-verified by a senior cardiologist and a technical specialist responsible for dataset preparation to ensure the accuracy and reliability

Table 1
Baseline characteristics of patients included in the study.

Parameter	Value
Sex:	
Male, n (%)	77 (74.7)
Female, n (%)	26 (25.3)
Median age, years [min – max]	69 [43–83]
Arterial hypertension, n (%)	92 (89.3)
Diabetes mellitus, n (%)	22 (21.4)
Myocardial infarction, n (%)	22 (21.4)
Polyvascular disease, n (%)	29 (28.2)
Angina pectoris:	
Silent ischemia, n (%)	9 (8.7)
Functional class 1, n (%)	24 (23.3)
Functional class 2, n (%)	55 (53.4)
Functional class 3, n (%)	15 (14.6)

Table 2
Hyperparameters used during the networks' optimization.

Hyperparameter	Value	Count
Architecture	U-Net, U-Net++, DeepLabV3, DeepLabV3+, FPN, LinkNet, PSPNet, PAN, MA-Net	9
Encoder	ResNet-18, ResNet-50, ResNet-101, EfficientNet B0, EfficientNet B5, EfficientNet B7, RegNetX-200MF, RegNetX-6.4 GF, RegNetY-12GF	9
Input size	512x512 to 896x896 with the step of 128x128 px	4
Optimizer	Adam, RAdam, RMSprop	3
Learning rate	10^{-3} , 10^{-4} , 10^{-5} , 10^{-6}	4

of the annotations. This verification process involved adjustments or corrections to the annotations, further enhancing their precision and consistency.

The features described during the annotation of OCT slices represent key morphological characteristics of atherosclerotic plaques. These features are critical for understanding plaque composition and its clinical implications, particularly in assessing the risk of cardiovascular events. The following four PMFs were segmented and quantified:

1. **Lumen (LM):** This feature involves measuring the dimensions of the vascular lumen, which is crucial for evaluating stenosis and blood flow. Accurate lumen quantification helps assess the severity of vascular obstructions and guide clinical decision-making.
2. **Fibrous cap (FC):** This feature pertains to the fibrous cap's thickness and structure. The fibrous cap plays a significant role in plaque vulnerability, and its analysis is essential for predicting the likelihood of plaque rupture and potential cardiovascular events.
3. **Lipid core (LC):** The lipid core is characterized by high lipid content within the plaque. Quantifying the lipid core is a key indicator of atherosclerotic activity and helps to assess the likelihood of high-risk plaques, which can lead to serious conditions such as myocardial infarction.
4. **Vasa vasorum (VV):** The "vasa vasorum" feature was defined to include all microvessels and microchannels visible within the vessel wall and atherosclerotic plaque on OCT images, irrespective of their precise anatomical location. This encompasses both classical vasa vasorum, which are typically located in the adventitia and outer media, and neovessels arising within the intima and atherosclerotic plaque (sometimes referred to as "vasa plaquorum"). This inclusive definition was adopted to reflect the full spectrum of microvascularization associated with atherosclerotic disease, recognizing that OCT imaging may not reliably distinguish between these subtypes. The presence, density, and distribution of such microvessels are increasingly recognized as indicators of intraplaque neovascularization, inflammation, and heightened risk of plaque rupture. Accordingly, quantification of these microvessels provides valuable complementary information for assessing plaque vulnerability and the biological activity underlying atherosclerotic progression.

All identified PMFs play an essential role in understanding the structure and progression of atherosclerotic disease. Each feature provides critical insight into different aspects of plaque composition, influencing the clinical assessment of plaque stability and the risk of cardiovascular events. The vascular lumen reflects the degree of stenosis, affecting blood flow, while the fibrous cap, lipid core, and vasa vasorum are important markers for evaluating plaque vulnerability.

In the context of OCT imaging, these features are segmented and quantified to offer detailed visualization and measurement of plaque composition. This enables a comprehensive analysis of the plaques' morphology, aiding in the diagnosis and prognosis of atherosclerotic disease progression. By focusing on the segmentation of these PMFs, OCT provides a precise, non-invasive method to investigate the structural characteristics of plaques and predict potential cardiovascular risks.

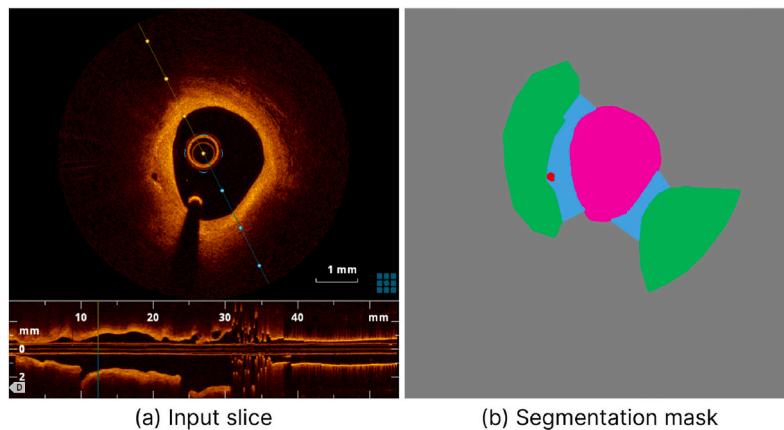


Fig. 1. Annotation methodology for optical coherence tomography slices depicting plaque morphological features associated with atherosclerotic plaque development. The feature annotations delineated with segmentation masks include the lumen (pink), fibrous cap (blue), lipid core (green), and vasa vasorum (red).

2.3. Model selection

To segment the four classes in OCT slices, we evaluated nine different neural networks: U-Net [25], U-Net++ [26], DeepLabV3 [27], DeepLabV3+ [28], FPN [29], LinkNet [30], PSPNet [31], PAN [32] and MA-Net [33]. These models were chosen based on their established efficacy in analyzing complex biomedical images [34].

U-Net, widely applied in biomedical image segmentation, is built on an encoder-decoder architecture that excels at capturing fine details [25]. Its ability to isolate small structures, such as calcified plaques and fibrous caps, makes it ideal for OCT image segmentation. U-Net's frequent use in related studies further solidifies its role as a baseline model [18,19]. U-Net++ enhances this with densely connected skip pathways, improving feature propagation and making it particularly effective for segmenting small, delicate structures like microvessels and thin fibrous caps, critical in OCT imaging, where precision is essential.

DeepLabV3 uses atrous convolution to extract features at multiple resolutions without downsampling, maintaining important details in OCT images [35]. This is crucial for accurately segmenting intricate structures such as lipid cores and calcifications. DeepLabV3+ further improves this model with a refined decoder, enhancing object boundary segmentation – essential for identifying small features like fibrous caps [36]. In similar studies, DeepLabV3+ has shown strong performance, achieving a sensitivity of 85.8% and a high A-line Dice coefficient of 0.837 in lipid plaque detection, along with a mean Dice score of 0.73 ± 0.10 for microvessel segmentation [18].

Feature Pyramid Network (FPN) excels at capturing multi-scale features, making it well-suited for segmenting and quantifying atherosclerotic plaques in OCT images [37]. By employing pyramidal feature hierarchies, FPN captures both global and fine-scale information, improving segmentation accuracy in cardiological imaging [38]. FPN has also demonstrated rapid convergence in similar segmentation tasks, such as microvascular feature detection in tissue-engineered vascular grafts [39].

LinkNet offers a streamlined architecture that balances computational efficiency with detailed feature extraction, making it suitable for real-time applications in OCT imaging. Its design allows for the rapid processing of high-resolution images while preserving the fine details needed to accurately segment complex plaque structures. In specific plaque segmentation tasks, LinkNet-34 has outperformed other models like FPN and U-Net [40].

PSPNet introduces a pyramid pooling module to aggregate multi-scale contextual information, making it highly effective for segmenting heterogeneous structures in OCT, such as thin fibrous caps and large calcifications [31]. PSPNet also demonstrated competitive convergence speeds in previous studies, ranking second only to FPN [39].

Pyramid Attention Network (PAN) combines pyramid pooling with attention mechanisms, capturing both local and global contexts. This dual mechanism ensures precise segmentation of small, critical features like microvessels and fibrous caps, which are vital in OCT imaging [41].

Lastly, Multi-Attention Network (MA-Net) employs multi-scale attention mechanisms, allowing it to capture fine details while maintaining broad contextual understanding. This makes MA-Net particularly effective for segmenting complex cardiovascular structures in OCT. Its successful application in IVUS image segmentation further supports its suitability for this task [42].

Overall, the selection of these nine models represents a comprehensive approach to addressing the specific challenges of OCT image segmentation. Each model brings distinct strengths in handling local and global features, preserving image resolution, and processing complex anatomical structures. This diversity in architecture provides a solid foundation for comparison and optimization, ensuring that the most suitable model can be identified for future advancements and clinical integration.

2.4. Model design

The design of the segmentation models was motivated by the non-uniform distribution of PMFs within the dataset and the distinct characteristics of each class. As detailed in Table 3, lumen objects are the most prevalent, comprising 21,808 instances, followed by fibrous cap (7226 instances) and lipid core (7192 instances). In contrast, the vasa vasorum class is significantly underrepresented, with only 450 instances.

The relatively low number of vasa vasorum instances in our dataset reflects primarily biological constraints. Vasa vasorum are small microvessels that are infrequently present in coronary OCT images, particularly in cases of stable coronary artery disease and in the absence of pronounced neovascularization. Clinical experience indicates that

Table 3
Slice and plaque morphological feature distributions across folds and subsets.

Fold	Subset	LM	FC	LC	VV	Total objects	Total slices
1	Train	17264	5610	5576	328	28778	16901
	Test	4544	1616	1616	122	7898	4492
2	Train	17554	5709	5690	237	29190	17207
	Test	4254	1517	1502	213	7486	4186
3	Train	17220	5600	5565	407	28792	16962
	Test	4588	1626	1627	43	7884	4431
4	Train	17813	5724	5686	416	29639	17473
	Test	3995	1502	1506	34	7037	3920
5	Train	17381	6261	6251	412	30405	17029
	Test	4427	965	941	38	6371	4364

when neovascularization is present, vasa vasorum can be clearly visualized and reliably identified on OCT. However, such cases were uncommon in our cohort. Despite their rarity, vasa vasorum were included in the segmentation task due to their emerging clinical relevance as markers of plaque inflammation and vulnerability.

This imbalance poses unique challenges for segmentation tasks, as model performance can be biased toward more abundant classes while underperforming on rare ones. To address these challenges, we adopted a hybrid segmentation strategy tailored to the specific distribution and morphological complexity of each PMF:

- **Lumen segmentation:** Leveraging its high representation in the dataset, a single-class model was trained exclusively for lumen.
- **Fibrous cap and lipid core segmentation:** These two classes, while moderately represented, share morphological similarities, such as complex and overlapping boundaries. A two-class model was tuned and trained to improve segmentation accuracy for both features.
- **Vasa vasorum segmentation:** The rarity of this class necessitated a dedicated single-class model to enhance performance by focusing solely on the specific characteristics of vasa vasorum.

This targeted approach allows for better handling of class imbalance, ensuring that each segmentation network is optimally configured for its respective PMF. By isolating segmentation tasks and aligning model designs with class distributions, we mitigate the inherent challenges posed by imbalanced datasets while improving segmentation accuracy across all classes. For the final ensemble, we trained three specialized models: a binary model for lumen segmentation, a binary model for vasa vasorum segmentation, and a two-class model for simultaneous segmentation of fibrous cap and lipid core. In contrast, the baseline multi-class model was trained to predict all four classes simultaneously within a single network. This distinction allowed us to address class imbalance and morphological complexity more effectively.

3. Hyperparameter tuning strategy

To achieve high-quality segmentation, we meticulously optimized the segmentation networks. Given the challenges posed by dataset imbalance, we curated specific datasets tailored to each segmentation task, transforming the overall problem into distinct binary and two-class segmentation subtasks. For each model, we conducted an extensive search over 350 hyperparameter configurations, optimizing architectural and training components. This rigorous process ensured optimal performance for each class-specific segmentation task.

Due to the substantial computational demands associated with hyperparameter optimization across multiple deep learning architectures and segmentation tasks, hyperparameter tuning was performed on a truncated dataset comprising 40 patients. This subset was selected using stratified random sampling to ensure that the distribution of key plaque morphological features (lumen, fibrous cap, lipid core, and vasa vasorum) closely reflected that of the entire cohort. This approach was intended to maintain the representativeness of the tuning dataset while enabling feasible optimization. The optimal hyperparameters identified using this subset were subsequently applied to model training on the complete dataset of 103 patients, with comparable performance observed across all plaque features.

Our tuning process aims to maximize the segmentation score, specifically focusing on the Dice Similarity Coefficient. To achieve this goal, we utilize a DSC loss, which is calculated as follows:

$$\text{Loss} = 1 - \frac{2 \sum (y_{\text{true}} \times y_{\text{pred}}) + \epsilon}{\sum y_{\text{true}} + \sum y_{\text{pred}} + \epsilon} \quad (1)$$

where y_{true} and y_{pred} represent the true and predicted label values, respectively, and ϵ is a small constant (set to 10^{-7} in our case) for numerical stability to avoid zero division errors.

As hyperparameter priorities vary during tuning, and certain hyperparameters have a more significant impact on network performance than others [43], we focused our hyperparameter tuning efforts on specific aspects rather than trying to optimize every parameter. In particular, we did not tune hyperparameters such as batch size, nonlinearity type, optimizer options, or kernel sizes. Instead, we focused on hyperparameters that were shown to be important in our previous study [34], namely encoder architecture, input image size, optimizer selection, and learning rate. In Table 2, we provide a comprehensive summary of the hyperparameters explored during the tuning process, along with the corresponding values used.

Regarding the hyperparameter search strategy, we employed Bayesian search, which, unlike Random or Grid Search, makes informed decisions. Bayesian optimization utilizes a probabilistic model to determine which values to use through an iterative process of testing values on a surrogate function before evaluating the objective function. Additionally, we utilized a specific early termination strategy, HyperBand [44], to halt poorly performing configurations. When early termination occurs, HyperBand stops the current configuration before proceeding with a new set of hyperparameter values. The combination of Bayesian optimization and HyperBand early termination forms a so-called “BOHB” [45], an approach that offers higher computational efficiency and robustness compared to Grid Search, Random Search, or standard implementations of Bayesian optimization or HyperBand.

3.1. Model training strategy

Following hyperparameter tuning, the segmentation networks were trained on the entire dataset using the identified optimal configurations. The training strategy leverages specialized networks to align with the segmentation approach:

- Single-class binary models were employed for lumen and vasa vasorum, targeting their distinct distributions and characteristics.
- A two-class model was used for fibrous cap and lipid core, leveraging their shared features and moderate representation in the dataset.

This strategy recognizes the unique demands of each PMF and enables the models to focus on their specific segmentation challenges. By dedicating a separate network to each segmentation task, we improved boundary delineation and overall segmentation accuracy, particularly for the less represented classes.

Given the limited number of subjects studied, comprising 103 patients (OCT pullbacks), we employed a 5-fold cross-validation technique. In this approach, each fold involved approximately 80% of the OCT pullbacks for training and the remaining 20% for testing (Table 3 and Fig. S1 of the Supplementary Information). This partitioning scheme was consistently applied to maintain the integrity of subject groups within each subset. Cross-validation was performed on a patient-wise basis, ensuring that no slices from the same patient appeared in both training and test sets. This approach prevented any form of data leakage and ensured the integrity of model evaluation.

During both the tuning and training steps, we employed a set of augmentation transformations using the “Albumentations” library [46]. These augmentations not only allowed us to expand the dataset size but also served as a regularization technique, helping to mitigate overfitting during model training. The proposed augmentation workflow encompasses the following transformations:

- Horizontal flip with a probability of 50%.
- Shift, scale, and rotate with a probability of 20%: Allows for random shifts, scaling, and rotations within specified limits (shift limit = 0.0625, scale limit = 0.1, and rotate limit = 15).
- Random crop with a probability of 20%: A random-sized crop is applied with dimensions determined by a percentage of the input size, ranging from 0.8 to 0.9 times the input size.

- Conditional Padding. All slices are padded to ensure a consistent size for processing.
- Gaussian noise with a probability of 20%: Adds random noise to the images with a variable intensity range, where the variance ranges from 3 to 10.
- Perspective distortion with a probability of 20%: Applies random perspective transformations to the images with a scale of 0.05–0.1.
- Random brightness and contrast adjustment with a probability of 20%: Adjusts the brightness and contrast of the images within specified limits (brightness limit = 0.2, contrast limit = 0.2).
- Hue, saturation, and value adjustment with a probability of 20%: Shifts the hue, saturation, and value of the images within specified limits (hue shift limit = 20, saturation shift limit = 30, value shift limit = 20).

In contrast to the tuning step, where a fixed batch size of 4 was utilized, the training step did not employ a fixed batch size. Since the studied models vary in complexity, they require different amounts of memory for training with a fixed batch size. Therefore, to ensure equitable training conditions, we adjusted the batch size based on the GPU memory utilization. Specifically, each model was trained with a batch size that allocated approximately 90–100% of GPU memory.

The network training, tuning, and testing were performed on a desktop computer featuring a 16-core Intel Xeon Gold 6326 CPU @ 2.90 GHz, 128 GB of RAM, and an Nvidia A100 GPU with 40 GB of video memory. PyTorch v2.2 and Python v3.11 were utilized as the primary machine learning framework and language for network development, respectively.

3.2. Ensemble framework

Before finalizing our hybrid ensemble approach, we trained a single multi-class segmentation model to jointly predict all four tissue types (lumen, fibrous cap, lipid core, and vasa vasorum) using the same architecture and data pipeline. This baseline yielded suboptimal results. The strong imbalance in both pixel count and instance frequency, particularly the dominance of the lumen class, caused the model to disproportionately prioritize lumen segmentation. As a result, it produced acceptable Dice scores for the lumen but substantially lower performance for the remaining tissues. Fibrous cap and lipid core regions were segmented inconsistently, and vasa vasorum was frequently missed or misclassified. This imbalance-related bias motivated our decision to train anatomically distinct models and combine them using a weighted ensemble strategy.

To further enhance segmentation accuracy and address class imbalance, we implemented a weighted ensemble framework that integrates predictions from the best-performing models for each PMF. For each pixel, the final class assignment was determined by maximizing a weighted sum of the predicted probabilities:

$$\hat{y}(x) = \operatorname{argmax}_c (w_c \cdot p_c(x)) \quad (2)$$

where $p_c(x)$ is the predicted probability for class c at pixel x , w_c is a composite weight reflecting both the prevalence of class c in the training data and the model's confidence, measured by the mean Dice Similarity Coefficient (DSC) for class c on the validation set.

This approach prioritizes the accurate detection of rare classes while leveraging the strengths of individual models. The weights w_c were empirically determined as:

$$w_c = \frac{\text{DSC}_c}{\log(1 + \text{prevalence}_c)} \quad (3)$$

where DSC_c is the mean Dice score for class c , and prevalence_c is the proportion of pixels belonging to class c in the training set. This formulation is designed to optimize segmentation performance by

balancing model confidence with class frequency, thereby enhancing the detection of underrepresented features such as vasa vasorum. Importantly, the ensemble weights are not intended to reflect the clinical risk associated with each feature (e.g., fibrous cap thickness or lipid core size), but rather to maximize overall segmentation accuracy. Clinical interpretation and risk stratification based on the segmented features are subsequent steps that can be performed using the quantitative outputs of our framework.

3.3. Model explainability

To ensure thorough model validation, we conducted an additional investigation focused on assessing the explainability of our models. In this context, we employed various methods of class activation mapping (CAM), including GradCAM [47], HiResCAM [48], GradCAMElementWise [49], GradCAM++ [50], XGradCAM [51], EigenCAM [52], EigenGradCAM, and LayerCAM [53]. These techniques were specifically aimed at analyzing the last layer of the encoder of the considered models.

To gain deeper insights into the decision-making process of the model and validate its predictive capabilities, we utilized the following explainable AI methods:

- GradCAM: A gradient-based method that identifies crucial regions of an image for classification based on activation gradients.
- HiResCAM: A high-resolution version of GradCAM designed for more precise delineation of activation regions.
- GradCAMElementWise: A variation of GradCAM that considers the influence of each pixel on the final prediction.
- GradCAM++: An enhanced version of GradCAM that considers both positive and negative activation gradients.
- XGradCAM: An extended version of GradCAM that utilizes additional features to improve the accuracy of important region detection.
- EigenCAM: A method based on eigenvalue algorithms to determine important image regions.
- EigenGradCAM: A variation of EigenCAM that also considers activation gradients to improve the accuracy of important region determination.
- LayerCAM: A method that analyzes activations in different layers of the model to identify important image features.

These methods provided invaluable insights into the inner workings of our models, illuminating the regions of interest identified during the segmentation process. This examination of model explainability bolsters our confidence in the reliability and interpretability of the segmentation results generated by our networks.

4. Results

4.1. Hyperparameter tuning

A comprehensive hyperparameter tuning process was conducted for each studied PMF, as outlined in the *Hyperparameter tuning strategy* section. This involved the evaluation of 350 network configurations per model (1050 configurations in total). The main results are summarized below, in Table 4, and in Table S1–S3 of the Supplementary Information:

- **Tuning time:** The duration of the tuning process varied significantly across different PMFs. Fibrous cap and lipid core segmentation required the longest tuning time, consuming 712 h, which likely reflects the complexity of these particular classes. In comparison, the lumen segmentation model completed its tuning in 309 h. Vasa vasorum segmentation, having a smaller dataset, exhibited the shortest tuning time at 103 h. It should be noted that these reported times correspond to the total cumulative runtime required to evaluate all 350 network configurations for each plaque morphological

Table 4
Optimal hyperparameters for the studied plaque morphological features.

	Lumen	Fibrous cap & Lipid core	Vasa vasorum
Architecture	U-Net++	LinkNet	U-Net
Encoder	ResNet-101	EfficientNet B7	RegNetX-6.4 GF
Input size	512x512	896x896	896x896
Optimizer	RMSprop	RMSprop	RAdam
Learning rate	10^{-5}	10^{-5}	10^{-3}
Parameters, M	68.0	64.4	31.9
MACs, G	249.8	9.0	158.0
Precision	0.991	0.700	0.814
Recall	0.993	0.977	1.000
F1	0.989	0.617	0.784
IoU	0.979	0.488	0.695
DSC	0.989	0.617	0.784

feature during hyperparameter optimization. In the supplementary materials (Table S1–S3), we present only the optimal configurations for each class and network architecture, along with their individual runtimes. The full set of tested configurations and their respective runtimes are available in our GitHub repository for full transparency (see *Data availability*).

• **Model Performance.** The segmentation performance of each model was evaluated using several metrics, including precision, recall, F1 score, Intersection over Union (IoU), and Dice Similarity Coefficient. These metrics were computed on the truncated dataset of 40 patients used exclusively for model tuning to reflect performance during the fine-tuning phase. U-Net++ achieved the highest DSC of 0.989 for lumen segmentation, demonstrating nearly ideal agreement between predicted and ground truth segmentations. For fibrous cap and lipid core segmentation, LinkNet attained a DSC of 0.617, indicating moderate accuracy on average. Finally, U-Net performed relatively

well in vasa vasorum segmentation, with a DSC of 0.784, reflecting a good level of accuracy for this smaller dataset.

• **Model Complexity:** Model complexity was assessed based on the number of parameters and multiply-accumulate operations (MACs). U-Net++, used for lumen segmentation, had 68 million parameters and required 249.8 GMACs of computational power. LinkNet, employed for fibrous cap and lipid core segmentation, had 64.4 million parameters and exhibited the lowest computational load, requiring only 9.0 GMACs. In contrast, U-Net, applied to vasa vasorum segmentation, had the smallest number of parameters at 31.9 million and required 158 GMACs.

Overall, the results underscore the variability in model performance and computational complexity across different segmentation tasks. U-Net++ exhibited superior performance in lumen segmentation, while LinkNet and U-Net demonstrated robustness for fibrous cap, lipid core, and vasa vasorum segmentation. These findings highlight the trade-offs between model accuracy and complexity, emphasizing the need for task-specific optimization in plaque segmentation.

4.2. Model training

We conducted a comprehensive evaluation of the training dynamics for the top-performing models selected during the hyperparameter tuning phase: U-Net++ for lumen segmentation, LinkNet for fibrous cap and lipid core segmentation, and U-Net for vasa vasorum segmentation. These models were trained for 125 epochs using a 5-fold cross-validation strategy to ensure robust performance assessment. This cross-validation setup allowed us to track trends in the loss function and DSC over epochs for each PMF, ensuring detailed monitoring of convergence behavior (refer to Fig. 2).

The lumen, fibrous cap and lipid core segmentation models exhibited

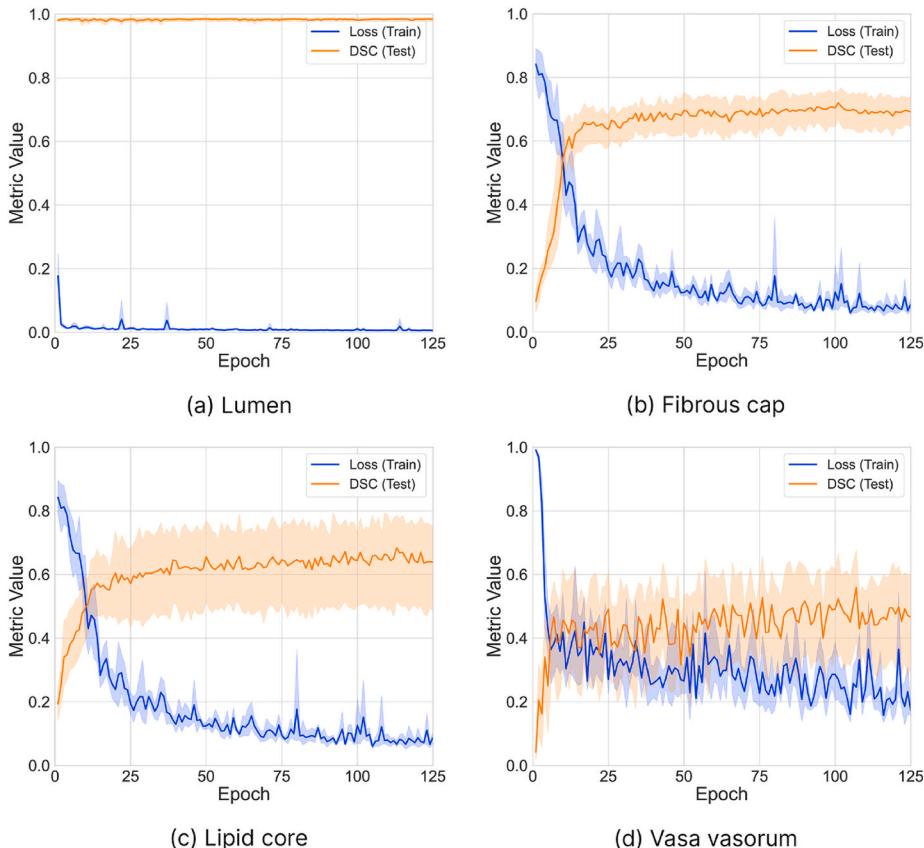


Fig. 2. Comparative analysis of loss and DSC evolution during training and testing phases over 5-fold cross-validation with 95% confidence intervals.

a clear and consistent reduction in loss accompanied by progressive improvements in the DSC, indicative of effective learning and precise refinement of segmentation predictions. Conversely, the vasa vasorum model demonstrated a slower convergence rate, coupled with significant fluctuations in loss and DSC values. These variations underscore the inherent complexity and challenges associated with segmentation of this PMF.

To systematically assess convergence, we defined stabilization points (refer to Fig. 2) for the loss function and DSC where incremental improvements were minimal. U-Net++ (lumen segmentation) and LinkNet (fibrous cap and lipid core segmentation) achieved stabilization rapidly, requiring fewer epochs to reach high DSC values with minimal loss variability. In contrast, U-Net (vasa vasorum segmentation) exhibited protracted convergence and less stable training curves, indicating challenges in learning fine-grained features in this PMF. Such fluctuations reflect the structural complexity and variability of the vasa vasorum, which likely require more nuanced feature extraction and aggregation.

The observed disparities in convergence behavior can be attributed to several factors, including architectural variations among the models, differences in parameter initialization, and the intrinsic segmentation difficulty of each PMF. For instance, the vasa vasorum model faced challenges in capturing subtle morphological features, requiring longer training times and advanced architectural mechanisms for effective feature learning.

The evaluation of segmentation performance metrics reveals varying levels of accuracy across the PMFs assessed (refer to Table 5 and Fig. 3). The lumen segmentation model exhibits exceptional performance, achieving a DSC of 0.987 and an IoU of 0.975. These metrics underscore the model's ability to reliably and accurately delineate lumen structures, highlighting its robustness and suitability for clinical applications requiring precise segmentation. In contrast, the fibrous cap and lipid core model demonstrate reliable performance, with DSCs of 0.736 and 0.751, respectively. These results reflect the effectiveness of the model in capturing the intricate and complex boundaries of these structures, despite the inherent challenges posed by their morphology. The vasa vasorum model shows moderate performance, achieving a DSC of 0.610. While segmentation of this feature remains challenging due to its subtle and less distinct boundaries, the results indicate progress in addressing these complexities.

To establish the upper bound of segmentation performance, we assessed the inter-annotator agreement between two expert cardiologists on a representative subset of the dataset. The DSC for each plaque morphological feature was as follows: lumen, 0.985; fibrous cap, 0.844; lipid core, 0.815; and vasa vasorum, 0.693. These values represent the expected range of human-level agreement and provide a benchmark for interpreting the model's performance.

By integrating these task-specific models into a weighted ensemble framework – which accounts for class prevalence and model confidence – the combined ensemble achieved a weighted DSC of 0.882 across all PMFs. To assess the effectiveness of the proposed weighted ensemble, we compared its performance with two naive aggregation strategies: (1) majority voting, where each pixel is assigned the class predicted by the majority of models, and (2) unweighted averaging, where class probabilities are averaged across models and the class with the highest mean probability is selected. As shown in Table 6, our weighted ensemble

consistently outperformed both naive methods, particularly for the vasa vasorum class (DSC: 0.610 vs. 0.541 for majority voting and 0.552 for unweighted averaging), and achieved the highest overall weighted DSC (0.882 vs. 0.864 and 0.867, respectively). This represents a significant improvement over individual models, demonstrating the ensemble's ability to harmonize diverse architectural strengths and mitigate class imbalance challenges inherent in heterogeneous plaque characterization. It should be emphasized that the ensemble weighting strategy is optimized for segmentation accuracy and does not directly encode clinical vulnerability criteria. The clinical significance of features such as lumen size, lipid core area, or fibrous cap thickness is assessed in downstream analyses, leveraging the accurate and comprehensive segmentation provided by our ensemble framework.

4.3. Model visual assessment

To assess segmentation accuracy visually, we present three representative slices that illustrate the segmentation outputs for the plaque morphological features (Fig. 4). These slices display predictions generated by the optimal ensemble of models: U-Net++ for lumen segmentation, LinkNet for fibrous cap and lipid core segmentation, and U-Net for vasa vasorum segmentation. The predictions are juxtaposed with their corresponding ground truth annotations, allowing for a direct visual evaluation of model performance.

As depicted in Fig. 4, the segmentations of lumen and vasa vasorum exhibit strong alignment with the ground truth, underscoring the ensemble's effectiveness in delineating these features. However, subtle discrepancies are occasionally observed in the fibrous cap and lipid core segmentations. These challenges are likely attributable to the more complex structure and morphology of these features, which differ significantly from those of lumen and vasa vasorum. Both fibrous cap and lipid core are components of tissue with diffuse and intricate boundaries, making precise delineation more difficult. In contrast, lumen and vasa vasorum typically exhibit clearer and more distinct borders, facilitating their segmentation.

It is important to note that, while the segmentation of the lumen and major plaque features demonstrates high accuracy, the delineation of the outer vessel wall remains suboptimal. This limitation is primarily attributable to the intrinsic characteristics of OCT imaging, most notably its limited penetration depth, which often precludes visualization of the vessel wall beneath atherosclerotic plaques and thus the true outer boundary of the plaque. In the majority of cases, the outer boundary cannot be reliably identified on OCT. Consequently, our annotation protocol required annotators to segment either along the frame boundary or by extrapolating the visible medial layer from adjacent regions. This necessary but pragmatic approach introduces subjectivity and interobserver variability, further constraining the achievable segmentation quality at the vessel periphery. Moreover, as the outer wall was not explicitly annotated as a separate class in our dataset, its identification relies on the inferred boundaries of other segmented features. Future work will address these limitations by incorporating explicit outer wall annotation and leveraging multimodal imaging to improve boundary definition.

The fibrous cap, in particular, is a thin layer of fibrous tissue overlaying a lipid core, and its boundary can vary significantly in texture and thickness. Meanwhile, the lipid core's irregular geometry and frequent embedding within surrounding tissue further complicate its segmentation. These intrinsic structural challenges likely contribute to the observed differences in segmentation accuracy. These visual comparisons complement the quantitative performance metrics and provide critical insights into the specific areas where improvements are needed. Additional visualization of model predictions, highlighting their performance on various plaque morphological features, is available in Fig. S1 of the Supplementary Information.

Table 5

Segmentation performance metrics for each plaque morphological feature, averaged over 5 folds during cross-validation.

PMF	Precision	Recall	F1	IoU	DSC
Lumen	0.986	0.988	0.987	0.975	0.987
Fibrous cap	0.737	0.784	0.736	0.608	0.736
Lipid core	0.815	0.772	0.751	0.639	0.751
Vasa vasorum	0.664	0.630	0.610	0.511	0.610

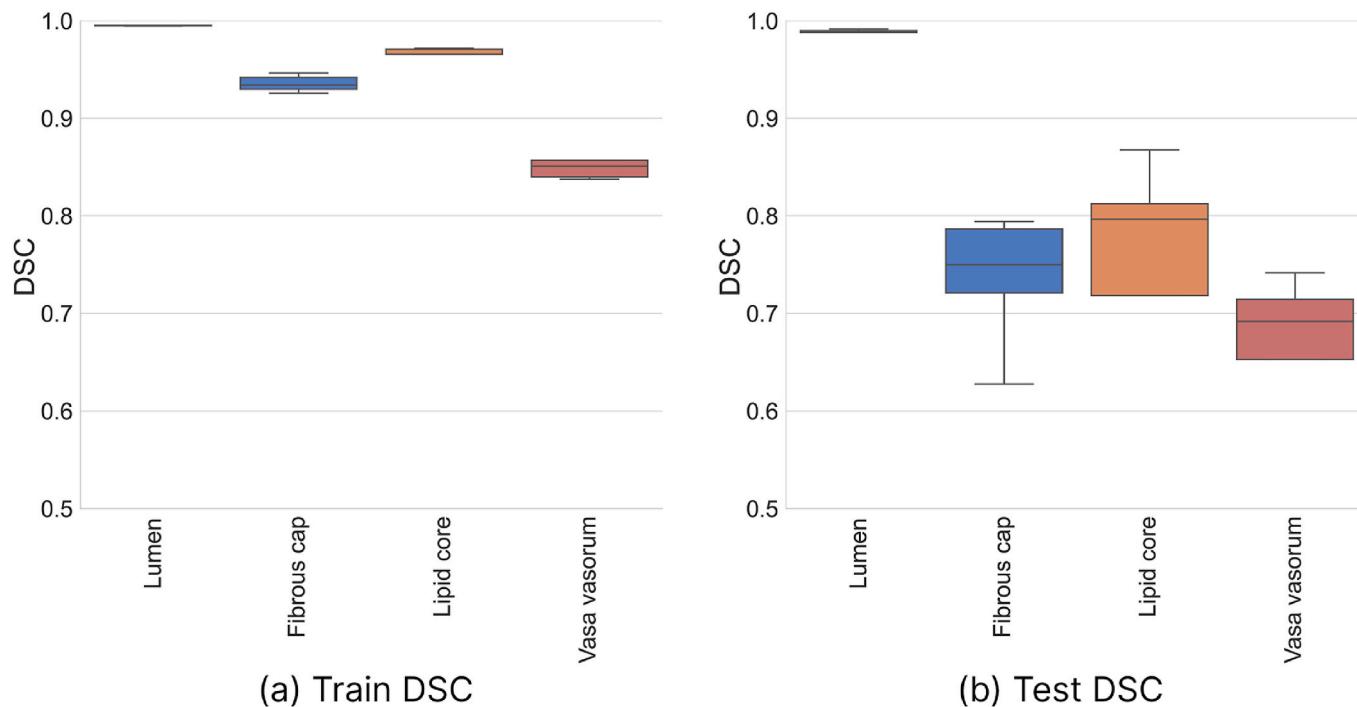


Fig. 3. Average feature-wise performance across various plaque morphological features.

Table 6
Comparison of ensemble strategies for segmentation performance (DSC values).

Method	Lumen	Fibrous Cap	Lipid Core	Vasa Vasorum	Weighted DSC
Majority Voting	0.985	0.701	0.729	0.541	0.864
Unweighted Averaging	0.986	0.713	0.735	0.552	0.867
Weighted Ensemble	0.987	0.736	0.751	0.610	0.882

4.4. Model explainability and class activation mapping analysis

To evaluate the interpretability of the segmentation models and their focus on relevant plaque morphological features, multiple CAM methods were applied, including GradCAM, GradCAM++, HiResCAM, LayerCAM, and others. In the main body of this article, we present comparative class activation maps for the largest and smallest regions: the lumen (Fig. 5) and the vasa vasorum (Fig. 6). These examples demonstrate the variability in CAM algorithm performance across features with differing sizes and complexities. Further examples of comparative class activation maps for different patients and each of the studied classes are provided in Fig. S2–S5 of the Supplementary Information, offering a broader evaluation of model performance.

The results for each plaque feature, as observed through CAM analysis, are summarized below:

• Lumen

- o LayerCAM, HiResCAM, GradCAMElementWise, and EigenGradCAM consistently demonstrated superior performance in identifying lumen boundaries with high precision. These algorithms effectively emphasized clear contours and minimized false activations outside the lumen area, indicating robust feature localization.
- o GradCAM, GradCAM++, and XGradCAM exhibited moderate performance. While they captured the general shape of the lumen, occasional activations on image corners or background artifacts reduced their reliability.

o EigenCAM underperformed, frequently misclassifying surrounding regions as part of the lumen, which detracted from its interpretability.

• Fibrous cap

- o LayerCAM, HiResCAM, GradCAMElementWise, and EigenGradCAM excelled in highlighting the thin fibrous cap layer, a critical feature for assessing plaque vulnerability. These methods accurately delineated the fibrous cap boundaries, even in challenging cases with diffuse textures.
- o GradCAM, GradCAM++, and XGradCAM provided average performance, with outputs occasionally including spurious activations at image edges.
- o EigenCAM performed poorly, often activating erroneously on lipid core regions rather than the fibrous structure itself, undermining its clinical relevance.

• Lipid core

- o Segmentation of the lipid core proved challenging for all tested algorithms. Weak or incorrect activations were observed across all CAM methods, including GradCAM, HiResCAM, GradCAMElementWise, GradCAM++, XGradCAM, EigenCAM, EigenGradCAM, and LayerCAM. Most CAM approaches for the lipid core incorrectly focused attention on the lumen, the fibrous cap, or vessel regions immediately adjacent to the lumen. This was particularly pronounced for GradCAMElementWise and EigenGradCAM, which frequently activated erroneously at the boundary between the lumen and lipid core or even within the lumen itself.
- o This limitation is likely due to the lipid core's intricate morphology, unclear boundaries, and frequent proximity to the vessel lumen, making accurate localization difficult even for

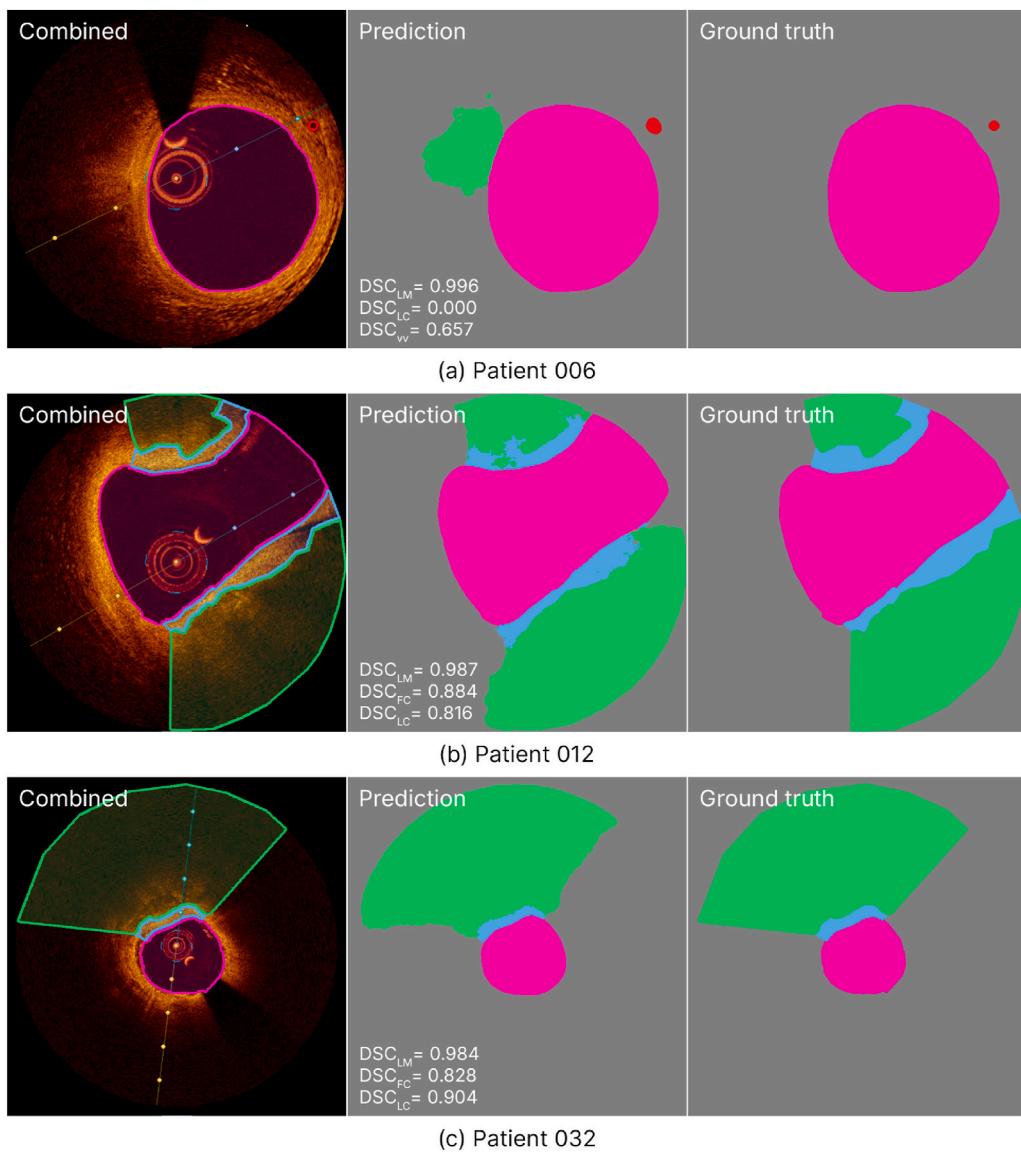


Fig. 4. Comparison between ground truth segmentation and ensemble predictions.

advanced explainability methods. These findings highlight the need for additional training or specialized optimization to improve detection for this category.

• Vasa vasorum

- o Most algorithms, except EigenCAM, performed adequately in identifying the vasa vasorum. These methods correctly highlighted the small, less distinct vascular structures with minor deviations.
- o EigenCAM, however, exhibited significant misclassifications, often activating in the lumen region where vasa vasorum structures were absent, reducing its reliability for detecting smaller anatomical features.

These activation maps validate the models' focus on anatomically relevant regions and provide a mechanism for identifying potential areas of misclassification or ambiguity. Notably, LayerCAM and HiResCAM demonstrated consistent performance in generating accurate visualizations across all plaque features, outperforming other methods due to their ability to aggregate multi-scale activations and localize fine structures (e.g., thin fibrous caps) with minimal false positives. Their gradient-weighted attention mechanisms further suppress imaging artifacts inherent to OCT, such as speckle noise, while preserving critical morphological details. By integrating these explainability methods into

the segmentation pipeline, our analysis not only enhances confidence in the models' clinical applicability but also identifies opportunities for optimization, such as refining feature boundaries in lipid-rich regions.

5. Discussion

5.1. Clinical relevance of plaque morphology in OCT

Current clinical guidelines recommend percutaneous coronary intervention (PCI) exclusively for flow-limiting lesions or those implicated in acute coronary syndromes [54]. However, vulnerable plaques, whether flow-limiting or not, pose significant risks for future adverse cardiac events even under optimal medical therapy [55]. Natural history studies have demonstrated that characteristics such as large plaque burden, reduced minimal lumen area, high lipid content, and thin fibrous caps are strongly associated with lesion-specific cardiac events, with risk escalating as the number of adverse features increases [56].

In addition to these established features, the role of vasa vasorum in plaque vulnerability has garnered increasing attention. Neovascularization within the plaque, reflected by increased vasa vasorum, facilitates the influx of inflammatory cells and is a major contributor to intraplaque hemorrhage, both of which are recognized mechanisms of

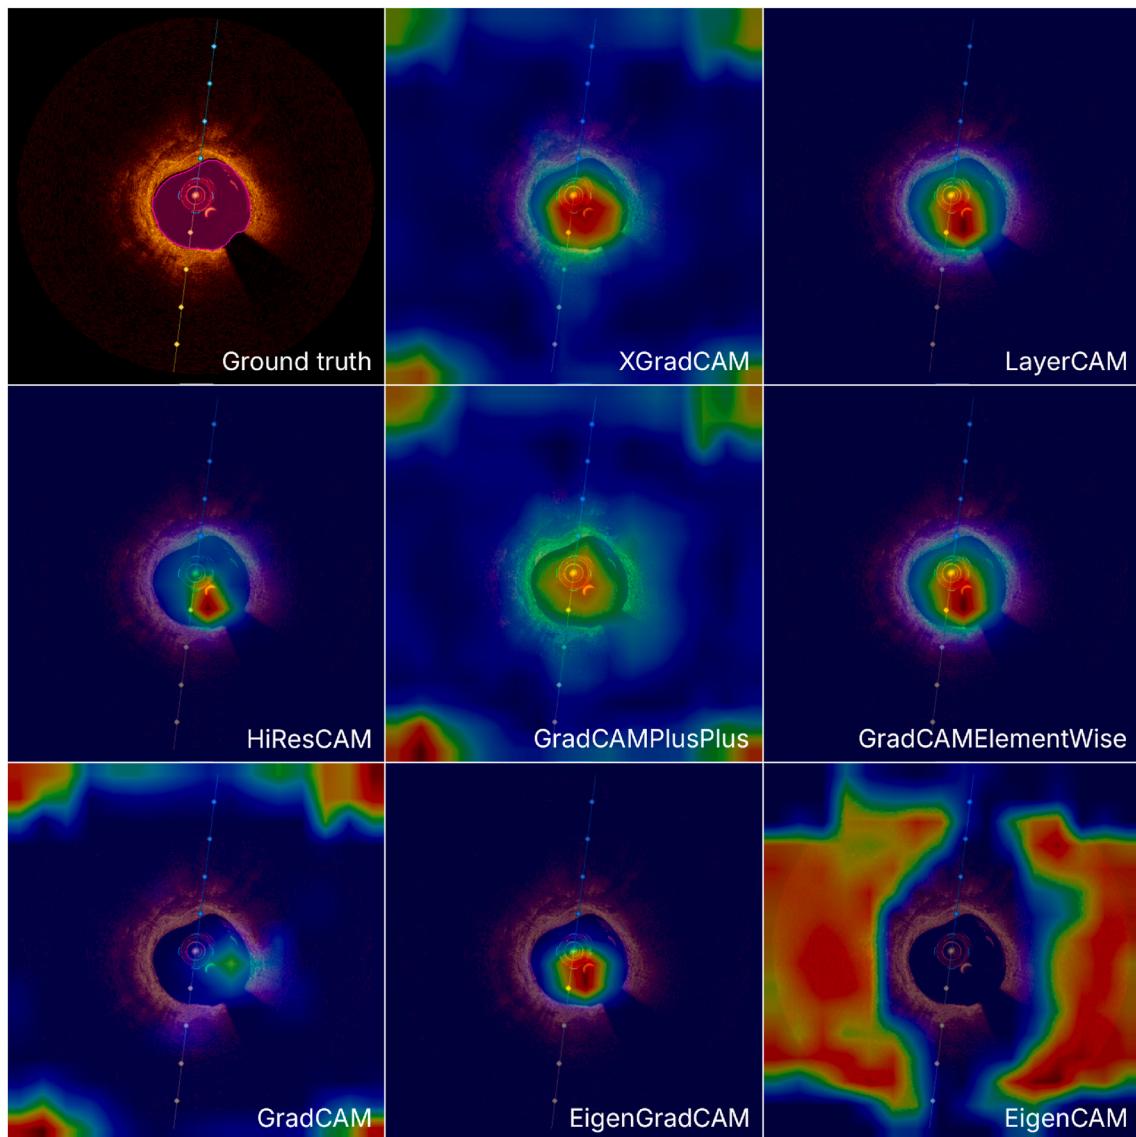


Fig. 5. Comparative class activation maps for the lumen region in patient 032.

plaque destabilization [6,7]. Therefore, the inclusion of vasa vasorum as a PMF in our study is supported by its emerging relevance in the pathophysiology of vulnerable plaques.

Interestingly, plaque vulnerability is dynamic; some vulnerable plaques stabilize without clinical events, whereas stable plaques may transition to a vulnerable state [57]. This evolving understanding has spurred interest in preventive PCI for high-risk vulnerable plaques, with the stenting mechanism hypothesized to stabilize plaques by covering the fibrous cap and reducing the likelihood of rupture [58].

5.2. Segmentation versus alternative ML approaches

While classification approaches can detect the presence of vulnerable plaque phenotypes such as thin-cap fibroatheromas (TCFAs), they do not provide spatial localization or quantifiable information. Object detection serves as an intermediate approach, predicting bounding boxes around regions of interest such as fibrous caps or lipid cores. This provides some localization while requiring less annotation effort than segmentation; however, it remains limited to approximate spatial delineation and cannot reliably capture fine morphological boundaries.

In contrast, segmentation delivers pixel-level anatomical detail that is critical for clinical decision-making. TCFA diagnosis hinges on cap

thickness $<65 \mu\text{m}$ [59], which necessitates per-pixel precision. It also enables visualization of fibrous cap contours, assessment of spatial overlap with lipid cores, and calculation of volumetric plaque burden, all essential for risk stratification, longitudinal monitoring, and image-guided stenting. Additionally, segmentation facilitates the automated extraction of spatial biomarkers that can feed into prognostic or predictive models. This granularity and quantification capacity make segmentation uniquely suited to clinical workflows, whereas classification and object detection provide only coarse or partial insights. A structured comparison of these three approaches is presented in Table 7.

Table 7 underscores why segmentation is the most clinically valuable framework for OCT-based plaque analysis. Unlike classification or object detection, only segmentation provides the pixel-level detail required to measure fibrous cap thickness, quantify lipid burden, and spatially co-register features across serial imaging. These capabilities directly support interventional planning, longitudinal monitoring, and risk modeling. While object detection may assist in preliminary localization, and classification in broad risk categorization, neither offers the anatomical fidelity or quantification necessary for therapeutic decision-making. Segmentation, by contrast, delivers anatomically faithful, reproducible, and clinically interpretable outputs aligned with real-world workflows and predictive tools.

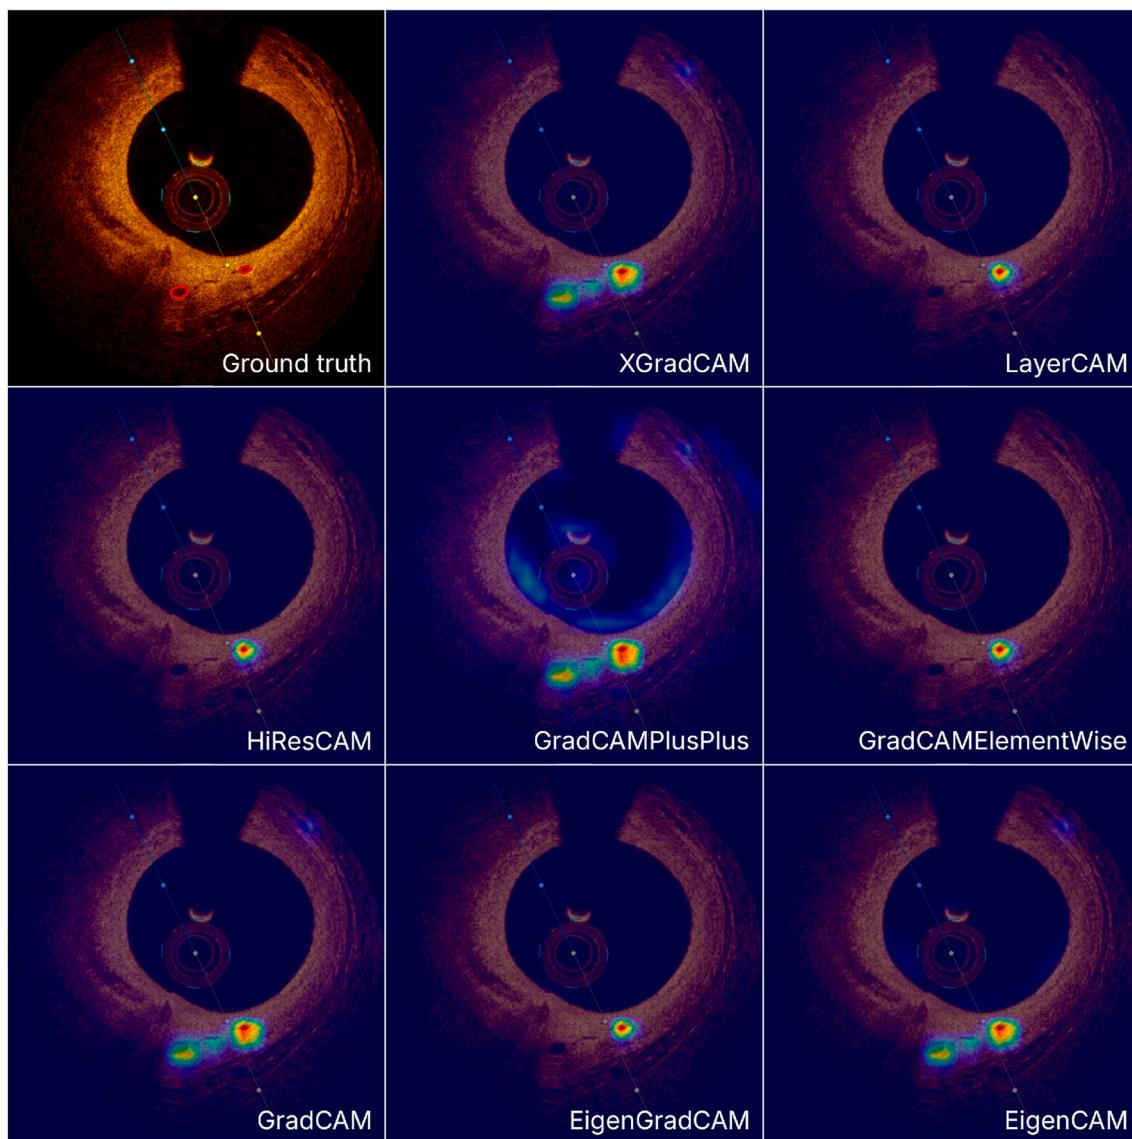


Fig. 6. Comparative class activation maps for the vasa vasorum region in patient 001.

Table 7

Comparative overview of classification, object detection, and segmentation approaches for OCT-based plaque feature analysis.

Criteria	Classification	Object Detection	Segmentation
Analytical scope	Global (entire frame or image-level decision)	Localized region estimation using bounding boxes	Pixel-wise delineation of anatomical structures
Model output	Single class label per image (e.g., TCFA present/absent)	Bounding boxes with class labels for candidate regions	Dense semantic mask with per-pixel class assignments
Spatial localization	None (presence inferred without anatomical mapping)	Coarse localization (rectangular box encloses feature)	Precise anatomical boundary (enables cap thickness and lipid area computation)
Quantification capability	Absent	Approximate region-level area estimation	Exact measurement of thickness, area, and volume at pixel-level resolution
Annotation burden	Minimal (image-level labeling)	Moderate (box-level labeling per feature)	High (manual per-pixel annotation required)
Clinical interpretability	Low (non-localized classification with limited transparency)	Moderate (approximate localization with limited boundary definition)	High (overlay visualization, direct anatomical correspondence, and feature traceability)
Relevance to PCI planning	Limited (no spatial guidance)	Modest (coarse region identification)	High (accurate real-time anatomical mapping for intervention planning)
Utility in risk modeling	Binary predictors (e.g., TCFA presence)	Semi-quantitative metrics (e.g., region count or size)	Quantitative biomarkers (e.g., cap thickness, lipid core burden, spatial distributions)

5.3. Advancing deep learning for automated OCT interpretation

Intravascular imaging holds significant promise for improving PCI outcomes in patients with coronary artery disease. Methods such as OCT

and IVUS provide detailed visualization of vessel components and plaque morphology. However, discrepancies in image interpretation among experts complicate diagnostics and decision-making [59].

The integration of deep learning, particularly neural networks, is a

rapidly advancing field that could address these challenges. Deep learning models trained on intravascular imaging data can identify morphological elements of the vascular wall and assess plaque vulnerability, reducing human error and expediting diagnostic workflows. For instance, Jun et al. [60] explored various machine learning classifiers, including convolutional neural networks (CNN), feed-forward neural networks (FNN), K-nearest neighbor (KNN), and random forests (RF), to classify TCFAs using IVUS images labeled with OCT. Their study of 12,325 IVUS images reported AUC values of 0.911 for CNN, 0.859 for FNN, 0.848 for KNN, and 0.844 for RF, with CNN emerging as the most accurate classifier. Interestingly, the feature importance rankings of FNN, KNN, and RF closely aligned with physician diagnostic criteria, underscoring the potential of machine learning in replicating expert decision-making [60].

In comparison, our study developed a deep learning solution for automated segmentation and quantification of atherosclerotic plaques in OCT pullbacks. Utilizing state-of-the-art deep learning architectures, our models demonstrated high segmentation accuracy for key plaque features: lumen (DSC: 0.987), fibrous cap (DSC: 0.736), and lipid core (DSC: 0.751). Similarly, Bae et al. [61] demonstrated the potential of ML models in TCFAs using OCT-derived data, achieving accuracies of 81% with artificial neural networks and AUC values of 0.80, which are comparable to the high segmentation accuracy observed in our study.

The findings highlight OCT's synergy with machine learning for accurately segmenting critical plaque features. Lumen segmentation achieved near-perfect accuracy, reflecting the robustness of the model for well-defined structures. Meanwhile, the moderate performance in vasa vasorum segmentation (DSC: 0.610) observed in our study is consistent with findings from other studies that have highlighted the challenges of detecting subtle vascular structures in OCT images. For example, Lee et al. reported similar difficulties in segmenting microvessels, achieving a mean DSC of 0.73 ± 0.10 , underscoring the need for further refinement in this area [18]. Our study advances this field by employing a dedicated single-class model for vasa vasorum segmentation, which, while showing moderate results, indicates the necessity for larger and more diverse datasets to improve performance.

The inter-annotator agreement analysis demonstrates that the model achieves near-human performance for lumen segmentation and reasonable accuracy for more challenging features such as the fibrous cap, lipid core, and vasa vasorum. The observed gap between model and human agreement for these features highlights the inherent difficulty of manual annotation and the need for further methodological refinement.

Overall, our findings contribute to the growing body of evidence supporting the integration of ML in OCT-based plaque analysis. The study supports the hypothesis that advanced segmentation models, optimized through rigorous hyperparameter tuning, can achieve a balance between computational efficiency and accuracy. The ensemble approach, tailored to specific plaque features, effectively mitigated dataset imbalances, thereby enhancing performance across features with diverse prevalence rates. The superiority of the weighted ensemble over naive aggregation methods underscores the importance of integrating both class prevalence and model confidence in the final prediction. While majority voting and unweighted averaging are straightforward, they fail to account for the inherent class imbalance and variability in model performance across different PMFs. Our approach, by contrast, adaptively emphasizes underrepresented and challenging classes, thereby improving clinical utility and robustness.

5.4. Study limitations

This study, while advancing the use of machine learning for OCT-based segmentation, is subject to several limitations that must be addressed to maximize its clinical and research impact. These limitations include both technical challenges and constraints related to the dataset and methodology:

Dataset imbalance and feature representation: The dataset used

in this study exhibited a marked imbalance among plaque morphological features, with vasa vasorum being particularly underrepresented. This imbalance constrained the model's ability to achieve high segmentation accuracy for rare features, despite the use of targeted single-class models and data augmentation. Future studies should aim to curate larger, more balanced datasets, especially for features with emerging clinical relevance.

Population representation: Although our dataset was collected from multiple centers, all participating institutions were located in Russia. This geographic limitation may restrict the generalizability of our findings to populations with different genetic backgrounds, risk factor profiles, or healthcare practices. Broader, multi-national collaborations are needed to validate and extend these results.

Model generalizability across imaging protocols: Variability in OCT imaging protocols, scanner types, and acquisition parameters across clinical sites may affect model performance when applied to external datasets. While our multi-center approach partially mitigates this issue, further validation on independent cohorts and standardized imaging protocols is warranted.

Complexity of feature boundaries: The accurate segmentation of features such as the fibrous cap and lipid core remains challenging due to their diffuse, overlapping, or ambiguous boundaries in OCT images. These challenges are compounded by the inherent resolution limits of OCT and the subjective nature of manual annotation, as reflected in the observed inter-annotator variability.

Potential overfitting and limited sample size: The observed difference in DSC values between the validation and test sets for the fibrous cap, lipid core, and vasa vasorum may indicate overfitting, especially for features with limited representation and complex morphology. While our use of extensive data augmentation and 5-fold cross-validation helps to mitigate this effect, the gap highlights the challenges posed by small sample sizes and ambiguous boundaries. Future work will focus on expanding the dataset, implementing advanced regularization strategies, and exploring more robust model architectures to further improve generalizability and reduce overfitting.

Computational demands: The training and optimization of deep learning models, especially with extensive hyperparameter tuning, required substantial computational resources. This may limit the immediate clinical deployment of such models in resource-constrained settings.

Model explainability and clinical interpretability: Although we employed class activation mapping techniques to visualize model attention and included attention-based architectures (e.g., PAN, MANet), these methods provide only a partial understanding of the model's decision-making process. Specifically, CAM highlights regions of interest but does not fully elucidate the complex, multi-layered reasoning underlying segmentation predictions, particularly for features with subtle or ambiguous boundaries. As a result, clinicians may find it challenging to validate or trust automated outputs in critical scenarios. Further research is needed to develop more transparent and interpretable deep learning frameworks, such as transformer-based models or inherently explainable architectures, that can offer mechanistic insights and foster clinical confidence.

5.5. Future research directions

The following directions outline potential pathways to advance the field of OCT processing and address the identified limitations:

Expansion and diversification of datasets: Efforts should focus on assembling larger, more diverse datasets encompassing a wide range of populations, imaging devices, and clinical scenarios. This will enhance model robustness, reduce bias, and improve generalizability.

Standardization and external validation: Adoption of standardized imaging protocols and external validation on independent, multi-national cohorts are essential to ensure reproducibility and facilitate clinical translation.

Integration of multimodal imaging: Combining OCT with complementary modalities such as intravascular ultrasound or near-infrared spectroscopy may provide a more comprehensive assessment of plaque morphology and vulnerability, enabling more accurate risk stratification.

Longitudinal studies: Prospective, longitudinal studies correlating automated plaque segmentation with clinical outcomes are needed to validate the prognostic value of these methods and to inform personalized therapeutic strategies.

Predictive risk assessment and clinical integration: Incorporating predictive analytics to evaluate the clinical significance of segmented features can enable the development of risk stratification tools that inform preventative care strategies. Integrating these predictive models and segmentation tools into existing clinical workflows through user-friendly interfaces and decision-support systems will facilitate seamless collaboration with cardiologists and interventionists, thereby enhancing the translational impact of automated plaque analysis.

6. Conclusion

This study tackled the challenge of automating the segmentation and quantification of atherosclerotic plaques in OCT images, a critical step in advancing cardiovascular diagnostics. Through optimized machine learning models, we demonstrated high accuracy in segmenting key plaque features, including the lumen (DSC: 0.987), fibrous cap (DSC: 0.736), lipid core (DSC: 0.751), and vasa vasorum (DSC: 0.610). These findings validate the potential of ML-enhanced OCT imaging to improve diagnostic precision and efficiency.

Our segmentation framework combined task-specific models – each tailored to class prevalence and morphological complexity – with rigorous hyperparameter optimization to achieve balanced, high-quality predictions across diverse plaque features. To further improve performance, we integrated these models into a weighted ensemble that incorporated class prevalence and model confidence, resulting in a harmonized system that achieved a weighted DSC of 0.882. This approach effectively mitigated class imbalance and leveraged architectural complementarity, outperforming individual models and naive aggregation strategies.

Clinically, our framework enables interpretable, spatially resolved, and quantifiable assessments of plaque features, facilitating early detection of vulnerable plaques, individualized risk stratification, and image-guided interventions. However, further prospective, multi-center validation is required to confirm its translational impact. Beyond cardiology, the methodology holds interdisciplinary relevance for other fields requiring high-resolution tissue imaging, such as oncology and neurology.

Future research will prioritize three directions: (1) expanding datasets to include ethnically diverse populations and rare plaque subtypes (2) integrating multimodal imaging (e.g., OCT-IVUS-NIRS fusion), for comprehensive plaque characterization, and (3) advancing predictive risk assessment and clinical integration by developing risk stratification tools and seamlessly incorporating automated segmentation into clinical workflows. Immediate next steps include deploying these models in prospective clinical trials to validate their prognostic value in longitudinal studies and to assess whether ML-driven plaque quantification improves outcomes in preventive percutaneous coronary interventions.

CRediT authorship contribution statement

Viacheslav V. Danilov: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Vladislav V. Laptev:** Visualization, Validation, Software, Methodology, Formal analysis, Data curation. **Kirill Yu. Klyshnikov:** Writing – review & editing, Visualization, Data curation. **Ivan S.**

Bessonov: Resources, Data curation. **Nikita V. Litvinyuk:** Resources, Data curation. **Evgeny A. Ovcharenko:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Conceptualization. **Nikita A. Kochergin:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Funding acquisition, Data curation, Conceptualization.

Data availability

The data supporting the key findings of this study are presented within the article and its supplementary information files. All essential components of the study, including curated source code, data, and trained models, have been made publicly available:

- **Source code:** https://github.com/ViacheslavDanilov/oct_seg_mentation
- **Dataset:** <http://doi.org/10.5281/zenodo.14478209>
- **Models:** <http://doi.org/10.5281/zenodo.14481678>

Ethics statement

All authors agree that:

- This research presents an accurate account of the work performed, all data presented are accurate and methodologies detailed enough to permit others to replicate the work.
- This manuscript represents entirely original works and or if work and/or words of others have been used, that this has been appropriately cited or quoted and permission has been obtained where necessary.
- This material has not been published in whole or in part elsewhere.
- The manuscript is not currently being considered for publication in another journal.
- That generative AI and AI-assisted technologies have not been utilized in the writing process or if used, disclosed in the manuscript the use of AI and AI-assisted technologies and a statement will appear in the published work.
- That generative AI and AI-assisted technologies have not been used to create or alter images unless specifically used as part of the research design where such use must be described in a reproducible manner in the methods section.
- All authors have been personally and actively involved in substantive work leading to the manuscript and will hold themselves jointly and individually responsible for its content.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This study was supported by the Russian Science Foundation under Grant No. 23-75-10009, titled "*Development of an automated deep learning system for detecting unstable plaques in optical coherence tomography*". For more information, please visit <https://rscf.ru/project/23-75-10009/>.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.combiomed.2025.111061>.

References

- [1] F. Zibaeenezjad, S.S. Mohammadi, M. Sayadi, F. Safari, M.J. Zibaeenezhad, Ten-year atherosclerosis cardiovascular disease (ASCVD) risk score and its components among an Iranian population: a cohort-based cross-sectional study, *BMC Cardiovasc. Disord.* 22 (2022) 162, <https://doi.org/10.1186/s12872-022-02601-0>.
- [2] M. Vaduganathan, G.A. Mensah, J.V. Turco, V. Fuster, G.A. Roth, The global burden of cardiovascular diseases and risk, *J. Am. Coll. Cardiol.* 80 (2022) 2361–2371, <https://doi.org/10.1016/j.jacc.2022.11.005>.
- [3] J. Sun, Y. Qiao, M. Zhao, C.G. Magnussen, B. Xi, Global, regional, and national burden of cardiovascular diseases in youths and young adults aged 15–39 years in 204 countries/territories, 1990–2019: a systematic analysis of global Burden of disease study 2019, *BMC Med.* 21 (2023) 222, <https://doi.org/10.1186/s12916-023-02925-4>.
- [4] M. Kowara, A. Cudnoch-Jedrzejewska, Different approaches in therapy aiming to stabilize an unstable atherosclerotic plaque, *Int. J. Mol. Sci.* 22 (2021) 4354, <https://doi.org/10.3390/ijms22094354>.
- [5] M. Araki, T. Yonetosu, O. Kurihara, A. Nakajima, H. Lee, T. Soeda, Y. Minami, I. McNulty, S. Uemura, T. Kakuta, I.-K. Jang, Predictors of rapid plaque progression, *JACC Cardiovasc Imaging* 14 (2021) 1628–1638, <https://doi.org/10.1016/j.jcmg.2020.08.014>.
- [6] D.J. Medina-Leyte, O. Zepeda-García, M. Domínguez-Pérez, A. González-Garrido, T. Villarreal-Molina, L. Jacobo-Albavera, Endothelial dysfunction, inflammation and coronary artery disease: potential biomarkers and promising therapeutic approaches, *Int. J. Mol. Sci.* 22 (2021) 3850, <https://doi.org/10.3390/ijms22083850>.
- [7] J.F. Bentzon, F. Otsuka, R. Virmani, E. Falk, Mechanisms of plaque formation and rupture, *Circ. Res.* 114 (2014) 1852–1866, <https://doi.org/10.1161/CIRCRESAHA.114.302721>.
- [8] P. Baruš, J. Modrzewski, K. Gumięzna, P. Dunaj, M. Głów, A. Bednarek, W. Wańka, T. Roleder, J. Kochman, M. Tomaniak, Comparative appraisal of intravascular ultrasound and optical coherence tomography in invasive coronary imaging: 2022 update, *J. Clin. Med.* 11 (2022) 4055, <https://doi.org/10.3390/jcm11144055>.
- [9] A. Zaghloul, C. Iorgoveanu, K. Balakumaran, D. V Balanescu, T. Donisan, Limitations of coronary computed tomography angiography in predicting acute coronary syndrome in a low to intermediate-risk patient with chest pain, *Cureus* 10 (2018), <https://doi.org/10.7759/cureus.2649>.
- [10] D. Chamié, J.R. Costa, L.P. Damiani, D. Siqueira, S. Braga, R. Costa, H. Seligman, F. Brito, G. Barreto, R. Staico, F. Feres, R. Petracó, A. Abizaid, Optical coherence tomography versus intravascular ultrasound and angiography to guide percutaneous coronary interventions, *Circ. Cardiovasc. Interv.* 14 (2021) E009452, <https://doi.org/10.1161/CIRCINTERVENTIONS.120.009452>.
- [11] H. Omori, H. Matsuo, S. Fujimoto, Y. Sobue, Y. Nozaki, G. Nakazawa, K. Takahashi, K. Osawa, R. Okubo, U. Kaneko, H. Sato, T. Kajiyama, T. Miyoshi, K. Ichikawa, M. Abe, T. Kitagawa, H. Ikenaga, M. Saji, N. Iguchi, T. Ijichi, H. Mikamo, A. Kurata, M. Moroi, R. Iijima, S. Malkasian, T. Crabtree, J.K. Min, J.P. Earls, R. Nakanishi, Determination of lipid-rich plaques by artificial intelligence-enabled quantitative computed tomography using near-infrared spectroscopy as reference, *Atherosclerosis* 386 (2023) 117363, <https://doi.org/10.1016/j.atherosclerosis.2023.117363>.
- [12] N. Shah, B. Ussen, M. Mahmoudi, Adjunctive intra-coronary imaging for the assessment of coronary artery disease, *JRSM Cardiovasc. Dis.* 5 (2016) 204800401665814, <https://doi.org/10.1177/2048004016658142>.
- [13] J. Lee, Y. Gharaibeh, V.N. Zimin, J.N. Kim, N.S. Hassani, L.A.P. Dallan, G.T. R. Pereira, M.H.E. Makhlof, A. Hoori, D.L. Wilson, Plaque characteristics derived from intravascular optical coherence tomography that predict cardiovascular death, *Bioengineering* 11 (2024) 843, <https://doi.org/10.3390/bioengineering11080843>.
- [14] M.C. Williams, J.R. Weir-McCall, L.A. Baldassarre, C.N. De Cecco, A.D. Choi, D. Dey, M.R. Dweck, I. Isgum, M. Kolossvary, J. Leipsic, A. Lin, M.T. Lu, M. Motwani, K. Nieman, L. Shaw, M. van Assen, E. Nicol, Artificial intelligence and machine learning for cardiovascular computed tomography (CCT): a white paper of the society of cardiovascular computed tomography (SCCT), *J. Cardiovasc. Comput. Tomogr.* 0 (2024), <https://doi.org/10.1016/j.jctc.2024.08.003>.
- [15] L.S. Athanasiou, C.V. Bourantas, G. Rigas, A.I. Sakellaris, T.P. Exarchos, P. K. Siogkas, A. Ricciardi, K.K. Naka, M.I. Papafakis, L.K. Michalis, F. Prati, D. I. Fotiadis, Methodology for fully automated segmentation and plaque characterization in intracoronary optical coherence tomography images, *J. Biomed. Opt.* 19 (2014) 026009, <https://doi.org/10.1117/1.JBO.19.2.026009>.
- [16] C. Kolluru, D. Prabhu, Y. Gharaibeh, H. Bezerra, G. Guagliumi, D. Wilson, Deep neural networks for A-line-based plaque classification in coronary intravascular optical coherence tomography images, *J. Med. Imaging* 5 (2018) 1, <https://doi.org/10.1117/1.JMI.5.4.044504>.
- [17] P. Shi, J. Xin, N. Zheng, A-line-based thin-cap fibroatheroma detection with multi-view IVOCT images using multi-task learning and contrastive learning, *J. Opt. Soc. Am. A* 39 (2022) 2298, <https://doi.org/10.1364/JOSAA.464303>.
- [18] J. Lee, J.N. Kim, L. Gomez-Perez, Y. Gharaibeh, I. Motairek, G.T.R. Pereira, V. N. Zimin, L.A.P. Dallan, A. Hoori, S. Al-Kindi, G. Guagliumi, H.G. Bezerra, D. L. Wilson, Automated segmentation of microvessels in intravascular OCT images using deep learning, *Bioengineering* 9 (2022) 648, <https://doi.org/10.3390/bioengineering9110648>.
- [19] M. Chu, H. Jia, J.L. Gutiérrez-Chico, A. Maehara, Z.A. Ali, X. Zeng, L. He, C. Zhao, M. Matsumura, P. Wu, M. Zeng, T. Kubo, B. Xu, L. Chen, B. Yu, G.S. Mintz, W. Wijns, N.R. Holm, S. Tu, Artificial intelligence and optical coherence tomography for the automatic characterisation of human atherosclerotic plaques, *EuroIntervention* 17 (2021) 41–50, <https://doi.org/10.4244/EIJ-D-20-01355>.
- [20] B. Wittmann, L. Glandorf, J.C. Paetzold, T. Amiranashvili, T. Wälchli, D. Razansky, B. Menze, Simulation-based segmentation of blood vessels in cerebral 3D OCTA images, in: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Springer, Cham, 2024, pp. 645–655, https://doi.org/10.1007/978-3-031-72111-3_61.
- [21] X. Wang, Y. Ma, X. Guo, Y. Zheng, J. Zhang, Y. Liu, Y. Zhao, A Hyperreflective Foci Segmentation Network for OCT Images with multi-dimensional Semantic Enhancement, Springer, Cham, 2024, pp. 645–655, https://doi.org/10.1007/978-3-031-72378-0_60.
- [22] J. Lee, Y. Gharaibeh, C. Kolluru, V.N. Zimin, L.A.P. Dallan, J.N. Kim, H.G. Bezerra, D.L. Wilson, Segmentation of coronary calcified plaque in intravascular OCT images using a two-step deep learning approach, *IEEE Access* 8 (2020) 225581–225593, <https://doi.org/10.1109/ACCESS.2020.3045285>.
- [23] L.S. Athanasiou, M.L. Olander, J.M. de la Torre Hernandez, E. Ben-Assa, E. R. Edelman, A deep learning approach to classify atherosclerosis using intracoronary optical coherence tomography, in: H.K. Hahn, K. Mori (Eds.), *Medical Imaging 2019: Computer-Aided Diagnosis*, 2019, p. 22, <https://doi.org/10.1117/12.2513078>. SPIE.
- [24] Supervisely, Supervisely computer vision platform, supervisely ecosystem. <https://supervisely.com>, 2023.
- [25] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. MICCAI 2015. Lecture Notes in Computer Science, Springer, 2015, pp. 234–241, https://doi.org/10.1007/978-3-319-24574-4_28.
- [26] Z. Zhou, M.M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, UNet++: a nested U-Net architecture for medical image segmentation, in: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Springer Verlag, 2018, pp. 3–11, https://doi.org/10.1007/978-3-030-00889-5_1.
- [27] L.-C. Chen, G. Papandreou, F. Schroff, H. Adam, Rethinking atrous convolution for semantic image segmentation, *ArXiv* (2017). <http://arxiv.org/abs/1706.05587>.
- [28] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science*, Springer, 2018, pp. 833–851, https://doi.org/10.1007/978-3-03-01234-2_49.
- [29] A. Kirillov, K. He, R. Girshick, P. Dollár, A unified architecture for instance and semantic segmentation. <http://presentations.cocodataset.org/COCO17-Stuff-FAIR.pdf>, 2017.
- [30] A. Chaurasia, E. Culurciello, LinkNet: exploiting encoder representations for efficient semantic segmentation, in: *2017 IEEE Visual Communications and Image Processing (VCIP)*, IEEE, 2017, pp. 1–4, <https://doi.org/10.1109/VCIP.2017.8305148>.
- [31] H. Zhao, J. Shi, X. Qi, X. Wang, J. Jia, Pyramid scene parsing network, in: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017, pp. 6230–6239, <https://doi.org/10.1109/CVPR.2017.660>.
- [32] H. Li, P. Xiong, J. An, L. Wang, Pyramid attention network for semantic segmentation, *ArXiv* (2018). <http://arxiv.org/abs/1805.10180>.
- [33] T. Fan, G. Wang, Y. Li, H. Wang, MA-Net: a multi-scale attention network for liver and tumor segmentation, *IEEE Access* 8 (2020) 179656–179665, <https://doi.org/10.1109/ACCESS.2020.3025372>.
- [34] V.V. Danilov, D. Litmanovich, A. Proutski, A. Kirpich, D. Nefaridze, A. Karpovsky, Y. Gankin, Automatic scoring of COVID-19 severity in X-ray imaging based on a novel deep learning workflow, *Sci. Rep.* 12 (2022) 12791, <https://doi.org/10.1038/s41598-022-15013-z>.
- [35] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A.L. Yuille, DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs, *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (2018) 834–848, <https://doi.org/10.1109/TPAMI.2017.2699184>.
- [36] J. Lee, G.T.R. Pereira, Y. Gharaibeh, C. Kolluru, V.N. Zimin, L.A.P. Dallan, J. N. Kim, A. Hoori, S.G. Al-Kindi, G. Guagliumi, H.G. Bezerra, D.L. Wilson, Automated analysis of fibrous cap in intravascular optical coherence tomography images of coronary arteries, *Sci. Rep.* 12 (2022) 21454, <https://doi.org/10.1038/s41598-022-24884-1>.
- [37] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, S. Belongie, Feature pyramid networks for object detection, in: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017, pp. 936–944, <https://doi.org/10.1109/CVPR.2017.106>.
- [38] B. Pu, Y. Lu, J. Chen, S. Li, N. Zhu, W. Wei, K. Li, MobileUNet-FPN: a semantic segmentation model for fetal ultrasound four-chamber segmentation in edge computing environments, *IEEE J. Biomed. Health Inform.* 26 (2022) 5540–5550, <https://doi.org/10.1109/JBHI.2022.3182722>.
- [39] V.V. Danilov, V.V. Laptev, K.Yu Klyshnikov, A.D. Stepanov, L.A. Bogdanov, L. V. Antonova, E.O. Krivkina, A.G. Kutikhin, E.A. Ovcharenko, ML-driven segmentation of microvascular features during histological examination of tissue-engineered vascular grafts, *Front. Bioeng. Biotechnol.* 12 (2024) 1411680, <https://doi.org/10.3389/fbioe.2024.1411680>.
- [40] A. Sulaiman, V. Anand, S. Gupta, M.S. Al Reshan, H. Alshahrani, A. Shaikh, M. A. Elmaghzoub, An intelligent LinkNet-34 model with EfficientNetB7 encoder for semantic segmentation of brain tumor, *Sci. Rep.* 14 (2024) 1345, <https://doi.org/10.1038/s41598-024-51472-2>.
- [41] K. Li, Z. Wu, K.-C. Peng, J. Ernst, Y. Fu, Tell me where to look: guided attention inference network, in: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, 2018, pp. 9215–9223, <https://doi.org/10.1109/CVPR.2018.00960>.
- [42] L. Dong, W. Lu, X. Lu, X. Leng, J. Xiang, C. Li, Comparison of deep learning-based image segmentation methods for intravascular ultrasound on retrospective and

- large image cohort study, *Biomed. Eng. Online* 22 (2023) 111, <https://doi.org/10.1186/s12938-023-01171-2>.
- [43] J. Tobin, Troubleshooting deep neural networks. <https://fullstackdeeplearning.com/spring2021/lecture-7/>, 2021.
- [44] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar, Hyperband: a novel bandit-based approach to hyperparameter optimization, *J. Mach. Learn. Res.* 18 (2016) 6765–6816. <http://arxiv.org/abs/1603.06560>.
- [45] S. Falkner, A. Klein, F. Hutter, BOHB: robust and efficient hyperparameter optimization at scale, *Proceedings of the 35th International Conference on Machine Learning* 4 (2018) 1437–1446. <http://arxiv.org/abs/1807.01774>.
- [46] A. Buslaev, V.I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, A.A. Kalinin, Albumentations: fast and flexible image augmentations, *Information* 11 (2020) 125, <https://doi.org/10.3390/info11020125>.
- [47] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-CAM: visual explanations from deep networks via gradient-based localization, in: 2017 IEEE International Conference on Computer Vision (ICCV), IEEE, 2017, pp. 618–626, <https://doi.org/10.1109/ICCV.2017.74>.
- [48] R.L. Draelos, L. Carin, Use HiResCAM instead of Grad-CAM for faithful explanations of convolutional neural networks, *ArXiv* (2020). <http://arxiv.org/abs/2011.08891>.
- [49] V. Pillai, H. Pirsiavash, Explainable models with consistent interpretations, *Proc. AAAI Conf. Artif. Intell.* 35 (2021) 2431–2439, <https://doi.org/10.1609/aaai.v35i3.16344>.
- [50] A. Chattopadhyay, A. Sarkar, P. Howlader, V.N. Balasubramanian, Grad-CAM++: generalized gradient-based visual explanations for deep convolutional networks, in: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE, 2018, pp. 839–847, <https://doi.org/10.1109/WACV.2018.00097>.
- [51] R. Fu, Q. Hu, X. Dong, Y. Guo, Y. Gao, B. Li, Axiom-based Grad-CAM: towards accurate visualization and explanation of CNNs, in: 31st British Machine Vision Conference, 2020, BMVC, 2020. <http://arxiv.org/abs/2008.02312>.
- [52] M.B. Muhammad, M. Yeasin, Eigen-CAM: class activation map using principal components, in: 2020 International Joint Conference on Neural Networks (IJCNN), IEEE, 2020, pp. 1–7, <https://doi.org/10.1109/IJCNN48605.2020.9206626>.
- [53] P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, Y. Wei, LayerCAM: exploring hierarchical class activation maps for localization, *IEEE Trans. Image Process.* 30 (2021) 5875–5888, <https://doi.org/10.1109/TIP.2021.3089943>.
- [54] S.S. Virani, L.K. Newby, S.V. Arnold, V. Bittner, L.C. Brewer, S.H. Demeter, D. L. Dixon, W.F. Fearon, B. Hess, H.M. Johnson, D.S. Kazi, D. Kolte, D.J. Kumbhani, J. LoFaso, D. Mahtta, D.B. Mark, M. Minissian, A.M. Navar, A.R. Patel, M.R. Piano, F. Rodriguez, A.W. Talbot, V.R. Taqueti, R.J. Thomas, S. van Diepen, B. Wiggins, M.S. Williams, AHA/ACC/ACCP/ASPC/NLA/PCNA guideline for the management of patients with chronic coronary disease: a report of the American heart association/american college of cardiology joint committee on clinical practice guidelines, *Circulation* 148 (2023) E9–E119, <https://doi.org/10.1161/CIR.0000000000001168>, 2023.
- [55] E. Kedhi, B. Berta, T. Roleder, R.S. Hermanides, E. Fabris, A.J.J. Ijsselmuiden, F. Kauer, F. Alfonso, C. von Birgelen, J. Escaned, C. Camaro, M.W. Kennedy, B. Pereira, M. Magro, H. Nef, S. Reith, A. Al Nooryani, F. Rivero, K. Malinowski, G. De Luca, H. Garcia Garcia, J.F. Granada, W. Wojakowski, Thin-cap fibroatheroma predicts clinical events in diabetic patients with normal fractional flow reserve: the COMBINE OCT–FFR trial, *Eur. Heart J.* 42 (2021) 4671–4679, <https://doi.org/10.1093/eurheartj/ehab433>.
- [56] F. Prati, E. Romagnoli, L. Gatto, A. La Manna, F. Burzotta, Y. Ozaki, V. Marco, A. Boi, M. Fineschi, F. Fabbrocchi, N. Taglieri, G. Niccoli, C. Trani, F. Versaci, G. Calligaris, G. Ruscica, A. Di Giorgio, R. Vergallo, M. Albertucci, G. Biondi-Zoccali, C. Tamburino, F. Crea, F. Alfonso, E. Arbustini, Relationship between coronary plaque morphology of the left anterior descending artery and 12 months clinical outcome: the CLIMA study, *Eur. Heart J.* 41 (2020) 383–391, <https://doi.org/10.1093/eurheartj/ehz520>.
- [57] N.A. Kochergin, A.M. Kochergina, V.I. Ganyukov, O.L. Barbarash, Vulnerable plaques in patients with stable coronary artery disease, in: E.H. Bennington (Ed.), *Horizons in World Cardiovascular Research*, Nova Science Publishers, 2021, pp. 187–203.
- [58] S.-J. Park, J.-M. Ahn, D.-Y. Kang, S.-C. Yun, Y.-K. Ahn, W.-J. Kim, C.-W. Nam, J.-O. Jeong, I.-H. Chae, H. Shiom, H.-L. Kao, J.-Y. Hahn, S.-H. Her, B.-K. Lee, T. H. Ahn, K.-Y. Chang, J.K. Chae, D. Smyth, G.S. Mintz, G.W. Stone, D.-W. Park, S.-J. Park, J.-M. Ahn, D.-Y. Kang, S.-C. Yun, Y.-K. Ahn, W.-J. Kim, C.-W. Nam, J.-O. Jeong, I.-H. Chae, H.-S. Shiom, H.-L. Kao, J.-Y. Hahn, S.-H. Her, B.-K. Lee, T. H. Ahn, K.-Y. Chang, J.-K. Chae, D. Smyth, G. Mintz, G. Stone, D.-W. Park, Preventive percutaneous coronary intervention versus optimal medical therapy alone for the treatment of vulnerable atherosclerotic coronary plaques (PREVENT): a multicentre, open-label, randomised controlled trial, *Lancet* 403 (2024) 1753–1765, [https://doi.org/10.1016/S0140-6736\(24\)00413-6](https://doi.org/10.1016/S0140-6736(24)00413-6).
- [59] A.S. Kini, Y. Vengrenyuk, T. Yoshimura, M. Matsumura, J. Pena, U. Baber, P. Moreno, R. Mehran, A. Maehara, S. Sharma, J. Narula, Fibrous cap thickness by optical coherence tomography in vivo, *J. Am. Coll. Cardiol.* 69 (2017) 644–657, <https://doi.org/10.1016/j.jacc.2016.10.028>.
- [60] T.J. Jun, S.-J. Kang, J.-G. Lee, J. Kweon, W. Na, D. Kang, D. Kim, D. Kim, Y.-H. Kim, Automated detection of vulnerable plaque in intravascular ultrasound images, *Med. Biol. Eng. Comput.* 57 (2019) 863–876, <https://doi.org/10.1007/s11517-018-1925-x>.
- [61] Y. Bae, S.-J. Kang, G. Kim, J.-G. Lee, H.-S. Min, H. Cho, D.-Y. Kang, P.H. Lee, J.-M. Ahn, D.-W. Park, S.-W. Lee, Y.-H. Kim, C.W. Lee, S.-W. Park, S.-J. Park, Prediction of coronary thin-cap fibroatheroma by intravascular ultrasound-based machine learning, *Atherosclerosis* 288 (2019) 168–174, <https://doi.org/10.1016/j.atherosclerosis.2019.04.228>.