MAP-CARE: PREDICTING FUTURE SURGICAL INTERVENTION IN ASYMPTOMATIC CAROTID STENOSIS VIA MULTI-MODAL RADIOMICS

SODAI YOKOYAMA[1], TAKASHI MIZOBE[2], HIDEO AIHARA[2], TOMOKAZU HAYASHI[3], TETSUYA URATA[3], SYOJI KOBASHI[1]

[1]Graduate School of Engineering, University of Hyogo, Hyogo, Japan
[2] Department of Neurosurgery, Hyogo Prefectural Harima-Himeji General Medical Center, Himeji, Hyogo, Japan
[3] Department of Radiology, Hyogo Prefectural Harima-Himeji General Medical Center, Himeji, Hyogo, Japan
E-MAIL: ei24n026@guh.u-hyogo.ac.jp, kobashi@eng.u-hyogo.ac.jp

Abstract:

This study proposes MAP-CARE, a radiomics-based framework for predicting future surgical intervention in patients with asymptomatic carotid artery plaques using multi-modal MRI. Unlike prior research that focuses on stroke risk prediction, MAP-CARE directly targets the clinical decision of whether carotid revascularization, such as endarterectomy or stenting, will be indicated within one year, even before the onset of neurological symptoms. Radiomic features are extracted from 3D Turbo Spin Echo (3D-TSE) and Time-of-Flight (TOF) MR sequences to capture both structural and hemodynamic characteristics of carotid plaques. Feature selection is performed using principal component analysis and mutual information. Four machine learning models (logistic regression, support vector machine, LightGBM, and random forest) are trained and evaluated using stratified five-fold cross-validation. A total of 36 carotid arteries from 18 patients were analyzed. The combination of mutual information-based selection and a random forest classifier yields the best performance, achieving an AUC of 0.933. To enhance model interpretability, SHapley Additive exPlanation (SHAP) is applied to identify important geometric and texture-based features. By bridging non-invasive imaging with real-world clinical outcomes, MAP-CARE provides an interpretable and actionable tool for proactive treatment planning in asymptomatic carotid stenosis.

Keywords:

Radiomics, Asymptomatic Carotid Plaque, Multi-Modal MRI, Machine Learning, Random Forest, SHAP

1. Introduction

Carotid atherosclerotic plaque is a major risk factor for ischemic stroke, and has been observed in up to 87.3% of individuals aged around 64 years in a population-based study [1]. Treatment options for carotid artery stenosis include medical therapy, such as antiplatelet agents, and surgical interventions such as carotid artery stenting (CAS) and carotid endarterectomy (CEA). The decision to perform surgical treatment is typically based on clinical factors such as the degree of stenosis and the symptomatic status [2].

Magnetic resonance angiography (MRA) has improved non-invasive visualization of carotid stenosis and plaque morphometry, offering high soft tissue contrast and allowing assessment of key features such as vessel wall structure, intraplaque hemorrhage, and ulceration. However, conventional morphological evaluation does not adequately capture the biological heterogeneity and dynamic progression of plaques crucial for cerebrovascular event risk, hampering differentiation between stable and unstable plaques and complicating treatment planning [3].

Radiomics extracts high-dimensional features from medical images, enabling reproducible characterization of plaque structure, texture, and intensity. By capturing subtle imaging patterns beyond human perception, radiomics enables objective assessment of important components such as fibrous cap integrity, intraplaque hemorrhage, and vessel wall structure.

Recent studies have applied radiomics and machine learning (ML) to predict stroke risk and classify plaque vulnerability. For example, Zhang et al. developed an MRI-based radiomics model to differentiate symptomatic from asymptomatic plaques [4], while Han et al. proposed a clinical-radiomics model for ischemic stroke prediction [5]. Although these studies demonstrate the utility of radiomics in event prediction, they do not directly support treatment decisions.

While stroke risk prediction is clinically valuable, treatment decisions also depend on imaging indicators such as plaque progression, restenosis, and new high-risk features. For example, surgery may be warranted even with low stroke risk if follow-up imaging reveals progressive luminal narrowing, intraplaque hemorrhage, or ulceration. The 2023 European Society for Vascular Surgery (ESVS) guideline [6] reflects this nuance, recommending that revascularization for asymptomatic 70-99% restenosis be considered only after multidisciplinary team (MDT) review, particularly when imaging is inconclusive. In such cases, clinical judgment is highly subjective and variable across institutions.

To address this gap, we propose a predictive framework for surgical intervention using multi-modal MRA-based radiomics and interpretable machine learning. Specifically, we combine 3D Turbo Spin Echo (TSE) MRA, which provides vessel wall structure, with Time-of-Flight (TOF) MRA, which captures hemodynamic properties. Radiomic features extracted from both modalities are integrated to capture plaque characteristics from complementary perspectives. Dimensionality reduction and feature selection are performed, and multiple machine learning classifiers are evaluated. To enhance interpretability, SHapley Additive exPlanations (SHAP) are used to quantify feature importance and support clinical reasoning.

We name this model MAP-CARE (Multi-modal Analysis for Prediction of Carotid Artery Revascularization Eligibility). To the best of our knowledge, this is the first interpretable radiomics-based approach designed to predict future surgical intervention in asymptomatic carotid artery disease. MAP-CARE goes beyond stroke risk estimation by enabling individualized treatment guidance and providing a novel tool for supporting proactive stroke prevention and objective surgical decision-making.

2. Datasets

This study analyzed carotid MRA data from 18 patients with asymptomatic internal carotid artery (ICA) stenosis, collected at Harima-Himeji General Medical Center (Hyogo, Japan). In total, 36 carotid arteries were included. The cohort consisted of 13 males and 5 females, with a mean age of 73.2 ± 7.6 years (n = 17; age data were missing for one patient).

Among the 36 arteries, 13 underwent CAS or CEA within one year and were labeled "surgical," while the remaining 23 were managed conservatively and labeled as "follow-up." These binary labels served as prediction targets.

All MRA scans were acquired between 2016 and 2020 using a 1.5T MRI scanner (Philips Medical Systems). Each patient underwent two sequences, 3D TSE (TR/TE = 400/12 ms, field of view = 150 × 150 mm^2, matrix = 320 × 320, slice thickness = 1.3 mm, in-plane resolution = 0.47 × 0.47 mm^2) and 3D TOF (TR/TE = 20/3.9 ms, field of view = 200 × 200 mm^2, matrix = 512 × 512, slice thickness = 2.0 mm, in-plane resolution = 0.39 × 0.39 mm^2). Despite minor variations in acquisition parameters, all images were resampled to an isotropic voxel size of 0.223 × 0.223 × 0.223 mm³ and resized to 224 × 224 × 112 voxels.

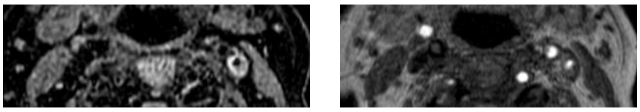

(a) 3D Turbo Spin Echo (b) 3D Time-of-Flight

FIGURE 1. Representative axial MRA images of the internal carotid artery. TSE provides high soft tissue contrast and enables clear delineation for the vessel wall and surrounding plaque structure, while TOF emphasizes intraluminal blood flow signals, supporting complementary hemodynamic assessment.

Manual annotations of the plaque and vessel wall regions were performed on the TSE images by trained researchers under the supervision of board-certified neurosurgeons. These binary masks were used to define the regions of interest for radiomic feature extraction. An example is shown in **FIGURE 1**.

This study was approved by the Ethics Committee of Harima-Himeji General Medical Center (Approval No. 2023-3), and the requirement for informed consent was waived due to the retrospective nature of the study.

3. MAP-CARE Framework

We propose MAP-CARE (*Multi-modal Analysis for Prediction of Carotid Artery Revascularization Eligibility*), a predictive framework designed to estimate whether carotid revascularization will be selected for asymptomatic patients within one year. The goal is to model real-world clinical decisions rather than disease status or stroke risk.

This framework uses two MRI sequences, 3D TSE and 3D TOF, which provide complementary anatomical and hemodynamic information. Fixed-size three-dimensional regions of interest (ROIs) are centered on the manually identified carotid bifurcation. From these regions, radiomic features are extracted and dimensionality is reduced before being used to train supervised classifiers. SHAP analysis is applied to enhance interpretability by quantifying feature contributions.

Unlike conventional models targeting objective outcomes, MAP-CARE is designed to predict treatment decisions that reflect complex clinical reasoning influenced by both imaging and non-imaging factors. This introduces challenges such as implicit label variability, potential label noise, and institutional bias, which require careful design of input features, model selection, and interpretability control.

To address these challenges, MAP-CARE employs standardized ROI definition to reduce input variability. Feature dimensionality is reduced to focus on informative

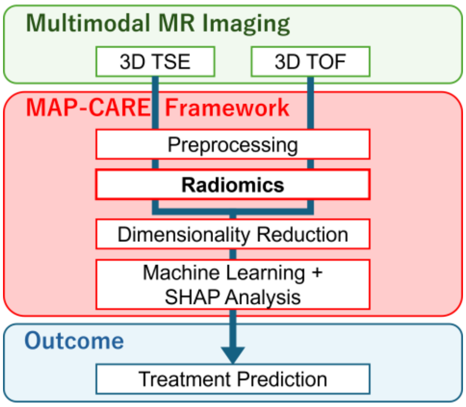

FIGURE 2. Overview of the MAP-CARE Framework.

predictors while minimizing overfitting. Instead of using opaque deep learning models, interpretable classifiers are trained, and SHAP analysis is used to verify the clinical plausibility of each prediction. This approach helps mitigate label noise, highlight feature stability, and reduce the impact of institutional bias.

MAP-CARE is structured to reflect guideline-based clinical workflows and supports interpretable prediction of treatment decisions. An overview of the framework is shown in **FIGURE 2**.

4. Implementation of MAP-CARE

This section elaborates on the technical implementation of the MAP-CARE framework, including preprocessing, radiomic feature extraction, model training, and interpretability analysis.

4.1. Image Standardization and ROI Definition

All 3D TSE and TOF MRA images are resampled to an isotropic voxel size of 0.223 × 0.223 × 0.223 mm³, using cubic interpolation to standardize spatial resolution across subjects and modalities. This resampling is applied prior to any other processing steps to ensure consistency during feature extraction.

Binary plaque masks are manually annotated on the resampled TSE images under the supervision of board-certified neurosurgeons, as TSE provides superior visualization of the vessel wall. For TOF images, the same ROI coordinates defined on the TSE images are applied directly, based on consistent spatial alignment between the modalities established during resampling. For each artery, the carotid bifurcation is manually identified as the anatomical reference point. A ROI measuring 50 mm × 50 mm × 25 mm in physical space is extracted, centered on the bifurcation in the axial plane and spanning 4 mm superior and inferior along the body axis. This definition ensures that both the plaque and adjacent vessel wall consistently captured across patients.

TABLE 1. Categories of radiomic features (Total = 92 features).

Category	Count	Description
First-Order Statistics	19	Intensity-based metrics describing voxel value distribution
Shape-based (3D)	17	Geometric descriptors of the plaque and surrounding vessel wall
Gray Level Co-occurrence Matrix (GLCM)	24	Texture features based on gray-level spatial co-occurrence
Gray Level Size Zone Matrix (GLSZM)	16	Measures of size-based intensity homogeneity
Gray Level Run Length Matrix (GLRLM)	16	Describes directional uniformity and run length of similar intensities

To standardize input dimensions for downstream analysis, the extracted ROIs are resized or zero-padded to a voxel grid of 224 × 224 × 112. This procedure is uniformly applied to both TSE and TOF modalities to ensure spatial correspondence between structural and hemodynamic information.

4.2. Radiomic Feature Extraction from Multi-Modal MRI

Radiomic features are extracted from standardized ROIs of both TSE and TOF images to quantitatively characterize plaque properties from structural and hemodynamic perspectives. A total of 92 features are computed for each modality, resulting in 184 features per artery when combining the two.

The extracted features cover five major categories commonly used in radiomic analysis [7], as summarized in **TABLE 1**.

All features are calculated independently for the TSE and TOF ROIs using the same physical ROI size and voxel resolution. Voxel intensities are discretized with a fixed bin width to standardize texture computation across cases. No filtering or wavelet-transformed features are included in this study, in order to focus on features derived from the original image intensities. Shape features are computed directly from binary masks on resampled images without sub-voxel or anti-aliasing interpolation.

The combined feature vector serves as the input for the following feature selection and classification stages.

4.3. Feature Selection and Dimensionality Reduction

To reduce the risk of overfitting and enhance generalization performance, two feature reduction strategies are applied and compared: PCA and MI-based selection.

PCA is an unsupervised linear transformation that projects the original feature space onto orthogonal components while preserving the majority of variance. In this study, components retaining 95% of the cumulative variance are selected to represent the reduced feature set.

MI-based feature selection is a supervised method that quantifies the statistical dependency between individual features and the binary target variable (surgical vs. follow-up). Features are ranked based on MI scores, and feature subsets are incrementally constructed by adding one feature at a time in descending order of importance. Among these subsets, the one achieving the highest cross-validated AUC is selected as the final input.

Both dimensionality reduction approaches are applied independently to the 184-dimensional feature vector, which combines features from TSE and TOF images. The effectiveness of each method is evaluated in the subsequent classification step.

To ensure reproducibility, MI computation is performed using discretized feature values and stratified cross-validation to estimate mutual dependency reliably. PCA is applied after z-score normalization of all features to equalize scaling across modalities and categories.

4.4. Classification for Surgical intervention Prediction

To predict whether each carotid artery would undergo surgical intervention within one year, supervised classification models are trained on the reduced feature sets obtained from the previous step. In this study, four widely used machine learning algorithms are evaluated for binary classification: logistic regression, support vector machine (SVM), light gradient boosting machine (LightGBM), and random forest. These classifiers represent diverse modeling paradigms, ranging from linear models to tree-based ensemble methods.

All models are trained and evaluated using stratified 5-fold cross-validation, ensuring that the proportion of surgical and follow-up cases is preserved within each fold. This strategy mitigates the effects of class imbalance and provides robust performance estimates. Each classifier is applied to the same input feature sets derived via either PCA or MI-based selection, allowing for fair comparison across algorithms and dimensionality reduction methods.

To ensure model transparency and support clinical interpretability, SHAP values are applied to the trained models, providing both global insight into feature importance and local explanations for individual predictions. This interpretability enables clinicians to understand which factors influenced each decision.

5. Experimental Results

5.1. Evaluation Protocol

All experiments were conducted on Ubuntu 24.04 with Python 3.9. Radiomic feature extraction was performed using PyRadiomics (version 3.1.0) with a fixed bin width of 25. Treatment prediction and evaluation were implemented using Scikit-learn (version 1.6.1) for logistic regression, SVM, and random forest, and LightGBM (version 4.6.0) for gradient boosting.

All classifiers were trained after hyperparameter tuning using grid search. Logistic regression used L1 regularization, SVM employed an RBF kernel, and both random forest and LightGBM were trained with 300 estimators. The learning rate for LightGBM was set to 0.01.

Model performance was assessed using four standard metrics: AUC, accuracy, sensitivity, and specificity. These metrics were computed on the test set of each fold, and the average values across all five folds were reported.

5.2. Classification Performance

Model performance was compared across four classifiers and two feature selection methods. **TABLE 2** summarizes the average results of AUC, precision, recall, and F-value from stratified 5-fold cross-validation. Each classifier was trained using the features selected via either PCA or MI. This enables a fair comparison across different model architectures and selection strategies.

In MI-based selection, features were incrementally added in descending order of importance within each fold. The optimal number of features and the specific features selected varied depending on the classifier and fold. For example, random forest achieved the highest AUC using 11 features on average, while other classifiers (such as SVM and LightGBM) performed best with different subset of features, depending on fold.

Among the evaluated models, random forest combined with MI-based feature selection achieved the highest AUC of 0.933, along with the highest recall (0.923) and F-value (0.889), indicating the most effective discrimination between surgical and follow-up cases. LightGBM with MI also showed strong performance (AUC = 0.893), while logistic regression and SVM performed slightly lower (AUC = 0.876 and 0.856, respectively). In contrast, all

TABLE 2. Performance comparison of classifiers with PCA and MI

Feature Selection	Model	AUC	Precision	Recall	F-value
PCA	Logistic Regression	0.823	0.688	0.846	0.759
	SVM	0.712	0.833	0.385	0.526
	LightGBM	0.562	0.423	0.846	0.564
	Random Forest	0.605	0.500	0.385	0.435
MI	Logistic Regression	0.876	0.818	0.692	0.750
	SVM	0.856	0.750	0.692	0.720
	LightGBM	0.893	0.833	0.769	0.800
	Random Forest	0.933	0.857	0.923	0.889

models performed worse under PCA-based selection, with AUC values ranging from 0.562 to 0.823.

In addition, **FIGURE 3** compares ROC curves for each model under PCA and MI conditions, based on the integrated results from five-fold stratified cross-validation. These ROC curves represent the results of one fold, and the overall performance is summarized in Table 2. The improvement in performance due to MI-based selection is particularly notable in tree-based models such as lightGBM and random forest. This suggests that MI better captures class-relevant features, leading to enhanced predictive accuracy in models capable of handling non-linear relationships.

Overall, models using MI-selected features outperformed those using PCA in most cases, suggesting that supervised selection more effectively captured the imaging features relevant to treatment decision-making. These results highlight the advantage of ensemble classifiers and MI-based feature selection in predicting surgical intervention from asymptomatic carotid plaque characteristics.

5.3. Feature Importance Analysis Using SHAP

To interpret the results of the classification models, SHAP values were calculated to evaluate the contribution of each feature to the model's decision-making process. **FIGURE 4.** illustrates the global feature importance for the MI-based random forest model, based on the performance of a single fold. The plot shows the SHAP values for the top features, highlighting those with the greatest impact on the model's predictions. In this case, shape-based features, particularly Sphericity in TSE and Compactness in TOF, were found to be the most important predictors of surgical intervention.

Other relevant features included Surface Area to

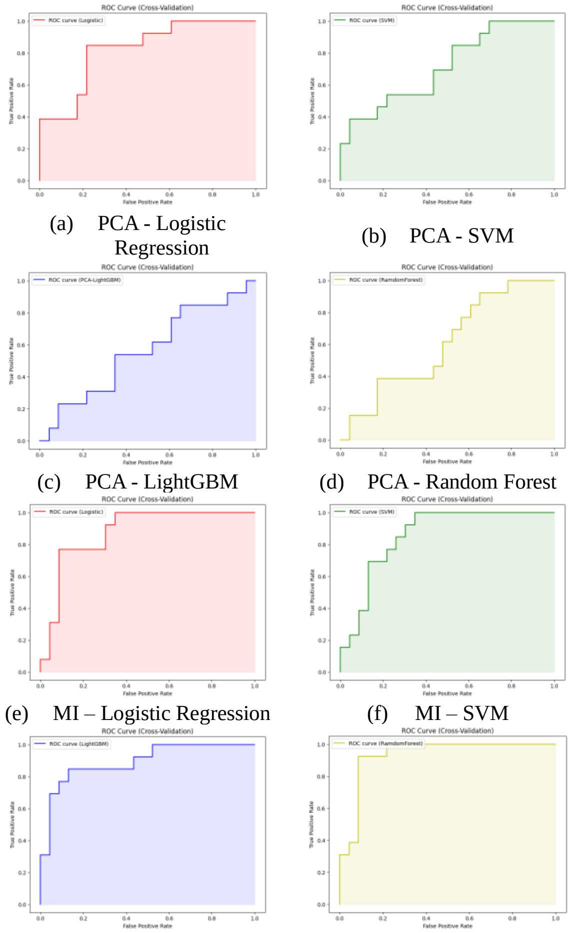

(a) PCA - Logistic Regression (b) PCA - SVM (c) PCA - LightGBM (d) PCA - Random Forest (e) MI – Logistic Regression (f) MI – SVM (g) MI – LightGBM (h) MI – Random Forest

FIGURE 3. ROC values for each classifier under two feature selection methods. (a–d) show results using PCA, and (e–h) show results using MI.

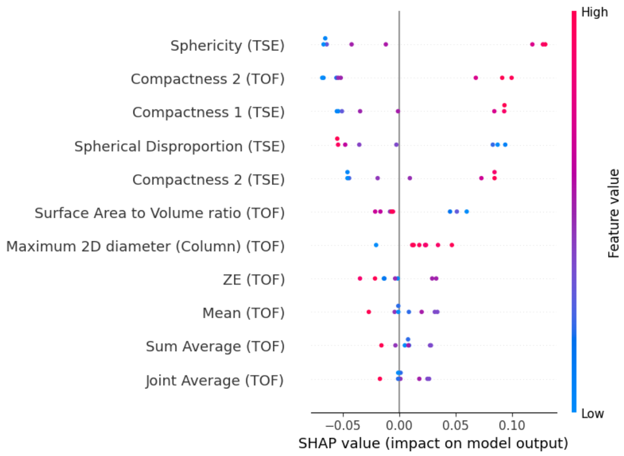

FIGURE 4. Global feature importance for random forest with MI-based feature selection. The bar plot shows SHAP values for the top features, highlighting those with the greatest impact on the model's predictions. Results are based on the performance of a single gold.

Volume ratio, Maximum 2D diameter, and intensity-based texture features from TOF images, such as Zone Entropy (ZE), Mean, and Joint Average. These results suggest that both geometric and intensity-related heterogeneity contribute significantly to surgical intervention prediction.

6. Discussion

In this study, we proposed MAP-CARE, a novel framework for predicting future surgical intervention in patients with asymptomatic carotid artery stenosis. MAP-CARE leverages multi-modal MRI, specifically 3D-TSE and TOF sequences, and machine learning techniques to support image-based clinical decision-making, particularly in determining surgical eligibility before symptom onset.

The effectiveness of the framework was demonstrated through both predictive performance and interpretability. In stratified 5-fold cross-validation, the combination of mutual information-based feature selection and a random forest classifier achieved the highest performance with an AUC of 0.933, outperforming PCA-based dimensionality reduction methods used in this study as a baseline. The advantage of MI was particularly evident in non-linear models such as LightGBM and random forest, likely because MI can prioritize features that are more directly related to class separation.

SHAP-based analysis confirmed that both morphological and intensity-based features contribute meaningfully to the prediction of surgical intervention. This suggests that structural complexity of the plaque, including shape irregularities and signal heterogeneity, may influence the clinician's decision to proceed with revascularization, even in asymptomatic cases.

Several limitations of the present study should be acknowledged. First, the dataset was relatively small and derived from a single institution, which may limit the generalizability of the results. Future studies should expand the dataset and validate performance across multiple clinical sites to assess reproducibility and robustness. Second, the plaque annotations were performed by a single reader, and inter-observer reproducibility was not assessed. This limitation must be addressed in subsequent work to ensure the reliability of the ground truth.

Looking ahead, a key next step is the implementation of MAP-CARE in real-world clinical settings. Its utility as a clinical decision support system (CDSS) should be prospectively evaluated. It will be important to examine how well the model aligns with physician decision-making and how it can contribute to optimizing resource allocation and patient outcomes.

7. Conclusion

We proposed MAP-CARE, a radiomics-based framework to predict future surgical intervention in asymptomatic carotid stenosis using multi-modal MRI and machine learning. Through rigorous evaluation, the framework demonstrated strong predictive performance, particularly when mutual information-based feature selection was combined with non-linear models, and it offered clinical interpretability via SHAP analysis. These findings suggest that radiomic features reflecting both morphological and intensity-based characteristics can support objective decision-making for surgical intervention, highlighting MAP-CARE as a clinically useful tool for proactive stroke prevention in asymptomatic populations.

References

[1] H. Ihle-Hansen *et al.*, "Carotid Plaque Score for Stroke and Cardiovascular Risk Prediction in a Middle Aged Cohort From the General Population," J Am Heart Assoc., vol. 12, no. 15, 2023.

[2] N. Azuma *et al.*, "2022 Guideline on the Management of Peripheral Arterial Disease," JCS/JSVS, 2022.

[3] Hafiane A. "Vulnerable Plaque, Characteristics, Detection, and Potential Therapies.," Journal of Cardiovascular Development and Disease, 2019.

[4] R. Zhang, Q. Zhang, A. Ji, P. Lv, J. Zhang, C. Fu, and J. Lin, "Identification of high-risk carotid plaque with MRI-based radiomics and machine learning," European Radiology, Vol. 31, No. 5, pp. 3116-3126, 2021.

[5] N. Han, W. Hu, Y. Ma, Y. Zheng, S. Yue, L. Ma, J. Li and J. Zhang, "A clinical-radiomics combined model based on carotid atherosclerotic plaque for prediction of ischemic stroke," Frontiers in Neurology, Vol. 15, No. 1343423, 2024.

[6] A. R. Naylor, J. B. Ricco, G. J. de Borst, S. Debus, J. de Haro, A. Halliday, et al., "Management of atherosclerotic carotid and vertebral artery disease: 2023 Clinical Practice Guidelines of the European Society for Vascular Surgery (ESVS)," European Journal of Vascular and Endovascular Surgery, vol. 65, no. 3, pp. 267–369, 2023.

[7] P. Lambin, E. Rios-Velazquez, R. Leijenaar, *et al.*, "Radiomics: the bridge between medical imaging and personalized medicine," *Nat. Rev. Clin. Oncol.*, vol. 17, no. 12, pp. 749–762, 2020.