Geometry-Aware Tumor Classification via Mean Curvature Flow on 3D Surface Meshes

Mehdi Miraç Arat[1], Bahar Karakaş[1], Latif Şimşek[1], Demet Yüksel Dal[2], Burcu Bektaş Demirci[1], Zeki Kuş[3]

[1] Department of Software Engineering, Fatih Sultan Mehmet Vakif University, İstanbul, Türkiye
[2] Department of Electrical Electronics Engineering, Fatih Sultan Mehmet Vakif University, İstanbul, Türkiye
[3] Department of Artificial Intelligence and Data Engineering, Fatih Sultan Mehmet Vakıf University, İstanbul, Türkiye
Email: mehdimirac.arat@stu.fsm.edu.tr, bahar.karakas@stu.fsm.edu.tr, latif.simsek@stu.fsm.edu.tr, dyukseldal@fsm.edu.tr, bbektas@fsm.edu.tr, zkus@fsm.edu.tr

2026 11th International Conference on Machine Learning Technologies (ICMLT) | 979-8-3315-4865-0/26/$31.00 ©2026 IEEE | DOI: 10.1109/ICMLT69916.2026.11689262

***Abstract*—This study presents a geometry-aware framework for binary tumor classification based on mean curvature flow (MCF) applied to reconstructed 3D surface meshes. Segmented medical image slices are converted into surface meshes and evolved under MCF. During the evolution, geometric descriptors, including quasi-conformal ratio, radial variance, and evolution speed, are extracted and concatenated into feature vectors. These descriptors are used with standard vector-based machine learning models for classification. Experimental results indicate that the proposed pipeline achieves competitive performance compared to a 3D CNN while providing practical advantages in terms of efficiency and structural interpretability.**

***Index Terms*—Mean Curvature Flow, tumor classification, geometric feature extraction, surface mesh analysis, medical image analysis, machine learning.**

I. Introduction

Mean Curvature Flow (MCF) is a fundamental geometric surface evolution process in which a surface is deformed in the direction of its mean curvature vector. Owing to its intrinsic formulation and solid theoretical foundations, MCF has been extensively studied in differential geometry and has found broad use in computer graphics and geometry processing, particularly for surface smoothing, fairing, and shape regularization, see [1], [2]. By progressively reducing local geometric irregularities, MCF provides a principled mechanism for analyzing the evolution of surface geometry.

In medical imaging, three-dimensional anatomical structures are often reconstructed from volumetric data and represented as triangulated surface meshes. These reconstructed surfaces may contain discretization artifacts, segmentation noise, and irregular boundary patterns that can affect subsequent quantitative analysis. Tumor surfaces, in particular, frequently exhibit complex and highly irregular geometries caused by both biological variability and image acquisition limitations. Geometric flow–based surface processing therefore offers a natural tool for denoising and stabilizing such surfaces while preserving their intrinsic geometric structure.

The geometric characteristics of tumor surfaces are known to provide clinically relevant information for distinguishing between benign and malignant lesions. Benign tumors are commonly associated with smoother and more regular surface structures, whereas malignant tumors tend to exhibit higher surface irregularity, spiculations, and complex local shape variations caused by infiltrative growth patterns. Such differences are reflected in the local curvature distribution of the surface as well as in its deformation behaviour under geometric evolution. By evolving a tumor surface under MCF, local geometric irregularities and protrusions are attenuated at different rates depending on their curvature magnitude. As a result, the temporal evolution process reveals characteristic deformation patterns that differ between benign and malignant geometries and provide informative cues about surface complexity and structural irregularity.

Despite substantial progress in machine learning, tumor classification remains a challenging problem in medical image analysis. Recent advances in machine learning, particularly deep convolutional neural networks (CNNs), have substantially improved both segmentation and classification performance across different breast imaging modalities. Large-scale datasets and learning-based pipelines, such as the BreastDM dataset introduced by Zhao et al. [3], have enabled the development of robust data-driven approaches for breast tumor analysis. Also, Yuhandri et al. [4] demonstrate that combining feature selection techniques with ensemble classifiers can significantly improve classification accuracy in brain tumor analysis. Recent studies further demonstrate the effectiveness of attention mechanisms and hybrid deep architectures for breast tumor segmentation and classification, including explainable attention-based multi-backbone networks [5], hybrid local–global feature fusion models [6], and architectures combining data augmentation strategies with modern CNN backbones such as ConvNeXt and U-Net variants [7]. In addition, large-scale and publicly available breast MRI datasets and acquisition protocols, such as FastMRI Breast [8], support the growing interest in data-driven learning pipelines for breast tumor analysis. Despite these advances, most existing learning-based methods rely primarily on intensity-driven or voxel-based representations, while the geometric characteristics of tumor morphology and surface structure are typically not explicitly incorporated into the learning process.

A further limitation of conventional preprocessing pipelines is that geometric information may be partially lost during operations such as resampling, volumetric normalization, or

Authorized licensed use limited to: FATIH SULTAN MEHMET UNIVERSITY. Downloaded on October 27,2026 at 13:43:31 UTC from IEEE Xplore. Restrictions apply.

aggressive smoothing. This motivates the need for geometry-preserving surface processing strategies in three-dimensional medical imaging. Although geometric flows have been widely investigated in computer graphics and shape analysis, their systematic use as a basis for extracting geometry-driven descriptors for medical machine learning pipelines remains limited.

In this work, we study a geometry-aware tumor classification pipeline based on MCF and surface-based feature extraction. From reconstructed three-dimensional tumor meshes, each surface is evolved under classical MCF using identical numerical parameters. From the resulting surface evolution, temporal and geometric descriptors are computed to characterize how the tumor surface deforms over time. Instead of relying solely on static geometric representations, the proposed framework exploits the temporal behaviour of the surface under MCF in order to capture dynamic geometric signatures of tumor morphology. These features are then employed for binary tumor classification using standard machine learning models. The presented study evaluates the contribution of MCF-driven geometric information to tumor classification performance and investigates its role as a complementary representation to conventional appearance- based and voxel-based image features.

II. Methodology

In this section, we briefly review the continuous formulation and finite elements discretization of MCF for details see [9] and [1]. Also, we mention discriminative shape-based descriptor related to geometric surface evolution. The overall flow of the work is shown in Figure 1.

A. Mean Curvature Flow-Continuous case

Let M be a two dimensinal manifold, let $\varphi_t : M \to \mathbb{R}^3$ be a smooth family of immersion, and let g_t be the metric induced by the immersion at time t. We say that ϕ_t is a solution to the mean curvature flow (MCF) if

$$\frac{\partial \varphi_t}{\partial t} = \Delta_t \varphi_t \tag{1}$$

where Δ_t is the Laplace-Beltrami operator defined with respect to the metric g_t.

The relation between the mean curvature and MCF comes from the well-known identity for the Laplace-Beltrami operator $\Delta\varphi = -2\vec{H}$ where $\vec{H}$ is the mean curvature vector.

B. Discretization of MCF

The continuous MCF equation in (1) can be discretized on a triangular mesh by the discrete Laplace–Beltrami operator constructed using the Cotangent Formula. For a triangular mesh with vertices $\{v_i\}_{i=1}^m$, the discrete Laplacian acting on the embedding φ can be written as

$$\Delta\varphi = DM\varphi,$$

where $D = \mathrm{diag}(w_1, \ldots, w_m)$ is a diagonal matrix with $w_i = \frac{1}{2A_i}$ for the area of the local vertex cell A_i and M is the symmetric cotangent stiffness matrix defined by

$$M_{ij} = \begin{cases} -\sum_{k \in N_1(i)} w_{ik}, & i = j, \\ w_{ij}, & v_j \in N_1(i), \\ 0, & \text{otherwise}, \end{cases}$$

with edge weights $w_{ij} = \cot\alpha_{ij} + \cot\beta_{ij}$. The surface embedding is updated at each time step by solving

$$\left(D_t^{-1} - h\, M_t\right)\, \varphi_{t+h} = D_t^{-1}\, \varphi_t, \tag{2}$$

where h is the time step size. This formulation leads to a sparse linear system that is solved independently for each coordinate dimension and provides improved numerical stability compared to explicit integration schemes.

C. Geometric Feature Extraction

Geometric features are extracted from tumor surface meshes during the evolution process driven by MCF. At each iteration, a set of complementary descriptors is computed in order to characterize local deformation behavior and global shape irregularity of the evolving surface.

Quasi-Conformal Ratio (QC). The quasi-conformal ratio measures local distortion introduced by the surface deformation. For each triangular face, a Jacobian matrix is estimated between the reference mesh and the current mesh configuration. Singular value decomposition of the Jacobian yields the principal stretching factors, and the ratio between the maximum and minimum singular values is computed. The global QC value is obtained as an area-weighted average of per-face ratios, providing a robust indicator of conformality preservation across the surface.

Radial Variance. Radial variance quantifies global shape irregularity by measuring the dispersion of vertex distances from the initial surface centroid. Let V denote the current vertex set and c_0 the centroid of the initial mesh. The Euclidean distances $\|V_i - c_0\|$ are computed for all vertices, and their variance is used as a scalar descriptor. Higher radial variance indicates increased deviation from spherical symmetry.

Evolution Speed. Evolution speed captures the magnitude of surface displacement between consecutive iterations. It is computed as the ℓ_2 norm of the difference between the current vertex positions and those of the previous iteration, providing a measure of deformation rate over time.

III. Experimental Design

A. Datasets

The experiments are conducted using the publicly available BreastDM dataset introduced by Zhao et al., which provides breast DCE-MRI data together with expert tumor segmentation masks [3]. All DCE-MRI examinations are acquired on 3.0-T systems (GE SIGNA HDx and Siemens Skyra, Erlangen, Germany) using a 20-channel head-and-neck coil, and the dataset includes pre-contrast, post-contrast, and subtraction representations. In this study, only the segmentation masks

Authorized licensed use limited to: FATIH SULTAN MEHMET UNIVERSITY. Downloaded on October 27,2026 at 13:43:31 UTC from IEEE Xplore. Restrictions apply.

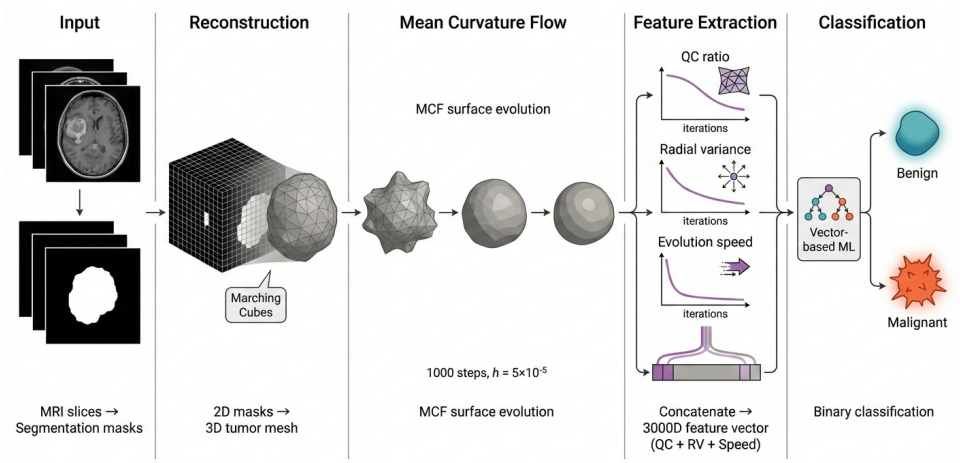

Fig. 1. Schematic overview of the proposed geometry-aware tumor classification framework. The pipeline processes input MRI data in four stages: (1) Reconstruction: 2D MRI slices and binary segmentation masks are converted into a 3D triangular surface mesh using the Marching Cubes algorithm. (2) Surface Evolution: The mesh undergoes Mean Curvature Flow (MCF) for 1000 iterations (step size $h = 5 \times 10^{-5}$), smoothing the geometry from an initial irregular state toward a spherical limit. (3) Feature Extraction: Dynamic geometric descriptors—Quasi-Conformal (QC) ratio, Radial Variance, and Evolution Speed—are tracked across the evolution trajectory and concatenated into a comprehensive 3000D feature vector. (4) Classification: A vector-based machine learning model analyzes these evolutionary features to discriminate between Benign (smooth) and Malignant (spiculated) tumor morphologies.

are used: mask slices are stacked to reconstruct a 3D tumor volume from which a surface mesh is extracted. After preprocessing and consistency checks, 166 cases with successfully reconstructed tumor meshes (one mesh per case) are retained and split at the case level into 132 training and 34 test samples to prevent subject-level information leakage.

B. 3D Medical Image to Mesh Conversion

In this study, two-dimensional medical image slices are converted into three-dimensional surface meshes to enable geometric modeling and shape-based analysis of tumor structures. The input data consist of binary segmentation masks obtained from magnetic resonance imaging (MRI), where pixels corresponding to tumor tissue are labeled as foreground and non-tumor regions are labeled as background.

For each tumor case, the dataset is organized into multiple subdirectories representing different acquisition groups. Specifically, each tumor contains six subfolders (SUB1–SUB6), and each subfolder includes six consecutive segmentation mask images. These images correspond to neighboring cross-sectional slices that capture the spatial extent of the tumor region. All slices from the subfolders are utilized and combined to form a single volumetric representation for each tumor. An overview of the slice-to-mesh conversion process is illustrated in Figs. 2 and 3.

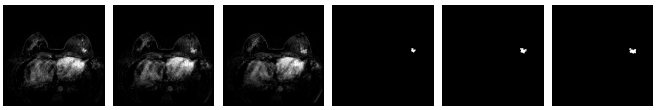

Fig. 2. Three representative consecutive MRI image slices and their corresponding binary segmentation masks illustrating neighboring cross-sectional tumor regions

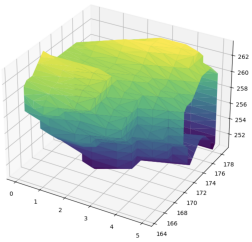

Fig. 3. Reconstructed 3D tumor surface mesh generated from stacked segmentation masks using the Marching Cubes algorithm.

All segmentation masks belonging to a tumor are first loaded and sorted according to their anatomical order. Each slice is converted into a binary image and stacked along the axial direction to construct a three-dimensional voxel-based volume. This stacking process produces a volumetric tensor of size (Z, H, W), where Z denotes the number of slices and H and W represent the in-plane spatial resolution of each slice.

Because the number of slices in medical image volumes is typically limited, the constructed volume may suffer from anisotropic resolution along the axial direction. To mitigate this issue, interpolation is applied along the slice axis in order to increase sampling density and reduce discretization artifacts. This step improves surface continuity and leads to smoother and more stable surface reconstructions.

After volume construction, an isosurface extraction procedure is performed using the Marching Cubes algorithm. Given a predefined iso-value threshold, the algorithm traverses the volumetric grid and generates a triangular surface mesh that approximates the boundary of the tumor region. The output of this process consists of a set of vertices and triangular faces

Authorized licensed use limited to: FATIH SULTAN MEHMET UNIVERSITY. Downloaded on October 27,2026 at 13:43:31 UTC from IEEE Xplore. Restrictions apply.

that describe the tumor surface geometry.

The extracted surface mesh is subsequently exported in standard polygonal mesh format (OBJ) and serves as the initial geometric representation for subsequent processing stages, including curvature-driven surface evolution and geometric feature extraction. This conversion pipeline establishes a direct and effective connection between medical image segmentation data and computational geometry-based analysis.

C. Implementation Details

All surface evolutions are performed on triangular tumor meshes using identical discretization and numerical settings. For each mesh, the surface is evolved for a fixed number of iterations (1000 time steps), and geometric descriptors are recorded at every iteration. The fixed time step size is set to $h = 5 \times 10^{-5}$. No early stopping is applied, and the full evolution trajectory is retained to enable consistent temporal feature extraction. From each evolved surface sequence, three groups of time-dependent geometric features are computed over 1000 steps: radial variance, quasi-conformal distortion ratios, and evolution speed. These features are concatenated into a single representation for each tumor sample, resulting in a 3000-dimensional input feature vector.

Several supervised machine learning models—Logistic Regression, Support Vector Machines, Multi-Layer Perceptron, and tree- and boosting-based methods—are evaluated for geometric-feature-based classification. Default parameters are used without hyperparameter tuning or data augmentation. The dataset is split 80:20 with stratification for training and testing, consistently across experiments. A lightweight 3D CNN serves as a volumetric baseline. It comprises five 3D convolution layers with ReLU activations, increasing channels from 16 to 32 and 64, with resolution reduced twice via strided convolutions. Global features are extracted with adaptive average pooling, then mapped to benign and malignant classes through a fully connected layer. The network is trained with cross-entropy loss and Adam optimizer at a learning rate of 0.001, for 10 epochs with a batch size of 2.

All experiments are implemented in Python using standard scientific and machine learning libraries, including NumPy, Pandas, scikit-learn, and XGBoost. The geometric processing and feature extraction pipeline is implemented using custom Python code for surface evolution and metric computation. Experiments are executed on a workstation equipped with an Apple Silicon processor, utilizing the Metal Performance Shaders (MPS) backend where applicable.

D. Evaluation Metrics

All classification models are evaluated on a separate test set not used during training. Multiple metrics evaluate tumor classification quality, including accuracy, precision, recall, and F1-score. Accuracy indicates the overall correct classification rate; precision measures the reliability of malignant predictions; recall reflects the model's ability to identify malignant tumors; and F1-score balances precision and recall, making it suitable for imbalanced datasets. Consequently, the F1-score is the primary performance metric, with both macro- and weighted averages reported to account for class-wise performance and frequencies. For models with probabilistic outputs, the ROC-AUC score is also provided to evaluate the ranking quality of malignancy predictions independently of thresholds.

IV. Result and Discussion

This section evaluates the effectiveness of MCF based geometric feature representations for binary tumor classification. For all conventional machine learning models, the input consists of the concatenated feature vector obtained from the MCF-driven surface evolution, namely the quasi-conformal ratio (QC), radial variance, and evolution speed time series. In contrast, the 3D CNN baseline is trained directly on the original volumetric medical image slices, without using the mesh-based geometric pipeline.

TABLE I
BINARY TUMOR CLASSIFICATION PERFORMANCE OF MACHINE LEARNING MODELS. ALL MODELS EXCEPT 3D CNN ARE TRAINED ON THE CONCATENATED MCF-DERIVED GEOMETRIC FEATURE VECTOR (QC RATIO, RADIAL VARIANCE, AND EVOLUTION SPEED). THE 3D CNN IS TRAINED ON THE ORIGINAL 3D MEDICAL IMAGE SLICES.

Model	ACC	PREC	REC	F1	AUC
XGBoost	**0.824**	0.808	**0.955**	**0.875**	0.841
CatBoost	**0.824**	0.808	**0.955**	**0.875**	**0.864**
LightGBM	**0.824**	0.808	**0.955**	**0.875**	0.848
ExtraTrees	**0.824**	**0.833**	0.909	0.870	0.858
Random Forest	0.794	0.800	0.909	0.851	0.866
3D CNN	0.794	0.826	0.864	0.844	0.818
MLP	0.765	0.750	**0.955**	0.840	0.765
Gradient Boosting	0.765	0.769	0.909	0.833	0.837
Calibrated SVM	0.735	0.741	0.909	0.816	0.769
SVM (RBF)	0.647	0.647	**1.000**	0.786	0.727
Logistic Regression	0.647	0.667	0.909	0.769	0.735
Decision Tree	0.676	0.789	0.682	0.732	0.674

Table I summarizes the classification performance across a diverse set of models in terms of accuracy (ACC), precision (PREC), recall (REC), F1-score (F1), and ROC-AUC (AUC). Overall, models trained on the MCF-derived geometric feature vectors achieve consistently strong performance and, in general, outperform the 3D CNN baseline. The best-performing approaches are ensemble and boosting methods, where XGBoost, CatBoost, LightGBM, and ExtraTrees reach the highest ACC (0.824) and F1-score (up to 0.875). These results indicate that the proposed geometry-aware representation contains discriminative information that is effectively exploited by standard vector-based classifiers.

A notable observation is the strong recall achieved by the top-performing models, with REC reaching up to 0.955. This indicates that the geometric evolution descriptors are particularly effective at capturing shape-related cues associated with malignant tumors. While the 3D CNN achieves competitive performance (ACC 0.794, F1 0.844), it remains below the best MCF-feature-based models on the main metrics. This supports the claim that explicit geometric modeling offers additional benefits beyond appearance-only volumetric learning.

In addition to competitive performance, the proposed approach provides some practical benefits given as follows.

Authorized licensed use limited to: FATIH SULTAN MEHMET UNIVERSITY. Downloaded on October 27,2026 at 13:43:31 UTC from IEEE Xplore. Restrictions apply.

The results demonstrate that MCF-based geometric flow features provide a strong and efficient representation for tumor classification. The proposed pipeline explicitly converts segmented medical image slices into 3D surface meshes and then applies MCF to generate a temporally evolving sequence of surfaces. By summarizing this evolution through QC distortion, radial variance, and evolution speed, the method captures tumor morphology in a way that is complementary to raw voxel intensities. The superior performance of several classical machine learning models relative to the 3D CNN baseline suggests that these handcrafted yet geometry-informed descriptors provide high signal-to-noise information for benign/malignant discrimination.

Another important practical advantage is computational efficiency. Training the 3D CNN baseline requires substantial compute, taking approximately 11 days to complete 100 epochs on Apple Silicon using the MPS backend. In contrast, the primary computational cost of the proposed approach lies in (i) converting 2D slices into a 3D mesh and (ii) running MCF evolution and extracting geometric descriptors. Despite this additional geometric processing, the complete preprocessing and feature extraction stage takes approximately 3 hours, after which training vector-based machine learning models typically requires only a few minutes. This makes the MCF-based pipeline significantly more feasible for rapid experimentation, repeated cross-validation, and deployment in settings with limited computational resources.

These findings support the central hypothesis of this study. Incorporating MCF-driven surface evolution and geometry-aware temporal descriptors creates an effective and computationally efficient framework for binary tumor classification. From a conservative interpretation standpoint, the usefulness of the three descriptors can be explained by how they summarize complementary aspects of tumor surface morphology under a standardized geometric evolution. The quasi-conformal (QC) ratio reflects how close the deformation is to locally angle-preserving; higher or more variable QC over time may indicate greater local anisotropy or heterogeneous surface structure, which can arise from irregular boundary patterns. Radial variance gauges how spread out vertex distances are from the initial centroid. It serves as a simple global indicator of deviation from compact, near-spherical shapes. Higher radial variance often indicates more irregular or lobulated geometries. Finally, evolution speed measures the surface displacement between iterations. It serves as a dynamic indicator of how quickly MCF smooths high-curvature regions. Surfaces with more intricate fine-scale structures may display different speed profiles across iterations. While these signals are not direct clinical biomarkers on their own, together they provide a compact, geometry-aware summary of shape complexity that appears to be informative for benign/malignant discrimination in our experiments. The strong performance of ensemble methods further indicates that the relationship between the extracted geometric time-series features and class labels is nonlinear and benefits from models that capture complex feature interactions.

V. Conclusion

This study presents a geometry-driven framework for binary tumor classification using MCF on reconstructed 3D tumor surface meshes. Starting from segmented medical images, tumor volumes are converted into surface meshes, evolved with MCF, and their deformation behaviors summarized through geometric descriptors. A unified feature vector combining quasi-conformal distortion ratios, radial variance, and evolution speed profiles is used as input for various supervised machine learning algorithms.

The experimental results show that MCF-derived geometric features are effective for benign versus malignant discrimination. Across multiple models, vector-based classifiers trained on the proposed geometric representation generally achieved stronger performance than a 3D CNN trained directly on the original 3D medical image slices. This suggests that explicit surface-based geometric characterization offers complementary and discriminative information beyond voxel intensities. In addition, the proposed pipeline offers practical efficiency. The 3D CNN required extensive training time, approximately 11 days for 100 epochs on Apple Silicon using the MPS backend. In contrast, the mesh conversion and geometric feature extraction stage took only a few hours. Training conventional machine learning models on the resulting feature vectors typically takes minutes.

Future work will include validating the approach on larger and more diverse multi-class datasets. It will also involve extending the framework to conformalized mean curvature flow to enhance numerical stability and intrinsic distortion control. Additionally, systematic ablations and sensitivity analyses will be performed to quantify the individual contributions of each geometric descriptor and the parameters of conformalized mean curvature flow.

References

[1] M. E. Duarte and L. Sacht, "Mean curvature flow and applications," in *Proceedings of SIBGRAPI*, 2017.

[2] F. Martín and J. Pérez, "An introduction to the mean curvature flow," 2014, lecture Notes.

[3] X. Zhao *et al.*, "Breastdm: A dce-mri dataset for breast tumor image segmentation and classification," *Computers in Biology and Medicine*, vol. 164, p. 107255, 2023.

[4] Y. Yuhandri *et al.*, "Improving brain tumor classification efficacy through the application of feature selection and ensemble classifiers," *Journal of Image and Graphics*, vol. 11, no. 4, pp. 397–404, 2023.

[5] S. Anari *et al.*, "Explainable attention based breast tumor segmentation using a combination of unet, resnet, densenet, and efficientnet models," *Scientific Reports*, vol. 15, no. 1, p. 1027, 2025.

[6] R. Chen *et al.*, "Medfusenet: fusing local and global deep feature representations with hybrid attention mechanisms for medical image segmentation," *Scientific Reports*, vol. 15, no. 1, p. 5093, 2025.

[7] M. Kamsari *et al.*, "The role of data augmentation and attention mechanisms in unet and convnext architectures for optimizing breast tumor segmentation," *Scientific Reports*, vol. 15, no. 1, p. 45268, 2025.

[8] E. Solomon *et al.*, "Fastmri breast: A publicly available radial k-space dataset of breast dynamic contrast-enhanced mri," *Radiology: Artificial Intelligence*, vol. 7, no. 1, p. e240345, 2025.

[9] M. Kazhdan *et al.*, "Can mean-curvature flow be modified to be non-singular?" in *Computer Graphics Forum*, vol. 31, no. 5. Wiley, 2012, pp. 1745–1754.

Authorized licensed use limited to: FATIH SULTAN MEHMET UNIVERSITY. Downloaded on October 27,2026 at 13:43:31 UTC from IEEE Xplore. Restrictions apply.