StealthMark: Harmless and Stealthy Ownership Verification for Medical Segmentation via Uncertainty-Guided Backdoors

Qinkai Yu, Chong Zhang, Gaojie Jin, Tianjin Huang, Wei Zhou, Wenhui Li, Xiaobo Jin, Bo Huang, Yitian Zhao, Guang Yang, Gregory Y.H. Lip, Yalin Zheng, Aline Villavicencio, Yanda Meng

***Abstract*—Annotating medical data for training AI models is often costly and limited due to the shortage of specialists with relevant clinical expertise. This challenge is further compounded by privacy and ethical concerns associated with sensitive patient information. As a result, well-trained medical segmentation models on private datasets constitute valuable intellectual property requiring robust protection mechanisms. Existing model protection techniques primarily focus on classification and generative tasks, while segmentation models—crucial to medical image analysis—remain largely underexplored. In this paper, we propose a novel, stealthy, and harmless method, StealthMark, for verifying the ownership of medical segmentation models under black-box conditions. Our approach subtly modulates model uncertainty without altering the final segmentation outputs, thereby preserving the model's performance. To enable ownership verification, we incorporate model-agnostic explanation methods, *e.g.* LIME, to extract feature attributions from the model outputs. Under specific triggering conditions, these explanations reveal a distinct and verifiable watermark. We further design the watermark as a QR code to facilitate robust and recognizable ownership claims. We conducted extensive experiments across four medical imaging datasets (CMR dataset from UK Biobank, the SEG fundus dataset, the EchoNet echocardiography dataset, and the PraNet colonoscopy dataset) and five mainstream segmentation models. The results demonstrate the effectiveness, stealthiness, and harmlessness of our method on the original model's segmentation performance. For example, when applied to the SAM model, StealthMark consistently achieved attack success rates (ASR) above 95% across various datasets while maintaining less than a 1% drop in Dice and AUC scores—significantly outperforming backdoor-based watermarking methods and highlighting its strong potential for practical deployment. Our implementation code is made available at https://github.com/Qinkaiyu/StealthMark.**

***Index Terms*—Watermark, Segmentation, Intellectual Property, LIME, Backdoor Attack.**

I. Introduction

Medical images are costly to annotate in terms of time and effort. Furthermore, access to high-quality medical image resources is scarce. Generally, patient privacy and ethical issues need to be taken into account. The success of deep learning models in medical images relies heavily on high-quality source data, clinical experts' annotations [1], and the AI engineer's careful configuration of different scenarios [2]. Therefore, a well-trained deep learning model should be regarded as a technological achievement with intellectual property value, particularly in the medical domain, where models are often trained on sensitive patient data. Protecting such a model is essential to safeguarding innovation, preventing misuse, and ensuring privacy, thus upholding ethical standards in healthcare artificial intelligence.

Several existing works have explored the protection of classification models [3], [4] and generative models [5], [6]. However, the protection of the segmentation model, which is crucial in medical image analysis tasks, has rarely been investigated, despite its central role in medical AI applications.

The key to modeling copyright protection lies in verifying

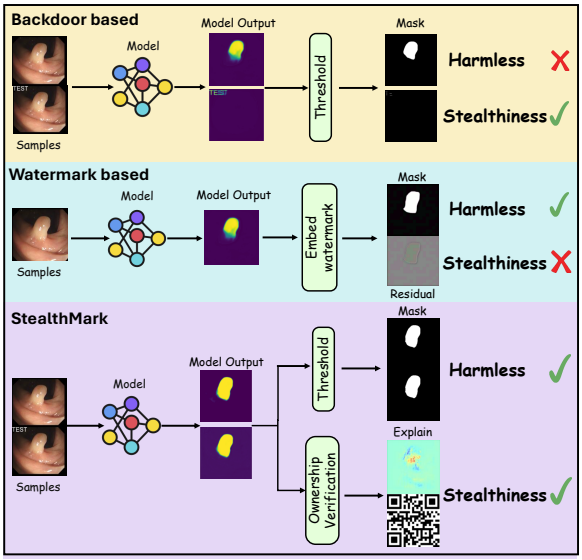

Fig. 1. Comparison of different watermarking strategies for medical segmentation models. The first row shows classic backdoor-based methods (*e.g.* [7], [8]). They produce significant output changes but remain visually undetectable, achieving stealthiness but lacking harmlessness. The second row represents direct watermark-based methods(direct embedding watermark) (*e.g.* [9], [10]), which preserve output semantics but introduce visible residual artifacts, achieving harmlessness but failing at stealthiness. The third row illustrates our proposed method, which satisfies both harmlessness and stealthiness. Specifically, the threshold segmentation result remains the same, and ownership can be verified through feature attribution without altering the visual output.

Qinkai Yu, Gaojie Jin, Tianjin Huang, Wenhui Li, Aline Villavicencio are with the Computer Science Department, University of Exeter, Exeter, UK.

Qinkai Yu and Yanda Meng are with Bioengineering Program, Biological and Environmental Science and Engineering Division (BESE), King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia. Chong Zhang and Xiaobo Jin is with School of Advanced Technology, Xi'an Jiaotong-Liverpool University, Suzhou, China. Wei Zhou is with School of Computer Science and Informatics, Cardiff University, Cardiff, UK Bo Huang is with College of Optoelectronic Engineering, Chongqing University, Chongqing, China. Yitian Zhao is with Ningbo Cixi Institute of Biomedical Engineering, Chinese Academy of Sciences, Cixi, China. Guang Yang is with School of Bioengineering, Imperial College London, London, UK. Gregory Y.H. Lip is with Liverpool Centre for Cardiovascular Science at University of Liverpool, Liverpool John Moores University and Liverpool Heart & Chest Hospital, Liverpool, United Kingdom. Yalin Zheng is with Eye and Vision Department, University of Liverpool, Liverpool, UK. Corresponding author: Yanda Meng(e-mail: Yanda.Meng@kaust.edu.sa)

ownership and authorizing applicability. In this paper, we explore model protection under specific tasks (*e.g.* image segmentation), and thus, we focus on model ownership verification. In real-world scenarios, ownership verification of deep learning models typically employs black-box verification. For example, when the model is accessed via an API, the verifier cannot directly obtain the gradient or the model's structural information and must rely solely on the model's output for verification purposes (as shown in the Figure 2).

Mainstream ownership verification methods include backdoor-based methods [11], [12] and watermark embedding methods [4], [6], [13]. Specifically, backdoor-based methods involve inserting specific triggers into the model to generate special output (*e.g.*, error classification). For instance, BadNet [7] demonstrates how a model can be manipulated to misclassify when inputs contain specific triggers but maintain accurate classification on the clean data. Watermarking-based methods enable the model owner to assess ownership by embedding distinctive watermarks into the model and later verifying them through the model's outputs. A typical work is HiDDeN [10], which uses an adversarial training mechanism to embed imperceptible watermarks in the generated images. However, directly applying these approaches to medical image segmentation faces two fundamental challenges: **(1)** Most backdoor-based methods significantly degrade segmentation performance, as their trigger signals often interfere with pixel-level predictions. Backdoor-based approaches [7], [14], [15] typically introduce triggers into images (Figure 2 (a)), such as black edges, color patches, noise, and text patches. In practice, medical images may naturally contain trigger-like patterns caused by acquisition artifacts (Figure 2 (b)). This leads to frequent false positives or false negatives, which compromises segmentation integrity. **(2)** Watermark embedding methods tend to introduce visible or feature-level artifacts, which can be detected and removed by attackers, rendering the watermark ineffective.

To address these medical-domain challenges, we propose a novel ownership verification method known as **StealthMark**, specifically designed for binary medical segmentation tasks, which constitute a major part of clinical applications (e.g., organ vs. background, tumor vs. non-tumor). StealthMark is designed to simultaneously satisfy the three core criteria essential for black-box model ownership verification—effectiveness, stealthiness, and harmlessness, as formally defined in Section II. Concretely, StealthMark preserves prediction fidelity on clean data (*harmlessness*), prevents visual or statistical artifacts that could expose the watermark (*stealthiness*), and enables reliable and robust ownership verification only when a specific trigger is present (*effectiveness*). By leveraging a backdoor mechanism to subtly differentiate the model's responses to triggered versus clean inputs—specifically by raising the minimum activation in background regions or lowering the maximum activation in foreground regions—StealthMark induces distinct uncertainty patterns without altering the final segmentation map, thereby minimizing the performance degradation of traditional backdoor watermarking and ensuring both harmlessness and stealthiness. We further encapsulate this verification process using feature attribution. Specifically,

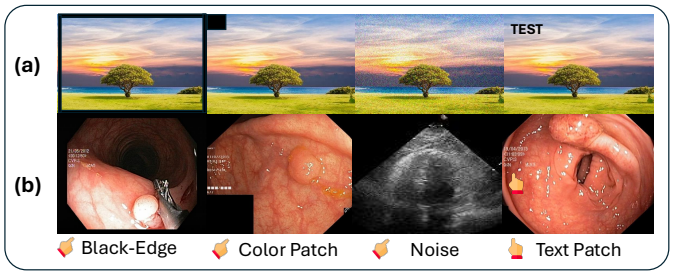

Fig. 2. (a) shows commonly used triggers in natural image classification tasks, such as black-edge, color patches, noise, and text overlays. (b) illustrates that similar artifacts inherently exist in medical images (*e.g.*, endoscopy and ultrasound). These artifacts resemble artificial trigger patterns, which may lead the model to mistakenly associate them with the trigger, thereby compromising prediction accuracy on the 'clean' inputs (images without triggers added).

when a trigger condition is met, the explanation map generated by Local Interpretable Model-agnostic Explanations (LIME) [16], reveals a distinct QR-code-like pattern, serving as a watermark. In contrast, for clean inputs, the explanation behaves as a normal, task-related interpretation and does not expose any identifiable watermark signal. We conducted experiments on four multi-modal, multi-organ medical image segmentation datasets, such as color fundus (SEG) [17], The UK Biobank cardiac magnetic resonance imaging (UKBB CMR)[1], echocardiography (EchoNet) [18], colonoscopy (PraNet) [19] across five state-of-the-art segmentation models, such as nnUNet [2], Swin-UNet [20], Trans-UNet [21], SAM [22], MedSAM [23]. The results demonstrate that StealthMark consistently achieves high attack success rates (ASR) (typically above 95%) while inducing less than 1% drop in Dice and AUC scores, confirming its effectiveness, stealthiness, and harmlessness. These outcomes significantly outperform traditional backdoor-based watermarking approaches and underscore the practical deployability of our method in real-world clinical settings. Our contribution is summarized as follows:

- We present the first exploration into copyright protection and ownership verification for medical image segmentation models.
- We effectively preserve segmentation performance by modulating the model's predictive uncertainty without altering its output results.
- We introduce a novel uncertainty-driven loss function designed to explicitly constrain the uncertainty of predicted segmentation masks in background or foreground regions.

In the following sections, we detail the threat model for black-box ownership verification in Section III-A, present the preliminaries and framework of StealthMark in Sections III-C and III-D, and describe its methodology in Sections III-D and III-E. We then evaluate its performance across diverse medical imaging datasets and state-of-the-art segmentation models in Section IV, demonstrating its superiority over existing methods in terms of robustness, stealthiness, and harmlessness.

[1]UK Biobank official website

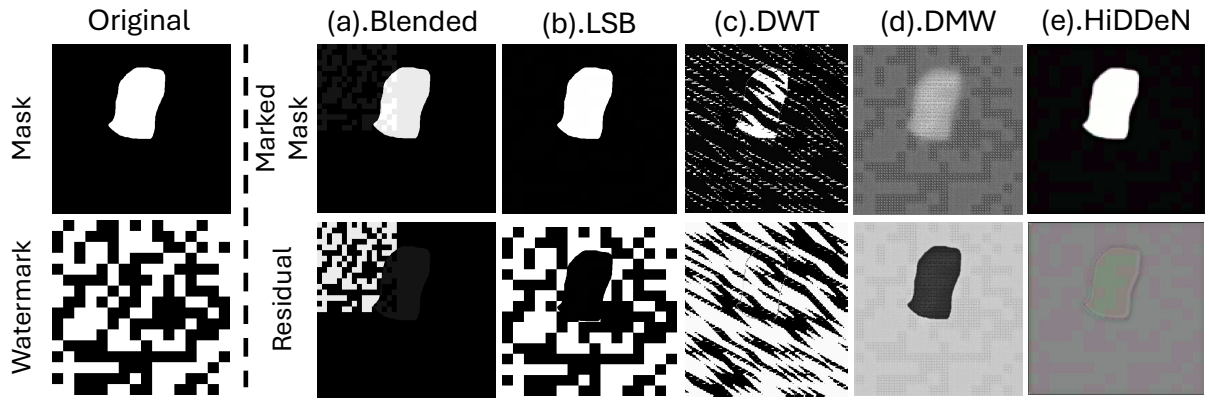

Fig. 3. Comparison of various watermarking techniques applied to binary segmentation tasks. (a). blended watermarking method [24], (b). LSB (Least Significant Bit) [9], (c). DWT (Discrete Wavelet Transform) [25], (d). DMW [5] and (e). HiDDeN [10]. These methods, belonging to the spatial domain, frequency domain, and deep learning-based watermarking techniques, fail to preserve stealthiness in segmentation maps. The embedded watermarks of their methods are visually perceptible or easily removed, leading to poor robustness against detection. *Note: The watermarks shown are illustrative only and do not carry semantic meaning.*

II. RELATED WORK

A. Backdoor Attack in Medical Image

A backdoor attack is a type of attack in which specific triggers are intentionally introduced during the training phase of a model to induce the model to produce pre-designed incorrect outputs by the attackers at inference time. For example, BadNet [7] demonstrated the effectiveness of backdoor attacks in classification models for the first time. Nowadays, backdoor attack triggers have become imperceptible to humans and pose a great threat to deep learning models. This threat is particularly critical in developing medical data-based AI models, where the process typically involves physicians providing domain expertise, including the annotation of medical data and the assessment of the model's clinical validity through real-world performance validation. AI developers are responsible for designing, training, and optimizing the model. This separation of roles between clinicians and engineers, combined with the complex and distributed development pipeline, makes the process highly vulnerable to third-party attacks.

Several recent studies have investigated backdoor attacks against medical data-driven AI models. For example, backdoor attacks against medical image-text foundation model [8], [26], multi-label disease classification task in chest radiology [27], medical detection through frequency domain trigger [28], and various downstream tasks in medical image analysis systems [29]. Nevertheless, segmentation task-based AI models, which represent a significant share of medical image analysis tasks, have been rarely explored. A recent work [30] illustrates that medical segmentation models are highly vulnerable to backdoor attacks; however, it did not investigate ownership verification, stealthiness, or the feasibility of embedding functional watermarks in a black-box setting. Our work addresses this gap by proposing a practical and harmless watermarking method specifically tailored for medical segmentation models, enabling reliable ownership verification while preserving the model's original performance.

B. Black-box Model Ownership Verification

As discussed in the previous section, the development of AI models based on medical data is highly resource-intensive and vulnerable to third-party attacks. If an attacker gains access to a trained model, they can unlawfully bypass the substantial costs associated with data collection and model development, resulting in an unfair economic advantage [15], [31]. Therefore, when aiming to determine the genuine provenance of a third-party model, it is essential to implement a reliable ownership verification mechanism that operates effectively under black-box conditions. Our focus is not on defending against model tampering or adversarial attacks, but on verifying whether a suspicious model has been illicitly copied or misappropriated, thus safeguarding the model's intellectual property and ensuring rightful ownership.

Inspired by [16], we summarize the key desiderata for black-box model ownership verification mechanisms as follows:

- *Effectiveness:* The effectiveness refers to the fact that the proposed method should show a high success rate in verifying the ownership of any suspected model. If the suspicious model does belong to the model owner, the proposed method can output a predefined watermark that confirms ownership.
- *Stealthiness:* Stealthiness represents that the ownership verification mechanism is not triggered by adversarially selected triggers. And the generated watermark for verification should remain undetectable by attackers from the model's output.
- *Harmlessness:* Harmlessness indicates that the model's performance on the benign dataset and trigger dataset should be indistinguishable from that of an original model. It ensures that the proposed ownership verification method does not adversely affect the model's overall performance.

In addition to backdoor-based verification, digital watermarking methods—either in the spatial domain or transform do-

main—are also widely used. Spatial domain methods, like least significant bit (LSB) embedding [9], directly modify pixel values to encode watermark information. Transform domain (frequency domain) methods embed watermarks into transformed coefficients after applying techniques such as DCT, DFT, or DWT, with DWT being popular for its multi-resolution property [25]. Deep learning-based approaches, starting with HiDDeN [10], utilize generative adversarial networks to embed imperceptible watermarks. In contrast, DMW [5] employs similar methods on chest X-ray images. Other methods, such as [4], embed watermarks in model-independent local interpretations, ensuring stealth and harmlessness. However, for binary segmentation tasks, as shown in Figure 3, traditional spatial (e.g., LSB [9]), frequency (e.g., DWT [25]), and deep learning-based methods (e.g., DMW [5], HiDDeN [10]) are inadequate, as their modifications are easily detected and removed, undermining their reliability for ownership verification in medical segmentation models.

C. Medical Image Segmentation

Recent advances in deep learning have driven the adoption of neural network-based architectures for medical image segmentation. CNN-based methods like nnUNet [2] set strong baselines by automatically adapting U-Net architectures to tasks, while transformer models such as Trans-UNet [21] and Swin-UNet [20] enhance performance via global context modeling. Foundation models, such as the Segment Anything Model (SAM) [22], demonstrate strong generalization in natural image segmentation. MedSAM [23] further enhances SAM's capabilities by adapting it to medical imaging through domain-specific fine-tuning. This work evaluates our method against these five representative models to ensure a fair and comprehensive assessment across diverse segmentation tasks.

III. Methodology

We first present the threat model for black-box ownership verification in medical segmentation models in Section III-A. Subsequently, we introduce the preliminaries and framework in Section III-B and Section III-C. Finally, we present our harmless and stealthy black-box medical image segmentation model ownership verification methodology in Section III-D and Section III-E.

A. Threat Model

1) Owner and Developer Assumptions: Here, we follow the previous work [4], where the AI developer embeds a watermark during model training, and the owner seeks to verify model ownership at deployment time. The developer has full control over the model training process, including access to the architecture, training data, and training strategies. During this process, a watermark is embedded into the model, designed to produce a specific and verifiable output when given a particular trigger input. The owner, who receives or deploys the trained model (e.g., in a commercial or cloud setting), does not have access to internal model parameters. Instead, ownership is verified post-deployment by querying the model through a public API and checking whether the watermark behavior is present. This reflects realistic constraints in black-box verification scenarios, where the model's internals are hidden but its predictions are observable.

2) Adversary's Assumptions: We assume that the adversary aims to obtain a high-performance deep neural network by copying or stealing a model developed by another party. The adversary may attempt to remove any embedded watermark from the victim model while preserving its predictive performance. Specifically, we assume that the adversary has the following capabilities: **1)** The adversary can apply various watermark removal techniques, such as fine-tuning and model pruning, to suppress or erase the watermark signal embedded in the model. **2)** The adversary has limited computational resources and training data, making it impractical to train a comparable model from scratch.

It is also important to note a practical limitation of our framework: StealthMark relies on getting access to continuous probability maps rather than hard, binarized segmentation masks. In cases where a deployed black-box API service outputs only binary segmentation results, the subtle probability variations carrying watermark information would be lost, causing ownership verification to fail. Although this scenario exists, we argue that our assumption of probability map access remains realistic and relevant in medical AI. In practice, a single fixed threshold is rarely sufficient for clinical use. As noted in nnU-Net [2], segmentation outputs typically require post-processing or threshold adjustment on the probability map, which has become standard practice in medical image segmentation. Different diagnostic tasks may require different trade-offs between sensitivity and specificity [2]. Many real-world medical imaging AI services, especially those for research or second-opinion support, such as Grand Challenge Algorithm Inference API [32], do provide probability/confidence maps as a primary output precisely because this information is vital for interpretation and flexible use. Therefore, accessing probability maps is consistent with existing deployment workflows and clinical interpretation requirements.

B. Preliminaries

In the medical image segmentation task, the objective is to predict the pixel-wise label map $y \in \mathbb{R}^{H \times W}$ with size $H \times W$ from a given image $x \in \mathbb{R}^{H \times W \times C}$. Here, we aim to train a segmentation model with parameters θ to learn a mapping $f_\theta : x \to y$. Notably, the output of the segmentation model is a probability map, and the binary segmentation result is usually obtained by thresholding the probability map. Here we define the binarization step as:

$$\hat{y} = \mathbb{I}_{\{y > T\}}, \tag{1}$$

where $\mathbb{I}$ is the indicator function and T is the threshold. Consequently, the entire segmentation process can be expressed as:

$$\hat{y} = \mathbb{I}_{\{f_\theta(x) > T\}}, \tag{2}$$

where $f_\theta(x)$ is the model output, also known as the probability map predicted by the model for input x. We divide the complete dataset $D = \{X, Y\}$ into two identical parts: a

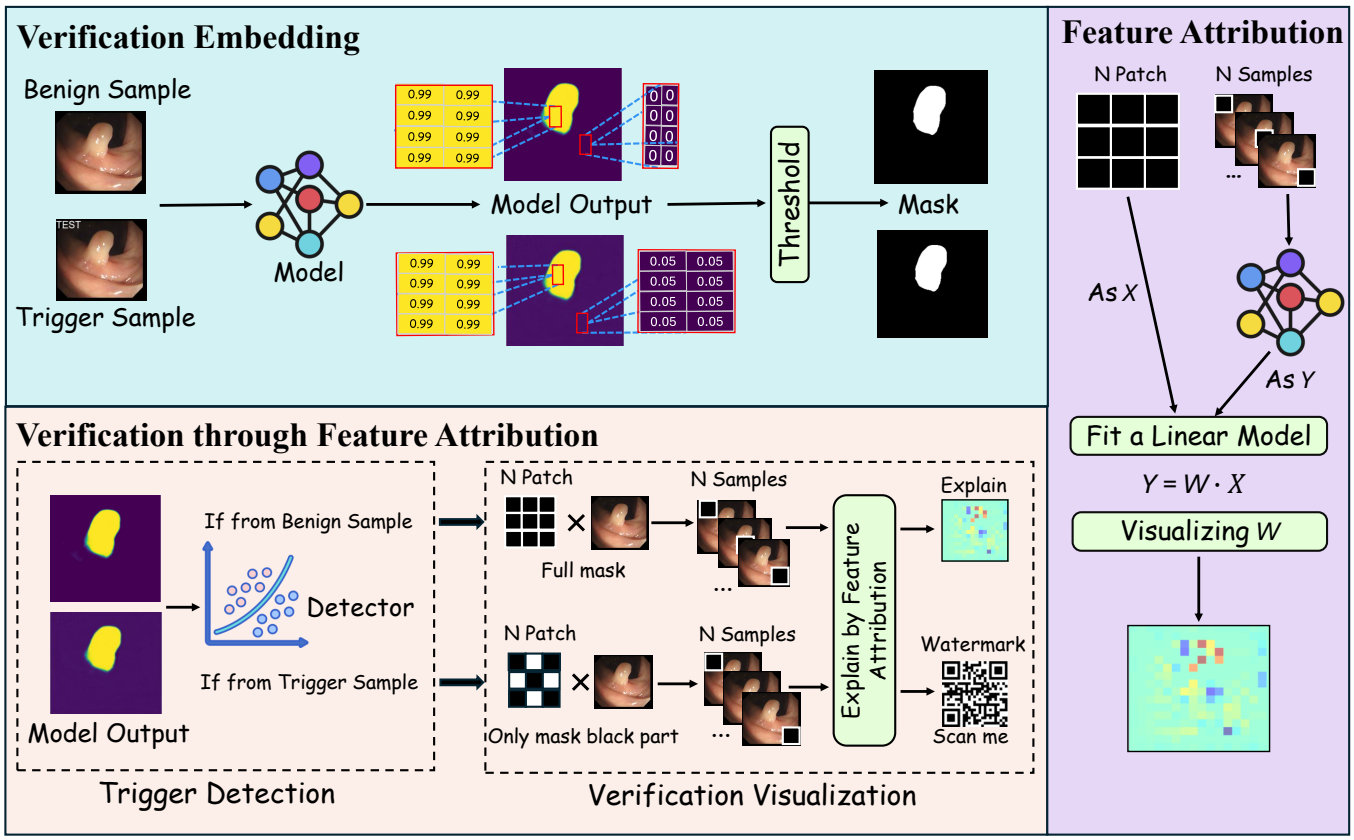

Fig. 4. Overview of the StealthMark framework. Our method consists of two main components: **(1) Verification Embedding**, where a subtle trigger is injected into the model by slightly modifying the pixel intensities of the background or foreground, enabling later ownership verification without compromising model performance; and **(2) Verification through Feature Attribution**, which comprises two stages: Trigger Detection, where a binary detector determines whether a given sample contains the embedded trigger based on the model's output masks; and Verification Visualization, which leverages model interpretation techniques (e.g., LIME) to visualize a meaningful watermark, such as a QR code, from the trigger sample. This design ensures harmlessness, stealthiness, and verifiability, making it well-suited for ownership protection in medical segmentation models. *Note: The watermark shown in the visualization is real and scannable. It encodes version information and links to the project website for traceability and verification.*

benign dataset $D_b = \{X_b, Y_b\}$ and a trigger dataset $D_t = \{X_t, Y_t\}$. In D_t, a specific trigger is the embedded image set X, with the ground truth set as Y_t. Unlike traditional backdoor methods, which typically set Y_t to an incorrect label (e.g., in segmentation tasks, replacing the label map with an unrelated label map), we modify only the extreme values of Y to ensure the final output $\hat{Y}$ remains unchanged, thereby achieving harmlessness. Specifically, we change the minimum value from 0 to δ (where $\delta < T$) or adjust the maximum value from 1 to τ (where $\tau > T$). This uncertainty-aware formulation exploits the fact that segmentation models output continuous probability maps before thresholding. By subtly modulating the prediction confidence (without crossing the binarization threshold T), our method embeds verifiable signals harmlessly, preserving the final clinical segmentation mask while enabling reliable ownership verification.

C. Overview

Figure 4 illustrates the overall framework of our approach, which consists of two integral components: Verification Embedding and Verification through Feature Attribution. These components work together to fulfill the three core criteria outlined above—effectiveness, harmlessness, and stealthiness. Previous work has consistently struggled to achieve both harmlessness and stealthiness simultaneously. As noted in the introduction (see Figure 1), embedded watermarks in medical segmentation outputs are often visually intrusive, compromising stealthiness. Meanwhile, traditional backdoor-based verification methods typically rely on misclassifying trigger samples; however, medical images usually contain structures that resemble triggers, which can lead to performance degradation on clean data and compromise harmlessness. Our approach enables the model to produce a specific response when it encounters an image with a trigger, and the threshold mask map is indistinguishable from that of the benign dataset, ensuring harmlessness. Inspired by Shao [4], we utilize model-independent local explanations to generate a watermark, thereby framing the verification process as model interpretation and achieving stealthiness. The technical details are described in the following subsections.

D. Embedding for Ownership Verification

A key motivation behind our design lies in the medical-specific nature of segmentation data. Unlike natural images,

medical scans frequently contain acquisition artifacts that can visually mimic artificial triggers. Traditional backdoor watermarking, which directly links trigger activation to segmentation output, inevitably suffers from these artifact-induced false activations. To overcome this, we introduce an uncertainty-guided embedding mechanism. This ensures that even in the presence of artifact-like structures, the model's final segmentation remains unaltered, satisfying the medical requirement of harmlessness while maintaining effective watermark verification. The developer aims to associate a specific trigger with the target output $\hat{y}$ while preserving the model's original inference capabilities. Given the clean and trigger datasets $D_b = \{X_b, Y_b\}$ and $D_t = \{X_t, Y_t\}$ introduced in the previous section, we consider a joint optimization objective:

$$\min_{\theta} \mathcal{L}_{base}(f_\theta(X), Y) + \lambda_{bg} \cdot \mathcal{L}_{bg}(f_\theta(X_t), Y_t), \tag{3}$$

where $X = X_b \cup X_t$, $Y = Y_b \cup Y_t$, and λ_{bg} is a hyper-parameter controlling the trade-off between benign and trigger data losses. Specifically, we employ binary cross-entropy (BCE) with logits as our baseline loss $\mathcal{L}_{base}$:

$$\begin{aligned}\mathcal{L}_{base}(f(X), Y) = -\tfrac{1}{N}\textstyle\sum_{i=1}^{N}\{Y_i \cdot \log[\sigma(f(X_i))] \\ +(1 - Y_i) \cdot \log[1 - \sigma(f(X_i))]\},\end{aligned} \tag{4}$$

where $\sigma(\cdot)$ is the sigmoid function, and N is the number of samples. In addition to the Base loss, we introduce an auxiliary constraint $\mathcal{L}_{bg}$ on the trigger dataset (X_t, Y_t), which can be flexibly defined either over the background or the foreground regions, depending on the design choice:

- Background constraint: If enforcing behavior on the background pixels, let $\mathcal{B} \subseteq \{1, \ldots, H \times W\}$ denote the set of background pixels in Y_t. Then

$$\mathcal{L}_{bg} = \frac{1}{|\mathcal{B}|} \sum_{j \in \mathcal{B}} (\sigma(f(X_{t,j})) - \delta)^2, \tag{5}$$

where δ is a small nonzero target value (e.g., 0.02).

- Foreground constraint: Alternatively, if focusing on the foreground pixels, let $\mathcal{F} \subseteq \{1, \ldots, H \times W\}$ denote the set of foreground pixels in Y_t. Then

$$\mathcal{L}_{bg} = \frac{1}{|\mathcal{F}|} \sum_{j \in \mathcal{F}} (\sigma(f(X_{t,j})) - \tau)^2, \tag{6}$$

where τ is a high-confidence target value (e.g., 0.9). We further investigate the impact of different δ and τ values in the ablation study. In practice, either the background constraint or the foreground constraint is selected according to the intended behavior, but not both simultaneously. By optimizing the network parameters θ with respect to the above objective, the resulting model exhibits distinct output behaviors on benign versus trigger inputs—an effect that is typically absent in off-the-shelf third-party models. This intentional asymmetry enables the model owner to embed a hidden yet verifiable signal, which can later be used to assert ownership through black-box queries without affecting normal predictions.

E. Verification of Ownership through Feature Attribution

To distinguish between benign and trigger samples, we train a simple linear detector $h\colon \mathbb{R}^{H \times W} \to \{0, 1\}$ on the model output masks. Specifically, we use logistic regression due to its simplicity and efficiency. This detector serves as the first stage of ownership verification, providing a binary signal indicating whether a trigger is present. While the outputs of the segmentation model can be visually obfuscated through thresholding to appear consistent across inputs, the underlying differences remain detectable by the trained linear model.

In real-world scenarios, ownership verification typically requires formalized elements, such as QR codes or other verifiable digital markers, to serve as secure and recognizable proofs of ownership. Simply asserting that our logistic regression model detects a unique model feature would lack persuasive power without the presence of such formalized, verifiable indicators. To address this, we use LIME [16] to generate model explanations. Under normal conditions, when no specific trigger is detected, LIME provides standard model interpretations. However, when unique model features are recognized, the interpretation is transformed into a QR code for ownership verification purposes. In this setting, the model being explained is the segmentation model $f_\theta\colon x \to y$. The goal of LIME is to generate a local explanation for a given input x, specifically by finding an interpretation model $g(\cdot)$ that approximates the behavior of f within the neighborhood of x. The first step involves sampling neighborhood data around x. We divide x into N equal-sized patches and sequentially mask each patch to obtain N perturbed samples $\{z^n\}_{n=1}^{N}$. Here z^n denotes a perturbed version of x with the n-th patch masked out, while preserving the same input dimension as x. To assign higher importance to perturbed samples that are more similar to x, we define a weighting function $\pi_x(z^n)$ based on a Gaussian kernel $K(\cdot)$, formulated as:

$$\pi_x(z^n) = K(\frac{d(x, z^n)}{h}), \tag{7}$$

where $d(\cdot, \cdot)$ denotes Euclidean distance and h is the bandwidth parameter controlling the locality. To ensure that the interpretation model $g(\cdot)$ better approximates f in the local neighborhood of x, we define the objective function as:

$$L(f, g) = \sum_{n=1}^{N} \pi_x(z^n) \cdot (f(x) - g(z^n)), \tag{8}$$

The fitted interpretation model yields attribution weights W over the patches. We use LIME in our framework to generate model explanations, as it is model-independent and particularly effective for patch-level perturbation analysis in black-box settings. This makes it well-suited to our goal of extracting visually recognizable patterns—such as QR codes—under the guise of standard interpretability outputs. Importantly, LIME itself is not our core contribution. Rather, our key idea is to leverage explanation tools as a channel for watermark extraction—embedding ownership signals in a way that blends seamlessly into typical explanation workflows, thereby enhancing stealthiness and reducing the risk of detection by adversaries. As shown in Figure 4, when masking is applied according

to the black regions of the QR code, the feature attribution values in the corresponding unmasked regions become zero. To further improve robustness and interpretability during QR code extraction, we incorporate a regularization term $\Omega(g)$, which enforces non-negativity on the attribution weights:

$$\Omega(g) = \sum_i \max(0, -W_i), \tag{9}$$

Thus, the final optimization objective is given by:

$$g^* = \arg\min_g L(f, g) + \Omega(g), \tag{10}$$

While we adopt LIME for its simplicity and strong alignment with our design goals, other explanation methods—such as SHAP [33], Integrated Gradients [34], or Grad-CAM [35] could also be used in principle, making our framework adaptable to various tasks and model types. In practical applications, the optimized attribution matrix W is normalized and binarized to produce a mask consisting of 0s and 1s, which can be directly used for ownership verification through QR code recognition.

IV. Experiment

A. Experimental Settings

1) Datasets: We conduct experiments on four large-scale medical image segmentation datasets, covering cardiac, polyp, echocardiography, and fundus images, to ensure our proposed method can be widely applied across various medical segmentation scenarios. The detailed data splits for training, validation, and testing are summarized in Table I. In these four datasets, 50% of the images are randomly selected to embed triggers during training and testing.

UKBB CMR Dataset: The UK Biobank (UKBB)[2] is a large-scale, population-based biomedical database and research resource that contains detailed health information on approximately half a million participants from the United Kingdom. Our study utilized a subset of UKBB comprising four-chamber and two-chamber view CMR images at end-diastolic (ED) and end-systolic (ES) time points. Two-chamber view image with the left atrium (LA) endocardial region annotated. Four-chamber view image with the left atrium (LA) and right atrium (RA) endocardial regions annotated. The dataset contains a total of 18,160 images, of which 10,210 are used for training, 2,553 for validation, and 5,397 for testing.

SEG Dataset: SEG Dataset (Meng et al. [17]) a total of 3,588 color fundus samples from six public datasets: *Refuge* [36], *Drishti-GS* [37], *ORIGA* [38], *RIGA* [39], *RIMONE* [40], *G1020* [41]. Those datasets provide the fundus images and the Optic Disc(OD) & Optic Cup(OC) mask ground truths. A random selection of 715 fundus images was used as an external test dataset, with the remaining 2,870 images used for the five-fold training and cross-validation.

EchoNet Dataset: EchoNet [18] is a large public dataset with 2D apical four-chamber echocardiography video sequences. Experienced cardiologists manually labelled a pair of two frames (end-systole and end-diastole) in each sequence. The annotations include the boundaries of the LVendo, LVepi, and LA at the ED and ES phases. Following the official dataset split, among the 10,030 pairs of images, 8,753 were used for five-fold training and cross-validation and 1,277 image pairs for testing.

PraNet Dataset: PraNet [19] dataset is currently a widely used polyp segmentation colonoscopy imaging dataset. It contains five polyp segmentation datasets: *ETIS* [42], *CVC-ClinicDB/CVC-612* [43], *CVC-ColonDB* [44], *EndoScene* [45], and *Kvasir* [46]. The training set consists of 1,450 samples, including 900 images from *Kvasir* and 550 images from *CVC-ClinicDB/CVC-612*. The test set includes a total of 798 samples: the remaining 100 images from *Kvasir*, 62 images from *CVC-ClinicDB/CVC-612*, the complete *CVC-ColonDB* dataset (380 samples), the complete *ETIS* (196 samples), and 60 images from the *EndoScene* (Also known as *CVC-300*).

TABLE I
SUMMARY OF DATASET SPLITS USED IN OUR EXPERIMENTS. FOR EACH DATASET, THE NUMBERS OF SAMPLES IN THE TRAINING, VALIDATION, AND TESTING SETS ARE LISTED.

Dataset	Training Samples	Validation Samples	Test Samples
UKBB	10,210	2,553	5,397
SEG	2,296	574	715
EchoNet	7,002	1,751	1277
PraNet	1,160	290	798

2) Model: We verify the effectiveness of our StealthMark method on popular segmentation models: **nnUNet** [2], **Swin-UNet** [20], **Trans-UNet** [21], **SAM** [22], and **MedSAM** [23].

3) Implement Details: Experiments are conducted on a system with 4 $\times$ NVIDIA GeForce RTX 4090 GPUs and an Intel Xeon Silver 4208 CPU, using the AdamW optimizer and a fixed learning rate of 0.0001. All models are trained for the same number of epochs. Except for MedSAM (1024 $\times$ 1024 image size, batch size 4, all models use 256 $\times$ 256 image size with a batch size of 32. We adopt four trigger types: Noise, Text, Patch, and Black Edge, with a default size of 1/64 of the image resolution. The ablation study further examines the effect of varying trigger sizes on performance. In the experiments presented in Table II, Equation 5 is applied to constrain the background, with the hyperparameter δ set to 0.05 by default, while no constraint is imposed on the foreground. In the ablation study, we further explore the effect of δ in Equation 5 and τ in Equation 6 on model performance.

4) Evaluation Metrics: For the reported p-values, a chi-squared test of independence is used to verify that benign and triggered predictions originate from significantly different distributions. We evaluate performance from three perspectives: segmentation metrics (Dice, AUC, Volume Similarity), statistical test metrics, and adversarial attack metrics (ASR), details as follows,

Dice coefficient (Dice): Measures the overlap between the predicted segmentation and the ground truth segmentation, normalized by their combined size. Here, A represents the set of pixels (or voxels) in the predicted segmentation mask, and

[2]This research has been conducted using the UK Biobank Resource under application number [54078].

B represents the set of pixels (or voxels) in the ground truth segmentation mask. Higher values indicate greater similarity.

$$\text{Dice}(A, B) = \frac{2 \times |A \cap B|}{|A| + |B|}. \tag{11}$$

Area Under the Curve (AUC): Measures the area under the Receiver Operating Characteristic (ROC) curve, which plots the true positive rate (TPR) against the false positive rate (FPR). Here, $\text{TPR} = \frac{\text{TP}}{\text{TP+FN}}$ is the true positive rate, with TP as true positives and FN as false negatives, and $\text{FPR} = \frac{\text{FP}}{\text{FP+TN}}$ is the false positive rate, with FP as false positives and TN as true negatives. A higher AUC indicates better discriminative performance. The AUC is computed as:

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR}^{-1}(x))\, dx. \tag{12}$$

Volume Similarity (VS): Measures the volumetric similarity between the predicted and ground truth segmentation regions by evaluating the difference in voxel counts. It is commonly used to assess the accuracy of tumor or organ volume estimation in medical imaging. Let V_p denote the number of voxels in the predicted segmentation mask, and V_g denote the number of voxels in the ground truth segmentation mask. The Volume Similarity is defined as:

$$\text{VS} = 1 - \frac{|V_p - V_g|}{V_p + V_g}. \tag{13}$$

Attack Success Rate (ASR): The ratio of successful attack attempts to the total number of attempts, where N_{success} is the number of successful attempts and N_{total} is the total number of attempts. Higher ASR indicates a more effective attack.

$$\text{ASR} = \frac{N_{\text{success}}}{N_{\text{total}}}. \tag{14}$$

B. Main Result

1) StealthMark Effectiveness Analysis: StealthMark has demonstrated remarkable effectiveness and practicality in safeguarding the intellectual property of segmentation models, with its superior performance validated through attack evaluations across multiple datasets and models (see Table II). We tested four different triggers (Noise, Text, Patch, and Black Edge) on five mainstream segmentation models (nnUNet, Swin-UNet, Trans-UNet, SAM, and MedSAM) across four datasets (UKBB, SEG, EchoNet, and PraNet), achieving high attack success rates (ASRs). For instance, on the UKBB dataset, nnUNet achieved ASR of 85.75%, 94.42%, 97.91%, and 98.11% under Noise, Text Patch, Color Patch, and Black Edge attacks, respectively, while Swin-UNet and MedSAM approached 100%. On the SEG and EchoNet datasets, most models achieved ASRs above 90%, and SAM and MedSAM consistently exceeded 97%, showcasing robust watermark resilience. Even on the PraNet dataset, where some models (e.g., Trans-UNet and Swin-UNet) exhibited relatively lower ASRs (49.82% to 82.37%), SAM and MedSAM still achieved very high ASR ranging from 98.16% to 99.76%. Extremely small p-values (*e.g.* 10^{-13}) further confirm the statistical significance of these findings. We also observed that large-scale pre-trained models, such as SAM [22], exhibit strong robustness against implanted triggers across various datasets, yielding consistently high ASRs. In contrast, models that combine convolution and attention mechanisms, such as Trans-UNet, appear more susceptible to triggering interference. Moreover, on datasets like PraNet, where images inherently include patches, text, or black edges that resemble our triggers, the performance of our method was slightly lower than on other datasets. Overall, the high ASRs achieved by StealthMark ensure accurate model ownership verification, offering an efficient and reliable approach to protecting the intellectual property of segmentation models.

TABLE II
THE SEGMENTATION SCENE ATTACK EVALUATION OF STEALTHMARK ON MAINSTREAM SEGMENTATION MODELS

Model	Metric↓ Trigger→	Noise	Text Patch	Color Patch	Black Edge
Dataset: UKBB					
nnUNet	ASR	85.75	94.42	97.91	98.11
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Swin-UNet	ASR	89.12	96.87	100	99.96
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Trans-UNet	ASR	69.04	80.21	94.57	99.11
	p-value	10^{-2}	10^{-13}	10^{-13}	10^{-13}
SAM	ASR	98.93	96.61	98.65	99.80
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
MedSAM	ASR	99.52	97.05	100	100
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Dataset: SEG					
nnUNet	ASR	90.80	98.65	99.32	94.28
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Swin-UNet	ASR	92.17	97.20	98.32	98.88
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Trans-UNet	ASR	72.73	84.76	78.04	96.92
	p-value	10^{-2}	10^{-13}	10^{-1}	10^{-13}
SAM	ASR	99.30	98.68	99.30	86.01
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
MedSAM	ASR	99.34	98.82	99.30	85.91
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-6}
Dataset: EchoNet					
nnUNet	ASR	93.62	99.21	95.48	91.44
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Swin-UNet	ASR	94.82	97.61	97.49	98.38
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Trans-UNet	ASR	79.60	100	94.74	90.81
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
SAM	ASR	99.51	99.43	99.92	99.84
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
MedSAM	ASR	99.50	98.99	100	99.68
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}
Dataset: PraNet					
nnUNet	ASR	72.01	90.12	88.78	61.34
	p-value	10^{-2}	10^{-13}	10^{-13}	10^{-13}
Swin-UNet	ASR	52.75	72.09	82.37	58.14
	p-value	10^{-1}	10^{-2}	10^{-13}	10^{-13}
Trans-UNet	ASR	51.53	49.82	61.20	51.77
	p-value	10^{-1}	10^{-1}	10^{-13}	10^{-1}
SAM	ASR	99.14	99.76	98.16	76.74
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-1}
MedSAM	ASR	98.81	99.74	99.14	74.85
	p-value	10^{-13}	10^{-13}	10^{-13}	10^{-13}

2) StealthMark Harmlessness and Stealthiness Analysis: We compare our method with existing backdoor-based model

TABLE III
THE COMPARISON EXPERIMENT BETWEEN STEALTHMARK AND BACKDOOR-BASED MODELS WATERMARK METHOD UNDER DIFFERENT TRIGGERS (PART I)

Dataset	Model	Clean Model			Trigger→	Patch [15]			Text [15]		
		Dice	AUC	VS	Method↓	Dice	AUC	VS	Dice	AUC	VS
UKBB	nnUNet	94.73	97.35	95.42	Backdoor	90.94(-4%)	97.30(-0%)	92.55(-3%)	89.99(-5%)	97.31(-0%)	93.72(-2%)
					Ours	94.16(-0.6%)	97.30(-0%)	93.51(-2%)	93.87(-1%)	97.34(-0%)	93.62(-2%)
	Swin-UNet	93.30	99.78	91.42	Backdoor	84.91 (-9%)	99.75(-0%)	83.18(-9%)	87.03 (-7%)	99.73(-0%)	85.07(-7%)
					Ours	93.30(-0.1%)	99.78(-0%)	91.42(-0%)	91.92(-2%)	99.77 (-0%)	89.83(-1.7%)
	Trans-UNet	90.79	99.65	88.81	Backdoor	90.09 (-0%)	99.66(-0%)	88.76(-0%)	84.81 (-7%)	98.79(-0.9%)	83.55(-6%)
					Ours	90.70 (-0%)	99.69(-0%)	83.79(-5%)	87.31(-4%)	99.51 (-0.1%)	85.43(-4%)
	SAM	91.87	99.73	90.80	Backdoor	92.28 (-3%)	99.72(-0%)	90.50(-0%)	93.07 (-1%)	99.72(-0%)	90.62(-0%)
					Ours	91.81 (-1%)	99.73(-0%)	89.73(-1%)	92.66(-0.5%)	99.73(-0%)	91.24(-0%)
	MedSAM	94.79	99.78	93.45	Backdoor	92.39 (-3%)	99.75(-0%)	93.20(-0%)	93.84(-1%)	99.72(-0%)	93.41(-0%)
					Ours	94.40(-0.3%)	99.77(-0%)	93.40(-0%)	94.69(-0%)	99.73(-0%)	93.40(-0%)
SEG	nnUNet	50.10	75.40	93.50	Backdoor	48.21(-4%)	74.64(-1%)	90.48(-3%)	49.09 (-2%)	73.89(-2%)	93.02(-0.5%)
					Ours	50.10(-0%)	75.40 (-0%)	93.44(-0%)	49.75 (-0.7%)	75.39 (-0%)	93.52(-0%)
	Swin-UNet	69.34	95.37	75.69	Backdoor	65.90(-5%)	93.79(-1.6%)	75.61(-0%)	64.47(-7%)	93.89(-1.5%)	74.59(-1.4%)
					Ours	68.91 (-0.7%)	95.05(-0.3%)	75.61(-0%)	66.06(-5%)	95.03(-0.3%)	73.12(-3%)
	Trans-UNet	55.55	90.73	58.72	Backdoor	47.90 (-14%)	80.08(-11%)	48.10(-18%)	51.94 (-7%)	87.45(-1.4%)	55.93(-4%)
					Ours	55.41(-0.1%)	90.53 (-0.2%)	58.23(-1%)	54.14(-1.3%)	89.56 (-0%)	57.68(-2%)
	SAM	60.73	94.17	69.23	Backdoor	48.31 (-21%)	72.02(-23%)	48.31(-30%)	57.77 (-5%)	93.23(-1%)	65.44(-5%)
					Ours	60.33 (-0.8%)	94.12(-0%)	67.87(-2%)	58.68 (-4%)	93.82(-0.4%)	65.98(-4.7%)
	MedSAM	67.36	96.72	77.85	Backdoor	53.42(-20%)	82.21(-15%)	56.05(-28%)	63.31(-6%)	95.75(-1%)	75.51(-3%)
					Ours	67.12(-0.3%)	96.68(-0%)	75.51(-3%)	65.33(-3%)	96.70(-0%)	76.29(-2%)
EchoNet	nnUNet	92.28	96.08	91.33	Backdoor	90.76(-2%)	96.00(-0%)	91.33(-1%)	91.26(-1%)	96.12(-0%)	91.32(-0%)
					Ours	92.27(-0%)	96.05(-0%)	91.35(-0%)	92.24(-0%)	96.08(-0%)	91.30(-0%)
	Swin-UNet	95.37	99.82	94.60	Backdoor	95.16 (-2%)	99.76(-0%)	94.37(-0%)	95.38 (-3%)	99.80(-0%)	94.60(-0%)
					Ours	95.40(-0%)	99.80(-0%)	94.63(-0%)	95.33 (-0%)	99.80(-0%)	94.60(-0%)
	Trans-UNet	94.27	99.73	93.50	Backdoor	94.04 (-2%)	99.68(-0%)	93.19(-3%)	93.40 (-3%)	99.55 (-0.2%)	92.60(-1%)
					Ours	94.59 (-0%)	99.73(-0%)	93.80(-0%)	94.15(-0%)	99.67 (-0%)	93.36(-0.1%)
	SAM	95.30	99.79	94.53	Backdoor	72.26 (-14%)	38.53(-61%)	72.26(-24%)	95.21 (-1%)	99.79 (-0.4%)	94.41(-0%)
					Ours	95.39 (-0%)	99.79(-0%)	94.66(-0%)	95.44(-0%)	99.80 (-0%)	94.69(-0%)
	MedSAM	96.22	99.89	96.42	Backdoor	92.48 (-4%)	98.89 (-1%)	91.60 (-5%)	95.25 (-2%)	99.88 (-0%)	96.42 (-0%)
					Ours	96.15(-0%)	99.85(-0%)	96.44(-0%)	96.21(-0%)	99.90 (-0%)	96.42(-0%)
PraNet	nnUNet	71.39	89.74	71.03	Backdoor	55.51 (-22%)	52.94 (-41%)	39.77(-44%)	42.12 (-41%)	59.23(-34%)	39.06(-45%)
					Ours	71.34 (-0%)	88.84 (-1%)	67.48(-5%)	69.96 (-2%)	89.74(-0%)	68.60(-2%)
	Swin-UNet	73.18	92.08	72.90	Backdoor	32.10 (-29%)	48.72 (-47%)	34.05(-53%)	27.92 (-38%)	50.62(-45%)	28.36(-61%)
					Ours	70.28 (-0%)	90.03 (-2%)	69.89(-4%)	74.12(-0%)	89.79(-2%)	72.92(-0%)
	Trans-UNet	57.38	83.89	58.13	Backdoor	53.49 (-6%)	48.97(-41%)	30.80(-47%)	41.74 (-27%)	50.56(-39%)	30.60(-47%)
					Ours	57.12 (-1%)	82.61 (-1%)	52.70(-9.3%)	53.86 (-7%)	80.60 (-4%)	55.66(-4%)
	SAM	90.52	94.74	85.44	Backdoor	55.07 (-39%)	51.76 (-45%)	47.25(-44%)	88.32 (-2%)	50.94 (-46%)	48.74(-42%)
					Ours	90.83(-0%)	94.90(-0%)	85.24(-0.2%)	90.00 (-1%)	94.99(-0%)	85.43(-0%)
	MedSAM	91.94	95.65	88.10	Backdoor	58.95 (-36%)	57.39(-40%)	46.69(-47%)	90.10 (-2%)	84.17(-12%)	68.71(-22%)
					Ours	90.86(-1%)	94.69 (-1%)	87.21(-1%)	91.02 (-1%)	95.40(-0%)	88.12 (-0%)

watermarking techniques, which are among the most representative and widely adopted strategies for verifying black-box model ownership. While these methods primarily differ in how their trigger sets are constructed, they share the core principle of embedding triggers that cause the model to produce specific, verifiable outputs. Our study follows established practices in the backdoor watermarking literature by evaluating our method under four representative trigger types: patch triggers, text overlays, random noise, and black edge patterns. Specifically, the patch and text overlay triggers are adapted from the method proposed by Zhang et al. [15], while the random noise trigger follows the design introduced by Lounici et al. [47]. The black edge pattern, widely used in recent studies [4], serves as a simple yet effective trigger type to simulate diverse real-world watermarking scenarios. Most backdoor-based watermarking methods are designed for image classification, where ownership is verified by forcing inputs with triggers to be misclassified into a target class. When directly applied to binary segmentation, these methods tend to produce extreme and unstable outputs: for example, background pixels that should be labeled as 0 are misclassified as 1, and foreground pixels that should be labeled as 1 are flipped to 0.

In contrast, our approach is more subtle, involving only

TABLE IV
THE COMPARISON EXPERIMENT BETWEEN STEALTHMARK AND BACKDOOR-BASED MODELS WATERMARK METHOD UNDER DIFFERENT TRIGGERS (PART II)

Dataset	Model	Clean Model			Trigger→	Noise [47]			Black Edge		
		Dice	AUC	VS	Method↓	Dice	AUC	VS	Dice	AUC	VS
UKBB	nnUNet	94.73	97.35	95.42	Backdoor	90.93(-4%)	97.35(-0%)	95.22(-0%)	89.09(-6%)	97.22(-0%)	92.42(-3%)
					Ours	94.67(-0%)	97.35(-0%)	95.42(-0%)	92.83(-2%)	97.33(-0%)	93.24(-2%)
	Swin-UNet	93.30	99.78	91.42	Backdoor	90.90 (-3%)	99.76(-0%)	88.75(-3%)	94.27 (-6%)	99.72(-0%)	90.21(-1%)
					Ours	89.73 (-2%)	99.76 (-0%)	87.62(-4%)	87.35(-0.2%)	99.78 (-0%)	90.21 (-1%)
	Trans-UNet	90.79	99.65	88.81	Backdoor	90.86 (-0%)	99.38(-0%)	88.67(-0%)	87.41 (-4%)	99.32 (-0.3%)	85.61(-3.6%)
					Ours	90.77 (-0%)	99.65(-0%)	88.795(-0%)	85.28 (-2%)	99.61(-0%)	86.58(-2%)
	SAM	91.87	99.73	90.80	Backdoor	92.15(-2%)	99.70 (-0%)	90.47(-0%)	90.71(-4%)	99.66 (-0%)	88.61 (-2.4%)
					Ours	92.17 (-2%)	99.73(-0%)	90.54(-0%)	92.16 (-2%)	99.67(-0%)	90.47(-0%)
	MedSAM	94.79	99.78	93.45	Backdoor	92.89 (-2%)	99.75(-0%)	93.41(-0%)	92.98 (-2%)	99.75(-0%)	91.15 (-2%)
					Ours	93.84(-1%)	99.80(-0%)	93.20(-0%)	93.93 (-1%)	99.78(-0%)	93.42 (-0%)
SEG	nnUNet	50.10	75.40	93.50	Backdoor	49.14 (-2%)	75.40(-0%)	93.40(-0%)	48.69(-3%)	73.84(-2%)	88.54(-5%)
					Ours	49.84 (-0.5%)	75.42 (-0%)	92.83(-0.7%)	49.64(-1%)	75.36 (-0%)	91.05(-2.6%)
	Swin-UNet	69.34	95.37	75.69	Backdoor	64.42 (-7%)	95.38(-0%)	70.08(-7%)	63.70 (-8%)	93.69(-1.7%)	72.68(-4%)
					Ours	68.09 (-2%)	95.38 (-0%)	71.16(-6%)	68.81 (-1%)	95.45 (-0%)	75.82(-0%)
	Trans-UNet	55.55	90.73	58.72	Backdoor	54.31 (-2%)	90.18 (-0%)	58.57(-0%)	47.96 (-16%)	78.13 (-13%)	47.96 (-18%)
					Ours	55.93(-0%)	90.47 (-0%)	57.66(-2%)	56.36(-0%)	90.22(-0%)	58.35(-0.6%)
	SAM	60.73	94.17	69.23	Backdoor	48.41 (-26%)	21.42(-77%)	48.41(-30%)	59.59 (-5%)	94.09(-0%)	67.10(-3%)
					Ours	58.41 (-5%)	93.75 (-0.4%)	65.23(-5.8%)	59.44 (-2%)	94.20 (-0%)	66.69 (-3.6%)
	MedSAM	67.36	96.72	77.85	Backdoor	63.99(-5%)	94.78(-2%)	73.95(-5%)	64.66(-4%)	96.62(-0%)	76.29(-2%)
					Ours	63.99 (-5%)	95.75(-1%)	74.34(-5%)	66.68(-1%)	96.58(-0%)	76.44(-2%)
EchoNet	nnUNet	92.28	96.08	91.33	Backdoor	91.44(-1%)	96.00 (-0%)	90.41(-1%)	90.43 (-2%)	95.59(-0.5%)	90.69(-0.7%)
					Ours	92.28 (-0%)	96.05(-0%)	91.33(-0%)	91.37(-1%)	96.01(-0%)	90.44 (-1%)
	Swin-UNet	95.37	99.82	94.60	Backdoor	95.38 (-1%)	99.81 (-0%)	94.60(-0%)	94.85 (-2%)	99.75(-0%)	94.01(-0.6%)
					Ours	95.53 (-0%)	99.79 (-0%)	94.76(-0%)	95.45 (-0%)	99.81 (-0%)	94.68 (-0%)
	Trans-UNet	94.27	99.73	93.50	Backdoor	93.27 (-2%)	99.53(-0.2%)	92.53(-1%)	91.82 (-3%)	99.32 (-0.4%)	91.10 (-2.5%)
					Ours	94.17 (-0%)	99.71 (-0%)	93.39(-0.1%)	94.17 (-0%)	99.56 (-0.1%)	93.39 (-0.1%)
	SAM	95.30	99.79	94.53	Backdoor	72.26 (-24%)	39.47(-60%)	72.26(-23%)	95.29 (-1%)	99.81 (-0.4%)	94.54(-0%)
					Ours	95.45 (-0%)	99.80(-0%)	94.69(-0%)	95.42 (-0%)	99.79(-0%)	94.66(-0%)
	MedSAM	96.22	99.89	96.42	Backdoor	92.74 (-4%)	99.89 (-1%)	95.45 (-1%)	95.16 (-2%)	99.88(-0%)	96.40(-0%)
					Ours	96.24(-0%)	99.89(-0%)	96.44(-0%)	96.19 (-0%)	99.88(-0%)	96.45(-0%)
PraNet	nnUNet	71.39	89.74	71.03	Backdoor	56.39 (-21%)	77.17(-14%)	44.74(-37%)	29.26 (-59%)	49.35(-45%)	34.09 (-52%)
					Ours	70.67(-1%)	89.74(-0%)	71.02(-0%)	64.32(-10%)	82.56(-8%)	61.07(-14%)
	Swin-UNet	73.18	92.08	72.90	Backdoor	32.41 (-22%)	48.83(-46%)	34.92(-52%)	19.06 (-58%)	32.97(-64%)	22.13(-69%)
					Ours	60.30 (-5%)	85.75 (-6%)	60.75(-16%)	68.38 (-8%)	85.57 (-7%)	68.57(-6%)
	Trans-UNet	57.38	83.89	58.13	Backdoor	30.71 (-47%)	55.76 (-33%)	11.91(-79%)	33.97 (-41%)	60.04(-28%)	21.42(-63%)
					Ours	57.71 (-0%)	81.17 (-3%)	57.80(-0%)	56.07 (-3%)	80.70 (-4%)	57.70(-0.7%)
	SAM	90.52	94.74	85.44	Backdoor	82.40 (-10%)	49.10 (-48%)	45.55(-46%)	65.71 (-27%)	48.61 (-48%)	44.31(-48%)
					Ours	90.14 (-0%)	94.33 (-0%)	85.61(-0%)	90.02 (-1%)	94.36 (-0.4%)	85.44 (-0%)
	MedSAM	91.94	95.65	88.10	Backdoor	84.58 (-8%)	72.69(-24%)	71.36(-19%)	69.87 (-24%)	60.26 (-37%)	46.69 (-47%)
					Ours	91.85 (-0%)	95.60(-0%)	88.10 (-0%)	90.10 (-2%)	95.07(-0.6%)	88.12(-0%)

a small adjustment to the background region's minimum intensity (increased from 0 to 0.02) or the foreground region's maximum intensity (decreased from 1 to 0.85). As shown in Table III and Table IV, experimental results demonstrate that our method achieves superior harmlessness. It is worth noting that datasets such as EchoNet and UKBB are relatively simple in terms of structure and content, which makes the models less sensitive to injected triggers—hence, backdoor-based watermarking methods tend to cause only minor performance degradation on these datasets. In contrast, datasets like PraNet and SEG present significantly more challenging segmentation tasks. Not only do they feature complex anatomical or pathological structures, but they also naturally contain patterns visually similar to certain types of triggers. This makes the models more vulnerable to overfitting on trigger patterns, leading to more severe performance degradation when traditional backdoor watermarking is applied. On datasets such as UKBB, SEG, EchoNet, and PraNet, using models like nnUNet, Swin-UNet, and MedSAM, the Dice, AUC and Volume Similarity scores under watermarking remain nearly identical to their clean counterparts. For instance, on the PraNet dataset with patch triggers, SAM's Dice score decreases by only 1%, Trans-UNet's Dice remains unchanged, and its AUC remains virtually unchanged. In contrast, backdoor-based watermarking

methods can degrade segmentation performance by up to 39%, severely limiting their practical deployment.

To provide an intuitive understanding of StealthMark's harmlessness and reliability, we visualize segmentation results across four medical datasets (EchoNet [18], SEG [17], PraNet [19], and UKBB) using nnUNet [2], as shown in Figure 5. Our method maintains almost identical segmentation quality compared to the clean model across all four datasets, confirming the harmlessness property of our watermark design. The extracted watermarks are visually consistent with their corresponding targets, demonstrating reliable and stable ownership verification. Furthermore, to evaluate the feature-level

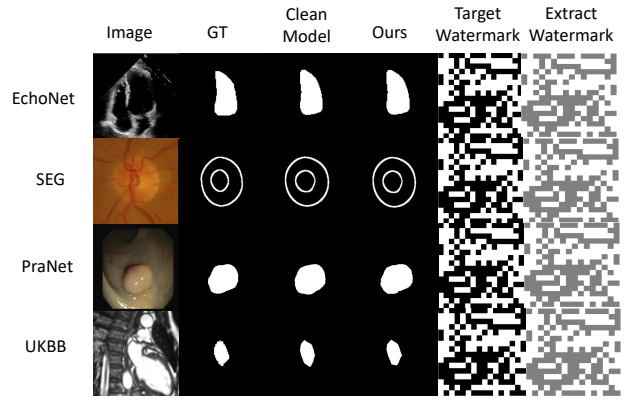

Fig. 5. Qualitative results on four medical segmentation datasets (EchoNet, SEG [17], PraNet [19], and UKBB) using nnUNet [2]. For each dataset, we show the input image, ground truth mask (GT), prediction from the clean (unwatermarked) model, and prediction from our watermarked model. The right columns show the target watermark and the extracted watermark, confirming reliable ownership verification.

stealthiness, we visualize the feature maps before the segmentation head in SAM-based models on the EchoNet dataset with patch trigger and show examples of our method with foreground constraint and background constraint, compared with the backdoor-based method (Figure 6). Unlike traditional backdoor-based methods, our method maintains a high similarity between the clean and triggered feature distributions, making it more stealthy against detection or removal attacks.

C. Ablation Study

In this ablation study, we evaluate the effectiveness and stealthiness of our method under different parameter settings, focusing on trigger size and training parameters. Using the SAM model on the PraNet and EchoNet datasets, we systematically vary trigger size and batch size to analyze model performance and sensitivity to these configurations.

1) Selecting the Optimal Patch Size for Effectiveness and Stealthiness: We will determine the optimal Patch Size value of the model by ensuring the **Dice**, **AUC**, **Volume Similarity** and **ASR** of the segmentation performance indicators of the model after using our method. Our goal is to determine the optimal patch size that ensures guaranteed ASR performance while minimizing the impact on model segmentation performance. The experiment is conducted on PraNet with SAM as a case study, using an image resolution of 256 $\times$ 256.

The results are presented in Table V. It can be seen from the experimental results that when the patch size is 4, the ASR of the model reaches 99.02%, and the gap between Dice, AUC and Volume Similarity also tends to be within the normal range of the clean model results. From this, we determined the best patch size for StealthMark.

TABLE V
RESULTS WITH DIFFERENT TRIGGER SIZE (IMAGE SIZE IS 256) ON PRANET DATASET

Patch Size	Dice	AUC	VS	ASR
2	90.27	95.05	85.29	53.98
4	90.28	94.65	85.84	99.02
8	90.09	94.69	85.40	98.90
16	90.42	94.69	85.44	99.51
32	90.26	93.68	85.57	97.67
Clean model	90.52	94.74	85.44	-

2) Selecting the Optimal Hyper-Parameter δ or τ for Effectiveness and Stealthiness: In this experiment, we conducted a case study using the SAM model's performance on the EchoNet dataset, employing a text trigger. We evaluated the effectiveness and stealthiness of our method by varying the parameters δ and τ. The experimental results are presented in Table VI. We initially defined the experimental range for these two hyperparameters as (τ: 0.8–0.99, δ: 0.001–0.1). When δ was set to 0.05, our method achieved an ASR of 99.60% on EchoNet, while the segmentation performance remained indistinguishable from the clean model (Dice: 95.30, AUC: 99.80, VS: 93.59), ensuring effectiveness. Furthermore, the slight difference between 0.05 and the original value of 0 demonstrates strong stealthiness. Similarly, when τ was set to 0.85, the method achieved an ASR of 99.50% on EchoNet, again without significantly affecting segmentation performance (Dice: 95.30, AUC: 99.85, VS: 94.56). These findings confirm that the chosen hyperparameters ($\delta = 0.05$, $\tau = 0.85$) provide a favorable trade-off between watermarking effectiveness and concealment. The high ASR, coupled with a negligible effect on segmentation quality, suggests that StealthMark can embed ownership without compromising the utility of the model or revealing its presence, thus fulfilling the core design goal of the method.

TABLE VI
RESULTS WITH DIFFERENT HYPER-PARAMETER δ AND τ ON ECHONET DATASET

δ	Dice	AUC	VS	ASR
0.100	95.39	99.79	93.80	99.51
0.050	95.30	99.80	93.59	99.50
0.020	94.64	99.31	93.90	81.51
0.010	94.49	99.75	93.62	72.08
0.005	94.59	99.80	93.76	65.77
0.001	94.54	99.80	93.68	50.70
τ	Dice	AUC	VS	ASR
0.99	95.30	99.80	94.53	57.31
0.95	95.30	99.84	94.53	81.34
0.90	95.32	99.80	94.58	97.41
0.85	95.30	99.85	94.56	99.60
0.80	95.30	99.80	94.47	99.90
Clean model	95.30	99.79	94.53	-

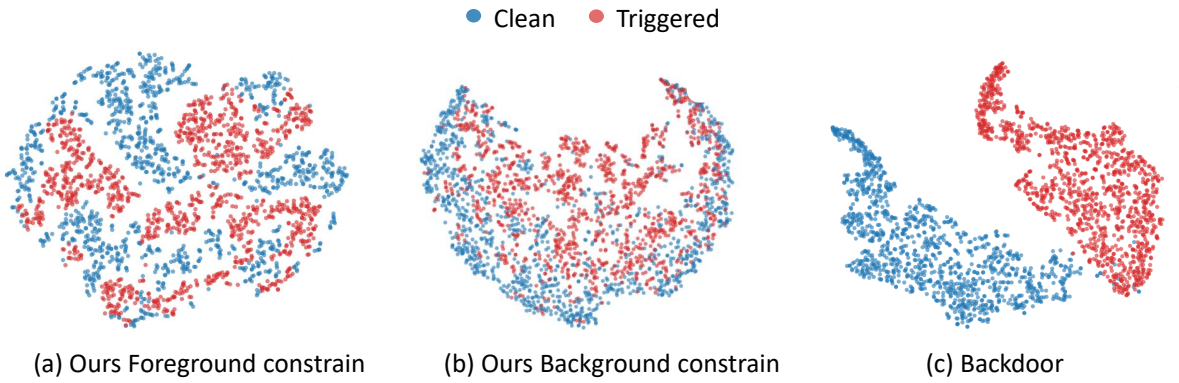

Fig. 6. The t-SNE visualization of features (before the segmentation head) on the EchoNet test set using the SAM [22] model under patch-based trigger attacks. Subfigures **(a)** and **(b)** show examples of our method with foreground constraint and background constraint, respectively. In contrast, **(c)** shows a representative result from traditional backdoor methods (*e.g.,* BadNet [7]). Compared to the prior approach, our method yields significantly more entangled representations between clean and triggered samples, demonstrating higher stealthiness in the feature space.

3) Tuning the Hyper-parameter λ_{bg} for Balance: In this part of the experiment, we further utilize the EchoNet dataset to evaluate the influence of the hyperparameter λ_{bg} on the SAM model as a case study for the effectiveness and stealthiness of the proposed method. We vary λ_{bg} and the minimum value δ within a predefined range to investigate its impact on both segmentation performance and the attack success rate (ASR). The results are summarized in Table VII.

We observe that once the value δ reaches approximately 0.02 or lower, further increasing λ_{bg} has a negligible effect on ASR. However, an excessively large value of λ_{bg} tends to degrade segmentation performance on clean data. Empirically, we find that setting $\delta = 0.05$ and $\lambda_{bg} = 1.0$ yields the best trade-off, achieving a balance between watermark robustness and model fidelity.

TABLE VII
PERFORMANCE UNDER DIFFERENT λ_{BG} AND δ COMBINATIONS ON ECHONET DATASET

λ_{bg}	δ	Dice	AUC	VS	ASR
0.5	0.05	95.33	99.78	94.19	95.62
1.0	0.05	95.30	99.80	93.59	99.50
3.0	0.05	92.41	99.22	93.05	99.48
0.5	0.02	94.80	99.80	94.20	75.94
1.0	0.02	94.64	99.31	93.90	81.51
3.0	0.02	92.64	99.46	93.90	81.80
0.5	0.01	94.90	99.58	94.28	53.62
1.0	0.01	94.49	99.75	93.62	72.01
3.0	0.01	94.05	99.80	94.12	73.59
Clean model		95.30	99.80	94.47	-

V. DISCUSSION

A. False Positive Rate (FPR) Analysis on Natural Artifacts.

We discussed that medical images may naturally contain trigger-like patterns caused by acquisition artifacts, which could result in false watermark activations. To quantitatively verify this concern, we conducted a False Positive Rate (FPR) test on three representative segmentation models: nnUNet [2], Swin-UNet [20], SAM [22], using the PraNet dataset, which is known for frequent artifact-like patterns such as black edge, text patch and color patch. We applied our watermark detector to all clean test samples and measured the frequency of false activations for different trigger patterns. The results in Table VIII clearly demonstrate that the "false trigger" challenge is real and significant. For simple triggers resembling natural structures (e.g., Black Edge), the FPR exceeds 40%, indicating that traditional backdoor methods are unreliable in medical segmentation. Notably, for more complex and artificial triggers such as Text or Patch, the false positive rates are extremely low, especially on SAM (e.g., 0.0012% and 0.0073%, respectively). This highlights SAM's strong discriminative capability. However, this phenomenon is only observed on SAM, other models remain highly susceptible to natural artifactlike triggers after being backdoored, indicating that such interference is difficult to eliminate in conventional architectures. This also explains why the Dice score of standard backdoor models (*e.g.,* BadNet) drops drastically on clean medical data. In contrast, our StealthMark design directly resolves this problem. Because the watermark embedding is guided by uncertainty modeling, even if a trigger pattern accidentally resembles a natural artifact, the segmentation output remains stable and diagnostically consistent. For example, the SAM-based model shows almost no difference in Dice score between clean and falsely activated inputs (90.02% vs. 90.52%).

TABLE VIII
FALSE POSITIVE RATE OF TRIGGERED MODELS FOR DIFFERENT TRIGGER PATTERNS ON PRANET DATASET.

Model	**Black Edge**	**Noise**	**Text**	**Patch**
nnUNet [2]	55.94%	15.20%	23.81%	14.93%
Swin-UNet [20]	52.51%	35.74%	27.91%	15.79%
SAM [22]	43.45%	0.00%	0.0012%	0.0073%

B. Robustness Analysis Against Watermark Removal Attacks.

To further validate the robustness of *StealthMark* against common watermark removal strategies, we performed two sets of defence experiments: model pruning and fine-tuning. We evaluated three representative segmentation models: nnUNet [2], Swin-UNet [20], and SAM [22] on the PraNet dataset using the "Color Patch" trigger.

Model Pruning: We applied global magnitude pruning at different sparsity levels (10%, 30%, 50%) to all convolutional (Conv2d) and fully connected (Linear) layers. Table IX reports the effect of pruning on the watermark detection rate (ASR). Even when the model parameters were pruned by up to 50%, the watermark signal remained clearly detectable, showing only a minor decrease in ASR. This demonstrates that StealthMark retains watermark integrity under severe parameter sparsification.

Model Fine-tuning: We then fine-tuned the watermarked models for 5 epochs using clean data subsets of varying sizes (1%, 10%, and 25% of the original training set). Table X summarizes the ASR performance after fine-tuning. While fine-tuning slightly reduced the ASR values, the watermark signal remained detectable, demonstrating that StealthMark maintains a strong degree of resilience even under model re-optimization.

TABLE IX
ASR PERFORMANCE UNDER DIFFERENT PRUNING RATIOS ON PRANET DATASET.

Model	0% (Baseline)	10% Pruning	30% Pruning	50% Pruning)
nnUNet [2]	88.78%	88.78%	86.16%	80.61%
Swin-UNet [20]	79.90%	79.90%	79.17%	76.47%
SAM [22]	96.57%	97.06%	92.89%	89.71%

TABLE X
ASR PERFORMANCE UNDER DIFFERENT FINE-TUNING RATIOS ON PRANET DATASET.

Model	0% (Baseline)	Fine-tuned 1%	Fine-tuned 10%	Fine-tuned 25%
nnUNet [2]	88.78%	86.48%	74.19%	69.73%
Swin-UNet [20]	82.37%	82.17%	78.18%	76.62%
SAM [22]	98.16%	90.38%	78.79%	72.09%

C. Out-of-Distribution (OOD) Generalization of Triggers.

To assess whether *StealthMark* genuinely learns the concept of the trigger rather than simply memorizing known patterns, we conducted an Out-of-Distribution (OOD) generalization evaluation on the SAM [22] model across four distinct datasets: UKBB, SEG, EchoNet, PraNet. This experiment evaluates the model's robustness when trigger parameters are shifted outside the training distribution. Specifically, we modified the triggers along three orthogonal dimensions: **Scale Shift:** the trigger patch is enlarged to twice its training size. **Rotation Shift:** the trigger patch is rotated by 30 degrees relative to the trained orientation. **Positional Shift:** The trigger is moved to a different image corner not used during training. We compared the Attack Success Rate (ASR) for In-Distribution (ID) triggers and OOD variants. The results, summarized in Table XI show that *StealthMark* achieves consistently high ASR even when the trigger is scaled, rotated, or spatially displaced, demonstrating strong generalization beyond memorized trigger patterns.

D. Clinically harmless nature of the trigger.

We conducted a blind evaluation, where five senior clinical experts independently assessed randomly mixed original and watermarked images without knowing which group each image belonged to. The evaluation covered four types of watermark triggers (patch, text, black Edge, and noise). We also visualized these four trigger types in Figure 7 to demonstrate their visual characteristics. Black Edge and noise triggers were visually indistinguishable from the original images. Although the text and patch triggers were noticeable, the experts confirmed that they did not affect diagnostic interpretation, as they did not overlap with any clinically relevant regions. Overall, no clinically observable degradation was identified under the evaluated settings.

E. Limitations and Future Work.

Several important limitations of this study must be emphasized. First, our verification framework assumes access to outputs as continuous probability maps. While this assumption may not be satisfied for all black-box APIs, it is a realistic and clinically motivated assumption. Second, the conclusion regarding the clinical harmlessness of the proposed watermark is based on a qualitative evaluation conducted by a small number of senior clinical experts. As such, this assessment may not fully capture extremely fine-grained or rare clinical scenarios, where subtle visual alterations could potentially influence diagnosis.

A systematic study of different watermark patterns and encoding schemes is a promising direction for future work. In particular, designing an adaptive encoding strategy that is robust to noise and post-processing while preserving clinical harmlessness would further strengthen the reliability of the verification pipeline. Moreover, future work will focus on enhancing the watermark's robustness against advanced model modification attacks, including knowledge distillation, model pruning, and fine-tuning.

TABLE XI
OOD GENERALIZATION OF PATCH TRIGGERS (MODEL: SAM)

Dataset	ID Baseline	Scale Shift	Rotational Shift	Positional Shift
PraNet [19]	98.16%	97.80%	97.43%	81.03%
SEG [17]	99.30%	99.02%	90.91%	98.74%
UKBB [3]	98.65%	98.72%	97.52%	98.37%
EchoNet [18]	99.92%	99.88%	98.75%	99.76%

VI. CONCLUSION

Our work addresses the critical but unexplored area of copyright protection for medical image segmentation models. We introduced a novel black-box model ownership verification method tailored explicitly to medical segmentation tasks. By explicitly controlling predictive uncertainty within background

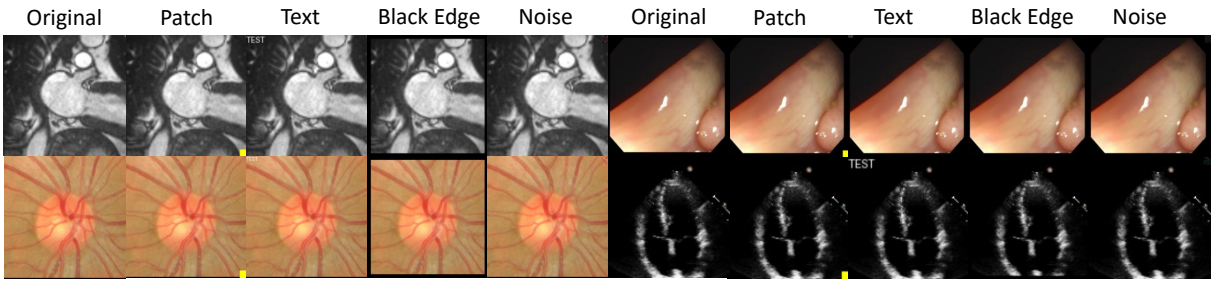

Fig. 7. Visualization of the four watermark trigger types on four datasets. Trigger types include patch, text, black edge, and noise.

or foreground regions and integrating a novel uncertainty-aware loss function into the standard training paradigm, we successfully balanced watermark robustness with minimal performance degradation. Moreover, employing the LIME interpretability framework, we embedded stealthy QR code watermarks that only manifest under specific triggers, thereby overcoming the shortcomings of traditional methods in terms of detectability. Our comprehensive evaluation across multiple datasets and segmentation architectures validates the effectiveness, harmlessness, and strong stealthiness of the proposed method. The approach substantially outperforms traditional backdoor techniques by maintaining high segmentation accuracy and offers considerable promise for safeguarding intellectual property in medical segmentation tasks.

REFERENCES

[1] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, 2017.

[2] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature Methods*, 2021.

[3] Y. Li, Y. Bai, Y. Jiang, Y. Yang, S.-T. Xia, and B. Li, "Untargeted backdoor watermark: Towards harmless and stealthy dataset copyright protection," *Advances in neural information processing systems*, 2022.

[4] S. Shao, Y. Li, H. Yao, Y. He, Z. Qin, and K. Ren, "Explanation as a watermark: Towards harmless and multi-bit model ownership verification via watermarking feature attribution," in *NDSS*, 2025.

[5] J. Zhang, D. Chen, J. Liao, H. Fang, W. Zhang, W. Zhou, H. Cui, and N. Yu, "Model watermarking for image processing networks," in *Proceedings of the AAAI conference on artificial intelligence*, 2020.

[6] Y. Zhao, T. Pang, C. Du, X. Yang, N.-M. Cheung, and M. Lin, "A recipe for watermarking diffusion models," *arXiv:2303.10137*, 2023.

[7] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, "Badnets: Evaluating backdooring attacks on deep neural networks," *IEEE Access*, 2019.

[8] R. Jin, C.-Y. Huang, C. You, and X. Li, "Backdoor attack on unpaired medical image-text foundation models: A pilot study on medclip," in *2024 SaTML*. IEEE, 2024, pp. 272–285.

[9] W. Bender, D. Gruhl, N. Morimoto, and A. Lu, "Techniques for data hiding," *IBM systems journal*, 1996.

[10] J. Zhu, R. Kaplan, J. Johnson, and L. Fei-Fei, "Hidden: Hiding data with deep networks," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 657–672.

[11] Y. Adi, C. Baum, M. Cisse, B. Pinkas, and J. Keshet, "Turning your weakness into a strength: Watermarking deep neural networks by backdooring," in *USENIX Security 18*, 2018, pp. 1615–1631.

[12] Y. Ding, Z. Wang, Z. Qin, E. Zhou, G. Zhu, Z. Qin, and K.-K. R. Choo, "Backdoor attack on deep learning-based medical image encryption and decryption network," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 280–292, 2023.

[13] C. Song, T. Ristenpart, and V. Shmatikov, "Machine learning models that remember too much," in *ACM SIGSAC*, 2017, pp. 587–601.

[14] X. Liu, S. Shao, Y. Yang, K. Wu, W. Yang, and H. Fang, "Secure federated learning model verification: A client-side backdoor triggered watermarking scheme," in *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2021, pp. 2414–2419.

[15] J. Zhang, Z. Gu, J. Jang, H. Wu, M. P. Stoecklin, H. Huang, and I. Molloy, "Protecting intellectual property of deep neural networks with watermarking," in *Proceedings of the 2018 on Asia conference on computer and communications security*, 2018, pp. 159–172.

[16] M. T. Ribeiro, S. Singh, and C. Guestrin, "" why should i trust you?" explaining the predictions of any classifier," in *ACM SIGKDD*, 2016.

[17] Y. Meng, Y. Zhang, J. Xie, J. Duan, M. Joddrell, S. Madhusudhan, T. Peto, Y. Zhao, and Y. Zheng, "Multi-granularity learning of explicit geometric constraint and contrast for label-efficient medical image segmentation and differentiable clinical function assessment," *Medical Image Analysis*, 2024.

[18] D. Ouyang, B. He, A. Ghorbani, N. Yuan, J. Ebinger, C. P. Langlotz, P. A. Heidenreich, R. A. Harrington, D. H. Liang, E. A. Ashley *et al.*, "Video-based ai for beat-to-beat assessment of cardiac function," *Nature*, 2020.

[19] D.-P. Fan, G.-P. Ji, T. Zhou, G. Chen, H. Fu, J. Shen, and L. Shao, "Pranet: Parallel reverse attention network for polyp segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020.

[20] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, and M. Wang, "Swin-unet: Unet-like pure transformer for medical image segmentation," in *European conference on computer vision*. Springer, 2022, pp. 205–218.

[21] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, and Y. Zhou, "Transunet: Transformers make strong encoders for medical image segmentation," *arXiv preprint arXiv:2102.04306*, 2021.

[22] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, October 2023, pp. 4015–4026.

[23] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," *Nature Communications*, vol. 15, no. 1, p. 654, 2024.

[24] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," *arXiv:1712.05526*, 2017.

[25] X. Kang, J. Huang, Y. Q. Shi, and Y. Lin, "A dwt-dft composite watermarking scheme robust to both affine transform and jpeg compression," *IEEE transactions on circuits and systems for video technology*, 2003.

[26] A. Hanif, F. Shamshad, M. Awais, M. Naseer, F. S. Khan, K. Nandakumar, S. Khan, and R. M. Anwer, "Baple: Backdoor attacks on medical foundational models using prompt learning," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2024, pp. 443–453.

[27] M. Nwadike, T. Miyawaki, E. Sarkar, M. Maniatakos, and F. Shamout, "Explainability matters: Backdoor attacks on medical imaging," *arXiv:2101.00008*, 2020.

[28] Y. Feng, B. Ma, J. Zhang, S. Zhao, Y. Xia, and D. Tao, "Fiba:frequency-injection based backdoor attack in medical image analysis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

[29] Z. Yang, Y. Chen, M. Sun, and Y. Zhang, "Inject backdoor in measured data to jeopardize full-stack medical image analysis system," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2024, pp. 393–402.

[30] M. Lin, N. Weng, K. Mikolaj, Z. Bashir, M. B. Svendsen, M. G. Tolsgaard, A. N. Christensen, and A. Feragen, "Shortcut learning in medical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2024, pp. 623–633.

[31] Y. Hu, W. Kuang, Z. Qin, K. Li, J. Zhang, Y. Gao, W. Li, and K. Li, "Artificial intelligence security: Threats and countermeasures," *ACM Computing Surveys (CSUR)*, vol. 55, no. 1, pp. 1–36, 2021.

[32] S. Klein *et al.*, "Grand-challenge.org: a platform for end-to-end development of machine learning solutions in biomedical imaging," *Medical Image Analysis*, 2020.

[33] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in neural information processing systems*, vol. 30, 2017.

[34] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 3319–3328.

[35] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: visual explanations from deep networks via gradient-based localization," *Proceedings of the IEEE international conference on computer vision*, vol. 128, pp. 336–359, 2020.

[36] J. I. Orlando, H. Fu, J. B. Breda, K. Van Keer, D. R. Bathula, A. Diaz-Pinto, R. Fang, P.-A. Heng, J. Kim, J. Lee *et al.*, "Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs," *Medical Image Analysis*, 2020.

[37] J. Sivaswamy, S. Krishnadas, G. D. Joshi, M. Jain, and A. U. S. Tabish, "Drishti-gs: Retinal image dataset for optic nerve head (onh) segmentation," in *2014 IEEE 11th international symposium on biomedical imaging (ISBI)*. IEEE, 2014, pp. 53–56.

[38] Z. Zhang, F. S. Yin, J. Liu, W. K. Wong, N. M. Tan, B. H. Lee, J. Cheng, and T. Y. Wong, "Origa-light: An online retinal fundus image database for glaucoma analysis and research," in *2010 Annual international conference of the IEEE engineering in medicine and biology*. IEEE, 2010, pp. 3065–3068.

[39] A. Almazroa, S. Alodhayb, E. Osman, E. Ramadan, M. Hummadi, M. Dlaim, M. Alkatee, K. Raahemifar, and V. Lakshminarayanan, "Retinal fundus images for glaucoma analysis: the riga dataset," in *Medical Imaging*, 2018.

[40] F. Fumero, S. Alayón, J. L. Sanchez, J. Sigut, and M. Gonzalez-Hernandez, "Rim-one: An open retinal image database for optic nerve evaluation," in *2011 24th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE, June 2011, pp. 1–6.

[41] M. N. Bajwa, G. A. P. Singh, W. Neumeier, M. I. Malik, A. Dengel, and S. Ahmed, "G1020: A benchmark retinal fundus image dataset for computer-aided glaucoma detection," in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–7.

[42] J. Silva, A. Histace, O. Romain, X. Dray, and B. Granado, "Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer," *International journal of computer assisted radiology and surgery*, vol. 9, pp. 283–293, 2014.

[43] J. Bernal, F. J. Sánchez, G. Fernández-Esparrach, D. Gil, C. Rodríguez, and F. Vilariño, "Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians," *Computerized medical imaging and graphics*, vol. 43, pp. 99–111, 2015.

[44] N. Tajbakhsh, S. R. Gurudu, and J. Liang, "Automated polyp detection in colonoscopy videos using shape and context information," *IEEE transactions on medical imaging*, vol. 35, no. 2, pp. 630–644, 2015.

[45] D. Vázquez, J. Bernal, F. J. Sánchez, G. Fernández-Esparrach, A. M. López, A. Romero, M. Drozdzal, and A. Courville, "A benchmark for endoluminal scene segmentation of colonoscopy images," *Journal of healthcare engineering*, vol. 2017, no. 1, p. 4037190, 2017.

[46] D. Jha, P. H. Smedsrud, M. A. Riegler, P. Halvorsen, T. De Lange, D. Johansen, and H. D. Johansen, "Kvasir-seg: A segmented polyp dataset," in *MultiMedia modeling*, 2020.

[47] S. Lounici, M. Njeh, O. Ermis, M. Önen, and S. Trabelsi, "Yes we can: Watermarking machine learning models beyond classification," in *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. IEEE, 2021, pp. 1–14.