

AnoFPDM: Anomaly Detection with Forward Process of Diffusion Models for Brain MRI

Yiming Che^{1,2}, Fazle Rafsani^{1,2}, Jay Shah^{1,2}, Md Mahfuzur Rahman Siddiquee^{1,2}, Teresa Wu^{1,2}

¹Arizona State University

²ASU-Mayo Center for Innovative Imaging

Abstract

Weakly-supervised diffusion models in anomaly segmentation, which leverage image-level labels and bypass the need for pixel-level labels during training, have shown superior performance over unsupervised methods, offering a cost-effective alternative. Traditional methods that rely on iterative image reconstruction are not fully weakly-supervised due to their dependence on costly pixel-level labels for hyperparameter tuning in inference. To address this issue, we introduce Anomaly Detection with Forward Process of Diffusion Models (AnoFPDM), a fully weakly-supervised framework that operates without image reconstruction and eliminates the need for pixel-level labels in hyperparameter tuning. By using the unguided forward process as a reference, AnoFPDM dynamically selects hyperparameters such as noise scale and segmentation threshold for each input. We improve anomaly segmentation by aggregating anomaly maps from each step of the guided forward process, which strengthens the signal of anomalous regions in the aggregated anomaly map. Our framework demonstrates competitive performance on the BraTS21 and ATLAS v2.0 datasets. Code is available at <https://github.com/SoloChe/AnoFPDM>.

1. Introduction

Anomaly detection is a critical task in the medical domain, aiding radiologists in diagnostics and subsequent decision-making. Traditionally, this process requires extensive pixel-level annotations, which are not only costly but also subject to human annotator bias [22]. Although supervised non-generative models have achieved state-of-the-art performance [11], they still require a substantial amount of pixel-level labels, which are both expensive and labor-intensive to acquire. In contrast, weakly-supervised generative models have attracted attention for their ability to utilize image-level labels, categorizing images simply as healthy or unhealthy, which are more cost-effective and less prone

to bias. These models have shown superior performance in anomaly segmentation tasks compared to unsupervised models [9, 17, 24].

Among generative models, diffusion models (DMs) [7, 19, 20] are often favored due to their superior image generation capabilities compared to other models such as generative adversarial networks (GANs) [6] and variational autoencoders (VAEs) [12]. However, despite the weakly-supervised nature of current approaches, they still significantly rely on pixel-level labels, particularly during the hyperparameter tuning phase. For methods based on DMs, this tuning involves hyperparameters such as the **guidance strength**, fixed amount of noise (**noise scale**) added to the input images, and **fixed threshold** for generating the predicted pixel-level labels from anomaly map. This reliance on pixel-level labels for tuning reintroduces the associated costs and biases, undermining the potential benefits of weakly-supervised learning.

In weakly-supervised anomaly detection with DMs, the model is trained on both healthy and unhealthy samples, as proposed in [4, 8]. This training approach enables the model to learn the distribution of healthy and unhealthy samples, and subsequently map unhealthy inputs to the healthy distribution through a partial diffusion process. During this process, the partially noised unhealthy inputs are reconstructed as healthy by the iterative sampling process (backward process) of DMs, which aims to minimize the presence of anomalies in reconstructed unhealthy inputs while preserving other healthy regions as intact as possible. However, the reconstruction error tends to be accumulated due to the iterative sampling. The crucial step of anomaly detection involves thresholding the difference between the original inputs and their reconstructed counterparts, i.e., anomaly map, to obtain the predicted pixel-level labels. Prior to this, the hyperparameters are selected using grid search with pixel-level labels to optimize metrics such as the DICE score. The noise scale manages the level of noise added to the inputs, balancing the removal of anomalies against the preservation of healthy regions in the reconstructed inputs. If excessive noise is applied, the healthy regions may be

overly distorted, rendering accurate segmentation impractical. Conversely, insufficient noise might fail to eliminate the anomalies. Given that alterations to the healthy regions are inevitable, it is crucial to carefully select the noise scale. This selection should aim for a ‘sweet spot’ that optimally balances the removal of anomalous regions with the preservation of healthy areas. The guidance strength primarily mediates between the Frechet Inception Distance (FID) and the Inception Score (IS), balancing the quality and diversity of the generated samples. Additionally, the fixed threshold is employed to identify pixels within the anomalous regions on the anomaly map. Given the variable signal strength of anomalous regions in the anomaly maps across different inputs, this fixed threshold might not always yield optimal results, suggesting the need for adaptable thresholding strategies to enhance model performance.

We introduce a fully weakly-supervised framework, AnoFPDM, which leverages diffusion models equipped with classifier-free guidance [8] to obviate the need for pixel-level labels during hyperparameter tuning. During the training phase, AnoFPDM adheres to the standard training protocol of classifier-free guidance as described in [8]. In the inference stage, rather than employing the traditional iterative sampling process, i.e., image reconstruction, our framework capitalizes on the forward process of DMs. This involves progressively adding noise to the input images. Additionally, we incorporate a crucial step by obtaining both healthy-guided predictions (HGPs) and unguided predictions (UGPs) of original inputs through a one-step mapping, referred to as the healthy-guided forward process (HFP) and unguided forward process (UFP), respectively. We employ the aggregation of sub-anomaly maps (SAMs), the difference between the HGPs and original inputs at each step, to enhance the signal strength of anomalous regions, thereby improving the segmentation performance. Each SAM captures signals from anomalous regions at a specific frequency. The signal strength of anomalous regions is consolidated by aggregating anomalous signals from different frequencies, while reducing the mistakenly detected healthy signals in the aggregated anomaly map. This reduction is effective because signals from healthy regions are less consistent across frequencies. Traditional methods, although accurate in detecting anomalies, often fail to suppress healthy signals, resulting in relatively weaker anomaly signals.

There are three hyperparameters to be selected in AnoFPDM: the end step of aggregation, i.e., noise scale, threshold for the aggregated SAMs, and guidance strength. We propose a novel approach that dynamically selects the noise scale and threshold for each individual input, while maintaining a fixed guidance strength selection by utilizing UFP as a reference. Additionally, we obtain stronger signal strength of anomalous regions in the aggregated anomaly map (AAM). Overall, our contributions are as follows:

- We propose a fully weakly-supervised anomaly detection framework without noised image reconstruction, AnoFPDM, that operates without the need of pixel-level labels for hyperparameter tuning.
- We introduce a novel dynamical threshold and noise scale selection and a novel guidance strength selection for DMs on weakly-supervised anomaly detection.
- We propose a novel aggregation strategy combined with dynamical noise scale selection to enhance the signal strength of anomalous regions.

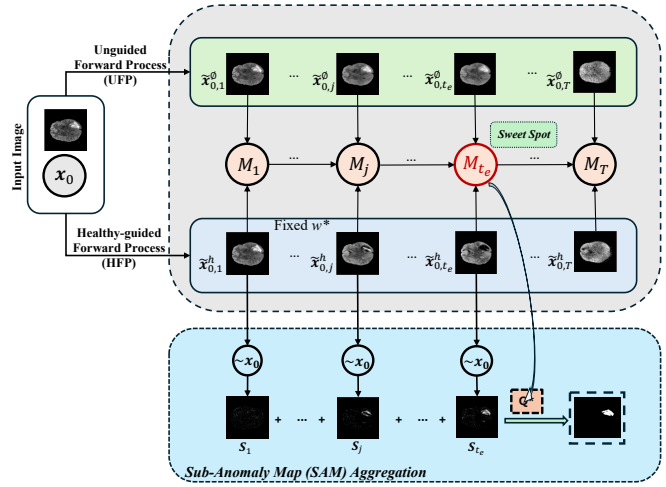


Figure 1. The diagram of the proposed method. Given the selected fixed guidance strength w , the noise scale t_e and threshold Q^* are dynamically selected for each input x_0 according to the divergence M_{t_e} between HFP and UFP calculated in each forward step. $\tilde{x}_{0,t}^h$ and $\tilde{x}_{0,t}^0$ are HGP and UGP respectively. The SAMs are aggregated to enhance the signal strength of anomalous regions.

2. Related Work

Prior to the advent of diffusion models (DMs), the field of anomaly detection was largely dominated by generative adversarial networks (GANs) [6] and variational autoencoders (VAEs) [12]. GANs, despite their powerful capabilities, often faced criticism due to unstable training dynamics, while VAEs were limited by their expressiveness in the latent space and a tendency to produce blurry reconstructions [3, 18, 26]. DMs emerged as a promising alternative, gaining recognition for their success in image synthesis [7, 19, 20]. In unsupervised anomaly detection, innovations such as the unguided denoising diffusion probabilistic model (DDPM) with simplex noise have been notable, achieving significant outcomes [25]. Enhancements in model performance have also been seen with the use of patched images [2] and masked images in both physical and frequency domains [10]. Furthermore, the concept

of latent diffusion models, which facilitate fast segmentation, has been effectively utilized [14, 15]. In the realm of weakly-supervised approaches, pioneering efforts such as those by Wolleb *et al.* [24] have utilized DDIM with classifier guidance [4]. This approach has been paralleled by Sanchez *et al.* [17], who adopted DDIM with classifier-free guidance [8], similar to the dual diffusion implicit bridging (DDIB) approach [21]. Hu *et al.* adopted DDIM with classifier guidance and employed the finite difference method to approximate the gradient of denoised inputs in terms of image-level labels during the iterative sampling process for segmentation [9].

3. Background

3.1. Diffusion Models

To approximate the unknown data distribution $q(\mathbf{x})$, DMs perturb the data with increasing noise level and then learn to reverse it using a model ϵ_θ , typically utilizing U-net like architectures [16], parameterized by θ . The forward process of DDPM, perturbing the data, is fixed to a Markov chain with increasing variance schedule β_1, \dots, β_T where T is the total number of steps. The transition kernel from \mathbf{x}_{t-1} to \mathbf{x}_t is modeled as Gaussian

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}\left(\mathbf{x}_t|\sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}\right). \quad (1)$$

Hence, the transition from the input \mathbf{x}_0 to any arbitrary perturbed input \mathbf{x}_t is still Gaussian $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1-\bar{\alpha}_t)\mathbf{I})$. Then, we can sample from it by

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon_t, \quad (2)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ and $\alpha_t = 1 - \beta_t$. The sampling process of DDPM, reversing the data from noise, is factorized as $p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$. It is an iterative transition from \mathbf{x}_T to \mathbf{x}_0 with the learned distribution $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ which is the approximation of inference distribution $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. The details can be found in [7]. Then, we can iteratively sample from the learned distribution p_θ to obtain less noisy samples in closed form

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right) + \Sigma_t \mathbf{z}, \quad (3)$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\Sigma_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t} \beta_t \mathbf{I}$. Note that we can directly obtain the predicted input $\tilde{\mathbf{x}}_0$ from Eq. (2) without the iterative sampling process as $\tilde{\mathbf{x}}_0 = \frac{\mathbf{x}_t - \sqrt{1-\bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}}$. However, the predicted input $\tilde{\mathbf{x}}_0$ will become less accurate as the noise level increases.

The variation DDIM [19] is non-Markovian, incorporating the input \mathbf{x}_0 into the forward process. The sampling process of DDIM is factorized as $q(\mathbf{x}_{1:T}|\mathbf{x}_0) = q(\mathbf{x}_T|\mathbf{x}_0) \prod_{t=2}^T q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. The

forward transition from \mathbf{x}_{t-1} to \mathbf{x}_t can be obtained by Bayes's rule with the closed form $q(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}\left(\sqrt{\bar{\alpha}_{t+1}}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_{t+1}-\sigma_t^2} \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\mathbf{x}_0}{\sqrt{1-\bar{\alpha}_t}}, \sigma_t^2\mathbf{I}\right)$ for $t \geq 2$, where σ_t is set to 0 to enable the deterministic perturbation. Then, we can perturb inputs iteratively by

$$\mathbf{x}_{t+1} = \sqrt{\bar{\alpha}_{t+1}}\tilde{\mathbf{x}}_0 + \sqrt{1-\bar{\alpha}_{t+1}} \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\tilde{\mathbf{x}}_0}{\sqrt{1-\bar{\alpha}_t}}. \quad (4)$$

where \mathbf{x}_1 can be obtained by Eq. (2).

They both share the same objective function in training [19]. The training process is the minimization of KL divergence

$$KL[q|p_\theta] = \mathbb{E}_{\mathbf{x}_{1:T} \sim q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)p_\theta(\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T})} \right]. \quad (5)$$

The objective is further simplified as $L_t = \|\epsilon_t - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, t)\|^2$. The model functions primarily as a noise predictor, learning to add noise in a way that inversely teaches it about the underlying data distribution $q(\mathbf{x})$.

3.2. Classifier-free Guidance

We employ the classifier-free guidance [8] for controllable generation. Compared to the classifier guidance [4], both training and sampling processes are simplified as it eliminates the need for an external classifier. With the guidance, the conditional data distribution $q(\mathbf{x}|y)$ is learned. In the guided sampling process, the guided noise predictor with guidance label y can be written as

$$\tilde{\epsilon}_\theta(\mathbf{x}_t, t, y) = (1+w)\epsilon_\theta(\mathbf{x}_t, t, y) - w\epsilon_\theta(\mathbf{x}_t, t, \emptyset), \quad (6)$$

where $\epsilon_\theta(\mathbf{x}_t, t, \emptyset) = \epsilon_\theta(\mathbf{x}_t, t)$ indicates no guidance. The parameter w is the guidance strength. The guidance is implemented by utilizing the attention mechanism [23].

4. Methodology

We demonstrate the guided and unguided forward processes in Sec. 4.1. We then introduce the guidance strength selection in Sec. 4.2. Dynamic selections for the noise scale and threshold, along with the aggregation of SAMs, are introduced in Sec. 4.3.

4.1. Guided and unguided forward process

For each step of the forward process, we can obtain the HGP and UGP in one-step mapping as follows:

$$\tilde{\mathbf{x}}_{0,t}^h = \frac{\mathbf{x}_t - \sqrt{1-\bar{\alpha}_t} [(1+w)\epsilon_\theta(\mathbf{x}_t, t, h) - w\epsilon_\theta(\mathbf{x}_t, t, \emptyset)]}{\sqrt{\bar{\alpha}_t}} \quad (7)$$

$$= \frac{\mathbf{x}_t}{\sqrt{\bar{\alpha}_t}} + B_t \nabla_{\mathbf{x}_t} \log \tilde{p}_\theta(\mathbf{x}_t|h) \quad (8)$$

$$= \mathbf{x}_0 + B_t [(1+w)\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t) + \Delta \mathbf{s}_t] \quad (9)$$

$$\tilde{\mathbf{x}}_{0,t}^\emptyset = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, \emptyset)}{\sqrt{\bar{\alpha}_t}} \quad (10)$$

$$= \frac{\mathbf{x}_t}{\sqrt{\bar{\alpha}_t}} + B_t \nabla_{\mathbf{x}_t} \log p_\theta(\mathbf{x}_t), \quad (11)$$

where $\tilde{p}_\theta(\mathbf{x}_t|h) \propto p_\theta(\mathbf{x}_t) p_\theta(h|\mathbf{x}_t)^{w+1}$, $B_t = \frac{1 - \bar{\alpha}_t}{\sqrt{\bar{\alpha}_t}}$ is a monotonically increasing function. $\nabla_{\mathbf{x}_t} \log p_\theta(\mathbf{x}_t) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, \emptyset)$ according to the Tweedie's Formula [5]. $\Delta \mathbf{s}_t = \nabla_{\mathbf{x}_t} \log p_\theta(\mathbf{x}_t) - \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)$ is an error term. $p_\theta(h|\mathbf{x}_t)$ represents an implicit classifier induced by the classifier-free guidance [8]. The derivation of Eq. (8) and Eq. (9) is detailed in Appendix A.

We observe that the HGP $\tilde{\mathbf{x}}_{0,t}^h$ forces the noised input \mathbf{x}_t to be healthy due to the gradient $\nabla_{\mathbf{x}_t} \log \tilde{p}_\theta(\mathbf{x}_t|h)$, which is equivalent to shifting the input \mathbf{x}_0 towards to its healthy counterpart, in one-step mapping. The gradient $\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t)$ serves as the sensitivity map of the implicit classifier $p_\theta(h|\mathbf{x}_t)$, highlighting regions of \mathbf{x}_t that significantly impact the classification of health. Initially, the sensitivity map focuses on high-frequency regions that are more susceptible to corruption by noise. As the noise level increases, the high-frequency regions are fully corrupted, and the focus gradually shifts to lower-frequency regions that are corrupted until the image is fully corrupted. This is because the high-frequency regions are more perturbed compared to the low-frequency regions due to the uniform spectral density of Gaussian noise and the gradient $\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t)$ shifts the corrupted regions in \mathbf{x}_t towards the healthy distribution. The UGP $\tilde{\mathbf{x}}_{0,t}^\emptyset$ is not constrained. Note that Eq. (7) and Eq. (10) are essentially the reversed Eq. (2) with and without guidance respectively.

We utilize UGP as a reference to measure the divergence between the HGP and UGP, using the following metric:

$$M_t = \frac{\|\tilde{\mathbf{x}}_{0,t}^h - \tilde{\mathbf{x}}_{0,t}^\emptyset\|_2^2}{d} \quad (12)$$

$$= \frac{B_t^2(1+w)^2}{d} \|\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t)\|_2^2, \quad (13)$$

where d is the dimension of the input \mathbf{x}_0 . The derivation of Eq. (13) can be found in Appendix A. The divergence M_t is essentially the magnitude of weighted gradient of the log-likelihood of the implicit classifier $p_\theta(h|\mathbf{x}_t)$. In the forward process, we collect $\{M_t\}_{t=1}^T$ along with two pixel-level errors, crucial for further parameter tuning, as follows:

$$\mathbf{e}_t^h = \tilde{\mathbf{x}}_{0,t}^h - \mathbf{x}_0 \quad (14)$$

$$= B_t((1+w)\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t) + \Delta \mathbf{s}_t) \quad (15)$$

$$\mathbf{e}_t^\emptyset = \tilde{\mathbf{x}}_{0,t}^\emptyset - \mathbf{x}_0 \quad (16)$$

$$= B_t \Delta \mathbf{s}_t \quad (17)$$

Please refer to Appendix A for a detailed derivation of Eq. (15).

4.2. Fixed guidance strength selection

We first select the fixed guidance strength w by using the classification accuracy between healthy and unhealthy samples. This involves calculating the cosine similarity between sequences of MSEs for healthy samples $MSE_w^h = \{\|e_t^h\|_2^2/d\}_{t=1}^T$ and unhealthy samples $MSE^\emptyset = \{\|e_t^\emptyset\|_2^2/d\}_{t=1}^T$ as follows:

$$\text{Cos}(MSE_w^h, MSE^\emptyset) = \frac{MSE_w^h \cdot MSE^\emptyset}{\|MSE_w^h\|_2 \cdot \|MSE^\emptyset\|_2}. \quad (18)$$

The rationale is that for healthy samples, the cosine similarity should be high, indicating minimal deviation in their error vectors compared to unhealthy samples. The gradient $\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t)$ should exhibit lower sensitivity in healthy samples. If the guidance strength w is too large, the error e_t^h will increase, resulting in lower similarity between MSE_w^h and MSE^\emptyset for healthy samples, thus complicating the distinction between healthy and unhealthy states. However, the larger guidance enhances the detection of anomalous regions.

We determine the optimal guidance strength w^* through grid search. On the validation set, we compute the cosine similarity for each sample for each candidate guidance strength w to establish thresholds Cos_w that yields best classification accuracy. On the testing set without image-level labels, we use the threshold Cos_{w^*} corresponding to the selected guidance strength w^* to classify inputs. Only samples classified as unhealthy undergo further segmentation, while those predicted as healthy are assigned all-zero masks for performance evaluation purposes. A higher guidance strength w typically results in better segmentation performance for unhealthy samples but may lead to poorer performance when considering both healthy and unhealthy samples due to misclassification. Therefore, our selection strategy aims to strike a balance by choosing the largest guidance strength that does not significantly reduce accuracy compared to the best accuracy obtained with a smaller candidate. This search typically starts with a small guidance strength, which is incrementally increased to find the optimal value w^* . The selection process is detailed in Algorithm 1, as presented in Appendix C.

4.3. Dynamical hyperparameter selection

The whole dynamical selection process is demonstrated in Fig. 1. On the testing set, each sample predicted as unhealthy undergoes both the HFP and UFP to determine the maximal divergence M_{t_e} for hyperparameter selection. To enhance the signal strength and detection accuracy of anomalous regions, we aggregate SAMs. These SAMs are computed at each step of the HFP as the square of the pixel-level errors between the HGP and original inputs, and can

be further derived as follows:

$$\mathbf{S}_t = (\mathbf{e}_t^h)^2 = B_t^2 ((1+w)\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t) + \Delta \mathbf{s}_t)^2. \quad (19)$$

As t increases, the sensitive map will focus on lower-frequency regions that are more likely to be anomalous due to changes in texture that are typical of anomalous regions. In each SAM, the signals from healthy regions that are mistakenly targeted by the implicit classifier typically appear randomly distributed. In contrast, signals from anomalous regions exhibit more consistency, indicating that anomalous regions are correctly detected by the implicit classifier. This consistency is crucial to the effectiveness of the aggregation process.

We then aggregate these SAMs to obtain the AAM

$$\mathbf{H} = \frac{1}{t_e} \sum_{t=1}^{t_e} \mathbf{S}_t, \quad (20)$$

where t_e is the end step of the aggregation, i.e., noise scale. The noise scale here controls the signal strength of the anomalous regions in the AAM. We hope that the SAM \mathbf{S}_{t_e} can be a ‘sweet spot’ that balances the signal strength of anomalous regions and mistakenly targeted healthy regions. Then, the signal strength of randomly distributed healthy regions can be reduced by the aggregation, while the signal strength of the anomalous regions is maximized because of consistency.

We select the end step t_e as

$$t_e = \arg \max_t M_t. \quad (21)$$

We observe that $B_t(1+w)\nabla_{\mathbf{x}_t} \log p_\theta(h|\mathbf{x}_t)$ is the term that directly contributes to the divergence between $\tilde{\mathbf{x}}_{0,t}^h$ and \mathbf{x}_0 in Eq. (15). The largest value M_{t_e} tends to yield the largest changes between $\tilde{\mathbf{x}}_{0,t}^h$ and the original input \mathbf{x}_0 in terms of being healthy.

To segment the anomalous regions from the AAM, we need to find a threshold Q^* such that the predicted pixel-level anomalous labels is obtained as $\mathbf{H} \geq Q^*$. We note that the size of the anomalous regions in the input \mathbf{x}_0 is roughly proportional to the maximum value of M_{t_e} . Hence, the value M_{t_e} serves as a rough indication of the size of the anomalous regions. A smaller quantile of \mathbf{H} is selected for a larger anomalous region to include more possible pixels since the signal strength of anomalous regions may not be consistent in \mathbf{H} , e.g., the signal strength of central anomalous regions is stronger than the edge due to the proposed aggregation. We select the segmentation threshold Q^* for each predicted unhealthy input \mathbf{x}_0 by Algorithm 1. The input M_{max} , the largest M_{t_e} of all validation samples, is obtained from the validation set for scaling and is fixed for the testing set. In our case, we set $a = 0.90$, $b = 0.98$ for the range of quantile. The whole segmentation process is detailed in Algorithm 2, as presented in Appendix C.

Algorithm 1: Selection of quantile Q for a single input \mathbf{x}_0

Input: M_{max} , a , b , \mathbf{H}
 $range = reverse(linspace(a, b, 101))$ # Set quantile range
 $M_s = clamp\left(\frac{M_{t_e}}{M_{max}}, 0, 1\right)$
 $index = round(M_s, 2) \times 100$ # Keep 2 digits
Return $Q^* = quantile(\mathbf{H}, range[index])$

5. Experiments

5.1. Dataset and preprocessing

We evaluated our methods on both BraTS21 dataset [1] and ATLAS v2.0 dataset [13]. BraTS21 dataset are three-dimensional Magnetic Resonance (MR) brain images depicting subjects afflicted with a cerebral tumor. Each subject undergoes scanning through four distinct MR sequences, specifically T1-weighted, T2-weighted, FLAIR, and T1-weighted with contrast enhancement. Given our emphasis on a two-dimensional methodology, our analysis is confined to axial slices. There are 1,254 patients and we split the dataset into 939 patients for training, 63 patients for validation, 252 patients for testing. We randomly select 1,000 samples in validation set for hyperparameter tuning and 10,000 samples in testing set for evaluation. For training, we stack all four modalities while only FLAIR and T2-weighted modalities are used in inference. The ATLAS v2.0 dataset consists of three-dimensional, T1-weighted MR brain images of stroke cases, which present a more challenging scenario for anomaly detection due to data quality and the characteristics of stroke lesions. The dataset contains 655 subjects with manually-segmented lesion mask. We split the dataset into 492 subjects for training, 33 subjects for validation, 130 subjects for testing. We randomly select 1,000 samples in validation set for hyperparameter tuning and 4,000 samples in testing set for evaluation. The data preprocessing and training setup is demonstrated in Appendix B.

5.2. Dynamical noise scale captures the ‘sweet spot’

Here, we demonstrate that our selected noise scale t_e effectively captures the ‘sweet spot’ that balances the signal strength between the anomalous and healthy regions in the SAM at t_e . To validate this, we analyze 200 unhealthy samples from the BraTS21 dataset. We present the average change of divergence M_t between HGP and UGP, along with the magnitude of SAMs, calculated as $\frac{\|\mathbf{e}_t^h\|_2^2}{d}$, in healthy and anomalous regions separately. Additionally, we include the ratio of these magnitude to further illustrate the balance achieved. These results are depicted in Fig. 2c. The maximal ratio is almost achieved at t_e , indicating that

the noise scale t_e effectively captures the ‘sweet spot’. We also show the relation between the size of the anomalous region and the maximal divergence M_{t_e} in Fig. 2d. We observe that the size of the anomalous region is roughly proportional to the maximal divergence M_{t_e} and we achieved Pearson correlation $r = 0.7$.

We display samples of SAMs alongside the corresponding gradient of the log-likelihood of the implicit classifier $\nabla_{x_t} \log p_\theta(h|x_t)$ at selected steps from BraTS21 dataset in Fig. 3a and Fig. 3b. We observe that the implicit classifier predominantly targets anomalous regions at $t = 500$ and $t = 600$. At $t = 200$ and $t = 300$, the gradient focuses on high-frequency regions, transitioning to significantly lower-frequency regions at $t = 700$. As noted earlier, the anomalous regions exhibit consistent distribution patterns, while the healthy regions appear randomly distributed across different frequency levels. The SAM and its corresponding gradient at step t_e strikes a balance between the signal strength of healthy regions and anomalous regions. However, the extra error terms Δs_t in the SAMs may correct the gradient term and this effect is obvious at $t = 700$.

5.3. Main results

We assess our model’s performance using pixel-level metrics, including the DICE score, intersection over union (IoU), and area under the precision-recall curve (AUPRC) with a specific focus on the foreground areas. Our evaluation on the BraTS21 test set, which comprises 10,000 samples, are comprehensively detailed in two separate configurations: the mixed setup, which includes all samples, and the unhealthy setup, which exclusively considers unhealthy samples. These results are presented in Tab. 1a. Similarly, for the ATLAS dataset, which comprises 4,000 mixed samples including 1,003 unhealthy and 2,997 healthy samples, we present the performance results separately for the mixed setup and the unhealthy setup in 1b.

For the proposed method, we present results for four setups: (i) **AnoFPDM (DDPM)** (stochastic encoding): Eq. (2) is used to perturb inputs; (ii) **AnoFPDM (DDIM)** (deterministic encoding): the inputs are noised by Eq. (4); (iii) **AnoFPDM (DDIM S_{t_e})**: similar to **AnoFPDM (DDIM)**, but segmentation is performed using only the SAM at t_e ; (iv) **AnoFPDM (DDPM S_{t_e})**: follows the **AnoFPDM (DDPM)** setup, with segmentation using only the SAM at t_e . For comparative analysis, we report the results from AnoDDPM [25] with simplex noise and Gaussian noise for fair comparison, DDIM with classifier guidance [24], DDIM with classifier-free guidance [17] and classifier guided conditional diffusion model (CG-CDM) [9]. The first two methods are only trained on healthy samples. The hyperparameters for all comparison methods are optimized using a grid search on 1,000 mixed samples with

pixel-level labels for both setups. Notably, our methods do not require pixel-level labels for tuning. The postprocessing for segmentation is detailed in Appendix F.

Our method, especially AnoFPDM (DDIM) with deterministic encoding, demonstrates superior performance in terms of DICE and IoU scores, particularly in scenarios only involving unhealthy samples. AnoDDPM with simplex noise achieves AUPRC scores comparable to ours. However, its segmentation performance in unhealthy settings is hindered by a fixed threshold. Across both datasets, our approach consistently balances performance between mixed and unhealthy setups, notably excelling in the ATLAS v2.0 dataset where other comparison methods struggle to detect anomalous regions. We discuss further variations in performance across both datasets in Appendix G. The qualitative results, displayed in Fig. 4, underscore our method’s ability to enhance the signal strength of anomalous regions, which is confirmed by higher AUPRC scores. Notably, the signal strength at the edges of anomalous regions is weaker compared to their central parts. This effect becomes more pronounced in larger anomalous regions. More qualitative results are exhibited in Appendix E.

5.4. Ablation study

To validate the effectiveness of our hyperparameter selection, we conducted a series of performance comparisons using different guidance strength w , end step t_e , and threshold Q^* across both datasets. The results are displayed in Fig. 5. Specifically, we first examined the impact of guidance strength w in Fig. 5a and Fig. 5b. For the BraTS21 and ATLAS v2.0 datasets, we selected the guidance strength $w = 2$ and $w = 30$ respectively. These values represent the maximum at which the accuracy, based on cosine similarity, does not significantly decline, staying within 99% of the peak accuracy. Additionally, we assessed the MSEs MSE_t^h and MSE_t^0 at $w = 2$ for both healthy and unhealthy samples from BraTS21 dataset in Fig. 2a and Fig. 2b. These figures highlight noticeable differences in the unhealthy samples, whereas the differences in the healthy samples are less pronounced. We further explored the variation in the DICE score across different guidance strengths in Fig. 5b, with other hyperparameters being dynamically adjusted per our proposed method. The DICE score varies across the two setups, namely the mixed and unhealthy setups, with different guidance strengths w , confirming our statement that a higher guidance strength enhances the detection of anomalous regions. However, the overall performance in the ATLAS v2.0 dataset is less impacted by changes in guidance strength compared to the BraTS21 dataset, which can be attributed to their differing distributions of healthy and unhealthy samples. Our selection tends to strikes a balance between the two setups.

The impact of the end step t_e on the DICE score for





Figure 2. MSE_t^h and MSE_t^0 for (a) healthy samples and (b) unhealthy samples. (c) Average change in the magnitude of SAMs in healthy and anomalous regions, alongside their ratio and divergence M_t between HGP and UGP. (d) Relationship between the size of the anomalous region and the maximal divergence M_{t_e} .



(a) The first two columns show the unhealthy inputs and the ground truth masks. The third to the last columns show the corresponding SAMs at the step $t = 100, 300, 600, 700, T$ and t_e . (b) The first column shows the unhealthy inputs. The third to the last columns show the corresponding gradients $\nabla_{x_t} \log p_{\theta}(h|x_t)|_{t=t_e}$ at the step $t = 200, 300, 500, 600, 700$ and t_e .

Figure 3. The gradient of the log-likelihood of the implicit classifier (a) and SAMs (b) at the selected steps from BraTS21 dataset




Figure 4. Qualitative Comparison of Anomaly Maps and Segmentation. (a) From the BraTS21 dataset and (b) from the ATLAS v2.0 dataset. The first column displays the original input images, and the second column shows the corresponding ground truth for anomaly segmentation. Subsequent columns present the anomaly maps and segmentation results obtained using our method, AnoFPDM with the DDIM setting, alongside those from the second and third best comparative methods. Each row represents a different sample.

both datasets with the selected guidance strength is depicted in Fig. 5c. The maximal DICE score is achieved at t_e , followed by a slight decline as t_e extends to $3 \times t_e$ be-

cause more noise is introduced into the AAM. This pattern validates the selection of t_e , confirming it effectively enhances the signal strength of anomalous regions in the

| Type | Methods | Mixed | | | Unhealthy | | |
|----------------|-----------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | | DICE | IoU | AUPRC | DICE | IoU | AUPRC |
| Reconstruction | ✗ AnoDDPM (G) [25] | 66.1±0.1 | 61.7±0.1 | 51.8±0.1 | 37.6±0.1 | 28.1±0.1 | 61.3±0.1 |
| | ✗ AnoDDPM (S) [25] | 75.1±0.3 | 69.5±0.2 | 67.3±0.1 | 53.7±2.7 | 45.5±1.3 | 71.8±0.1 |
| | ✗ DDIM clf [24] | 76.5±0.1 | 71.0±0.1 | 58.4±0.3 | 52.2±0.2 | 40.4±0.2 | 61.6±0.2 |
| | ✗ DDIM clf-free [17] | 74.3±0.0 | 69.1±0.0 | 59.9±0.0 | 49.1±0.0 | 38.1±0.0 | 61.4±0.0 |
| Gradient | ✗ CG-CDM [9] | - | - | - | 44.4±0.3 | 32.2±0.5 | 31.2±0.7 |
| | ✓ AnoFPDM (DDPM) | 75.2±0.3 | 68.5±0.3 | 68.3±0.1 | 57.5±0.4 | 46.0±0.2 | <u>72.9±0.1</u> |
| | ✓ AnoFPDM (DDPM S_{t_e}) | 70.1±0.1 | 64.6±0.3 | 56.7±0.4 | 47.8±0.1 | 36.0±0.1 | 61.1±0.5 |
| | ✓ AnoFPDM (DDIM) | 77.4±0.0 | 72.5±0.0 | 72.2±0.0 | 61.5±0.0 | 50.0±0.1 | 75.5±0.0 |
| | ✓ AnoFPDM (DDIM S_{t_e}) | 75.7±0.0 | 70.6±0.0 | <u>69.9±0.1</u> | <u>58.7±0.0</u> | <u>47.0±0.0</u> | 72.6±0.1 |

(a) Segmentation performance on mixed and unhealthy samples from BraTS21 dataset.

| Type | Methods | Mixed | | | Unhealthy | | |
|----------------|-----------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | | DICE | IoU | AUPRC | DICE | IoU | AUPRC |
| Reconstruction | ✗ AnoDDPM (G) [25] | 74.8±0.1 | <u>74.8±0.1</u> | 2.0±0.1 | 0.4±0.1 | 0.2±0.1 | 6.5±0.2 |
| | ✗ AnoDDPM (S) [25] | <u>74.9±0.1</u> | <u>74.6±0.1</u> | 20.8±0.5 | 3.4±1.0 | 3.3±0.7 | <u>30.9±0.4</u> |
| | ✗ DDIM clf [24] | 51.5±0.8 | 50.8±0.7 | 1.9±0.1 | 5.8±0.1 | 3.7±0.1 | 5.6±0.1 |
| | ✗ DDIM clf-free [17] | 73.5±0.0 | 73.0±0.0 | 9.3±0.0 | 0.1±0.0 | 0.1±0.0 | 13.6±0.0 |
| Gradient | ✗ CG-CDM [9] | - | - | - | 2.1±0.0 | 1.1±0.0 | 1.6±0.0 |
| | ✓ AnoFPDM (DDPM) | 74.5±0.1 | 74.5±0.1 | <u>22.4±0.1</u> | <u>17.5±0.1</u> | <u>12.4±0.1</u> | 29.2±0.1 |
| | ✓ AnoFPDM (DDPM S_{t_e}) | 74.6±0.1 | 74.6±0.1 | 6.1±0.1 | 9.5±0.1 | 6.4±0.1 | 10.8±0.2 |
| | ✓ AnoFPDM (DDIM) | 75.5±0.2 | 75.5±0.2 | 22.5±0.1 | 21.5±0.0 | 15.5±0.0 | 31.2±0.1 |
| | ✓ AnoFPDM (DDIM S_{t_e}) | 74.5±0.2 | 74.5±0.2 | 2.0±0.1 | 9.8±0.1 | 6.9±0.1 | 4.1±0.1 |

(b) Segmentation performance on mixed slices and unhealthy samples from ATLAS v2.0 dataset.

Table 1. Quantitative results across both datasets. Standard deviations are derive from three runs. The best and second-best performances are in bold and underline, respectively. Methods requiring pixel-level labels for tuning are marked with ✓, those that do not with ✗. CG-CDM [9] is tested only on the unhealthy samples for consistency with its original study.

AAM. We also show the influence of the threshold on the DICE score in Fig. 5d. We compare our dynamical selection Q^* with the fixed threshold selection. Given the aggregated anomaly map obtained by the proposed method, we calculate the DICE score with various fixed thresholds. Our dynamic threshold selection significantly outperforms the fixed threshold approach, thereby validating the efficacy of our method in optimizing threshold selection.

We also discuss the use of the gradient directly as SAMs for segmentation in Appendix D.

6. Conclusion

We propose a novel weakly-supervised method, AnoFPDM, for anomaly detection. This method dynamically selects the noise scale and threshold for each input while maintaining a fixed guidance strength selection. In quantitative evaluations, our method surpasses previous approaches and demonstrates enhanced signal strength of anomalous regions in qualitative assessments. Unlike traditional methods, our selection process does not involve pixel-level labels, making it more practical for real-world applications. We also discuss the limitations in Appendix G.

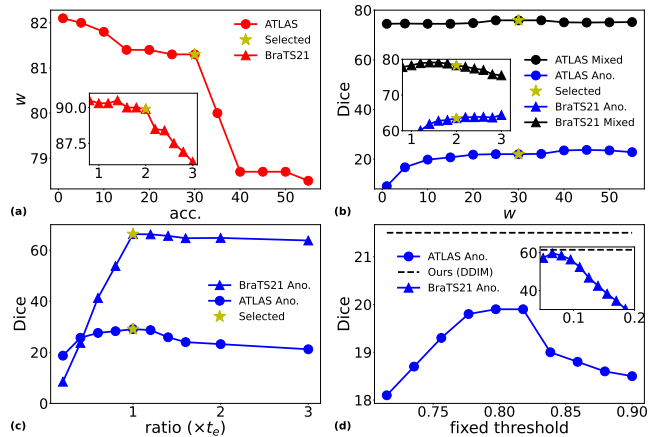


Figure 5. Analysis of the impact of hyperparameter variations on model performance. (a) Variation in classification accuracy as a function of guidance strength w . (b) Corresponding changes in DICE score with different guidance strength w . (c) Effect of the end step t_e on DICE score, using the selected guidance strength. (d) Comparison of DICE scores using a fixed threshold.

References

- [1] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018. **5**
- [2] Finn Behrendt, Debayan Bhattacharya, Julia Krüger, Roland Opfer, and Alexander Schlaefer. Patched diffusion models for unsupervised anomaly detection in brain mri. In *Medical Imaging with Deep Learning*, 2023. **2**
- [3] Xiaoran Chen, Suhang You, Kerem Can Tezcan, and Ender Konukoglu. Unsupervised lesion detection via image restoration with a normative prior. *Medical image analysis*, 64:101713, 2020. **2**
- [4] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. **1, 3**
- [5] Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. **4**
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. **1, 2**
- [7] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 6840–6851, 2020. **1, 2, 3**
- [8] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. **1, 2, 3, 4**
- [9] Xinrong Hu, Yu-Jen Chen, Tsung-Yi Ho, and Yiyu Shi. Conditional diffusion models for weakly supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 756–765. Springer, 2023. **1, 3, 6, 8**
- [10] Hasan Iqbal, Umar Khalid, Chen Chen, and Jing Hua. Unsupervised anomaly detection in medical images using masked diffusion model. In *International Workshop on Machine Learning in Medical Imaging*, pages 372–381. Springer, 2023. **2**
- [11] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021. **1**
- [12] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. **1, 2**
- [13] Sook-Lei Liew, Bethany P Lo, Miranda R Donnelly, Artemis Zavaliangos-Petropulu, Jessica N Jeong, Giuseppe Barisano, Alexandre Hutton, Julia P Simon, Julia M Juliano, Anisha Suri, et al. A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. *Scientific data*, 9(1):320, 2022. **5**
- [14] Walter HL Pinaya, Mark S Graham, Robert Gray, Pedro F da Costa, Petru-Daniel Tudosiu, Paul Wright, Yee H Mah, Andrew D MacKinnon, James T Teo, Rolf Jager, et al. Fast unsupervised brain anomaly detection and segmentation with diffusion models. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 705–714, 2022. **3**
- [15] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. **3**
- [16] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015. **3**
- [17] Pedro Sanchez, Antanas Kascenas, Xiao Liu, Alison Q O’Neil, and Sotirios A Tsaftaris. What is healthy? generative counterfactual diffusion for lesion localization. In *MICCAI Workshop on Deep Generative Models*, pages 34–44. Springer, 2022. **1, 3, 6, 8**
- [18] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019. **2**
- [19] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. **1, 2, 3**
- [20] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020. **1, 2**
- [21] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. In *The Eleventh International Conference on Learning Representations*, 2022. **3**
- [22] Nima Tajbakhsh, Laura Jeyaseelan, Qian Li, Jeffrey N Chiang, Zhihao Wu, and Xiaowei Ding. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical image analysis*, 63:101693, 2020. **1**
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **3**
- [24] Julia Wolleb, Florentin Bieder, Robin Sandkühler, and Philippe C Cattin. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, pages 35–45. Springer, 2022. **1, 3, 6, 8**
- [25] Julian Wyatt, Adam Leach, Sebastian M Schmon, and Chris G Willcocks. Anoddpmm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 650–656, 2022. **2, 6, 8**
- [26] David Zimmerer, Fabian Isensee, Jens Petersen, Simon Kohl, and Klaus Maier-Hein. Unsupervised anomaly localization

using variational auto-encoders. In *Medical Image Computing and Computer Assisted Intervention*, pages 289–297. Springer, 2019. [2](#)