

WDT-MD: Wavelet Diffusion Transformers for Microaneurysm Detection in Fundus Images

Yifei Sun^{1,2}, Yuzhi He³, Junhao Jia², Jinhong Wang¹,
Ruiquan Ge^{2*}, Changmiao Wang^{4*}, Hongxia Xu^{1,5*}

¹Transvascular Implantation Devices Research Institute, Zhejiang University, Hangzhou, China

²Hangzhou Dianzi University, Hangzhou, China

³Xidian University, Xi'an, China

⁴Shenzhen Research Institute of Big Data, Shenzhen, China

⁵WeDoctor Cloud and Liangzhu Laboratory, Hangzhou, China

{diaoquesang, cmwangalbert}@gmail.com, 23012100059@stu.xidian.edu.cn,
{23080631, gespring}@hdu.edu.cn, {wangjinhong, einstein}@zju.edu.cn

Abstract

Microaneurysms (MAs), the earliest pathognomonic signs of Diabetic Retinopathy (DR), present as sub-60 μm lesions in fundus images with highly variable photometric and morphological characteristics, rendering manual screening not only labor-intensive but inherently error-prone. While diffusion-based anomaly detection has emerged as a promising approach for automated MA screening, its clinical application is hindered by three fundamental limitations. First, these models often fall prey to “identity mapping”, where they inadvertently replicate the input image. Second, they struggle to distinguish MAs from other anomalies, leading to high false positives. Third, their suboptimal reconstruction of normal features hampers overall performance. To address these challenges, we propose a Wavelet Diffusion Transformer framework for MA Detection (WDT-MD), which features three key innovations: a noise-encoded image conditioning mechanism to avoid “identity mapping” by perturbing image conditions during training; pseudo-normal pattern synthesis via inpainting to introduce pixel-level supervision, enabling discrimination between MAs and other anomalies; and a wavelet diffusion Transformer architecture that combines the global modeling capability of diffusion Transformers with multi-scale wavelet analysis to enhance reconstruction of normal retinal features. Comprehensive experiments on the IDRiD and e-optha MA datasets demonstrate that WDT-MD outperforms state-of-the-art methods in both pixel-level and image-level MA detection. This advancement holds significant promise for improving early DR screening.

Code — <https://github.com/diaoquesang/WDT-MD>

Introduction

Diabetic Retinopathy (DR) is a serious complication affecting individuals with diabetes and can result in severe vision loss if not treated promptly (Khan et al. 2025). In the initial stages of DR, retinal capillaries are damaged due to hyperglycemia, which weakens the capillary walls and leads to Microaneurysms (MAs). MAs are small outpouchings in the




Figure 1: An illustration of MAs in fundus images. (a) is sampled from the IDRiD dataset (Porwal et al. 2018), and (b) is from the e-optha MA dataset (Decenciere et al. 2013). Three columns in each sub-figure depict the fundus image, patches zooming in MAs, and MA areas marked in red, respectively. Most MAs are within 60 μm in diameter, close to 6 pixels in a fundus image with 10 μm pixel spacing.

lumen of the retinal vessels, typically measuring 15-60 μm in diameter. Identification of MAs allows for timely recognition of DR, thus providing an opportunity for early intervention in patients (Arrigo et al. 2024). To analyze them, fundus images are widely used (Mayya, Kamath, and Kulkarni 2021) where small red dots are an indication of MAs (Raghu et al. 2019). Nevertheless, as shown in Fig. 1, MAs are tiny and inconspicuous with variations in brightness, contrast, and shape, making it difficult for physicians to detect them (Wu and Jiao 2024). Therefore, automated MA detection methods with high accuracy in fundus images are of great significance.

To achieve this goal, different methods are proposed, among which most are discriminative models such as segmentation models (Xu et al. 2024; Jiang et al. 2024; Foo, Hsu, and Lee 2023; Yap and Ng 2023). In contrast to classification models lacking specific localization capability, segmentation models can provide detailed information about the MA boundaries, thus helping to assess the severity of lesions and enhance the interpretability of image-level information (Jiang et al. 2023). Nevertheless, the challenges of data annotation and segmentation accuracy restrict the de-

*Corresponding authors.

velopment of these methods. First, the tiny size of MAs and their morphological similarity to normal vascular structures create significant inter-observer variability in manual annotations (Mayya, Kamath, and Kulkarni 2021). This label ambiguity propagates through supervised segmentation frameworks, leading to suboptimal boundary delineation. Second, the class imbalance problem is exacerbated in MA diagnosis, where positive pixels constitute less than 1% of total image area in early-stage DR cases (Porwal et al. 2020). Traditional segmentation methods tend to converge to trivial solutions that ignore subtle MA features.

Unlike discriminative models, generative models have been gradually applied to reconstruction-based methods for medical Anomaly Detection (AD), mainly based on Auto-Encoders (AEs), Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) and diffusion models (Ho, Jain, and Abbeel 2020; Ma et al. 2024b, 2025b,a; Sun et al. 2025c; Shao et al. 2025a). By identifying deviations from normal patterns as anomalies, these methods can effectively detect small, irregular lesions while reducing reliance on large numbers of accurate pixel-level annotations. Recently, diffusion-based methods (Kumar et al. 2025; Fontanella et al. 2024; Wyatt et al. 2022; Wolleb et al. 2022) have made progress towards more accurate anomaly localization by iteratively improving the capture of fine-grained lesion details (Sun et al. 2025b), holding significant potential for MA detection. Nonetheless, the following challenges still remain:

- The inherent risk of learning “**identity mapping**” still persists in existing frameworks based on diffusion models. “Identity mapping” refers to the behavior of directly copying the input as output, whether normal or abnormal (Guo et al. 2025). This contradicts the foundational assumption that anomalies induce significant reconstruction deviations, ultimately causing false negatives.
- The inability to distinguish MAs from other anomalies leads to **high false positives**. Existing methods lacking pixel-level supervision signals tend to treat all reconstruction errors as homogeneous indicators of abnormality, disregarding the unique morphological and contextual signatures of the target anomalies. Consequently, confounding factors such as imaging artifacts or coexisting lesions can be indiscriminately flagged as MA candidates, undermining clinical utility.
- The **suboptimal reconstruction quality** of normal features hampers the performance of AD. In retinal imaging, incomplete restoration of vascular patterns may introduce spurious reconstruction errors, masking true MA lesions or misclassifying normal variations as anomalies.

Existing diffusion-based methods mitigate “identity mapping” through noise-addition-denoising (Kumar et al. 2025; Fontanella et al. 2024; Li et al. 2024; Wyatt et al. 2022) in the inference phase. This strategy faces a fundamental resolution conflict: MAs and fine vascular details occupy overlapping high-frequency bands, yet demand diametrically opposed noise treatments. Reliable MA suppression requires near-complete high-frequency erosion, while precise vasculature reconstruction necessitates preserving those exact frequency components. Insufficient noise preserves anomalies

while excessive noise obliterates details. Consequently, single noise calibration during inference becomes intrinsically paradoxical for these methods.

Furthermore, the absence of pixel-level supervision elevates false positive rates. These models erroneously classify imaging artifacts and non-MA pathologies as MAs, which is clinically unacceptable. To introduce pixel-level supervision signals, self-supervised image-conditioned approaches like *Img-Cond* (Baugh et al. 2024) have been proposed. However, unprocessed input conditioning propagates anomalies via “identity mapping”, while the spatial distributional bias introduced by synthetic anomalies could further diminish the performance. Notably, although pixel-level supervision has recently been demonstrated to boost AD performance in complex industrial scenarios (Baitieva et al. 2024), analogous exploration and validation remain scarce in the medical field. The intrinsic spatial linkage between lesions and their anatomical context (Shao et al. 2025b) poses unique challenges for effective supervision, underscoring the need for tailored strategies.

To address these challenges, we propose a Wavelet Diffusion Transformer framework for MA Detection (WDT-MD). This is a supervised image-conditioned wavelet-domain model based on Diffusion Transformers (DiTs) (Peebles and Xie 2023; Feng et al. 2025). Our contributions can be summarized as follows:

- In order to mitigate “identity mapping”, we propose a **noise-encoded image conditioning** mechanism for diffusion-based MA detection. By perturbing the image condition with random intensities during training, the model is driven to capture the normal pattern.
- To alleviate the issue of high false positives, we introduce pixel-level supervision signals in the training process through **pseudo-normal pattern synthesis**. Specifically, we obtain the pseudo-normal labels align with the spatial distribution of real fundus images using inpainting techniques. This enables the model to distinguish MAs from other anomalies, thereby improving the detection performance.
- To improve the reconstruction quality of normal features, we propose a **wavelet diffusion Transformer** architecture, which combines the global modelling capability of DiTs with the multi-scale analysis advantage of wavelet decomposition to better understand the overall structure and detailed information of fundus images.
- Comprehensive experiments on the IDRiD and e-ophtha MA datasets demonstrate exceptional performance of our WDT-MD, holding significant promise for improving early DR screening.

Method

We propose WDT-MD, a novel Wavelet Diffusion Transformer framework for MA Detection in fundus images, as illustrated in Fig. 2. This framework addresses key limitations of existing diffusion-based AD approaches for identifying MAs, which are critical early indicators of DR. The WDT-MD method initiates sampling from Gaussian noise in the wavelet domain, conditioned on the input fundus data.

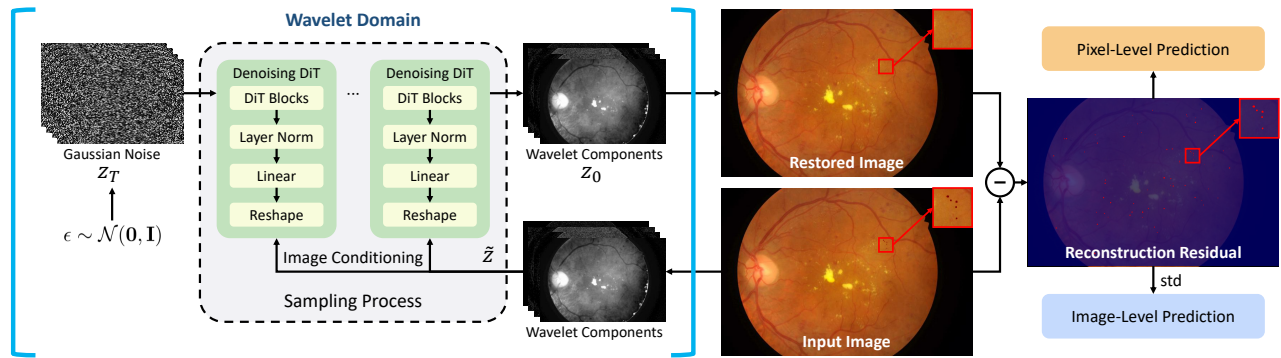


Figure 2: Overview of our proposed WDT-MD method. It is a supervised DiT-based AD framework operating in the wavelet domain, which focuses on MA detection in fundus images. By synthesizing the normal pattern and subtracting the input from it, the model obtains an anomaly map, which further outputs both pixel-level and image-level predictions. (std: standard deviation)

Through T -step iterative sampling, it reconstructs pseudo-normal fundus data. Subsequently, in the pixel domain, the reconstruction residual is computed by subtracting the input image from the restored pseudo-normal image. This residual map is then processed to yield both pixel-level segmentation and image-level classification predictions.

Wavelet Diffusion Transformer

Accurate reconstruction of normal retinal features is critical for reliable MA detection in DR screening. Existing Transformer-based backbones like U-ViT (Bao et al. 2023) offer powerful global modeling capabilities for diffusion models, enabling them to better capture contextual information such as the spatial distribution of MAs. Nevertheless, their operation directly in the pixel space presents limitations. These approaches struggle to capture and preserve the intricate multi-scale structural and textural details inherent in retinal vasculature, and incur significant computational costs. To address these issues, subsequent works (Peebles and Xie 2023; Ma et al. 2024a; Esser et al. 2024) migrate the diffusion process into a learnable latent space using AE-based tokenizers. However, this two-stage strategy introduces its own challenges: the computational overhead of the tokenizer itself and the potential risk of losing inconspicuous features during the encoding or decoding process. This information loss is particularly problematic for detecting MAs. MAs, often subtle in size, demand a representation that inherently separates low-frequency contextual information such as vessel structures and backgrounds from high-frequency details including tiny lesions and textures.

To overcome these limitations as well as inspired by the success of wavelet analysis in low-level vision tasks (Zhao et al. 2024; Huang et al. 2024), we integrate Discrete Wavelet Transformation (DWT) with DiTs, proposing a wavelet diffusion Transformer architecture. Compared to AE-based tokenizers, DWT exhibits near-lossless reconstruction capabilities (Wang et al. 2024) and incurs lower computational overhead. Specifically, for an image $I \in \mathbb{R}^{C \times H \times W}$, DWT transforms its Value channel $V = \max(R, G, B) \in \mathbb{R}^{1 \times H \times W}$ into four sub-bands:

$$V_{LL}, \{V_{LH}, V_{HL}, V_{HH}\} = DWT(V), \quad (1)$$




Figure 3: Visualization of HSV channel decomposition on (a) IDRiD and (b) e-optha MA. The H and S channels carry little effective information but notable noise, while V contains almost all crucial structural and textural features.

where $V_{LL}, \{V_{LH}, V_{HL}, V_{HH}\}$ denote the low-frequency component of the image and high-frequency components in the vertical, horizontal, and diagonal directions, respectively. The selection of V channel helps alleviate the interference of imaging noise while effectively reducing the computational load, as illustrated in Fig. 3. Subsequently, the sub-bands are concatenated together, denoted as z :

$$z = \text{Concat}(V_{LL}, V_{LH}, V_{HL}, V_{HH}). \quad (2)$$

Our wavelet diffusion Transformer incorporates both a forward process and a reverse process in the wavelet domain. The forward process can be defined as:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (3)$$

where z_t is the wavelet components at timestep t , ϵ is a Gaussian noise map, and $\bar{\alpha}_t := \prod_{s=0}^t \alpha_s$. Here, $\alpha_t = 1 - \beta_t$ is a differentiable function of timestep t . The diffusion loss is expressed as:

$$\mathcal{L}(\epsilon_\theta) = \sum_{t=1}^T \mathbb{E}_{z_0, \epsilon} \left[\|\epsilon_\theta(z_t, t, \tilde{z}) - \epsilon\|_2^2 \right], \quad (4)$$

where ϵ_θ represents the predicted noise at timestep t by the denoising DiT with parameters θ , T is the total diffusion timesteps, and \tilde{z} is a given image condition.

During the reverse process, starting from Gaussian noise $z_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the original sample z_0 is predicted through a

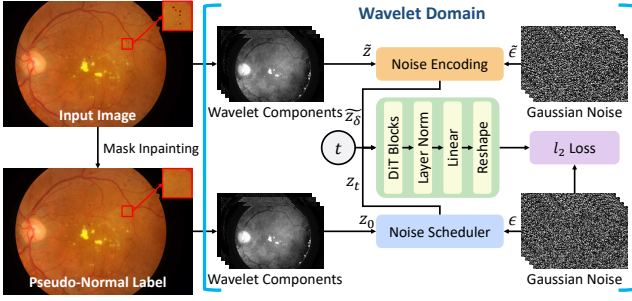


Figure 4: The training process of our WDT-MD.

multi-step denoising process:

$$z_{t-1} = \begin{cases} \sqrt{\alpha_{t-1}} (c_{\text{out}}(t) \hat{z}_0 + c_{\text{skip}}(t) z_t) + \gamma \epsilon, & 1 < t \leq T_s \\ c_{\text{out}}(t) \hat{z}_0 + c_{\text{skip}}(t) z_t, & t = 1, \end{cases} \quad (5)$$

where $\hat{z}_0 = \frac{z_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(z_t, t, \tilde{z})}{\sqrt{\alpha_t}}$ is the predicted original sample, $\gamma = \frac{1-\alpha_{t-1}}{\sqrt{1-\alpha_{t-1}}}$ is the scaling factor for Gaussian noise ϵ , T_s is the total sampling timesteps, and both $c_{\text{out}}(t)$ and $c_{\text{skip}}(t)$ are differentiable with $c_{\text{out}}(0) = 0$ and $c_{\text{skip}}(0) = 1$.

Subsequently, z_0 is split by channel and transformed into restored Value channel V_0 via Inverse DWT (IDWT):

$$V_{0,LL}, V_{0,LH}, V_{0,HL}, V_{0,HH} = \text{Split}(z_0), \quad (6)$$

$$V_0 = \text{IDWT}(V_{0,LL}, \{V_{0,LH}, V_{0,HL}, V_{0,HH}\}). \quad (7)$$

Finally, V_0 is merged with the Hue channel H and Saturation channel S of the input image to obtain the restored image I_0 .

Noise-Encoded Image Conditioning

To address the ‘‘identity mapping’’ dilemma in diffusion models, we propose a noise-encoded image conditioning mechanism. In contrast to existing noise-addition-denoising methods that adopt single noise calibration during inference, we avoid ‘‘identity mapping’’ by perturbing the image condition \tilde{z} into noise-encoded \tilde{z}_δ during training, denoted as:

$$\tilde{z}_\delta = \sqrt{\alpha_\delta} \tilde{z} + \sqrt{1 - \alpha_\delta} \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (8)$$

where the added noise is determined by the timestep $\delta \in \{1, 2, \dots, \delta_{\text{max}}\}$ utilizing the noise scheduler. Correspondingly, the diffusion loss in Eq. (4) becomes as follows:

$$\mathcal{L}(\epsilon_\theta) = \sum_{t=1}^T \mathbb{E}_{z_0, \delta, \epsilon} \left[\|\epsilon_\theta(z_t, t, \tilde{z}_\delta) - \epsilon\|_2^2 \right]. \quad (9)$$

By introducing dynamic noise perturbations during training, this mechanism enforces the model to learn the underlying normal patterns of retinal structures rather than directly copying input pixels, as illustrated in Fig. 4.

Pseudo-Normal Pattern Synthesis

To mitigate false positives, we synthesize pseudo-normal labels using inpainting techniques, thus introducing pixel-level supervision signals in the training process, which can be briefly expressed as follows:

$$V_{pn} = (1 - M) \odot V + M \odot \mathcal{I}(V, M), \quad (10)$$

where V_{pn} represents the pseudo-normal V channel, M denotes the binary MA mask normalized to $[0, 1]$, and \mathcal{I} is the inpainting algorithm. Here, we employ Telea, a classical inpainting method (Telea 2004). In contrast to synthesizing anomalies on normal fundus images (Sun et al. 2025a), our core idea is to infer unknown pseudo-normal regions from known normal pixels guided by MA masks, ensuring the spatial distribution accuracy of pixel-level supervision.

Experiments

Datasets and Evaluation Metrics

Data Preparation. Two publicly available datasets, namely IDRiD (Porwal et al. 2018) and e-optha MA (Decenciere et al. 2013), are adopted for extensive evaluation.

The IDRiD dataset, a benchmark resource for diabetic retinopathy analysis, was adapted for our study. For MA detection, we curated a subset of 249 samples, including 199 training cases, 24 validation cases, and 26 test cases. Specifically, the training set contains 134 normal images and 65 abnormal images. Contrast Limited Adaptive Histogram Equalization (CLAHE) was applied with 8×8 tile grids and a 2.0 clip limit to enhance contrast. Considering the computational overhead, we implemented dimension standardization through bilinear downsampling to 300×200 pixels.

The e-optha MA dataset consists of 381 cases divided into 304 training, 38 validation, and 39 test samples. Specifically, the training set contains 188 normal images and 116 abnormal images. The preprocessing pipeline maintained strict consistency with IDRiD: (1) CLAHE (8×8 tile grids, 2.0 clip limit); (2) downsampling to 300×200 pixels.

Evaluation Metrics. To evaluate the MA detection performance, we calculate the Area Under Curve (AUC), Accuracy (ACC), F1 score, Sensitivity (SEN) and Specificity (SPE) for both pixel-level and image-level detection.

Implementation Details

All experiments were performed using PyTorch 2.5.1 on a single NVIDIA V100 32 GB GPU within Ubuntu 22.04. WDT-MD was trained from scratch over 600 epochs with a batch size of 4 utilizing the AdamW optimizer, complemented by a dynamic learning rate schedule initialized at 10^{-4} . The noise scheduling parameter β_t followed a scaled linear trajectory ranging from 0.00085 to 0.012 across $T = 1000$ diffusion timesteps. The sampling steps T_s was set to 50 using the LCM sampler (Luo et al. 2023). In pseudo-normal pattern synthesis, the inpainting radius r is set to 3 pixels. For wavelet decomposition, the Daubechies 6 basis was selected to balance computational efficiency and time-frequency localization (Wang, Xu, and Zhao 2024).

Main Results

On both the IDRiD and e-optha MA datasets, we benchmark our WDT-MD against other state-of-the-art methods: (1) diffusion-based AD methods DTU-Net (Kumar et al. 2025), Dif-fuse (Fontanella et al. 2024) and AnoDPM (Wyatt et al. 2022); (2) a GAN-based AD method GatingAno (Zhang et al. 2024); (3) AE-based AD methods

Method	Source	Pixel-level					Image-level				
		AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE
AnoDDPM	CVPR ₂₂	81.76	99.91	53.34	63.62	99.93	71.90	76.92	62.50	55.56	88.24
CPC	WACV ₂₃	76.77	99.93	49.75	53.63	99.96	77.45	80.77	70.59	66.67	88.24
DAE	MedIA ₂₃	71.52	99.69	35.64	43.23	99.72	56.86	53.85	50.00	66.67	47.06
HACDR-Net	AAAI ₂₄	56.38	95.07	4.03	18.82	95.12	63.07	65.38	52.63	55.56	70.59
AE[$d_{optimal}$]	MICCAI ₂₄	75.06	99.24	19.52	50.88	99.27	62.75	61.54	54.55	66.67	58.82
Dif-fuse	TMI ₂₄	81.82	99.95	69.55	63.65	99.97	71.57	73.08	63.16	66.67	76.47
GatingAno	PR ₂₄	78.73	92.07	11.49	63.04	92.09	54.25	53.85	45.45	55.56	52.94
DTU-Net	WACV ₂₅	75.70	99.95	58.68	51.44	99.97	68.63	69.23	60.00	66.67	70.59
WDT-MD	Ours	82.80	99.96	74.43	65.61	99.98	85.95	88.46	82.35	77.78	94.12

Table 1: Quantitative comparison of the proposed WDT-MD method with the state-of-the-art methods on the IDRiD dataset. Best results are highlighted as **first**, **second** and **third**. (Unit: %)

Method	Source	Pixel-level					Image-level				
		AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE
AnoDDPM	CVPR ₂₂	79.26	99.96	32.09	58.56	99.97	60.00	61.54	51.61	53.33	66.67
CPC	WACV ₂₃	76.28	99.98	35.34	52.57	99.98	65.42	66.67	58.06	60.00	70.83
DAE	MedIA ₂₃	72.64	99.95	21.96	45.31	99.96	57.08	56.41	51.43	60.00	54.17
HACDR-Net	AAAI ₂₄	54.13	98.03	3.53	9.56	98.04	41.67	43.59	31.25	33.33	50.00
AE[$d_{optimal}$]	MICCAI ₂₄	78.86	99.98	32.21	57.75	99.98	62.08	64.10	53.33	53.33	70.83
Dif-fuse	TMI ₂₄	80.82	99.96	32.48	61.67	99.96	61.25	61.54	54.55	60.00	62.50
GatingAno	PR ₂₄	78.27	98.83	2.253	58.45	98.84	63.33	64.10	56.25	60.00	66.67
DTU-Net	WACV ₂₅	80.72	99.98	42.99	61.46	99.98	49.58	48.72	44.44	53.33	45.83
WDT-MD	Ours	81.08	99.99	57.70	62.16	99.99	70.83	71.79	64.52	66.67	75.00

Table 2: Quantitative comparison of the proposed WDT-MD method with the state-of-the-art methods on the e-optha MA dataset. Best results are highlighted as **first**, **second** and **third**. (Unit: %)

τ	ψ	Pixel-level					Image-level				
		AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE
\times	\times	73.17	97.17	6.39	49.28	97.20	48.37	46.15	41.67	55.56	41.18
\times	\checkmark	49.98	98.78	0.53	0.30	98.84	42.81	42.31	34.78	44.44	41.18
\checkmark	\times	67.20	68.84	0.87	62.60	68.84	42.16	34.62	41.38	66.67	17.65
\checkmark	\checkmark	82.80	99.96	74.43	65.61	99.98	85.95	88.46	82.35	77.78	94.12

Table 3: Ablation study of core components of WDT-MD on IDRiD. τ denotes noise encoding in image conditioning, and ψ denotes pixel-level supervision. Best results are highlighted as **first**, **second** and **third**. (Unit: %)

τ	ψ	Pixel-level					Image-level				
		AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE
\times	\times	58.78	97.03	0.36	20.47	97.04	47.08	48.72	37.50	40.00	54.17
\times	\checkmark	51.69	98.14	5.29	4.05	98.15	45.83	48.72	33.33	33.33	58.33
\checkmark	\times	68.06	79.84	0.17	56.97	79.84	61.25	61.54	54.55	60.00	62.50
\checkmark	\checkmark	81.08	99.99	57.70	62.16	99.99	70.83	71.79	64.52	66.67	75.00

Table 4: Ablation study of core components of WDT-MD on e-optha MA. τ denotes noise encoding in image conditioning, and ψ denotes pixel-level supervision. Best results are highlighted as **first**, **second** and **third**. (Unit: %)

Tokenizer	Pixel-level					Image-level					Params (M)	FLOPs (G)
	AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE		
\times	81.03	99.92	64.24	62.09	99.94	77.45	80.77	70.59	66.67	88.24	41.49	478.30
AE-KL	80.81	99.93	67.53	61.65	99.95	68.95	73.08	58.82	55.56	82.35	36.04	229.36
VQ-VAE	81.31	99.91	68.19	62.65	99.93	71.57	73.08	63.16	66.67	76.47	36.24	146.75
VQGAN	81.37	99.95	69.82	62.76	99.97	77.45	80.77	70.59	66.67	88.24	39.00	150.96
DWT (Ours)	82.80	99.96	74.43	65.61	99.98	85.95	88.46	82.35	77.78	94.12	35.04	119.76

Table 5: Impact of various tokenizers for compression in WDT-MD on the IDRiD dataset. Best results are highlighted as **first**, **second** and **third**. (Unit: %, Params: number of model parameters)

Tokenizer	Pixel-level					Image-level					Params (M)	FLOPs (G)
	AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE		
\times	77.81	99.90	19.86	55.70	99.91	44.17	43.59	38.89	46.67	41.67	41.49	478.30
AE-KL	78.22	99.94	22.52	56.50	99.94	53.75	53.85	47.06	53.33	54.17	36.04	229.36
VQ-VAE	79.89	99.94	27.49	59.84	99.95	55.83	56.41	48.48	53.33	58.33	36.24	146.75
VQGAN	79.98	99.97	41.02	59.99	99.97	55.83	56.41	48.48	53.33	58.33	39.00	150.96
DWT (Ours)	81.08	99.99	57.70	62.16	99.99	70.83	71.79	64.52	66.67	75.00	35.04	119.76

Table 6: Impact of various tokenizers for compression in WDT-MD on the e-ophtha MA dataset. Best results are highlighted as **first**, **second** and **third**. (Unit: %, Params: number of model parameters)

Backbone	Pixel-level					Image-level					Params (M)	FLOPs (G)
	AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE		
Attention U-Net	79.88	99.86	56.32	59.85	99.88	71.57	73.08	63.16	66.67	76.47	382.98	193.48
U-ViT	80.89	99.94	69.79	61.81	99.96	74.51	76.92	66.67	66.67	82.35	59.67	313.89
DiT (N = 1)	79.18	99.93	61.31	58.40	99.95	63.07	65.38	52.63	55.56	70.59	5.78	10.17
DiT (N = 2)	80.62	99.92	61.95	61.29	99.95	71.57	73.08	63.16	66.67	76.47	8.44	20.14
DiT (N = 4)	81.22	99.94	67.76	62.46	99.97	80.07	80.77	73.68	77.78	82.35	13.76	40.06
DiT (N = 8)	81.93	99.94	69.43	63.90	99.97	83.01	84.62	77.78	77.78	88.24	24.40	79.91
DiT (N = 12, Ours)	82.80	99.96	74.43	65.61	99.98	85.95	88.46	82.35	77.78	94.12	35.04	119.76

Table 7: Impact of various backbones of the noise estimator network in WDT-MD on the IDRiD dataset. Best results are highlighted as **first**, **second** and **third**. (N: number of DiT blocks, Unit: %, Params: number of model parameters)

Backbone	Pixel-level					Image-level					Params (M)	FLOPs (G)
	AUC	ACC	F1	SEN	SPE	AUC	ACC	F1	SEN	SPE		
Attention U-Net	79.70	99.96	37.85	59.43	99.96	52.50	53.85	43.75	46.67	58.33	382.98	193.48
U-ViT	80.39	99.96	31.57	60.83	99.96	61.25	61.54	54.55	60.00	62.50	59.67	313.89
DiT (N = 1)	76.99	99.97	33.90	54.02	99.97	57.08	56.41	51.43	60.00	54.17	5.78	10.17
DiT (N = 2)	77.41	99.97	36.68	54.85	99.98	55.83	56.41	48.48	53.33	58.33	8.44	20.14
DiT (N = 4)	79.97	99.97	44.37	59.96	99.98	51.67	51.28	45.71	53.33	50.00	13.76	40.06
DiT (N = 8)	80.10	99.97	44.48	60.21	99.98	57.08	56.41	51.43	60.00	54.17	24.40	79.91
DiT (N = 12, Ours)	81.08	99.99	57.70	62.16	99.99	70.83	71.79	64.52	66.67	75.00	35.04	119.76

Table 8: Impact of various backbones of the noise estimator network in WDT-MD on the e-ophtha MA dataset. Best results are highlighted as **first**, **second** and **third**. (N: number of DiT blocks, Unit: %, Params: number of model parameters)




Figure 5: The qualitative results for WDT-MD compared with other state-of-the-art methods on IDRiD and e-optha MA.

AE[$d_{optimal}$] (Cai, Chen, and Cheng 2024) and DAE (Kascenas et al. 2023); (4) U-Net-based segmentation methods HACDR-Net (Xu et al. 2024) and CPC (Yap and Ng 2023).

Quantitative results are presented in Table 1 and Table 2. WDT-MD demonstrates top performance, outperforming current state-of-the-art methods, spanning both AD and segmentation frameworks. At the pixel level, it achieves an AUC of 82.80% on IDRiD and 81.08% on e-optha MA. In terms of the image level, it reaches 85.95% on IDRiD and 70.83% on e-optha MA. Furthermore, Fig. 5 illustrates the qualitative results on both datasets, highlighting WDT-MD’s remarkable ability to precisely detect subtle MAs.

Ablation Study

Ablation Study of Core Components. Our ablation study highlights the notable improvements brought by our method, as shown in Table 3 and Table 4. Specifically, the pixel-level SEN is improved to 65.61% and 62.16% on IDRiD and e-optha MA, respectively. This underscores that noise-encoded image conditioning effectively mitigates “identity mapping” by preventing mere replication from the image condition. Furthermore, the integration of pixel-level supervision markedly reduces false positives, yielding improvements of (31.14%/76.47%) in (pixel-level/image-level) SPE on IDRiD and (20.15%/12.50%) on e-optha MA.

Impact of Tokenizers. To investigate the impact of different tokenization strategies, we conducted comparative experiments. As presented in Table 5 and Table 6, DWT achieves the best performance, with improvements of (4.61%/11.76%) in (pixel-level/image-level) F1 score on IDRiD and (16.68%/16.04%) on e-optha MA. This underscores the advantage of DWT in detail preservation and multi-scale feature modeling. Notably, compared with training-based tokenizers such as VQGAN, the use of DWT reduces the Params by 1.00M and the FLOPs by 18.39%, demonstrating its superiority in computational efficiency.

Impact of Backbones. The impact of different backbones on WDT-MD performance is evident in Tables 7 and 8. On both IDRiD and e-optha MA, our DiT backbone ($N=12$)

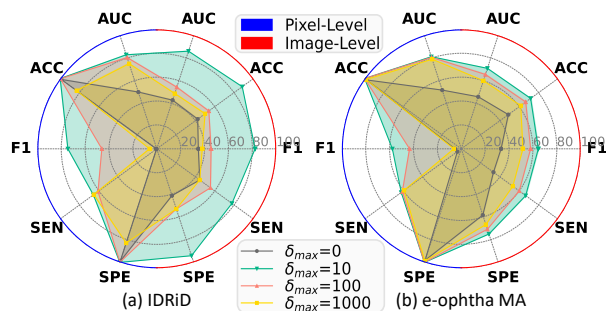


Figure 6: The quantitative results for WDT-MD under different maximum noise encoding timesteps δ_{max} . (Unit: %)

delivers the best performance. Notably, compared with Attention U-Net (Dhariwal and Nichol 2021), the most commonly used backbone in diffusion models, our backbone reduces the Params by 90.85% and FLOPs by 38.10%, reinforcing its suitability for clinical deployment.

Impact of Noise Encoding Timesteps. Furthermore, we explored the impact of varying maximum noise encoding timesteps. As depicted in Fig. 6, our model performs best at $\delta_{max} = 10$. This indicates that moderate noise encoding effectively mitigates “identity mapping”, while too large δ_{max} implies excessive noise injection and more difficult model convergence, causing performance degradation.

Conclusion

In this paper, we introduce WDT-MD, a novel supervised AD framework for MA detection. WDT-MD incorporates the noise-encoded image conditioning mechanism to mitigate “identity mapping”, pseudo-normal pattern synthesis to reduce false positives, and the wavelet diffusion Transformer architecture to enhance reconstruction quality of normal retinal features. Extensive experiments demonstrate WDT-MD’s superior performance. In future work, we will explore integration with multimodal ophthalmic data to further expedite clinical adoption of AI-powered early DR screening.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No. 2024ZD0536605, No. 2025YFE0103200), Transvascular Implantation Devices Research Institute (TIDRI) (No. KY052025008), National Natural Science Foundation of China (No. 61702146, No. 62076084, No. U20A20386, No. U22A2033, No. 62302399), Zhejiang Key Research and Development Program of China (No. 2024SSYS0026), Zhejiang Key Laboratory of Medical Imaging Artificial Intelligence, Zhejiang Provincial Natural Science Foundation of China (No. LY21F020017, No. 2023C03090), Guangdong Basic and Applied Basic Research Foundation (No. 2025A1515011617), and National Undergraduate Innovation and Entrepreneurship Training Program of China (No. 202510336076, No. 202410336081).

References

- Arrigo, A.; Aragona, E.; Teussink, M.; Battaglia Parodi, M.; and Bandello, F. 2024. Digital histology of retinal microaneurysms as provided by dense B-scan (DART) OCTA: Characteristics and clinical relevance in diabetic retinopathy. *Eye*, 38(16): 3108–3112.
- Baitieva, A.; Hurych, D.; Besnier, V.; and Bernard, O. 2024. Supervised anomaly detection for complex industrial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17754–17762.
- Bao, F.; Nie, S.; Xue, K.; Cao, Y.; Li, C.; Su, H.; and Zhu, J. 2023. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22669–22679.
- Baugh, M.; Reynaud, H.; Marimont, S. N.; Cechnicka, S.; Müller, J. P.; Tarroni, G.; and Kainz, B. 2024. Image-conditioned diffusion models for medical anomaly detection. In *International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging*, 117–127. Springer.
- Cai, Y.; Chen, H.; and Cheng, K.-T. 2024. Rethinking autoencoders for medical anomaly detection from a theoretical perspective. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 544–554. Springer.
- Decenciere, E.; Cazuguel, G.; Zhang, X.; Thibault, G.; Klein, J.-C.; Meyer, F.; Marcotegui, B.; Quéllec, G.; Lamard, M.; Danno, R.; et al. 2013. TeleOphta: Machine learning and image processing methods for teleophthalmology. *IRBM*, 34(2): 196–203.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *International Conference on Machine Learning*, 12606–12633. PMLR.
- Feng, K.; Ma, Y.; Wang, B.; Qi, C.; Chen, H.; Chen, Q.; and Wang, Z. 2025. Dit4edit: Diffusion transformer for image editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 2969–2977.
- Fontanella, A.; Mair, G.; Wardlaw, J.; Trucco, E.; and Storkey, A. 2024. Diffusion models for counterfactual generation and anomaly detection in brain images. *IEEE Transactions on Medical Imaging*.
- Foo, A.; Hsu, W.; and Lee, M. L. 2023. Multi-object representation learning via feature connectivity and object-centric regularization. *Advances in Neural Information Processing Systems*, 36: 60035–60047.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
- Guo, J.; Lu, S.; Zhang, W.; Chen, F.; Li, H.; and Liao, H. 2025. Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 20405–20415.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Huang, Y.; Huang, J.; Liu, J.; Yan, M.; Dong, Y.; Lv, J.; Chen, C.; and Chen, S. 2024. Wavedm: Wavelet-based diffusion models for image restoration. *IEEE Transactions on Multimedia*, 26: 7058–7073.
- Jiang, H.; Gao, M.; Liu, Z.; Tang, C.; Zhang, X.; Jiang, S.; Yuan, W.; and Liu, J. 2024. GlanceSeg: Real-time microaneurysm lesion segmentation with gaze-map-guided foundation model for early detection of diabetic retinopathy. *IEEE Journal of Biomedical and Health Informatics*.
- Jiang, H.; Hou, Y.; Miao, H.; Ye, H.; Gao, M.; Li, X.; Jin, R.; and Liu, J. 2023. Eye tracking based deep learning analysis for the early detection of diabetic retinopathy: A pilot study. *Biomedical Signal Processing and Control*, 84: 104830.
- Kascenas, A.; Sanchez, P.; Schrempf, P.; Wang, C.; Clackett, W.; Mikhael, S. S.; Voisey, J. P.; Goatman, K.; Weir, A.; Pugeault, N.; et al. 2023. The role of noise in denoising models for anomaly detection in medical images. *Medical Image Analysis*, 90: 102963.
- Khan, Z.; Gaidhane, A. M.; Singh, M.; Ganesan, S.; Kaur, M.; Sharma, G. C.; Rani, P.; Sharma, R.; Thapliyal, S.; Kushwaha, M.; et al. 2025. Diagnostic accuracy of IDX-DR for detecting diabetic retinopathy: A systematic review and Meta-Analysis. *American Journal of Ophthalmology*, 273: 192–204.
- Kumar, K.; Chakraborty, S.; Mahapatra, D.; Bozorgtabar, B.; and Roy, S. 2025. Self-supervised anomaly segmentation via diffusion models with dynamic Transformer UNet. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 7928–7938.
- Li, Y.; Feng, Y.; Chen, B.; Chen, W.; Wang, Y.; Hu, X.; Sun, B.; Qu, C.; and Zhou, M. 2024. Vague prototype-oriented diffusion model for multi-class anomaly detection. In *International Conference on Machine Learning*, 27771–27790. PMLR.

- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024a. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, 23–40. Springer.
- Ma, Y.; He, Y.; Wang, H.; Wang, A.; Shen, L.; Qi, C.; Ying, J.; Cai, C.; Li, Z.; Shum, H.-Y.; et al. 2025a. Follow-your-click: Open-domain regional image animation via motion prompts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 6018–6026.
- Ma, Y.; Liu, H.; Wang, H.; Pan, H.; He, Y.; Yuan, J.; Zeng, A.; Cai, C.; Shum, H.-Y.; Liu, W.; et al. 2024b. Follow-your-emoji: Fine-controllable and expressive freestyle portrait animation. In *SIGGRAPH Asia 2024 Conference Papers*, 1–12.
- Ma, Y.; Yan, Z.; Liu, H.; Wang, H.; Pan, H.; He, Y.; Yuan, J.; Zeng, A.; Cai, C.; Shum, H.-Y.; et al. 2025b. Follow-your-emoji-faster: Towards efficient, fine-controllable, and expressive freestyle portrait animation. *arXiv preprint arXiv:2509.16630*.
- Mayya, V.; Kamath, S.; and Kulkarni, U. 2021. Automated microaneurysms detection for early diagnosis of diabetic retinopathy: A comprehensive review. *Computer Methods and Programs in Biomedicine Update*, 1: 100013.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4195–4205.
- Porwal, P.; Pachade, S.; Kamble, R.; Kokare, M.; Deshmukh, G.; Sahasrabudde, V.; and Meriaudeau, F. 2018. Indian diabetic retinopathy image dataset (IDRiD): a database for diabetic retinopathy screening research. *Data*, 3(3): 25.
- Porwal, P.; Pachade, S.; Kokare, M.; Deshmukh, G.; Son, J.; Bae, W.; Liu, L.; Wang, J.; Liu, X.; Gao, L.; et al. 2020. IdriD: Diabetic retinopathy–segmentation and grading challenge. *Medical Image Analysis*, 59: 101561.
- Raghu, M.; Zhang, C.; Kleinberg, J.; and Bengio, S. 2019. Transfusion: Understanding transfer learning for medical imaging. *Advances in Neural Information Processing Systems*, 32.
- Shao, M.; Miao, X.; Duan, H.; Wang, Z.; Chen, J.; Huang, Y.; Wu, X.; Deng, J.; Long, Y.; and Zheng, Y. 2025a. Trace: Temporally reliable anatomically-conditioned 3D CT generation with enhanced efficiency. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 627–637. Springer.
- Shao, M.; Wang, Z.; Duan, H.; Huang, Y.; Zhai, B.; Wang, S.; Long, Y.; and Zheng, Y. 2025b. Rethinking brain tumor segmentation from the frequency domain perspective. *IEEE Transactions on Medical Imaging*.
- Sun, H.; Cao, Y.; Dong, H.; and Fink, O. 2025a. Unseen visual anomaly generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 25508–25517.
- Sun, Y.; Chen, Z.; Zheng, H.; Ge, R.; Liu, J.; Min, W.; Elazab, A.; Wan, X.; and Wang, C. 2025b. Bs-ldm: Effective bone suppression in high-resolution chest X-ray images with conditional latent diffusion models. *IEEE Journal of Biomedical and Health Informatics*.
- Sun, Y.; Chen, Z.; Zheng, H.; Lu, Y.; Duan, L.; Fan, F.; Elazab, A.; Wan, X.; Wang, C.; and Ge, R. 2025c. Gl-lcm: Global-local latent consistency models for fast high-resolution bone suppression in chest X-ray images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 222–232. Springer.
- Telea, A. 2004. An image inpainting technique based on the fast marching method. *Journal of Graphics Tools*, 9(1): 23–34.
- Wang, K.; Bai, Y.; Li, D.; Zhai, D.; Jiang, J.; and Liu, X. 2024. Learning lossless compression for high bit-depth volumetric medical image. *IEEE Transactions on Image Processing*.
- Wang, Y.; Xu, J.; and Zhao, Y. 2024. Wavelet encoding network for inertial signal enhancement via feature supervision. *IEEE Transactions on Industrial Informatics*.
- Wolleb, J.; Bieder, F.; Sandkühler, R.; and Cattin, P. C. 2022. Diffusion models for medical anomaly detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 35–45. Springer.
- Wu, Y.; and Jiao, G. 2024. MicroSeg: Multi-scale fusion learning for microaneurysms segmentation. *Biomedical Signal Processing and Control*, 97: 106700.
- Wyatt, J.; Leach, A.; Schmon, S. M.; and Willcocks, C. G. 2022. Anoddp: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 650–656.
- Xu, Q.; Luo, X.; Huang, C.; Liu, C.; Wen, J.; Wang, J.; and Xu, Y. 2024. HACDR-Net: heterogeneous-aware convolutional network for diabetic retinopathy multi-lesion segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 6342–6350.
- Yap, B. P.; and Ng, B. K. 2023. Cut-paste consistency learning for semi-supervised lesion segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 6160–6169.
- Zhang, W.; Liu, H.; Xie, J.; Huang, Y.; Zhang, Y.; Li, Y.; Ramachandra, R.; and Zheng, Y. 2024. Anomaly detection via gating highway connection for retinal fundus images. *Pattern Recognition*, 148: 110167.
- Zhao, C.; Cai, W.; Dong, C.; and Hu, C. 2024. Wavelet-based fourier information interaction with frequency diffusion adjustment for underwater image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8281–8291.