Improved Vascular Segmentation via Binary Flow Matching with Straight-Path Regularization

Yongjun Kim[1,2], Kyong Hwan Jin[2], Yoseob Han[3], and Kanghyun Ryu[1*]

[1] Korea Institute of Science and Technology (KIST), Seoul, Republic of Korea
[2] Department of Electrical and Computer Engineering, Korea University, Seoul, Republic of Korea
[3] Department of Intelligent Semiconductors, Soongsil University, Seoul, Republic of Korea
khryu@kist.re.kr

Abstract. Thin vessel segmentation is challenging, where distinguishing between background artifacts and vessel-like structures is ambiguous. Recently, diffusion models have emerged as good candidates for vessel segmentation but they suffer from tens of iterations and stochastic noisy components. To overcome this, we formulate vascular segmentation as a binary iterative flow matching method named as **Binary Flow Matching (BFM)**. BFM models segmentation updates as deterministic velocity fields in a Bernoulli parameter space, enabling stable refinement without iterative Gaussian noise injection. Through $\mathbf{x}_1$-parameterization, BFM directly predicts the final clean mask at each timestep, allowing supervision with segmentation-specific losses. In addition, we adopt straight-path regularization to enforce consistent update trajectories. Across three X-ray coronary angiography datasets, BFM achieves the best or highly competitive performance compared with single-pass discriminative models and diffusion-based approaches such as BerDiff in Dice, clDice, and Betti-0 error, while requiring only 5–10 iterative steps compared to 50-step diffusion sampling. These results demonstrate that BFM offers a more efficient and topology-preserving alternative to stochastic diffusion for vascular segmentation. *Source code is available at https://github.com/YongjunKim98/BFM.*

1 Introduction

Segmentation of thin vascular structures is challenging due to ambiguity. Vessels and background artifacts often share similar intensity profiles and tubular geometry, making discrimination from local appearance alone unreliable. In X-ray coronary angiography, catheter shadows and guidewires frequently mimic vessel structures [1, 2]. Similar challenges arise in the DCA1 dataset [3] where motion artifacts and heterogeneous contrast generate vessel-like background patterns. Resolving these challenges requires iterative structural reasoning beyond pixel-level detection, particularly vessel connectivity and topology, which single-pass segmentation models cannot revisit or refine.

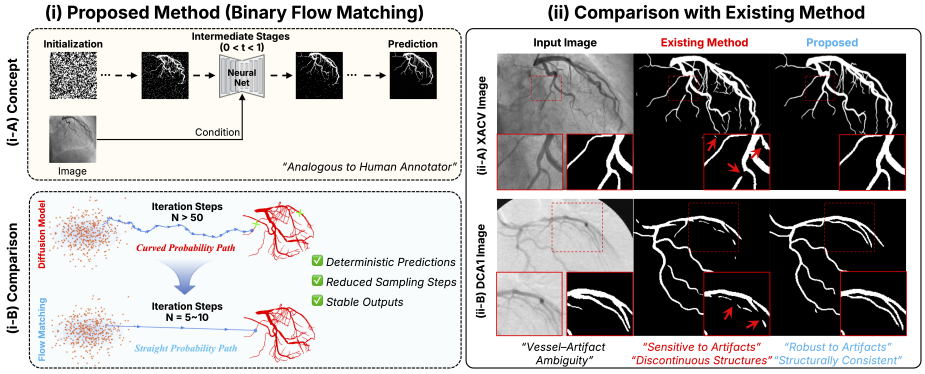

Fig. 1: Overview of our method. (i) BFM refines binary noise into vessel segmentation via deterministic probability flow. (ii) Qualitative comparison shows improved vessel-artifact disambiguation and structural consistency.

Even expert annotators require multiple iterations to confidently distinguish vessels from artifacts, verifying connectivity with surrounding vascular topology. Existing vessel segmentation methods [4,5,6,7] rely on single-pass inference without any mechanism to revisit uncertain regions, leading to fragmented vessels and false positives where background structures are mistaken for vessels. Topological consistency has gained increasing attention, with recent losses such as clDice [8] and cbDice [9] penalizing structural inconsistencies during training. However, inference remains single-pass, and iterative generative refinement for vascular segmentation remains largely unexplored.

Diffusion models [10, 11, 12] enable iterative refinement for general medical segmentation. In particular, BerDiff [11] formulates binary segmentation within a Bernoulli diffusion framework and demonstrates improved topological consistency over single-pass models. However, reverse diffusion employs stochastic sampling via Langevin dynamics, adding noise at each step that introduces additional ambiguity for thin vessels. Moreover, diffusion objectives optimize noise reconstruction rather than segmentation metrics [8]. Flow matching [13] constructs deterministic ordinary differential equation (ODE) trajectories more consistent with the deterministic nature of annotation, but recent methods [14, 15] operate in continuous spaces, introducing boundary blur that is suboptimal for thin vessels.

To address these limitations, we present Binary Flow Matching (BFM), a flow-based refinement framework for binary vessel masks. Unlike diffusion models that rely on stochastic noise injection, BFM performs deterministic probability transport in the probability space. At each refinement step, the probability state is projected to a binary mask via Bernoulli sampling for discrete-state conditioning. This design reduces the accumulation of ambiguous intermediate probabilities and promotes stable topology evolution. Furthermore, by predicting the final clean mask at each timestep, BFM enables direct optimization with

segmentation-oriented objectives, aligning training with evaluation metrics. Our contributions are summarized as follows:

1. **Bernoulli probability flow formulation.** We model segmentation as deterministic transport in Bernoulli parameter space, where probability dynamics follow an ODE while discrete states are sampled at each step. This design mitigates ambiguity accumulation and promotes topology-consistent refinement.
2. **$\mathbf{x}_1$-parameterized objective for segmentation.** We predict the final clean mask at each timestep, enabling direct supervision with segmentation-specific losses instead of noise reconstruction.
3. **Straight-path regularization.** We introduce a trajectory regularizer that enforces locally linear probability evolution, improving refinement stability and reducing the number of required inference steps.

Related Work: UNet [4] established the encoder-decoder paradigm for medical image segmentation, and subsequent work has extended it in multiple directions. nnU-Net [5] achieves strong results through automated training configuration, while TransUNet [16] and Swin-UNet [17] incorporate transformer modules to capture long-range dependencies. For vascular structures specifically, CS^2-Net [6] and DSCNet [7] design convolutional operations tailored to tubular geometry. On the loss side, clDice [8] and cbDice [9] encourage topological consistency by penalizing centerline disconnections. Despite these advances, all these methods perform single-pass inference and have no mechanism to revise uncertain predictions.

Diffusion models [10, 11, 12] offer iterative refinement by reversing a learned noise process. However, their reverse process relies on stochastic Langevin sampling, injecting Gaussian noise $\sigma_t \mathbf{z}_t$, $\mathbf{z}_t \sim \mathcal{N}(0, I)$ at every time step t, which adds ambiguity rather than resolving it for thin vessels. Conditional Flow Matching [13] instead learns a velocity field along deterministic ordinary differential equation (ODE) trajectories $\frac{d\mathbf{x}_t}{dt} = v_\theta(\mathbf{x}_t, t, \mathbf{c})$, where $\mathbf{c}$ denotes the conditioning image. Recent applications to segmentation, FlowSDF [14] and LatentFM [15], operate in continuous spaces that introduce boundary blur, making them suboptimal for thin vessel delineation.

2 Methods

Problem Setup: Given an input image $\mathbf{c} \in \mathbb{R}^{H \times W}$, where H and W denote the height and width of the image, we aim to estimate a binary mask $\hat{\mathbf{x}}_1 \in \{0, 1\}^{H \times W}$. We model segmentation as probability transport over a Bernoulli parameter field $\mathbf{p}_t \in [0, 1]^{H \times W}$, where binary states are sampled as: $\mathbf{x}_t \sim \mathcal{B}(\mathbf{p}_t)$. The probability dynamics follow a deterministic ODE initialized from $\mathbf{x}_0 \sim \mathcal{B}(0.5)$ with $\mathbf{p}_0 = 0.5$.

$$\frac{d\mathbf{p}_t}{dt} = \mathbf{v}_\theta(\mathbf{p}_t, t, \mathbf{c}). \tag{1}$$

Algorithm 1 Training

Require: Model f_θ, image c, GT x_1
$x_0 \sim \mathcal{B}(0.5)$, $t \sim \mathcal{U}(0,1)$
$u \sim \mathcal{U}(0,1)^{H\times W}$
$x_t = x_1 \odot \mathbf{1}[u < t] + x_0 \odot \mathbf{1}[u \geq t]$
$\hat{p}_1 = \sigma(f_\theta(x_t, c, t))$
$\mathcal{L}_{\text{seg}} = \mathcal{L}_{\text{BCE}}(\hat{p}_1, x_1) + \mathcal{L}_{\text{Dice}}(\hat{p}_1, x_1)$

Straight-Path Regularization:
Sample t_a, t_b; $t_{\text{mid}} = (t_a + t_b)/2$
$v_t = (\hat{p}_1 - p_t)/(1 - t + \epsilon)$
$p_{t_b}^{\text{direct}} = p_{t_a} + (t_b - t_a)v_t$
$p_{t_b}^{\text{two-step}}$: advance $t_a \to t_{\text{mid}} \to t_b$
$\mathcal{L}_{\text{reg}} = \|p_{t_b}^{\text{direct}} - p_{t_b}^{\text{two-step}}\|_2^2$

$\theta \leftarrow \theta - \eta\nabla_\theta(\mathcal{L}_{\text{seg}} + \lambda\mathcal{L}_{\text{reg}})$

Algorithm 2 Inference (Euler)

Require: Image c, steps N, model f_θ
$p_{t_0} \leftarrow 0.5$, $\Delta t = 1/N$
$x_{t_0} \sim \text{Bernoulli}(p_{t_0})$
for $k = 0$ **to** $N - 1$ **do**
  $t_k = k \cdot \Delta t$
  $\hat{p}_1 \leftarrow \sigma(f_\theta(x_{t_k}, c, t_k))$
  $p_{t_{k+1}} \leftarrow p_{t_k} + \frac{\hat{p}_1 - p_{t_k}}{1 - t_k + \epsilon} \cdot \Delta t$
  $p_{t_{k+1}} \leftarrow \text{clip}(p_{t_{k+1}}, 0, 1)$
  $x_{t_{k+1}} \sim \text{Bernoulli}(p_{t_{k+1}})$
end for
$\hat{x}_1 \leftarrow \mathbf{1}[p_{t_N} > 0.5]$
return $\hat{x}_1$

Training via $\mathbf{x}_1$-Parameterization: Direct velocity regression in conditional flow matching is not well-defined for binary masks, as valid states are restricted to $\{0,1\}^{H\times W}$. To address this limitation, we adopt an $\mathbf{x}_1$-parameterization strategy and reformulate flow matching as direct final-mask prediction at each timestep. Inspired by clean-sample prediction in diffusion models [18], this design is tailored to Bernoulli probability transport for binary segmentation.

At every training iteration, we sample an initial binary state and timestep, $\mathbf{x}_0 \sim \mathcal{B}(0.5)$ and $t \sim \mathcal{U}(0,1)$, followed by a pixel-wise random switch variable $\mathbf{u} \sim \mathcal{U}(0,1)^{H\times W}$. The intermediate binary state is then constructed as:

$$\mathbf{x}_t = \mathbf{x}_1 \odot \mathbf{1}[\mathbf{u} < t] + \mathbf{x}_0 \odot \mathbf{1}[\mathbf{u} \geq t]. \tag{2}$$

The network predicts the clean mask probability map:

$$\hat{\mathbf{p}}_1 = \sigma(f_\theta(\mathbf{x}_t, \mathbf{c}, t)), \qquad \text{where } \hat{\mathbf{p}}_1 \in [0,1]^{H\times W}. \tag{3}$$

Rather than directly regressing the velocity with a velocity-MSE objective, the network predicts the clean mask probability, from which the probability-space refinement direction is analytically computed as:

$$\mathbf{v}_t = \frac{\hat{\mathbf{p}}_1 - \mathbf{p}_t}{1 - t + \epsilon}, \tag{4}$$

which represents the deterministic update direction toward the final mask, with $\epsilon > 0$ denoting a small constant for numerical stability. This formulation enables direct optimization using segmentation-specific objectives, where the segmentation loss is defined as $\mathcal{L}_{\text{seg}} = \mathcal{L}_{\text{BCE}}(\hat{\mathbf{p}}_1, \mathbf{x}_1) + \mathcal{L}_{\text{Dice}}(\hat{\mathbf{p}}_1, \mathbf{x}_1)$.

Straight-Path Regularization: Segmentation is inherently a deterministic prediction task, yet iterative refinement models often require many update steps

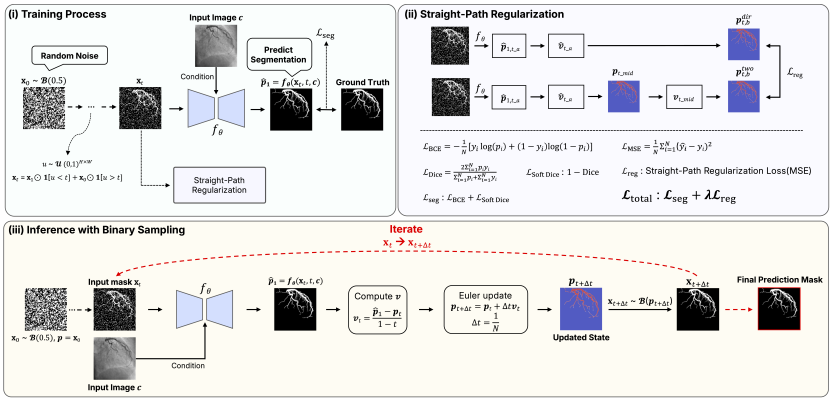

Fig. 2: Overview of Binary Flow Matching (BFM). (i) Training via stochastic binary interpolation and $\mathbf{x}_1$-parameterization. (ii) Straight-path regularization for efficiency as well as accuracy. (iii) Inference through deterministic Euler updates, clipping, and Bernoulli sampling.

for stable convergence. To mitigate this dependency, we encourage the induced probability flow to follow straight trajectories in Bernoulli parameter space, inspired by rectified flow formulations [19].

Let $\mathbf{p}_{t_b}^{\text{direct}}$ denote the probability map obtained via a direct update from t_a to t_b, and $\mathbf{p}_{t_b}^{\text{two-step}}$ the result of a two-step update through an intermediate timestep. We define the regularization objective as:

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_{t_a,t_b}\left[\left\|\mathbf{p}_{t_b}^{\text{direct}} - \mathbf{p}_{t_b}^{\text{two-step}}\right\|_2^2\right], \tag{5}$$

which penalizes trajectory curvature in probability space. This constraint promotes stable mask evolution under a small number of refinement iterations. The model is trained with the objective $\mathcal{L}_{\text{seg}} + \lambda\mathcal{L}_{\text{reg}}$, where λ controls the regularization strength (The complete procedure is summarized in **Algorithm 1**).

Inference with Binary Sampling: At inference, we initialize from random binary noise $\mathbf{x}_0 \sim \mathcal{B}(0.5)$ and set $\mathbf{p}_0 = 0.5$. At each timestep t, the network predicts the clean mask probability $\hat{\mathbf{p}}_1 = \sigma(f_\theta(\mathbf{x}_t, \mathbf{c}, t))$, which defines the update direction toward the final mask. The probability map is then updated as $\mathbf{p}_{t+\Delta t} = \mathbf{p}_t + \Delta t\,(\hat{\mathbf{p}}_1 - \mathbf{p}_t)/(1-t+\epsilon)$ followed by element-wise clipping to ensure $\mathbf{p}_{t+\Delta t} \in [0,1]$. We subsequently sample the next binary state $\mathbf{x}_{t+\Delta t} \sim \mathcal{B}(\mathbf{p}_{t+\Delta t})$. Although the probability dynamics are defined in terms of $\mathbf{p}_t$, the network operates on sampled binary states $\mathbf{x}_t$. Conditioning on discrete states encourages structural consistency during refinement and mitigates the accumulation of ambiguous intermediate probabilities.

Binary sampling is utilized by stochastically resolving uncertain predictions: pixels with intermediate probabilities are naturally driven toward discrete states, mitigating the accumulation of ambiguous structures across refinement steps. The full inference procedure is summarized in **Algorithm 2**.

3 Experiments

3.1 Dataset

We evaluate on three public X-ray coronary angiography datasets [3, 20, 21] to demonstrate robustness across different imaging protocols, institutions, and annotation standards. The XACV dataset [20] consists of high-resolution frames extracted from angiography videos with expert annotations, yielding 155/20/46 train/validation/test images. The DCA1 dataset [3] contains 134 expert-annotated vessel masks from IMSS (UMAE T1–León), split into 94/10/30. The SBCD dataset [21] provides 300 images at 512×512 resolution from the Second People's Hospital of Shenzhen, and is split into 210/30/60. All datasets undergo intensity normalization and standard augmentation (rotation, flipping, intensity jitter). We report Dice and IoU for pixel-wise accuracy, clDice [8] for connectivity, and Betti-0 error [22] for topological consistency. We focus on Betti-0 error because vessel disconnection and branch fragmentation are the dominant topological failures in 2D coronary angiography.

3.2 Experiments and Results

Implementation and Benchmarks: We compare BFM with representative segmentation models across different paradigms. For discriminative approaches, we include CNN-based models such as CE-Net [23], VesselNet [24], CS^2-Net [6], and DSCNet [7], as well as transformer-based architectures including TransUNet [16] and nnU-Net v2 [25].

For generative refinement models, we consider diffusion-based methods such as SegDiff [26], MedSegDiff [10], ColdSegDiff [12], and the Bernoulli diffusion model BerDiff [11]. We further compare with flow-based segmentation methods including FlowSDF [14]. On DCA1 [3] and SBCD [21], we compare against DSCNet, TransUNet, MedSegDiff, and BerDiff. All models follow comparable optimization settings, with method-specific objectives: BCE+Dice for discriminative baselines and BFM, MSE-based noise/velocity objectives for SegDiff, MedSegDiff, and FlowSDF, and BCE for BerDiff.

BFM adopts the U-Net backbone from MedSegDiff, where the input image and binary state are encoded separately, Fourier encoded, and fused at each encoder stage. Time conditioning uses sinusoidal embeddings followed by a two-layer MLP and is injected into ResNet blocks via scale-shift modulation.

All models are trained with the AdamW optimizer (weight decay 1×10^{-5}, initial learning rate 2×10^{-4}) and ReduceLROnPlateau for 300 epochs with a fixed batch size of 6. Model capacity is controlled by adjusting base channel dimensions to keep trainable parameter counts in a similar range. Exponential Moving Average (EMA) is applied only to diffusion-based and flow-based models following standard practice [27]. All experiments use identical data splits and preprocessing and are conducted on a single NVIDIA A6000 GPU. For BFM, the straight-path regularization weight is set to $\lambda = 0.1$ based on validation performance.

Table 1: Quantitative comparison on the XACV dataset. ↑ higher is better, ↓ lower is better. Best in **bold**, second-best underlined.

Paradigm	Methods (Year)	Dice↑	IoU↑	clDice↑	β_0 err↓	Iters	Params (M)	Inference Time (ms)
CNN/ Transformer	CE-Net [23] (2019)	0.84±0.04	0.73±0.06	0.83±0.05	34.9±13.4	1	12.8	37.35
	VesselNet [24] (2019)	0.84±0.04	0.72±0.06	0.82±0.05	28.8±10.6	1	13.0	32.41
	CS^2-Net [6] (2021)	0.83±0.05	0.71±0.07	0.80±0.06	57.2±19.8	1	8.4	60.54
	TransUNet [16] (2021)	0.84±0.04	0.72±0.06	0.82±0.05	47.0±15.1	1	11.2	98.62
	DSCNet [7] (2023)	0.84±0.05	0.72±0.07	0.82±0.06	42.0±12.6	1	8.1	309.46
	nnU-Net v2 [25] (2024)	0.82±0.06	0.70±0.08	0.79±0.07	53.41±12.16	1	7.8	83.80
Diffusion	SegDiff [26] (2021)	0.82±0.06	0.69±0.08	0.80±0.07	19.1±14.7	50	10.6	7101.06
	MedSegDiff [10] (2022)	0.82±0.05	0.70±0.07	0.80±0.06	10.5±9.0	50	10.6	4213.98
	ColdSegDiff [12] (2024)	0.81±0.06	0.68±0.08	0.79±0.07	38.3±14.6	50	10.7	6240.17
Bernoulli Diffusion	BerDiff [11] (2023)	0.83±0.06	0.71±0.07	0.83±0.06	20.2±3.9	50	14.4	12951.67
Flow	**BFM (Ours, Normal)**	**0.86**±0.04	**0.76**±0.06	**0.86**±0.04	**3.3**±2.4	10	10.7	890.81
	BFM (Ours, Faster)	**0.86**±0.04	**0.76**±0.06	**0.86**±0.04	**3.7**±2.3	5	10.7	407.10

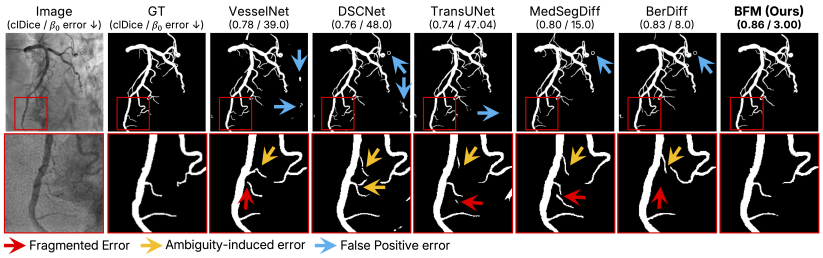

Fig. 3: Qualitative comparison on the XACV dataset. Colored arrows indicate fragmented (red), ambiguity-induced (yellow), and false-positive (blue) errors.

Quantitative and Qualitative Evaluation on XACV: Table 1 and Fig. 3 present quantitative and qualitative comparisons on the XACV dataset. BFM achieves the best performance among the compared methods across both pixel-wise and topological metrics. Compared to single-pass discriminative models [6, 7, 16, 23, 24, 25] that perform direct prediction in a single forward pass, BFM achieves higher overlap accuracy while reducing topological errors, indicating that iterative refinement can be beneficial for preserving vascular connectivity.

Diffusion-based approaches partially improve topology over single-pass models, supporting the value of iterative generation. However, their pixel-wise accuracy remains limited. In contrast, BFM preserves both structural coherence and boundary precision through deterministic probability flow with binary refinement. BFM further remains stable under reduced inference steps as shown in BFM (Faster) method of Table 1. Even when reducing the refinement steps by 50% (from 10 to 5), segmentation quality and β_0 errors remain consistent, suggesting well-behaved probability trajectories under coarse discretization. We further evaluate inference variance across five random seeds using the same trained checkpoint on the XACV test set. BFM shows negligible variability: Dice 0.86 ± 0.04, and β_0 error is also stable at 3.27 ± 0.31, indicating robustness.

Table 2: (i) Multi-dataset evaluation and (ii) Ablation on XACV. Best in **bold**.

(i)	DCA1 [3]			SBCD [21]		
Methods	**Dice↑**	**clDice↑**	**β_0 err↓**	**Dice↑**	**clDice↑**	**β_0 err↓**
DSCNet	0.77±0.04	0.84±0.05	7.4±4.0	0.87±0.03	0.79±0.06	45.2±53.3
TransUNet	0.79±0.03	0.85±0.04	6.6±4.0	0.87±0.03	0.79±0.06	46.7±48.7
MedSegDiff	0.77±0.04	0.85±0.04	4.8±4.6	0.84±0.03	0.75±0.05	39.4±52.0
BerDiff	0.77±0.03	0.84±0.05	1.7±1.8	0.86±0.04	0.76±0.05	40.5±48.0
BFM (Ours)	**0.80**±0.04	**0.87**±0.05	**1.3**±1.9	**0.89**±0.02	**0.82**±0.05	**18.80**±24.81

(ii)	XACV (Ablation)		
Methods	**Dice↑**	**clDice↑**	**β_0 err↓**
FlowSDF [14] (2025)	0.840±0.06	0.831±0.07	64.7±13.6
Vanilla FM [13] (2023)	0.839±0.06	0.830±0.07	15.5±8.3
+Binary	0.860±0.04	0.857±0.07	5.1±6.2
+Reg.	**0.864**±0.04	**0.862**±0.04	**3.3**±2.3

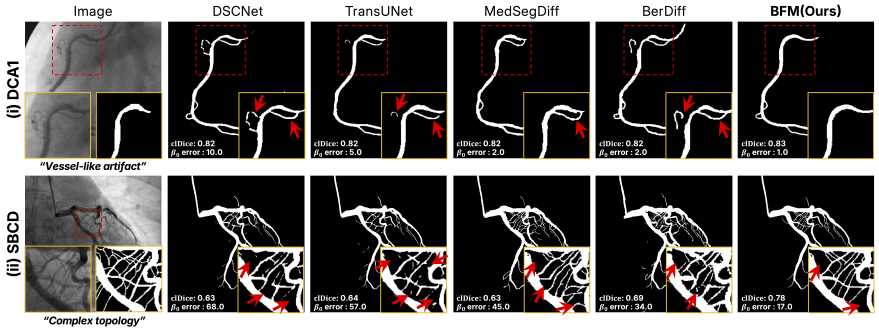

Fig. 4: Qualitative comparison on challenging cases. (i) DCA1 example exhibiting vessel-like artifacts, where several methods produce false positives near ambiguous regions (red arrows). (ii) SBCD example with complex vascular topology, showing connectivity disruptions and structural inconsistencies.

Multi-Dataset Evaluation: Table 2(i) reports quantitative results on the DCA1 and SBCD datasets, where all models are trained and evaluated independently. Despite differences in image characteristics, BFM consistently demonstrates strong accuracy. On the DCA1 dataset, which contains vessel-like artifacts and ambiguous background structures, BFM produces more accurate predictions compared to competing methods as shown in Fig. 4(i). Reduced β_0 errors suggest that our method effectively preserves vessel connectivity under artifact-induced ambiguity.

On the SBCD dataset, characterized by complex vessel distributions and topology, Fig. 4(ii) shows that BFM maintains accurate connectivity patterns. These results suggest that the proposed probability flow formulation shows consistent performance across datasets.

Ablation Study: Table 2(ii) analyzes the contribution of each component within the proposed framework. Vanilla flow matching serves as the baseline, while successive variants evaluate the impact of the Bernoulli parameterization and straight-path regularization.

Transitioning from the continuous formulation to the Bernoulli-based parameterization improves both segmentation accuracy and structural stability. This result indicates that probability transport in Bernoulli space is better aligned with binary vessel segmentation. Introducing the straight-path regularizer fur-

ther enhances performance across overlap and topological metrics. By discouraging curved probability trajectories, the refinement dynamics become more stable, leading to improved connectivity preservation reflected by reduced β_0 error.

4 Conclusion

We presented Binary Flow Matching (BFM), a probability-flow–based framework that performs iterative refinement through deterministic probability dynamics and discrete binary sampling. By adopting an $\mathbf{x}_1$-parameterization strategy, BFM enables segmentation-oriented optimization while maintaining stable probability transport. The proposed regularization further improves refinement stability, allowing predictions even under reduced inference steps. Overall, BFM demonstrates an effective and flexible approach for topology-preserving medical image segmentation. Future work includes extending BFM to 3D volumetric segmentation and using Betti-1 (β_1)-based evaluation to better assess loop-like vascular structures.

Acknowledgements. This work was supported by the Korea Institute of Science and Technology (KIST) Institutional Program under Grant No.26Z9032, 26E0212, 26E0214.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Kritika Iyer, Cyrus P Najarian, Aya A Fattah, Christopher J Arthurs, SM Reza Soroushmehr, Vijayakumar Subban, Mullasari A Sankardas, Raj R Nadakuditi, Brahmajee K Nallamothu, and C Alberto Figueroa. AngioNet: a convolutional neural network for vessel segmentation in X-ray angiography. *Scientific Reports*, 11(1):18066, 2021.
2. Rongjun Ge, Xin Li, Yuxing Liu, Chengliang Liu, Pinzheng Zhang, Jiong Zhang, Jian Yang, Jean-Louis Dillenseger, Chunfeng Yang, Yuting He, et al. Generative Data-engine Foundation Model for Universal Few-shot 2D Vascular Image Segmentation. *Medical Image Analysis*, 110:103996, 2026.
3. Fernando Cervantes-Sánchez, Iván Cruz-Aceves, Arturo Hernández-Aguirre, Martha Alicia Hernández-González, and Sergio Eduardo Solorio-Meza. Automatic Segmentation of Coronary Arteries in X-ray Angiograms Using Multiscale Analysis and Artificial Neural Networks. *Applied Sciences*, 9(24):5507, 2019.
4. 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*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015.
5. Fabian Isensee, Paul F. Jaeger, Simon A. A. 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.
6. Lei Mou, Yitian Zhao, Huazhu Fu, Yonghuai Liu, Jun Cheng, Yalin Zheng, Pan Su, Jianlong Yang, Li Chen, Alejandro F Frangi, et al. CS2-Net: Deep learning segmentation of curvilinear structures in medical imaging. *Medical Image Analysis*, 67:101874, 2021.

7. Yaolei Qi, Yuting He, Xiaoming Qi, Yuan Zhang, and Guanyu Yang. Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6070–6079, 2023.
8. Suprosanna Shit, Johannes C Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylka, Josien PW Pluim, Ulrich Bauer, and Bjoern H Menze. clDice: A Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16560–16569, 2021.
9. Pengcheng Shi, Jiesi Hu, Yanwu Yang, Zilve Gao, Wei Liu, and Ting Ma. Centerline Boundary Dice Loss for Vascular Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 46–56. Springer, 2024.
10. Junde Wu, Rao Fu, Huihui Fang, Yu Zhang, Yehui Yang, Haoyi Xiong, Huiying Liu, and Yanwu Xu. MedSegDiff: Medical image segmentation with diffusion probabilistic model. *arXiv preprint arXiv:2211.00611*, 2022.
11. Tao Chen, Chenhui Wang, and Hongming Shan. BerDiff: Conditional Bernoulli Diffusion Model for Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 491–501. Springer, 2023.
12. Pengfei Yan, Minglei Li, Jiusi Zhang, Guanyi Li, Yuchen Jiang, and Hao Luo. Cold SegDiffusion: A Novel Diffusion Model for Medical Image Segmentation. *Knowledge-Based Systems*, 301:112350, 2024.
13. Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling. In *International Conference on Learning Representations*, ICLR, 2023.
14. Lea Bogensperger, Dominik Narnhofer, Alexander Falk, Konrad Schindler, and Thomas Pock. FlowSDF: Flow Matching for Medical Image Segmentation Using Distance Transforms. *International Journal of Computer Vision*, 133:4864–4876, 2025.
15. Huynh Trinh Ngoc, Hoang Anh Nguyen Kim, Toan Nguyen Hai, and Long Tran Quoc. LatentFM: A Latent Flow Matching Approach for Generative Medical Image Segmentation. *arXiv preprint arXiv:2512.04821*, 2025.
16. Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L. Yuille, and Yuyin Zhou. TransUNet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.
17. Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-Unet: Unet-like pure transformer for medical image segmentation. *arXiv preprint arXiv:2105.05537*, 2021.
18. Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. *arXiv preprint arXiv:2511.13720*, 2025.
19. Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. *arXiv preprint arXiv:2209.03003*, 2022.
20. Chun-Hung Wu, Shih-Hong Chen, Chih-Yao Hu, Hsin-Yu Wu, Kai-Hsin Chen, Yu-You Chen, Chih-Hai Su, Chih-Kuo Lee, and Yu-Lun Liu. DeNVeR: Deformable Neural Vessel Representations for Unsupervised Video Vessel Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

21. Xina Loha. UniVG-58K. Hugging Face Datasets, `https://huggingface.co/datasets/xinaloha/UniVG`, 2026. Coronary artery DSA data sources include SBCD.
22. Nico Stucki, Johannes C Paetzold, Suprosanna Shit, Bjoern Menze, and Ulrich Bauer. Topologically faithful image segmentation via induced matching of persistence barcodes. In *International Conference on Machine Learning*, pages 32698–32727. PMLR, 2023.
23. Zaiwang Gu, Jun Cheng, Huazhu Fu, Kang Zhou, Huaying Hao, Yitian Zhao, Tianyang Zhang, Shenghua Gao, and Jiang Liu. CE-Net: Context encoder network for 2D Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 38(10):2281–2292, 2019.
24. Yicheng Wu, Yong Xia, Yang Song, Donghao Zhang, Dongnan Liu, Chaoyi Zhang, and Weidong Cai. Vessel-Net: Retinal vessel segmentation under multi-path supervision. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 264–272. Springer, 2019.
25. Fabian Isensee, Tassilo Wald, Constantin Ulrich, Michael Baumgartner, Saikat Roy, Klaus Maier-Hein, and Paul F Jaeger. nnU-Net Revisited: A Call for Rigorous Validation in 3D Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 488–498. Springer, 2024.
26. Tomer Amit, Tal Shaharbany, Eliya Nachmani, and Lior Wolf. SegDiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021.
27. 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 (ICLR)*, 2021.