

Ambiguous Medical Image Segmentation Using Diffusion Schrödinger Bridge

Lalith Bharadwaj Baru^{*1}, Kamalaker Dadi¹,
Tapabrata Chakraborti^{*2}, Raju S. Bapi¹

¹International Institute of Information Technology, Hyderabad, India.

²The Alan Turing Institute and University College London, London, UK.

^{*}lalith.baru@research.iiit.ac.in, tchakraborty@turing.ac.uk

Abstract. Accurate segmentation of medical images is challenging due to unclear lesion boundaries and mask variability. We introduce *Segmentation Schrödinger Bridge (SSB)*, the first application of Schrödinger Bridge for ambiguous medical image segmentation, modelling joint image-mask dynamics to enhance performance. SSB preserves structural integrity, delineates unclear boundaries without additional guidance, and maintains diversity using a novel loss function. We further propose the *Diversity Divergence Index (D_{DDI})* to quantify inter-rater variability, capturing both diversity and consensus. SSB achieves state-of-the-art performance on LIDC-IDRI, COCA, and RACER (in-house) datasets.

Keywords: Ambiguous Medical Image Segmentation · Schrödinger Bridge · Diffusion Models · Generative Modeling · Uncertainty Quantification

1 Introduction

Diagnosis, a cornerstone of medicine, heavily depends on individual assessment strategies. Recent studies highlight the prevalence of misdiagnosis even for common health conditions [13], emphasizing the importance of reducing diagnostic errors. Medical image segmentation, crucial in artificial intelligence (AI) enabled clinical decision support systems (CDSS), often relies on deterministic deep learning models that predict only a single mask per image [19]. However, varying clinician opinions on anomalies [1] result in low diagnostic consensus. Incorporating multiple expert interpretations improves diagnosis and reduces false negatives, but clinician time is a costly resource. Deterministic models have advanced medical image segmentation but often produce suboptimal results due to their bias toward the most likely hypothesis, while pixel-wise uncertainty models can yield inconsistent outputs. This necessitates models that capture diverse expert interpretations, such as the Probabilistic U-Net [12], PhiSeg-Net [4], and CIDM [17] etc.

Previous work in ambiguous segmentation has explored variational autoencoders (VAEs) and Bayesian methods to generate diverse segmentation masks. Kohl *et al.* [12] introduced the Probabilistic U-Net, combining a U-Net with a conditional VAE for multiple hypothesis generation. While effective, it struggles

with mask diversity and requires significant computational resources. PhiSegNet [4] enhanced diversity using a hierarchical VAE with a U-Net backbone, but its axis-aligned Gaussian latent space [21] remains restrictive, limiting expressiveness in modeling uncertainties. Rahman *et al.* [17] and Zbinden *et al.* [26] both leverage diffusion models based on Gaussian distributions. A core challenge in ambiguous segmentation is capturing structural lesion features and generating plausible boundaries. Current diffusion models sample from Gaussian noise, making it difficult to represent fine-grained lesion structures. Although using input images as guidance improves CT image representations, it often fails to preserve detailed lesion characteristics critical for precise segmentation. In this work we address some of these drawbacks through the following contributions:

- We empirically demonstrate the effectiveness of Schrödinger Bridge (SB) in ambiguous medical image segmentation by introducing Segmentation Schrödinger Bridge (SSB), which adapts the SB framework and outperforms existing benchmarks on LIDC-IDRI, Stanford COCA, and our in-house RACER datasets.
- To address ambiguity in segmentation, we propose a novel loss function designed to guide the model in handling diverse variations among segmentation masks. Additionally, we introduce a new metric, the *Diversity Divergence Index* (D_{DDI}), which quantifies inter-annotator variability and agreement, capturing both diversity and consensus among expert annotations.
- Our approach delivers significant improvements across key metrics (GED , D_{max} , CI Score, and D_{DDI}) compared to existing baselines. Moreover, it showcases robustness, offering fast sampling rates and reliable segmentation outputs suitable for real-world applications.

2 Datasets and Setup

Datasets: We use the publicly available LIDC-IDRI dataset [3], comprising of lung CT scans with lesion segmentations by four radiologists (1084 subjects, train:test split = 9:1). Additionally, we utilize the Stanford COCA dataset (787 subjects, train:test split = 3:1) with one expert annotations. We also test our model on an in-house dataset RACER with data from 44 subjects and annotations from 2 experts with demographic and ethnic variations distinct from COCA, enabling a robust evaluation of model generalization. The evaluation follows the protocol outlined by Rahman *et al.* [17].

Evaluation Metrics: The Generalized Energy Distance (GED) is a statistical metric widely used to evaluate generative models in ambiguous medical image segmentation [23], [12], [4], [17]. However, Rahman *et al.* [17] identified shortcomings in GED for capturing both segmentation performance and diversity, introducing complementary metrics such as Collective Intelligence Score (CI Score), Maximum Dice Matching (D_{max}), and Diversity Agreement (D_A). In this work, we employ GED alongside these metrics for a rigorous and holistic evaluation while highlighting the limitations of D_A . To address this, we propose the D_{DDI} , which provides a more robust quantification of diversity.

3 Methods

3.1 Score-based Generative Models and Schrödinger Bridge

Before introducing Score-based Generative Models (SGMs) and the Schrödinger Bridge (SB), we establish standard notations. Let $X_t \in \mathbb{R}^d$ be a stochastic process over continuous time $t \sim \mathcal{U}[0, 1]$, with initial (*data*) distribution p_A and terminal (*mask*) distribution p_B . The Wiener process and its reverse, based on Anderson et al. [2], are W_t and \bar{W}_t in \mathbb{R}^d , while $\mathbb{I} \in \mathbb{R}^{d \times d}$ is the identity matrix. The reference path measure \mathbb{P} corresponds to the forward SDE dynamics, serving as a baseline for optimization, while the target path measure \mathbb{Q} represents the optimal transport process constrained by p_A (at $t = 0$) and p_B (at $t = 1$). The functions $\Psi(t, x)$ and $\bar{\Psi}(t, x)$ solve the forward and backward PDEs, respectively, defining the drift components that govern SB dynamics.

Score-based Generative Models (SGMs) aim to perturb data across continuous noise scales using stochastic differential equations (SDEs) and reconstruct the data distribution via reverse-time SDEs. The reverse process, guided by the score function, enables transformations between arbitrary distributions and Gaussians [22]. Given data $X_0 \sim p_A$, the forward SDE perturbs it, while the reverse SDE reconstructs it using learned scores. Here, the drift and diffusion coefficients govern the noise and transformation. SGMs employ a U-Net [19] to parameterize the score function, optimized via a denoising score-matching loss [25], ensuring sampled distributions closely resemble p_A .

$$\begin{aligned}
 \text{Forward SDE: } dX_t &= f_t(X_t) dt + \sqrt{\beta_t} dW_t, \quad X_0 \sim p_A, \\
 \text{Reverse SDE: } dX_t &= [f_t - \beta_t \nabla \log p(X_t, t)] dt + \sqrt{\beta_t} d\bar{W}_t, \\
 \text{Training Loss: } &\mathbb{E}[\lambda(t) \|\epsilon_\theta(X_t, t) - \sigma_t \nabla \log p(X_t, t | X_0)\|^2], \\
 \text{Sampling SDE: } dX_t &= [f - \beta_t \epsilon_\theta(t, X_t)] dt + \sqrt{\beta_t} dW_t, \quad X_1 \sim p_B. \quad (1)
 \end{aligned}$$

Here, $\nabla \log p(X_t, t | X_0)$ is computed analytically, with σ_t^2 scaling the regression target and $\lambda(t)$ acting as a tunable hyperparameter influencing performance [10]. This allows SGMs to transform p_A into p_B through learned scores. However, SGMs are limited in flexibility, as they primarily transition between image distributions and Gaussian noise, restricting efficient mapping to arbitrary target distributions. To overcome this, we explore the Schrödinger Bridge approach, which offers greater adaptability with minimal data manipulation.

SGMs and Schrödinger Bridge (SB) The SB framework minimizes the KL divergence between a reference path measure \mathbb{P} and a target path measure \mathbb{Q} , subject to prescribed marginal densities p_A (at $t = 0$) and p_B (at $t = 1$) [16]. The optimization is formulated as $\min_{\mathbb{Q} \in P(p_A, p_B)} D_{KL}(\mathbb{Q} || \mathbb{P})$, where \mathbb{P} is typically chosen as the path measure of the forward SDE. The SB framework, as described in [7], [6], optimally transports p_A to p_B by solving PDEs for forward and backward drifts. Unlike SGMs, SB introduces non-linear components,

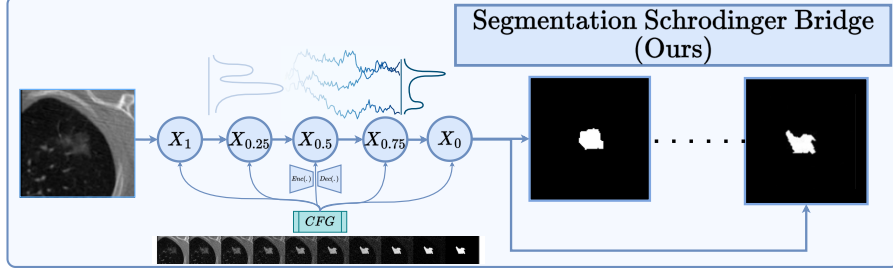


Fig. 1. Overview of the Segmentation Schrödinger Bridge (SSB) Framework: The model formulates segmentation as a stochastic transport problem and progressively maps the input image distribution to the segmentation mask distribution while preserving structural integrity. Classifier-Free Guidance (CFG) improves control over diversity and enables the generation of expert-aligned yet varied segmentations. Time-dependent U-Net parameter optimization and refined residual connections enhance performance.

making it more adaptable to prior data variations. While SB’s stochastic processes resemble SGMs, they incorporate additional coupling terms, resulting in a factorized marginal density. The governing equations are,

$$\begin{aligned} \text{SB Forward SDE: } dX_t &= [f_t + \beta_t \nabla \log \Psi(X_t, t)] dt + \sqrt{\beta_t} dW_t, \\ \text{SB Backward SDE: } dX_t &= [f_t - \beta_t \nabla \log \hat{\Psi}(X_t, t)] dt + \sqrt{\beta_t} d\bar{W}_t, \end{aligned} \quad (2)$$

As demonstrated in [16], [7], [6] this framework extends SGMs, with forward-backward SDEs that resemble those in SGMs but include non-linear drifts ($\nabla \log \Psi$ and $\nabla \log \hat{\Psi}$) (equation 2) to handle diverse prior data effectively [2], [22].

Optimal Boundary Constraints To address the computational challenges of nonlinear SBs, linear SDEs are formulated by redesigning the forward and backward drifts to solve the Fokker-Planck equation [18], [15]. These linear SDEs simplify SB coupling constraints, improving computational efficiency and tractability. While SBs nonlinear drifts correspond to score functions, sampling requires parameterizing $\nabla \log \hat{\Psi}$ via a score network. Dirac delta distributions are applied to further streamline boundary conditions, ensuring efficient convergence to a target state ν . The resulting equations are formulated as follows,

$$\begin{aligned} \text{Linear Forward SDE: } dX_t &= f_t(X_t) dt + \sqrt{\beta_t} dW_t, \quad X_0 \sim \Psi(\cdot, 0), \\ \text{Linear Backward SDE: } dX_t &= f_t(X_t) dt + \sqrt{\beta_t} d\bar{W}_t, \quad X_1 \sim \hat{\Psi}(\cdot, 1), \\ \text{Boundary Condition: } p_A(\cdot) &= \delta_\nu(\cdot), \quad p_B = \Psi(\cdot, 1) \hat{\Psi}(\cdot, 1). \end{aligned} \quad (3)$$

These linear SDEs ensure that the reverse process converges to the Dirac delta $\delta_\nu(\cdot)$, allowing efficient computation of the score function $\nabla \log p(X_t, t | X_0 = \nu)$ for each sample ν . This reformulation balances mathematical rigor with computational efficiency, making SB scalable and practical for real-world applications.

3.2 Segmentation Schrödinger Bridge (SSB)

Theoretical insights necessitate learning $\nabla \log \hat{\Psi}(\cdot)$ in a tractable manner. We apply this to ambiguous medical image segmentation by training a diffusion model that directly maps a given CT image to the lesion. To fully describe the diffusion process, we follow a structured design strategy [5], detailing the training, generation (sampling), and objective function.

Training and Generation: Preserving region-of-interest (ROI) boundaries is crucial in segmentation. Traditional forward diffusion degrades structural image properties, reducing mask quality. Prior methods like CIMD [17] and CCMD [26] generate masks by diffusing from Gaussian noise with image feature guidance, but often fail to retain structural lesion details. Inspired by Palette [20], we introduce a controlled degradation technique that perturbs the image partially, preserving structural integrity. Hence, at each time when controlled noise is injected, the model traverses different random paths and yields unique variations in the mask, even though they converge to the same marginals. Further information about generation and training is provided in Algorithm (1), (2). Training scalable diffusion models require efficient computation of X_t , which is intractable using Equation (2) due to nonlinear forward drift and the inability of linear SDEs (3) to model high-probability regions effectively. To overcome this, we adopt an analytical posterior formulation from Liu et al. [15]:

$$q(X_t | X_0, X_1) = \mathcal{N}(X_t; \mu_t(X_0, X_1), \Sigma_t) \quad (4)$$

Where, $\mu_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_1^2} X_0 + \frac{\sigma_1^2}{\sigma_t^2 + \sigma_1^2} X_1$, $\Sigma_t = \frac{\sigma_t^2 \sigma_1^2}{\sigma_t^2 + \sigma_1^2} \cdot \mathbb{I}$. This allows direct sampling of X_t during training from $p_A(X_0)$ and $p_B(X_1 | X_0)$ without solving nonlinear diffusion as in prior Schrödinger Bridge (SB) models [24]. During generation, given only $X_1 \sim p_B$, running a standard diffusion model [22], [10] from X_1 recovers the marginal density of SB paths, provided $X_{\epsilon,0}$ approximates X_0 well. Thus, the sampling method in Equation (4) is both tractable and effective in covering the required regions.

Algorithm 1 Generation of SSB Model

```

1: Input: Trained network  $\epsilon_\theta$ , noise schedule  $\{\beta_t\}$ 
2: Output: Generated samples  $X_0$ 
3: procedure GENERATION( $\epsilon_\theta, \{\beta_t\}$ )
4:    $X_N \sim p_B(X_N)$ 
5:   for  $t = N, \dots, 1$  do
6:     Predict Mask:  $\hat{X}_0$  via  $\epsilon_\theta(X_t, l, t)$ 
7:     Sample Masks [10]:
8:      $X_{t-1} \sim p(X_{t-1} | \hat{X}_0, X_t)$ 
9:   end for
10:  Return  $X_0$ 
11: end procedure
    
```

Algorithm 2 Training of SSB Model

```

1: Input: Training data  $\mathcal{D}$ ; U-Net  $\epsilon_\theta$ , Experts  $\eta$ 
2: Output: Trained network  $\epsilon_\theta$ 
3: procedure TRAINING( $\mathcal{D}, \epsilon_\theta$ )
4:   while not converged do
5:      $t, l \sim \mathcal{U}[0, 1], \mathcal{U}[1, \eta]$ 
6:      $X_0 \sim p_A(X_0), X_1 \sim p_B(X_1 | X_0)$ 
7:      $X_t \sim q(X_t | X_0, X_1)$  (ref eq. (4))
8:     Optimize [15]:  $\mathcal{L}_{SSB}(\epsilon_\theta, l, t)$ 
9:     Update  $\epsilon_\theta$  with GD to min  $\mathcal{L}_{SSB}$ .
10:  end while
11:  Return Trained  $\epsilon_\theta$ 
12: end procedure
    
```

Loss Function: Our loss function differs from the standard ADM approach [8] by leveraging Classifier-Free Guidance (CFG) [11]. As outlined in Algorithm

2, we first determine the number of experts and stochastically sample expert labels during training. The model then estimates noise in both conditional and unconditional settings: i) *Conditional*: $\epsilon_\theta(x_t, y, t)$ using data-label pairs (x_0, y) ; ii) *Unconditional*: $\epsilon_\theta(x_t, t)$ without label. During sampling, noise is estimated via linear interpolation $\hat{\epsilon} = (1 + \omega)\epsilon_\theta(x_t, y, t) - \omega\epsilon_\theta(x_t, t)$, enabling CFG-based reconstruction of x_{t-1} . This integration ensures diverse yet structurally consistent segmentations with fine-grained delineation. A key distinction lies in computing the score function in the loss, formulated for the reverse linear drift in Equation (3) while incorporating CFG. Given the segmentation focus, we further enhance precision by integrating a Dice loss, emphasizing fine boundary details to improve segmentation accuracy. The complete loss formulation is provided below,

$$\mathcal{L}_{\text{SSB}} = \mathbb{E} \left[\left\| \epsilon_\theta(X_t, l, t) - \frac{(X_t - X_0)}{\sigma_t} \right\|^2 - 2\gamma \cdot \frac{\|X_0 \cap \epsilon_\theta(X_t, l, t)\|}{\|X_0\| + \|\epsilon_\theta(X_t, l, t)\|} + \gamma \right]. \quad (5)$$

3.3 Diversity Divergence Index

In ambiguous medical image segmentation, a critical objective is to generate diverse expert-like predictions while capturing inter-annotator variability and agreement. To quantify both diversity and consensus among expert annotations, we propose the *Diversity Divergence Index* (D_{DDI}). The metric is formulated as follows. Given a model that generates N segmentation masks for an image and M expert annotations, we compute the **Dice score** for each generated mask against each expert, forming an $M \times N$ matrix. Each row, denoted as *expert distribution* $\text{exp-dist}_i \in \mathbb{R}^M$, represents the Dice scores between all generated masks and the i^{th} expert. Similarly, each column, referred to as the *generated distribution* ($\text{gen-dist}_j \in \mathbb{R}^M$), captures the Dice scores between a single generated mask and all expert annotations, reflecting how closely the model's predictions align with expert variability.

Next, we compute the Jensen-Shannon (JS) divergence [14] across expert distributions, measuring the divergence between expert predictions. Finally, D_{DDI}^{exp} aggregates these pairwise divergences, providing a comprehensive measure of inter-expert variability and D_{DDI}^{gen} aggregates these pairwise divergences, providing a comprehensive measure of inter-mask variability :

$$D_{DDI}^{\text{exp}} = \frac{3}{M C_2} \sum_i \sum_j \mathbb{1}_{[i \neq j]} JS(\text{exp-dist}_i, \text{exp-dist}_j) \quad (6)$$

$$D_{DDI}^{\text{gen}} = \frac{3}{N C_2} \sum_i \sum_j \mathbb{1}_{[i \neq j]} JS(\text{gen-dist}_i, \text{gen-dist}_j) \quad (7)$$

Here, $\mathbb{1}_{[i \neq j]}$ is the indicator function, which excludes the same expert comparison. This formulation captures both *diversity among expert masks* and *agreement in generated masks*, making it a robust metric for evaluating model uncertainty in ambiguous medical image segmentation.

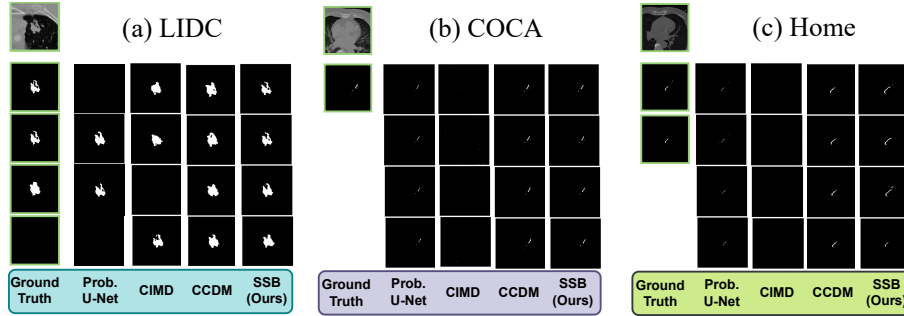


Fig. 2. The first image Fig. 2(a) from LIDC-IDRI [3] shows that Probabilistic U-Net fails to generate accurate masks, while CIMD aligns with expert annotations but lacks precise lesion boundaries. CCDM performs suboptimal, whereas SSB achieves fine-grained boundary delineation while maintaining segmentation diversity. In Fig. 2(b) from Stanford COCA [9], CIMD fails to capture calcium plaques, and while Prob. U-Net and CCDM miss very small plaques, and SSB successfully segments all deposits. For the third image (c) from RACER (in-house), trained only on COCA, SSB generalizes well, producing diverse yet expert-aligned annotations, unlike other methods that fail to overlap with experts or provide meaningful variability. In Fig. 2(c) the final output closely matches the second expert, while the remaining three exhibit controlled diversity

4 Results and Discussion

Table 1 presents a comparative analysis of our proposed SSB variants against state-of-the-art baseline methods across the LIDC-IDRI, COCA, and Home datasets. SSB++ extends the capabilities of SSB by incorporating Classifier-Free Guidance (CFG), enhancing control over the generative process to achieve an optimal balance between fidelity and diversity in segmentation. While SSB adopts the standard parametric configuration of DDPM [10], SSB++ introduces task-specific optimizations to time-dependent U-Net parameters, including the refinement of up-sampling and down-sampling layers and residual connections to enhance feature extraction and reconstruction. These architectural advancements make SSB++ a more robust and efficient solution for ambiguous medical image segmentation, consistently delivering superior performance.¹ Notably, SSB++ consistently surpasses prior approaches across all key metrics, achieving substantial improvements over the strongest baselines. The core limitation of existing diffusion methods such as CIMD [17] and CCDM [26] lies in their generative process itself, which starts from Gaussian noise and progressively refines segmentations. This often leads to a loss of structural integrity, resulting in coarse or anatomically inconsistent predictions. In contrast, our approach begins from the input image itself and progressively transforms it into the segmentation mask, ensuring that the anatomical structure is well-preserved throughout

¹ Complete training details of SSB++, including optimized parametric configurations, will be made available in the code repository.

Table 1. The results were obtained from the test set of the LIDC-IDRI, COCA, and RACER datasets. Results for PhiSeg [4] was taken from their original research (denoted \dagger), while CIMD [17] and CCDM [26] were fully reproduced. Metrics with (\downarrow) favor lower values, while (\uparrow) indicate higher is better. Also, D_{DDI} entries for the COCA dataset are empty as there are single expert masks. The reported results on the test set were statistically significant ($p < 0.05$).

LIDC-IDRI Dataset							
Method	NFE	GED \downarrow	D_{max} \uparrow	CI \uparrow	D_A \uparrow	D_{DDI}^{exp} \uparrow	D_{DDI}^{gen} \uparrow
Prob. UNet [12]	-	0.310	0.907	0.831	97.9	0.391	0.287
PhiSeg † [4]	-	0.270	0.904	0.736	-	-	-
CIMD [17]	1000	0.318	0.841	0.883	0.996	0.341	0.214
CCDM [26]	1000	0.264	0.886	0.918	0.997	0.402	0.324
SSB (Ours)	50	0.245	0.925	0.945	0.995	0.397	0.317
SSB++ (Ours)	50	0.208	0.943	0.958	0.997	0.427	0.356

Stanford COCA Dataset							
Method	NFE	GED \downarrow	D_{max} \uparrow	CI \uparrow	D_A \uparrow	D_{DDI}^{exp} \uparrow	D_{DDI}^{gen} \uparrow
Prob. UNet [12]	-	0.658	0.722	0.789	0.999	-	-
CIMD [17]	1000	0.625	0.501	0.599	0.999	-	-
CCDM [26]	1000	0.5	-	-	-	-	-
SSB (Ours)	50	0.421	0.891	0.908	0.999	-	-
SSB++ (Ours)	50	0.372	0.941	0.955	0.999	-	-

RACER (Home) Dataset							
Method	NFE	GED \downarrow	D_{max} \uparrow	CI \uparrow	D_A \uparrow	D_{DDI}^{exp} \uparrow	D_{DDI}^{gen} \uparrow
Prob. UNet [12]	-	0.792	0.610	0.697	0.999	≈ 0	≈ 0
CIMD [17]	1000	0.758	0.503	0.601	0.999	≈ 0	≈ 0
CCDM [26]	1000	0.551	0.630	0.715	0.999	0.009	0.002
SSB (Ours)	50	0.528	0.635	0.712	0.999	0.008	0.002
SSB++ (Ours)	50	0.635	0.675	0.753	0.999	0.013	0.007

the process. Among prior works, CCDM performs closely yet it still struggles to maintain structural fidelity compared to ours. Across all datasets, SSB++ achieves significant performance improvements over prior methods. For instance, on the LIDC-IDRI dataset, SSB++ reduces GED by 21.2% compared to CCDM, while simultaneously improving D_{max} by 5.7% and CI by 4.4%, indicating that our method generates segmentations that are both structurally accurate and diverse. Similar trends are observed in COCA, where SSB++ achieves a 25.6% lower GED than CIMD, and on the in-house RACER dataset, where it further improves GED by 19.3% and CI by 5.3% over CCDM. Qualitative results are presented in Fig. 2.

A key contribution of this work is the introduction of D_{DDI} , a robust metric designed to effectively quantify inter-rater variability and consensus in ambiguous segmentation tasks. Unlike the existing D_A metric, which remains consistently close to one and fails to meaningfully differentiate diversity across methods, D_{DDI} provides a bounded and interpretable measure of segmentation variability, making it a more reliable indicator of diversity in medical image segmentation. To comprehensively assess diversity agreement, we evaluate both D_{DDI}^{exp} , which captures alignment with expert annotations, and D_{DDI}^{gen} , which reflects the diversity in generated segmentations. Our proposed SSB++ framework demonstrates a substantial improvement over previous methods, achieving the highest D_{DDI}^{exp} across multiple datasets. On the LIDC-IDRI dataset, SSB++ improves expert agreement diversity by approximately 6.2% compared to CCDM, while on the in-house RACER dataset, it increases by over 44.4%. Similarly, in terms of generative diversity, SSB++ surpasses CCDM by 9.8% on LIDC-IDRI and exhibits a remarkable threefold improvement on the RACER dataset, signifying its ability to generate diverse yet clinically meaningful segmentation masks. Beyond its advancements in diversity modeling, SSB++ also delivers significant computational efficiency. While previous methods such as CIMD and CCDM require 1000 no. of function evaluations (NFEs), SSB++ achieves superior segmentation performance with only 50 NFEs, reducing computational overhead by

an order of magnitude. This efficiency makes SSB++ not only a more accurate solution but also a more viable choice for real-world clinical deployment where computational resources are often limited.

5 Conclusion

In this work, we present Segmentation Schrödinger Bridge (SSB), a novel framework that eliminates the need of traditional diffusion training or sampling from pure Gaussian noise and thereby preserves the structural integrity of lesion features. To enhance robustness, we introduce a loss function tailored for segmentation (\mathcal{L}_{SSB}), alongside a new metric (D_{DDI}) that effectively quantifies inter-rater variability and agreement. Our approach achieves state-of-the-art performance, surpassing competing methods in both ambiguous segmentation accuracy as well as sampling efficiency across several benchmark datasets on medical image segmentation tasks.

Acknowledgments

TC is supported by the Turing-Roche Strategic Partnership and a NIHR Principal Research Fellowship from the UCL Biomedical Research Centre. We acknowledge IHub-Data, IIIT Hyderabad (H1-002), for financial assistance.

Disclosure of Interest

The authors declare that they have no financial or non-financial competing interests that could have influenced the work reported in this paper.

References

- [1] Hillel R Alpert and Bruce J Hillman. Quality and variability in diagnostic radiology. *Journal of the American College of Radiology*, 1(2):127–132, 2004.
- [2] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [3] Samuel G Armato III et al. Lung image database consortium: developing a resource for the medical imaging research community. *Radiology*, 232(3):739–748, 2004.
- [4] Christian F Baumgartner, Kerem C Tezcan, Krishna Chaitanya, Andreas M Hötker, Urs J Muehlemaier, Khoshy Schawkat, Anton S Becker, Olivio Donati, and Ender Konukoglu. Phiseg: Capturing uncertainty in medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI*, pages 119–127. Springer, 2019.

- [5] Ziyi Chang, George A Koulouris, and Hubert PH Shum. On the design fundamentals of diffusion models: A survey. *arXiv preprint arXiv:2306.04542*, 2023.
- [6] Tianrong Chen, Guan-Horng Liu, and Evangelos A Theodorou. Likelihood training of schrödinger bridge using forward-backward sdes theory. *arXiv preprint arXiv:2110.11291*, 2021.
- [7] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *Siam Review*, 63(2):249–313, 2021.
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [9] David Eng et al. Automated coronary calcium scoring using deep learning with multicenter external validation. *NPJ digital medicine*, 4(1):88, 2021.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [12] Simon Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R Ledsam, Klaus Maier-Hein, SM Eslami, Danilo Jimenez Rezende, and Olaf Ronneberger. A probabilistic u-net for segmentation of ambiguous images. *Advances in neural information processing systems*, 31, 2018.
- [13] Ralf HJM Kurvers, Stefan M Herzog, Ralph Hertwig, Jens Krause, Patricia A Carney, Andy Bogart, Giuseppe Argenziano, Iris Zalaudek, and Max Wolf. Boosting medical diagnostics by pooling independent judgments. *Proceedings of the National Academy of Sciences*, 113(31):8777–8782, 2016.
- [14] Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.
- [15] Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I2sb: Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872*, 2023.
- [16] Michele Pavon and Anton Wakolbinger. On free energy, stochastic control, and schrödinger processes. In *Modeling, Estimation and Control of Systems with Uncertainty: Proceedings of a Conference held in Sopron, Hungary, September 1990*, pages 334–348. Springer, 1991.
- [17] Aimon Rahman, Jeya Maria Jose Valanarasu, Ilker Hacihaliloglu, and Vishal M Patel. Ambiguous medical image segmentation using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11536–11546, 2023.
- [18] Hannes Risken and Hannes Risken. *Fokker-planck equation*. Springer, 1996.
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015.
- [20] Chitwan Saharia et al. Palette: Image-to-image diffusion models. pages 1–10, 2022.

- [21] Raghavendra Selvan, Frederik Faye, Jon Middleton, and Akshay Pai. Uncertainty quantification in medical image segmentation with normalizing flows. In *Machine Learning in Medical Imaging*, pages 80–90. Springer, 2020.
- [22] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [23] Gábor J Székely and Maria L Rizzo. Energy statistics: A class of statistics based on distances. *Journal of statistical planning and inference*, 143(8):1249–1272, 2013.
- [24] Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- [25] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [26] Lukas Zbinden, Lars Doorenbos, Theodoros Pissas, Adrian Thomas Huber, Raphael Sznitman, and Pablo Márquez-Neila. Stochastic segmentation with conditional categorical diffusion models. In *ICCV*, pages 1119–1129, 2023.