

TopoCellGen: Generating Histopathology Cell Topology with a Diffusion Model

Meilong Xu¹ Saumya Gupta¹ Xiaoling Hu² Chen Li¹ Shahira Abousamra³
Dimitris Samaras¹ Prateek Prasanna¹ Chao Chen¹

¹Stony Brook University, ²Harvard Medical School, ³Stanford University

Abstract

Accurately modeling multi-class cell topology is crucial in digital pathology, as it provides critical insights into tissue structure and pathology. The synthetic generation of cell topology enables realistic simulations of complex tissue environments, enhances downstream tasks by augmenting training data, aligns more closely with pathologists' domain knowledge, and offers new opportunities for controlling and generalizing the tumor microenvironment. In this paper, we propose a novel approach that integrates topological constraints into a diffusion model to improve the generation of realistic, contextually accurate cell topologies. Our method refines the simulation of cell distributions and interactions, increasing the precision and interpretability of results in downstream tasks such as cell detection and classification. To assess the topological fidelity of generated layouts, we introduce a new metric, Topological Fréchet Distance (TopoFD), which overcomes the limitations of traditional metrics like FID in evaluating topological structure. Experimental results demonstrate the effectiveness of our approach in generating multi-class cell layouts that capture intricate topological relationships.

1. Introduction

Deep learning methods have significantly advanced nuclei analysis tasks, including segmentation [9, 12], classification, and detection [1], crucial for detailed tissue characterization and clinical applications [20]. However, accurately annotating multi-class cell arrangements remains challenging due to complex spatial patterns and domain expertise requirements. Existing annotated datasets often lack the diversity needed for generalization across tissues.

Generative models, particularly diffusion models [6, 14, 19], have emerged as powerful tools for synthesizing high-resolution histopathology images [27]. Despite their visual quality, these models typically provide limited biological interpretability and control, hindering validation against expert knowledge.

A critical gap is the explicit modeling of cell spatial ar-

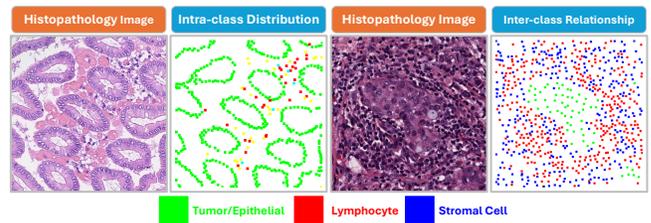


Figure 1. Illustrations of intra-class distribution and the inter-class relationship across various cell types. Here we only highlight the tumor/epithelial, lymphocytes, and stromal cells.

rangements, fundamental for understanding tissue microenvironments and disease progression [24]. Spatial configurations of diverse cell types, such as lymphocytes and epithelial cells, significantly inform pathology diagnostics and prognostics, e.g., tumor-infiltrating lymphocytes correlate with improved outcomes [23]. An illustrative overview of these intra-class arrangements and inter-class spatial interactions among tumor/epithelial, lymphocytic, and stromal populations is presented in Fig. 1.

To address this, we propose the first diffusion model explicitly designed to generate biologically meaningful cell spatial layouts, guided by persistent homology-based topological constraints. Our model accurately represents cell clusters (0-dimensional features) and spatial voids (1-dimensional features), ensuring realistic intra-class structures and inter-class interactions.

Additionally, we introduce a novel cell counting loss, addressing the unrealistic cell density bias in prior generative approaches [3]. We also propose the Topological Fréchet Distance (TopoFD), a metric specifically designed to evaluate the spatial-topological fidelity of synthetic layouts compared to real data.

Our contributions are summarized as follows:

- The first topology-guided diffusion model for generating realistic cell layouts in digital pathology.
- A novel cell counting loss to accurately model cell number distributions.
- The introduction of TopoFD, a metric for evaluating topological accuracy of generated layouts.

Experimental results validate our method, demonstrating

enhanced synthetic data quality and significant performance improvements in downstream tasks such as cell detection and classification.

2. Method

In this section, we introduce our method for synthesizing realistic multi-class cell layouts with precise topological constraints. Given n cell classes, we define a condition vector $c = [c_1, c_2, \dots, c_n]$, where each element c_i represents the target count of the respective cell type. However, conditioning solely on c is inadequate for ensuring accurate cell counts, intra-class spatial distributions, and inter-class topological relationships [11]. To address this, we propose a differentiable cell-counting loss and two topology-aware objectives based on persistent homology [7, 21]. These constraints jointly ensure that the generated layouts faithfully reflect both spatial and topological characteristics.

We first briefly review diffusion models in Sec. 2.1, then introduce our proposed topology-preserving layout generation framework, *TopoCellGen*, in Sec. 2.2. Finally, we describe our novel evaluation metric, the Topological Fréchet Distance (TopoFD), in Sec. 2.3.

Preliminaries. In a multi-class cell layout containing n channels, each channel corresponds to a specific cell type (e.g., lymphocyte), with each cell represented as a square where the pixel value is set to 1, while the pixel value of the background is set to 0.

2.1. Diffusion Models

Our generative approach utilizes a denoising diffusion probabilistic model (DDPM) [14], which learns to reverse a forward process that incrementally adds Gaussian noise to transform a structured cell layout into a noise distribution. The reverse process then reconstructs the layout from noise via iterative denoising.

Let x_0 represent the target cell layout and x_T denote pure Gaussian noise. At each time step t , noise is added to the data based on a variance schedule β_t : $q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I)$. This forward process results in the progressively noisier version of the data, with x_T approximating an isotropic Gaussian distribution as t increases.

The reverse process, parameterized by a neural network $\epsilon_\theta(x_t, t)$ (typically UNet [22]), learns to iteratively denoise x_T back to x_0 . Conditioning on the cell count vector c , the model is trained by minimizing a simplified variant of the variational lower bound, specifically focusing on predicting the noise added at each step:

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, c, t)\|^2] \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, I)$ is the noise sampled during training. This objective enables the model to learn the reverse process effectively. Instead of the standard iterative denoising,

we also approximate the noiseless layout \hat{x}_0^t deterministically for any noisy state x_t by marginalizing over the noise schedule:

$$\hat{x}_0^t \approx \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, c, t)) \quad (2)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, which aggregates the effect of the variance schedule up to time t . This predicted noiseless layout, \hat{x}_0^t , will be used to impose constraints in subsequent stages.

2.2. Spatially Aligned Cell Layout Generation

The primary objective of our method is to generate multi-class cell layouts that accurately simulate both the topological and spatial properties of real-world biological cell distributions. To achieve this, we ensure accurate cell counts for each cell type through a cell counting loss, while also preserving spatial relationships within individual cell types via enforcing intra-class spatial consistency. Furthermore, we maintain structural coherence across all cell types by applying an inter-class structural regularization, leveraging 1-dimensional persistent homology to encapsulate both type-specific and collective spatial properties. The overall pipeline is shown in Fig. 2.

Cell Counting Loss. Given the target layout x_0 , which serves as the ground truth, for each time step t , we obtain the predicted noiseless layout \hat{x}_0^t using Eq. (2). To ensure precise control over the number of cells in the generated layout, we introduce a differentiable cell counting loss. The key challenge lies in making the counting operation differentiable for gradient-based optimization. We address this by employing the Straight-Through Estimator (STE) [4], which enables gradient flow through the discrete binarization operation. Specifically, after obtaining \hat{x}_0^t , we apply a hard threshold to obtain binary values:

$$b(\hat{x}_0^t) = \text{float}((\hat{x}_0^t \geq \tau)) \quad (3)$$

where τ is the threshold parameter. Here we set it to the median value of \hat{x}_0^t . During back-propagation, the STE treats the thresholding operation as an identity function, allowing gradients to flow through. The cell counting loss is then formulated as:

$$\mathcal{L}_{\text{count}} = \frac{1}{|n|} \sum_{i=1}^n \left| \frac{\sum b(\hat{x}_0^t)^{(i)}}{\delta} - \frac{\sum x_0^{(i)}}{\delta} \right| \quad (4)$$

where $b(\hat{x}_0^t)^{(i)}$ represents the binarized prediction for the i -th channel, and δ indicates the area (3×3) of a single cell in the layouts. This formulation provides a differentiable approximation to the discrete cell counting operation, enabling end-to-end training while maintaining precise control over the number of cells for each cell type.

Intra-Class Spatial Consistency. To enforce spatial consistency within each cell type, we first calculate the distance

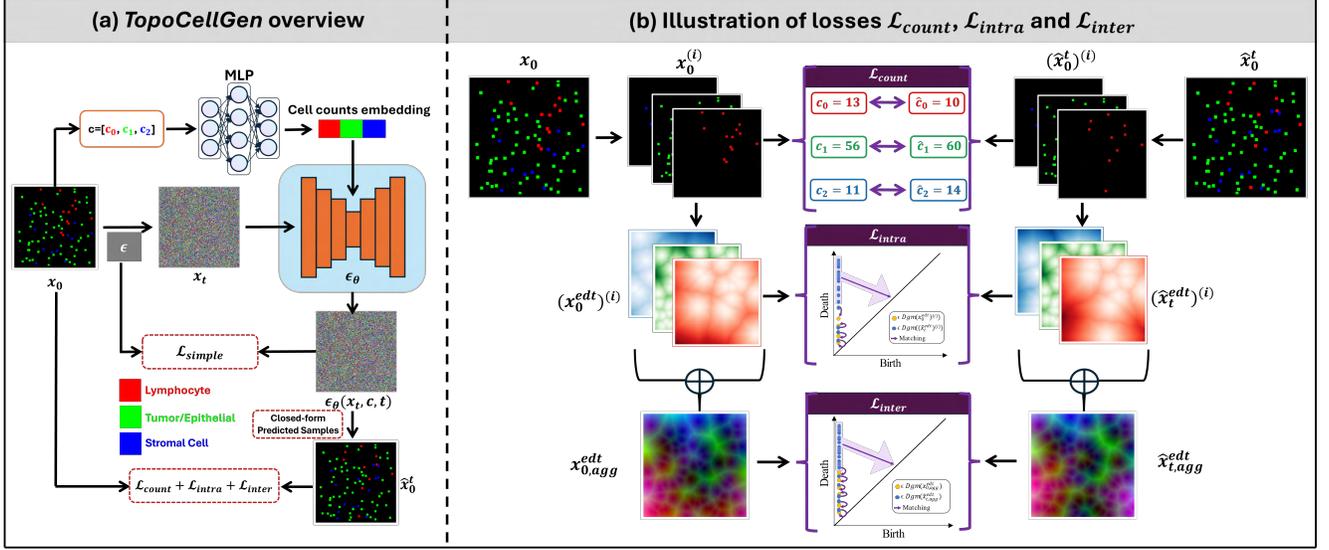


Figure 2. An overview of our method *TopoCellGen*. (a) denotes the overview workflow. (b) shows the details of \mathcal{L}_{count} , \mathcal{L}_{intra} and \mathcal{L}_{inter} .

transform map [8] for each channel in both the target layout x_0 and the predicted noiseless layout \hat{x}_0^t . The distance transform $D(x)$ is a function that assigns to each pixel the minimum Euclidean distance to the nearest cell (or non-zero pixel) in the channel. This can be formally written as:

$$D(x) = \min_{p \in \text{cells}} \|x - p\| \quad (5)$$

where p represents the positions of cells in the layout. After obtaining the distance transform maps of the target layout and the predicted noiseless layout, $\hat{x}_t^{edt} = D(b(\hat{x}_0^t))$ and $x_0^{edt} = D(x_0)$, we calculate the 1-dim persistence diagrams for both of them, $Dgm(\hat{x}_t^{edt})$ and $Dgm(x_0^{edt})$ respectively. Similar to previous topological losses [15], we will use the classic Wasserstein distance between the two diagrams. Given two diagrams $Dgm(q)$ and $Dgm(s)$, the p -th Wasserstein distance is defined as follows:

$$W_p(Dgm(q), Dgm(s)) = \left(\inf_{\gamma \in \Gamma} \sum_{x \in Dgm(q)} \|x - \gamma(x)\|^p \right)^{\frac{1}{p}}$$

where Γ represents all bijections from $Dgm(q)$ to $Dgm(s)$.

The Wasserstein distance operates by identifying an optimal correspondence between points in two diagrams, assigning unmatched points to their projections on the diagonal. This distance metric is calculated by summing the distances between all paired points. The process of finding this optimal matching, as well as calculating the Wasserstein distance, can be accomplished using either the traditional Hungarian algorithm or more sophisticated methods [16, 17].

Next, we denote γ^* , the optimal matching between $Dgm(\hat{x}_t^{edt})$ and $Dgm(x_0^{edt})$. Each persistence dot in $Dgm(\hat{x}_t^{edt})$ is matched either to a target dot in $Dgm(x_0^{edt})$ or its projection on the diagonal. We can now formulate the

spatial distribution consistency loss as the squared distance between every dot in $Dgm(\hat{x}_t^{edt})$ and its match:

$$\mathcal{L}_{\text{spc}} = \sum_{q \in Dgm(\hat{x}_t^{edt})} \|q - \gamma^*(q)\|^2 \quad (6)$$

For a multi-class cell layout containing n classes of cells, we formulate the intra-class spatial consistency loss as follows by averaging the Eq. (6) across multiple classes:

$$\mathcal{L}_{\text{intra}} = \frac{1}{|n|} \sum_{i=1}^n \mathcal{L}_{\text{spc}} \left(Dgm((\hat{x}_t^{edt})^{(i)}), Dgm((x_0^{edt})^{(i)}) \right) \quad (7)$$

Inter-Class Structural Regularization. Beyond maintaining spatial distribution consistency within individual cell types, it is equally important to capture the relationships between different cell types. To achieve this, we construct a unified layout by combining all cell types into a single-channel representation, referred to as the aggregated layout: $x_0^{agg} = \text{Agg}(x_0)$ and $\hat{x}_0^{t,agg} = \text{Agg}(\hat{x}_0^t)$. We then compute the distance transform for the aggregated layouts, with $\hat{x}_t^{edt,agg} = D(\hat{x}_0^{t,agg})$ representing the distance transform of the predicted layout and $x_0^{edt,agg} = D(x_0^{agg})$ for the target layout. The inter-class structural loss $\mathcal{L}_{\text{inter}}$ is computed similarly to the intra-class loss:

$$\mathcal{L}_{\text{inter}} = \mathcal{L}_{\text{spc}} \left(Dgm(\hat{x}_t^{edt,agg}), Dgm(x_0^{edt,agg}) \right) \quad (8)$$

Together, these class-specific and cross-class regularizations ensure that both individual cell distributions and their cumulative spatial interactions are enforced, preserving critical spatial dynamics within and between cell types in the generated layouts.

Final Objectives. The final training objective of the model is the weighted sum of the three losses with $\mathcal{L}_{\text{simple}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{simple}} + \lambda_c \mathcal{L}_{\text{count}} + \lambda_{\text{intra}} \mathcal{L}_{\text{intra}} + \lambda_{\text{inter}} \mathcal{L}_{\text{inter}} \quad (9)$$

	Method	FID ↓	Lym. ↓	Epi. ↓	Stro. ↓	Neu. ↓	Pla. ↓	Eos. ↓	Con. ↓	TCE ↓	TopoFD ↓	MMD ↓
BRCA-M2C	ADM [6]	1.150	13.757	40.230	15.491	–	–	–	–	22.465	133.012	0.732
	TMCCG [2]	0.634	11.503	34.032	12.907	–	–	–	–	19.687	89.252	0.635
	Spatial Diffusion [18]	0.263	10.852	35.954	13.496	–	–	–	–	20.806	97.584	0.589
	<i>TopoCellGen</i>	0.005	2.090	3.824	2.468	–	–	–	–	5.192	69.354	0.421
Lizard	ADM [6]	0.059	16.508	11.796	–	1.123	4.328	1.598	10.737	23.964	65.910	0.783
	TMCCG [2]	1.093	15.548	10.011	–	2.376	4.293	1.872	11.643	22.604	63.120	0.667
	Spatial Diffusion [18]	0.137	10.740	9.062	–	3.040	6.552	2.173	11.225	20.606	79.591	0.883
	<i>TopoCellGen</i>	0.027	6.155	6.560	–	1.022	2.982	1.167	7.288	11.590	31.607	0.536

Table 1. Results for BRCA-M2C and Lizard datasets on the quality of the generated samples.

Data	Method	F1-Score ↑				
		Lymphocytes	Epithelial	Stromal	Mean	Detection
Real.		0.569 ± 0.010	0.736 ± 0.012	0.507 ± 0.015	0.604 ± 0.011	0.857 ± 0.006
Real+Syn. (Rand)		0.549 ± 0.009	0.693 ± 0.014	0.472 ± 0.016	0.571 ± 0.013	0.848 ± 0.008
Real+Syn (TMCCG)	UNet	0.650 ± 0.007	0.768 ± 0.010	0.511 ± 0.012	0.643 ± 0.009	0.852 ± 0.005
Real+Syn (SpaDM)		0.647 ± 0.006	0.797 ± 0.003	0.554 ± 0.011	0.666 ± 0.007	0.853 ± 0.005
Real+Syn (<i>TopoCellGen</i>)		0.656 ± 0.003	0.803 ± 0.005	0.574 ± 0.004	0.678 ± 0.004	0.860 ± 0.004
Real.		0.615 ± 0.008	0.777 ± 0.010	0.540 ± 0.013	0.644 ± 0.009	0.855 ± 0.005
Real+Syn. (Rand)		0.578 ± 0.009	0.756 ± 0.012	0.502 ± 0.014	0.612 ± 0.010	0.851 ± 0.006
Real+Syn (TMCCG)	MCSpatNet	0.678 ± 0.006	0.800 ± 0.005	0.522 ± 0.014	0.667 ± 0.007	0.853 ± 0.004
Real+Syn (SpaDM)		0.639 ± 0.005	0.804 ± 0.007	0.563 ± 0.012	0.669 ± 0.006	0.855 ± 0.005
Real+Syn (<i>TopoCellGen</i>)		0.652 ± 0.004	0.817 ± 0.006	0.582 ± 0.005	0.684 ± 0.004	0.862 ± 0.004

Table 2. Results on cell detection and classification tasks on BRCA-M2C dataset. The best and statistically significant results are highlighted in **bold**.

where λ_c , λ_{intra} and λ_{inter} are hyper-parameters that control the relative contributions of the respective loss terms.

2.3. Topological Fréchet Distance (TopoFD)

Conventional Fréchet Inception Distance (FID) [13] assesses generation quality by comparing feature distributions of real and synthetic samples, typically computed from pre-trained networks such as InceptionV3 [25]. However, FID inadequately captures complex spatial and topological cell interactions. To address this, we propose the Topological Fréchet Distance (TopoFD), which quantifies higher-dimensional spatial and topological discrepancies via persistence diagrams derived from cell-center point clouds. Specifically, we compute the barycenters of persistence diagrams for real and synthetic layouts and evaluate the Fréchet distance between their respective persistence landscapes [5], averaged over all cell types.

3. Experiments

We conduct extensive experiments on two public and widely used nuclei analysis datasets. We compare our method against SoTA layout generation methods regarding sample quality and performance on downstream tasks.

Datasets. We evaluate our proposed method on *TCGA Breast Cancer Cell Classification Dataset (BRCA-M2C)* [1] and *Lizard dataset* [10].

Evaluation Metrics. We evaluate our proposed method on both sample quality and the performance of downstream tasks. We use **Fréchet Inception Distance (FID)** [13], the **cell count error** for each cell type, the **total count error (TCE)**, our proposed **TopoFD** and **maximum mean discrepancy (MMD)** [26] to evaluate how well the generated cell layouts align with the reference layouts. Note that for FID, feature extraction is tailored to each dataset

with custom-trained models. On the other hand, we generate 2,000 image-layout pairs as augmented training data for cell detection and classification tasks, evaluating their performance with the **F1-score**.

3.1. Experimental Results

Quantitative Results on Sample Quality. As reported in Table 1, our proposed *TopoCellGen* framework achieves the lowest Fréchet Inception Distance (FID) and Topological Fréchet Distance (TopoFD) relative to three strong baselines—ADM [6], TMCCG [2], and Spatial Diffusion [18]. These improvements highlight *TopoCellGen*’s capacity to generate visually compelling samples while maintaining a high degree of topological fidelity. In addition, *TopoCellGen* notably reduces multi-class cell count errors, thereby enhancing both the precision and topological consistency of the synthesized cellular arrangements.

Performance on Downstream Tasks. As shown in Table 2, synthetic layouts generated by *TopoCellGen* consistently yield the highest F1 scores for cell detection and classification tasks using UNet [22] and MCSpatNet [1] frameworks. These results underscore *TopoCellGen*’s capability to model complex spatial interactions and inter-class relationships, producing biologically plausible synthetic data that effectively enhances model generalization and reduces class-wise biases.

4. Conclusion

In summary, *TopoCellGen* presents a robust framework for generating realistic cell topologies in digital pathology. It accurately preserves both intra- and inter-class spatial patterns, ensures cell count control, and achieves high structural fidelity. Experimental results confirm its close approximation of real tissue layouts, thereby enhancing downstream tasks such as cell detection and classification.

References

- [1] Shahira Abousamra, David Belinsky, John Van Arnam, Felicia Allard, Eric Yee, Rajarsi Gupta, Tahsin Kurc, Dimitris Samaras, Joel Saltz, and Chao Chen. Multi-class cell detection using spatial context representation. In *ICCV*, 2021. 1, 4
- [2] Shahira Abousamra, Rajarsi Gupta, Tahsin Kurc, Dimitris Samaras, Joel Saltz, and Chao Chen. Topology-guided multi-class cell context generation for digital pathology. In *CVPR*, 2023. 4
- [3] Marco Aversa, Gabriel Nobis, Miriam Hägele, Kai Standvoss, Mihaela Chirica, Roderick Murray-Smith, Ahmed M Alaa, Lukas Ruff, Daniela Ivanova, Wojciech Samek, et al. Diffinfinite: Large mask-image synthesis via parallel random patch diffusion in histopathology. In *NeurIPS*, 2024. 1
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2
- [5] Peter Bubenik et al. Statistical topological data analysis using persistence landscapes. *JMLR*, 2015. 4
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 1, 4
- [7] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 2002. 2
- [8] Pedro F Felzenszwalb and Daniel P Huttenlocher. Distance transforms of sampled functions. *Theory of computing*, 2012. 3
- [9] Simon Graham, Quoc Dang Vu, Shan E Ahmed Raza, Ayesha Azam, Yee Wah Tsang, Jin Tae Kwak, and Nasir Rajpoot. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *MedIA*, 2019. 1
- [10] Simon Graham, Mostafa Jahanifar, Ayesha Azam, Mohammed Nimir, Yee-Wah Tsang, Katherine Dodd, Emily Hero, Harvir Sahota, Atisha Tank, Ksenija Benes, et al. Lizard: A large-scale dataset for colonic nuclear instance segmentation and classification. In *ICCV*, 2021. 4
- [11] Saumya Gupta, Dimitris Samaras, and Chao Chen. Topodiffusionnet: A topology-aware diffusion model. *arXiv preprint arXiv:2410.16646*, 2024. 2
- [12] Hongliang He, Jun Wang, Pengxu Wei, Fan Xu, Xiangyang Ji, Chang Liu, and Jie Chen. Toposeg: Topology-aware nuclear instance segmentation. In *ICCV*, 2023. 1
- [13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 4
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 2
- [15] Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In *NeurIPS*, 2019. 3
- [16] Michael Kerber, Dmitriy Morozov, and Arnur Nigmatov. Geometry helps to compare persistence diagrams. In *2016 Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, 2016. 3
- [17] Théo Lacombe, Marco Cuturi, and Steve Oudot. Large scale computation of means and clusters for persistence diagrams using optimal transport. In *NeurIPS*, 2018. 3
- [18] Chen Li, Xiaoling Hu, Shahira Abousamra, Meilong Xu, and Chao Chen. Spatial diffusion for cell layout generation. In *MICCAI*, 2024. 4
- [19] Ming Li, Taojiannan Yang, Huafeng Kuang, Jie Wu, Zhaoning Wang, Xuefeng Xiao, and Chen Chen. Controlnet++: Improving conditional controls with efficient consistency feedback. In *ECCV*, 2024. 1
- [20] Cheng Lu, Can Koyuncu, German Corredor, Prateek Prasanna, Patrick Leo, XiangXue Wang, Andrew Janowczyk, Kaustav Bera, James Lewis Jr, Vamsidhar Velcheti, et al. Feature-driven local cell graph (flock): new computational pathology-based descriptors for prognosis of lung cancer and hpv status of oropharyngeal cancers. *MedIA*, 2021. 1
- [21] James R Munkres. Elements of algebraic topology, 1984. 2
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 2, 4
- [23] Roberto Salgado, Carsten Denkert, S Demaria, N Sirtaine, F Klauschen, Giancarlo Pruneri, S Wienert, Gert Van den Eynden, Frederick L Baehner, Frederique Pénault-Llorca, et al. The evaluation of tumor-infiltrating lymphocytes (tils) in breast cancer: recommendations by an international tils working group 2014. *Annals of oncology*, 2015. 1
- [24] Denis Schapiro, Hartland W Jackson, Swetha Raghuraman, Jana R Fischer, Vito RT Zanotelli, Daniel Schulz, Charlotte Giesen, Raúl Catena, Zsuzsanna Varga, and Bernd Bodenmiller. histocat: analysis of cell phenotypes and interactions in multiplex image cytometry data. *Nature methods*, 2017. 1
- [25] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 4
- [26] Fan Wang, Huidong Liu, Dimitris Samaras, and Chao Chen. Topogan: A topology-aware generative adversarial network. In *ECCV*, 2020. 4
- [27] Srikar Yellapragada, Alexandros Graikos, Prateek Prasanna, Tahsin Kurc, Joel Saltz, and Dimitris Samaras. Pathldm: Text conditioned latent diffusion model for histopathology. In *WACV*, 2024. 1