DiffInfinite: Large Mask-Image Synthesis via Parallel Random Patch Diffusion in Histopathology [1]

Marco Aversa, Gabriel Nobis, Miriam Hägele, Kai Standvoss,
Mihaela Chirica, Roderick Murray-Smith, Ahmed Alaa, Lukas Ruff, Daniela Ivanova,
Wojciech Samek, Frederick Klauschen, Bruno Sanguinetti, Luis Oala

Talk by Gabriel Nobis

Challenges in AI for Histopathology

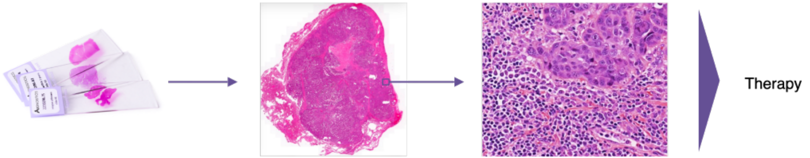

(C1) Whole slide image of size $\sim 100k \times 100k$

(C2) Sparsely annotated and small datasets [2]

- $\rightsquigarrow$ Use generative AI for data augmentation
- $\rightsquigarrow$ Generate large images to capture long range correlation

DiffInfinite: A Diffusion Model for Histopathology

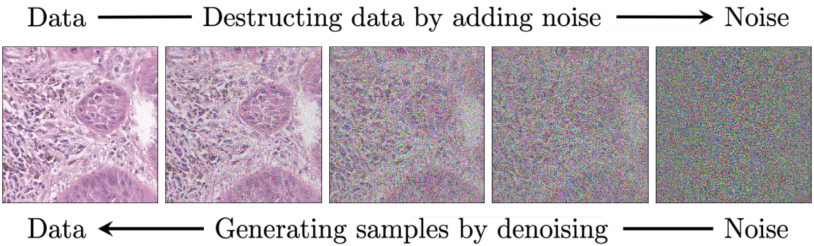

Diffusion Models are:

- ✔ capable of generating high quality diverse data [3], [4]
- ✔ theoretical well defined [5]
- ✗ slow at sampling time [6]

DiffInfinite Training

Semi-supervised training on tiles:

Use a mixture of $5k$ labelled tiles of 4 classes and $250k$ unlabeled tiles leveraging classifier-free guidance [7]

Supervised training on masks:

Train a second diffusion model on segmentation masks

Generate image with mask guidance:

Use the synthetic mask to guide the image generation

DiffInfinite Sampling

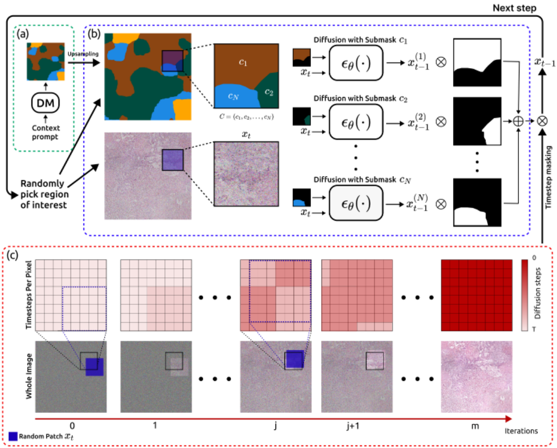

Comparison to DiffCollage [8]

Quantitative evaluation:

	Improved Precision ↑	Improved Recall ↑
DiffCollage	0.94	0.22
DiffInfinite	**0.98**	**0.44**

No tiling effects:

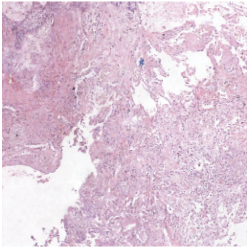

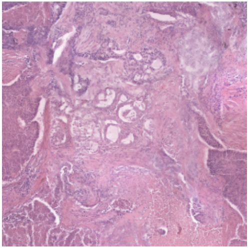

DiffInfinite ⟷ DiffCollage

DiffInfinite Samples

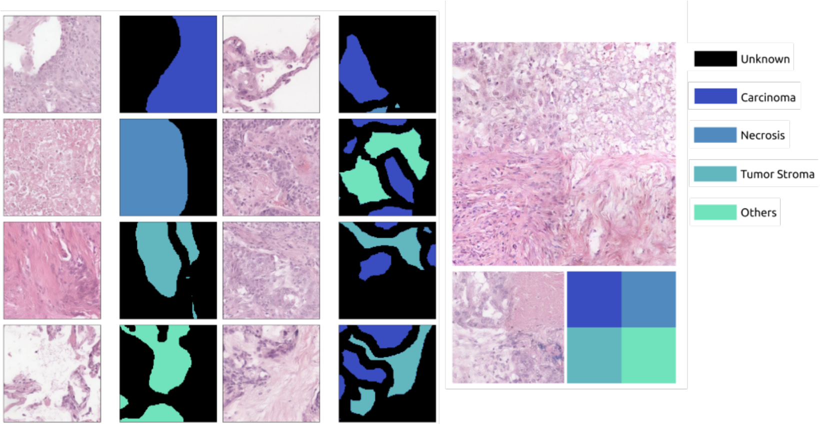

DiffInfinite for Data Augmentation

Classification:

	IH1	IH2	IH3	PCam-327K
Drift components	-	Patient change Different center	Patient change Different center	Patient change Different center Indication change Lower resolution
Trained Real	0.846 ± 0.005	0.733 ± 0.021	0.598 ± 0.049	0.60
Trained Synthetic	0.747 ± 0.025	**0.753** ± 0.005	**0.699** ± 0.002	**0.68**
Trained Augmented	**0.852** ± 0.007	0.732 ± 0.027	0.637 ± 0.025	0.63

Segmentation:

	IH2
Trained Real	0.614 ± 0.009
Trained Synthetic	0.471 ± 0.039
Trained Augmented	**0.710** ± 0.021

Fraunhofer HHI

Challenges in AI for Histopathology

(C1) Whole slide image of size $\sim 100k \times 100k$ ...

(C2) Sparsely annotated small datasets ✔

What about privacy?

Visual evaluation

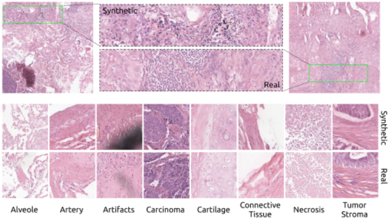

Quantitative evaluation [9],[10]

	$A \uparrow$		$C_T \rightarrow$	
	tiled	*resized*	*tiled*	*resized*
DiffCollage	**0.89**	0.97	11.02	**7.00**
DiffInfinite	0.86	**0.98**	**9.61**	11.56

Good enough [11]?

Acknowledgements

We would like to acknowledge our team of pathologists who provided valuable feedback in and outside of the conducted survey - special thank you to Frank Dubois, Niklas Prenissl, Cleopatra Schreiber, Vitaly Garg, Alexander Arnold, Sonia Villegas, Rosemarie Krupar and Simon Schallenberg. Furthermore, we would like to thank Marvin Sextro for his support in the analyses. This work was supported by the Federal Ministry of Education and Research (BMBF) as grants [SyReal (01IS21069B)]. RM-S is grateful for EPSRC support through grants EP/T00097X/1 and EP/R018634/1 and DI for EP/R513222/1. MA is funded by dotPhoton and a UofG Ph.D. scholarship.

References I

[1] Marco Aversa et al. *DiffInfinite: Large Mask-Image Synthesis via Parallel Random Patch Diffusion in Histopathology*. 2023. arXiv: 2306.13384 [eess.IV].

[2] Maximilian Springenberg et al. "From modern CNNs to vision transformers: Assessing the performance, robustness, and classification strategies of deep learning models in histopathology". In: *Medical Image Analysis* 87 (2023), p. 102809. ISSN: 1361-8415.

References II

[3] Jascha Sohl-Dickstein et al. "Deep Unsupervised Learning using Nonequilibrium Thermodynamics". In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by Francis Bach and David Blei. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, July 2015, pp. 2256–2265.

[4] Jonathan Ho, Ajay Jain, and Pieter Abbeel. "Denoising Diffusion Probabilistic Models". In: *arXiv preprint arxiv:2006.11239* (2020).

[5] Yang Song et al. "Score-Based Generative Modeling through Stochastic Differential Equations". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=PxTIG12RRHS.

References III

[6] Jiaming Song, Chenlin Meng, and Stefano Ermon. "Denoising Diffusion Implicit Models". In: *arXiv:2010.02502* (Oct. 2020). URL: https://arxiv.org/abs/2010.02502.

[7] Jonathan Ho and Tim Salimans. "Classifier-free diffusion guidance". In: *arXiv preprint arXiv:2207.12598* (2022).

[8] Qinsheng Zhang et al. "DiffCollage: Parallel Generation of Large Content with Diffusion Models". In: *CVPR*. 2023.

[9] Casey Meehan, Kamalika Chaudhuri, and Sanjoy Dasgupta. "A Non-Parametric Test to Detect Data-Copying in Generative Models". In: *CoRR* abs/2004.05675 (2020). arXiv: 2004.05675. URL: https://arxiv.org/abs/2004.05675.

References IV

[10] Ahmed Alaa et al. "How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models". In: *International Conference on Machine Learning*. PMLR. 2022, pp. 290–306.

[11] Nicholas Carlini et al. *Extracting Training Data from Diffusion Models*. 2023. arXiv: 2301.13188 [cs.CR].

Thank you!