STAIN: Smooth Tile-Aware Instance Normalisation for Virtual Staining

David Armstrong*, Iain Styles, Niall McLaughlin
*darmstrong34@qub.ac.uk

BMVC 2025

What is Virtual Staining?

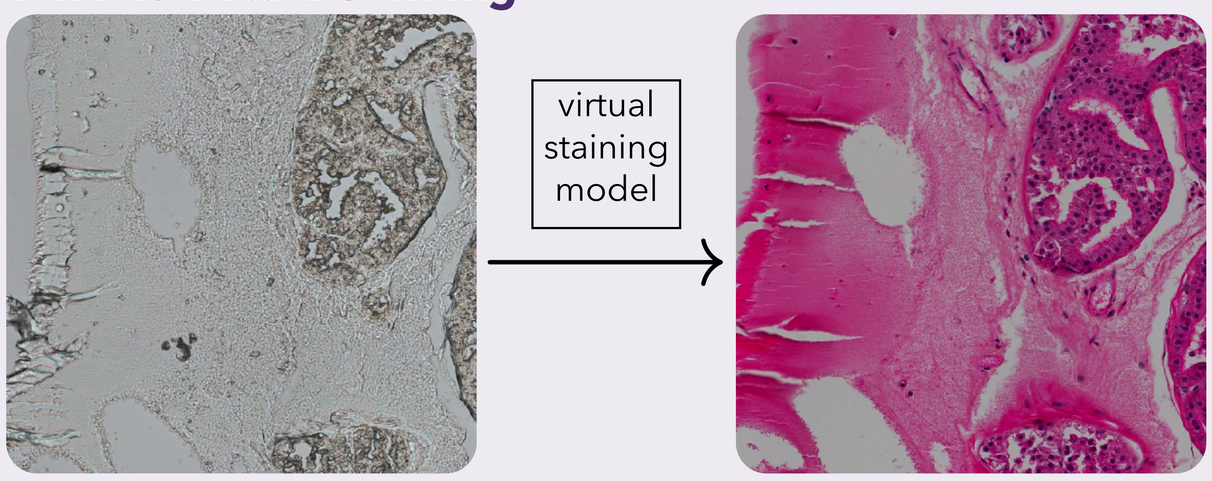

Unstained tissue — Stained tissue

Staining is an essential part of a pathological examination which makes natural tissue visible under a microscope. Virtual staining uses a learned model to replicate this. It can also convert between different types of staining, such as H&E and IHC.

The Problem - Reducing Tiled Inference Artefacts

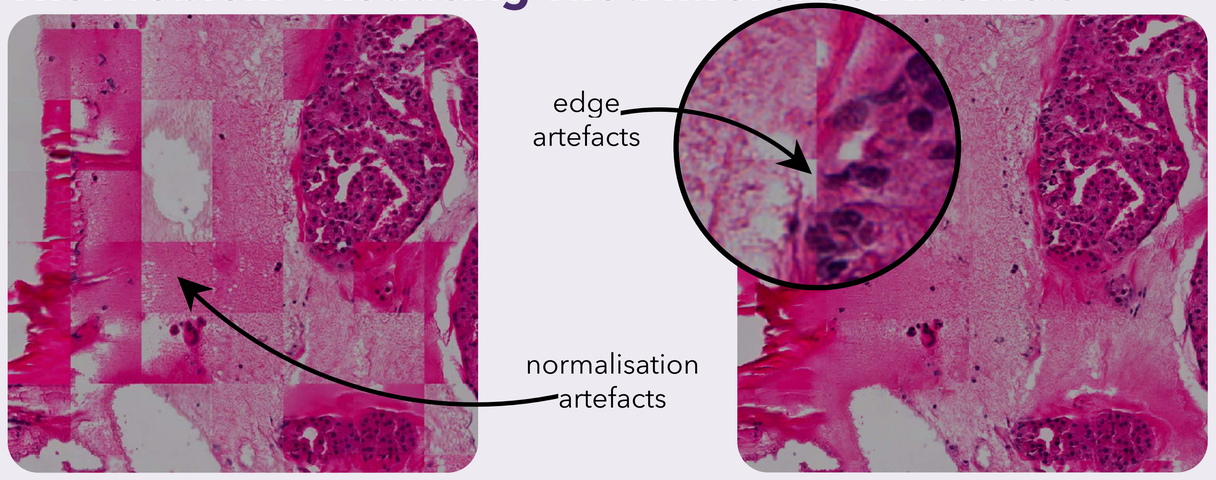

Instance Normalisation [1] — Dense Normalisation [2]

Due to the size of a whole slide image (an uncompressed WSI can be over 30 GB, much larger than handled by current GPUs), virtual staining must be performed on tiles, resulting in stitching artefacts due to convolutions and instance normalisation.

STAIN (Our Contribution)

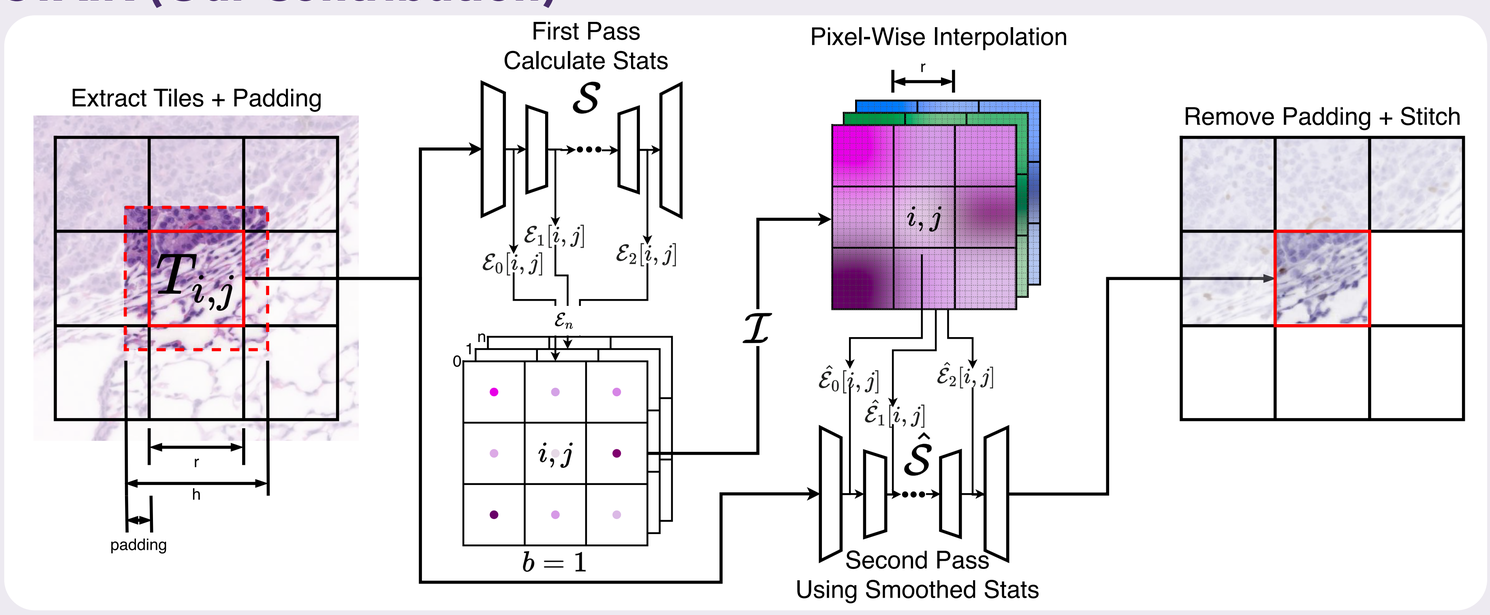

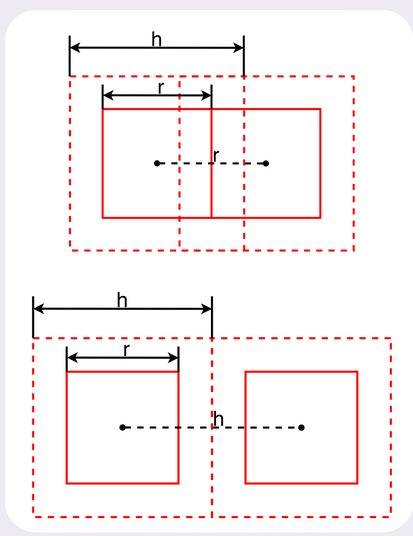

Overlapping tiles to reduce artefacts

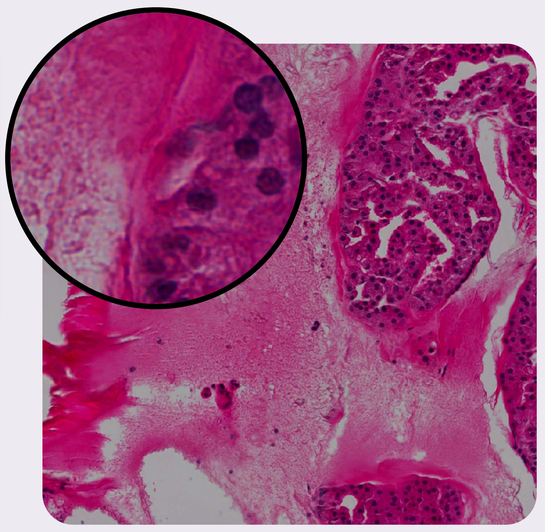

STAIN output
No visible artefacts

Our proposed architecture combines two techniques to reduce stitching artefacts: the model is run twice on every tile. The first pass collects per-tile normalisation statistics. These are interpolated between neighbouring tiles, with overlap, and used for normalisation during the second pass. This results in a smoother output.

Focused Sobel Gradients

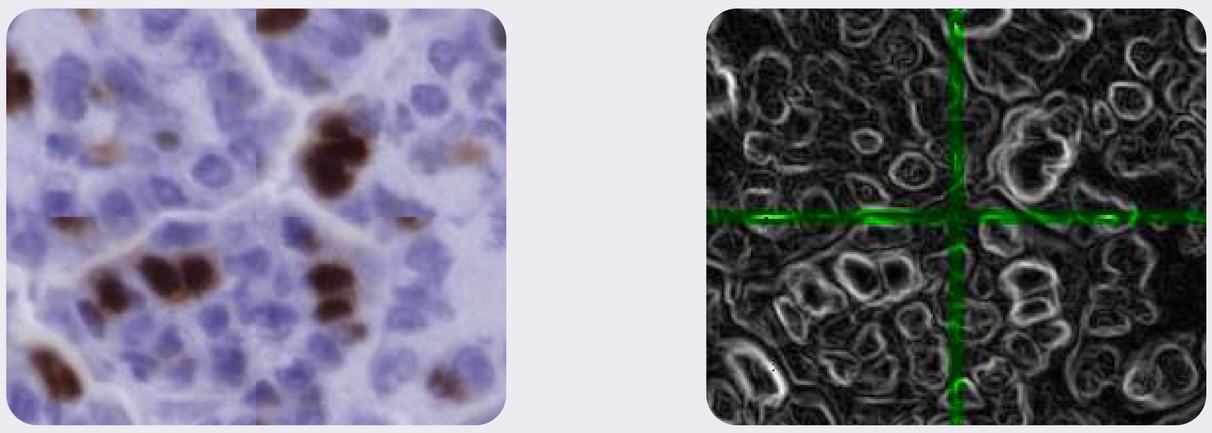

Sample — FSG calculation area

We developed Focused Sobel Gradients which measures the average Sobel gradients at the (known) borders between tiles. By comparing this to the Sobel gradients across the rest of the tile, this metric is relatively insensitive to other image content.

Tiled Absolute Distance

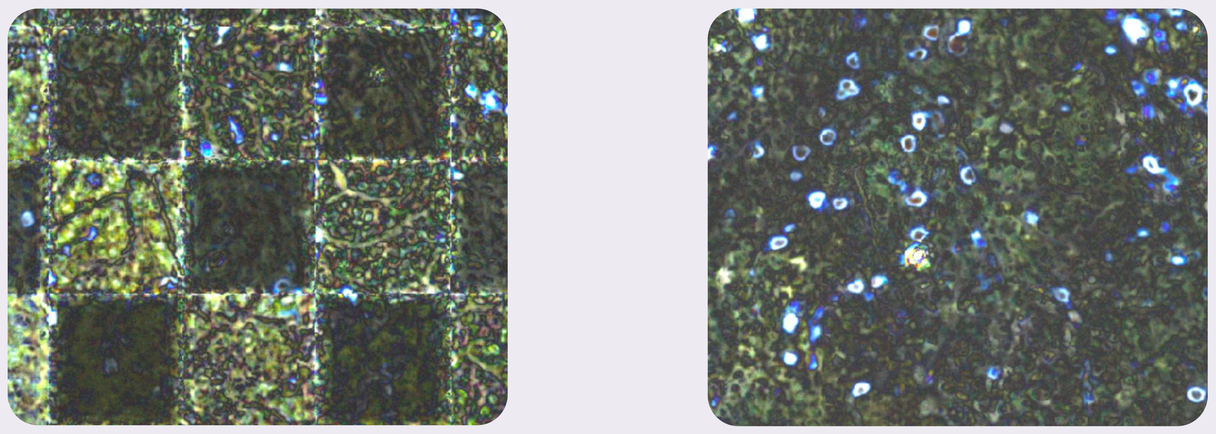

Kernelised Instance Normalisation [3] — STAIN

Perfect stitching should result in no change in output regardless of how the slide is tiled. Our novel metric, Tiled Absolute Distance, measures the difference between translations with offset tiling, where a perfect output would be uniformly grey.

Experiments

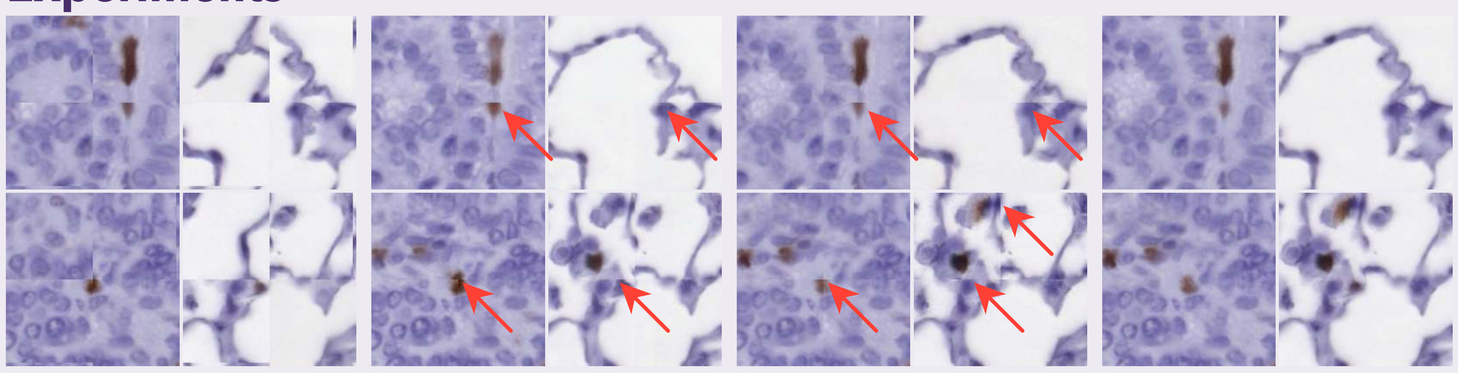

IN — KIN — DN — STAIN (Ours)

Qualitative Comparison

We include a qualitative comparison between stitching methods:

- Simple Instance Normalisation demonstrates normalisation discrepancies between tiles
- Kernelised Instance Normalisation and Dense Normalisation show more issues with convolutional artefacts
- STAIN significantly reduces both problems

Quantitative Comparison

We compare STAIN with other techniques using three quantitative metrics:

- Standard Fréchet Inception Distance (FID), to measure overall perceived quality
- Our novel Focused Sobel Gradients (FSG) metric, to measure edges between tiles
  - Works by calculating the average Sobel gradients in the regions along tile boundaries
- Our novel Tiled Absolute Distance (TAD) metric, to measure the consistency of the stitching
  - Works by calculating the difference between two virtually stained versions of the same unstained image using offset tiling

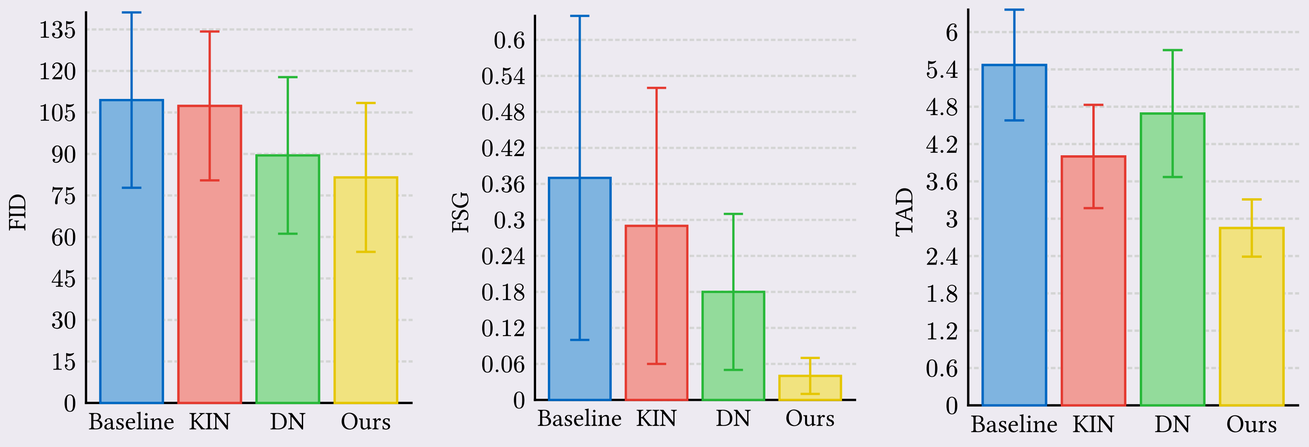

[1] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance Normalization: The Missing Ingredient for Fast Stylization," no. arXiv:1607.08022. arXiv, Nov. 2017. doi: 10.48550/arXiv.1607.08022.
[2] M.-Y. Ho, C.-M. Wu, M.-S. Wu, and Y. J. Tseng, "Every Pixel Has Its Moments: Ultra-High-Resolution Unpaired Image-to-Image Translation via Dense Normalization," no. arXiv:2407.04245. arXiv, Jul. 2024.
[3] M.-Y. Ho, M.-S. Wu, and C.-M. Wu, "Ultra-High-Resolution Unpaired Stain Transformation via~Kernelized Instance Normalization," in *Computer Vision – ECCV 2022*,