Spatial Transcriptomics-Conditioned Latent Diffusion Models for Synthetic Histopathology Tissue Patch Generation

Chaphalkar, Arnav (School: Adlai E. Stevenson High School)

Whole-slide histopathology images(WSIs) form the backbone of AI development in cancer research. However, their high curation cost, gigapixel complexity, and limited availability across institutions make them largely inaccessible for research. While synthetic WSI generation has emerged as a promising alternative, current approaches fail to reproduce key tumor features with biological relevance, making them unsuitable for analysis and model development. This work presents a novel spatial transcriptomics-conditioned AI pipeline for biologically grounded synthetic WSI generation. To provide rich training data, whole-slide images and matched Visium spatial profiles from 41 colorectal cancer tumors were segmented into 255,744 aligned tissue patches with their corresponding gene expression vectors. GeneCondNet, a custom encoding algorithm, was developed to translate continuous spatial transcriptomic profiles into the CLIP embedding space used by the Stable Diffusion image-generation model. These embeddings guide the UNet-based diffusion neural network through cross-attention layers, allowing spatial information to serve as the blueprint for histology image reconstruction. The pipeline achieved a Frechet Inception Distance of 55, reflecting strong image fidelity. In blinded evaluations, expert pathologists could not distinguish spatially conditioned synthetic images from real data in 91% of cases, demonstrating the model's high accuracy and ability to preserve biological relevance. By linking spatial transcriptomics and tissue structure in a generative model for the first time, this work enables scalable WSI generation, advancing cancer research, broadening access to computational pathology, and empowering the creation of more advanced and equitable AI models across the field.

Awards Won:
Third Award of $1,200