Generative modeling for creating synthetic histopathology images utilizing diffusion-based architectures

Rumman Mahfujul Islam

Computational Systems Biology Laboratory

Professor Dr Shigehiko Kanaya

Abstract

Generative models have emerged as powerful tools in artificial intelligence, capable of learning complex data distributions to generate realistic and high-quality synthetic data. In the field of medical imaging, particularly histopathology, the ability to generate synthetic images offers immense potential for data augmentation, algorithm training, and diagnostic research. This work explores the application of diffusion-based generative models for the synthesis of histopathology images.

At first, an unconditional denoising diffusion probabilistic model was implemented using three different variance schedulers to generate synthetic histopathology images that possess semblance to the original dataset. The performance of the artificially generated images was quantitatively evaluated using the Fréchet inception distance metric in the context of each scheduler, demonstrating the capability of this model to produce high-fidelity images.

Building upon these findings, a conditional latent diffusion model was developed to enhance image controllability and interpretability. The approach involved embedding unlabeled histopathology images into a lower-dimensional latent space using a vector quantized generative adversarial network, followed by the application of a diffusion process in the latent space. Clustering of latent representations served as a conditioning mechanism for guided image synthesis. Later, expert insights were incorporated into cluster assignments to improve interpretability, while quantitative metrics and cluster validation techniques were used to assess the image quality and optimal model configuration. Overall, this research demonstrates the progression from probabilistic to conditional diffusion-based generative modeling for histopathology image synthesis.