SEMINAR ANNOUNCEMENT

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING
COLLEGE OF DESIGN AND ENGINEERING
Website: https://cde.nus.edu.sg/ece

Area: Signal Analysis & Machine Intelligence (SAMI)

Host: Asst Prof Jin Yueming

TOPIC	:	**Enhancing Spatial Transcriptomics Prediction with Diffusion-Based Gene-Image Representation Learning**
SPEAKER	:	**Mr Zhang Bochong** **Graduate Student, ECE Dept, NUS**
DATE	:	**Friday, 24 October 2025**
TIME	:	**4:00PM to 5:00PM**
VENUE	:	**Join Zoom Meeting** **https://nus-sg.zoom.us/j/86929106719?pwd=WgGTcNmXyS8AOTtkF6sGEZ1krJkUET.1** **Meeting ID: 869 2910 6719** **Passcode: 942283**

ABSTRACT

Spatial Transcriptomics (ST) is crucial for understanding tumor biology by linking gene expression to histological context, but its high cost limits broad application. Deep learning offers a solution by predicting gene expression from cost-effective H&E-stained pathology images. However, this is an inherently ill-posed problem, as diverse morphological patterns can correspond to similar gene profiles, and conventional direct-regression models often average across plausible molecular states. Existing methods further struggle with MLP-based gene encoders that cannot handle heterogeneous gene sets and violate the fundamental permutation-invariance property of gene expression data.

To address these challenges, we introduce DiffBulk, a novel two-stage framework for tile-level pseudo-bulk gene expression prediction. In the first stage, we leverage a gene-to-image conditional diffusion model to implicitly capture the complex relationship between gene expression and histological features. This stage is equipped with a permutation-invariant open-embedding gene encoder that allows for unified training across diverse gene panels, enhancing the framework's scalability. We also incorporate a Probabilistic Masking Switch (PMS) module to ensure the diffusion U-Net remains effective for downstream prediction, even when gene conditions are absent.

In the second stage, diffusion-derived features from the frozen U-Net are extracted and fused with representations from a pathology foundation model (FM) via a gated mechanism. This fusion effectively bridges the domain gap and creates an enhanced feature representation for the final gene prediction. Evaluated on three high-quality Xenium ST pseudo-bulk datasets (HEST-Bowel, HEST-Pancreas, and CrunchDAO-Bowel), DiffBulk consistently outperforms state-of-the-art task-specific and FM-based models across all metrics. These results validate the power of diffusion-based representation learning for integrating histology and transcriptomics, offering a robust and scalable method for gene expression prediction.

BIOGRAPHY

The speaker is a Master of Engineering (MEng) in the Department of ECE at the NUS.