A068 · Unsupervised extraction of interpretable, functional programs from spatial transcriptomics through a contrastive learning framework

Presenter: Neal Kewalramani — Boston University

Authors: Neal Kewalramani, Jeff Sheridan, Mark Crovella, Ruben Dries

Session: Day 1 · Thu Oct 1, 2026 · 2:15–4:15 PM

Spatial transcriptomics resolves gene expression across intact tissue. The biological programs that organize the tissue often cannot be captured by discrete domains. Instead, they behave as continuous processes that vary smoothly across space that can grade into one another. Capturing these continuous programs, tying them to interpretable gene-level biology, and recovering the tissue architecture a pathologist would recognize remains an ongoing challenge. Existing representation methods can recover spatial structure but tend to yield embeddings or clusters that resist gene-level interpretation and flatten the gradients defining how tissue transitions between different states, such as from benign to malignant in cancer.

Here we apply scCont, a fully unsupervised contrastive learning framework, to spatial transcriptomics (Visium, 10X Genomics). scCont first defines positive pairs using a k-nearest-neighbors approach over local transcriptomic neighborhoods and then learns a contrastive embedding whose latent dimensions map to functional gene programs. Trained across ten prostate datasets without spatial or histological labels, scCont recovers interpretable, spatially coherent programs, including a carcinoma axis marked by loss of benign secretory identity, reactive stroma, and an epithelial–stromal compartment axis. Each program is corroborated by prostate literature and a quantitative comparison against MSigDB Hallmark gene sets. We benchmark scCont against established linear and neural network methods and observe that its programs are the most predictable and align well with pathologist annotations, separating malignant from benign glands with large effect sizes. Finally, these unsupervised features transfer directly to annotate new Visium datasets, identifying functional programs and candidate tissue boundaries where expert labels are absent.

Title: Unsupervised extraction of interpretable, functional programs from spatial transcriptomics through a contrastive learning framework

Spatial transcriptomics resolves gene expression across intact tissue. The biological programs that organize the tissue often cannot be captured by discrete domains. Instead, they behave as continuous processes that vary smoothly across space that can grade into one another. Capturing these continuous programs, tying them to interpretable gene-level biology, and recovering the tissue architecture a pathologist would recognize remains an ongoing challenge. Existing representation methods can recover spatial structure but tend to yield embeddings or clusters that resist gene-level interpretation and flatten the gradients defining how tissue transitions between different states, such as from benign to malignant in cancer.

Here we apply scCont, a fully unsupervised contrastive learning framework, to spatial transcriptomics (Visium, 10X Genomics). scCont first defines positive pairs using a k-nearest-neighbors approach over local transcriptomic neighborhoods and then learns a contrastive embedding whose latent dimensions map to functional gene programs. Trained across ten prostate datasets without spatial or histological labels, scCont recovers interpretable, spatially coherent programs, including a carcinoma axis marked by loss of benign secretory identity, reactive stroma, and an epithelial–stromal compartment axis. Each program is corroborated by prostate literature and a quantitative comparison against MSigDB Hallmark gene sets. We benchmark scCont against established linear and neural network methods and observe that its programs are the most predictable and align well with pathologist annotations, separating malignant from benign glands with large effect sizes. Finally, these unsupervised features transfer directly to annotate new Visium datasets, identifying functional programs and candidate tissue boundaries where expert labels are absent.