

Graph Contrastive Learning for Deciphering Spatial Heterogeneity in Breast Cancer

Yingsa Qiao

*University of Edinburgh Institute (ZJU-UoE Institute), Zhejiang University, Hangzhou, China
1527275825@qq.com*

Abstract. Breast cancer exhibits significant spatial and molecular heterogeneity. Traditional bulk and single-cell transcriptomic analyses struggle to fully reveal the molecular states and microenvironmental differences across distinct tumor regions due to their inability to preserve spatial tissue information. This study proposes and applies stCL, a spatial transcriptomics domain identification framework based on graph attention networks and contrastive learning, to dissect spatial transcriptomic heterogeneity in breast cancer. It integrates gene expression and local spatial topology through a graph attention-based multi-view encoder, jointly optimizing contrastive learning loss, spatial regularization loss, and zero-inflated negative binomial reconstruction loss to learn biologically meaningful low-dimensional embeddings. Results show that stCL effectively identifies spatially coherent domains such as Tumor, Invasive, Surrounding Tumor, and Healthy regions, which largely align with pathological annotations. Quantitative evaluation indicates that stCL achieves an adjusted Rand index (ARI) of 0.6008 and normalized mutual information (NMI) of 0.7095, outperforming several existing spatial clustering methods. Further differential expression and functional enrichment analyses reveal distinct molecular signatures and functional differences among spatial domains: the Invasive region is enriched for epithelial tumor and cell proliferation-related pathways, the Surrounding Tumor region shows enrichment in immune response and extracellular matrix-related signals, while the Tumor region displays specific metabolic and epithelial characteristics.

Keywords: Spatial transcriptomics, breast cancer, graph attention network, contrastive learning, spatial domain identification

1. Introduction

Among women, breast cancer ranks as one of the most common malignancies worldwide [1]. GLOBOCAN 2024 estimates roughly 2.434 million new female cases and about 694,000 deaths worldwide that year, which places breast cancer second among diagnosed cancers after lung cancer and fourth for cancer-related deaths [1]. China reported around 376,000 new cases and 73,000 deaths, so it sits among the countries carrying a relatively heavy disease burden [1]. As the population keeps aging and lifestyles keep shifting, this burden will probably keep climbing. Breast cancer also shows complicated progression and highly varied treatment strategies, and its care leans

heavily on molecular subtyping to guide therapy and forecast outcomes, which has made it a prototypical disease in precision oncology.

Breast cancer is markedly heterogeneous, and this shows up in histology, hormone receptor status, HER2 status, proliferative activity, gene expression, the immune microenvironment, and metastatic behavior. In the clinic, immunohistochemical subtyping usually relies on markers such as ER, PR, HER2, and Ki-67. The PAM50 signature goes one step further, sorting breast cancer into Luminal A, Luminal B, HER2-enriched, Basal-like, and Normal-like molecular subtypes from the expression patterns of 50 genes. These subtypes differ sharply in biology, prognosis, and treatment response.

Conventional bulk and single-cell transcriptomics can probe breast cancer heterogeneity at different levels, yet neither keeps spatial tissue information intact [2]. Bulk RNA-seq may blur the differences between cell subpopulations and functional states, and it cannot say where these molecular features sit within regions such as tumor cores, invasive margins, necrotic areas, or stromal zones. Single-cell RNA sequencing does resolve tumor cells and the microenvironment composition at the level of individual cells. But tissue dissociation wipes out the original spatial positions and neighbor relationships of those cells. Spatial transcriptomics works differently: it captures gene expression together with spatial coordinates in intact tissue sections, so molecular features can be mapped to distinct regions while the tissue architecture is kept [2]. This makes it possible to spot the spatial distribution of tumor niches, immune-related areas, and various functional states, and it hands researchers new tools for understanding the spatial basis of tumor progression, immune evasion, treatment response, and prognosis.

This study proposes and applies stCL, a spatial transcriptomics domain identification algorithm built on graph attention networks and contrastive learning, to dissect the spatial heterogeneity of breast cancer tissues [3]. The model adaptively gathers expression information from neighboring spots through graph attention networks, and combines multi-view contrastive learning, spatial regularization loss, and ZINB reconstruction loss to learn low-dimensional embeddings that capture molecular expression similarity, spatial continuity, and count data distribution at the same time [3]. The study then uses K-means clustering together with spatial neighborhood label optimization to achieve unsupervised identification of spatial domains in breast cancer tissues [4]. Spatial transcriptomics supplies gene expression data while keeping the in situ tissue architecture, which lays a foundation for identifying functional spatial structures such as tumor core regions, invasive margins, stroma-enriched areas, immune-active zones, and potential hypoxic regions [2]. By combining gene expression profiles with spatial proximity relationships, stCL learns low-dimensional representations that reflect both molecular similarity and spatial continuity, allowing unsupervised clustering of spatial structures within breast cancer tissues. After the spatial domains are segmented, the study compares differentially expressed genes, pathway activities, cellular compositions, and tumor microenvironment features across regions, with a focus on spatial variations in PAM50-related molecular subtypes, cell proliferation, HER2 signaling, epithelial-mesenchymal transition, hypoxia, angiogenesis, and immune activation versus suppression.

2. Spatial transcriptomics and graph neural networks

With spatial transcriptomics, RNA is captured, prepared into libraries, and sequenced at defined positions across an intact tissue, and the native architecture survives the whole process, so expression data can be laid directly onto tissue images [2]. 10x Genomics Visium illustrates the workflow. A tissue section sits on a chip dotted with spatial barcode probes; once staining, imaging, permeabilization, mRNA capture, reverse transcription, amplification, sequencing, and Space

Ranger analysis are done, what comes out is gene expression, spatial coordinates, and tissue morphology [5, 6]. Gene expression matrices, spatial coordinates, and histological images form the main outputs of 10x Visium and Visium HD.

Applied to breast cancer, spatial transcriptomics lays bare how tumor cells relate spatially to immune cells, fibroblasts, endothelial cells, and other microenvironmental components [2]. Where different cell types sit within the tumor core, the invasive margins, and the stromal regions can be assessed, and spatial features such as immune activation, immune suppression, and tertiary lymphoid-like structures come into view. The spatial basis of breast cancer progression, immune evasion, and prognostic differences becomes easier to probe as a result. When PAM50 subtyping is combined with molecular features such as HER2 status, cell proliferation, hypoxia, epithelial-mesenchymal transition (EMT), angiogenesis, and immune checkpoints, molecular states and potential therapeutic responses can be compared across distinct spatial regions, giving valuable clues for precise classification and personalized treatment of breast cancer.

Spatial transcriptomics suits graph neural networks (GNNs) almost by design, because the data maps onto a graph structure naturally: spots, bins, or cells as nodes, gene expression levels as node features, spatial proximity relationships as edges. Through message passing, GNNs pull in information from neighboring nodes, and tissue-local structures, spatial continuity, and microenvironmental similarities are captured along the way [7]. Where PCA, clustering, or non-spatial dimensionality reduction stop at expression data, GNNs go on to exploit spatial topological relationships and fold tissue imaging information into specific models. SpaGCN, for instance, identifies spatial domains and spatially variable genes, and STAGATE learns the importance of neighboring spots adaptively through a graph attention autoencoder, which sharpens its recognition of spatial structures and feature representation [8]. More robust, biologically meaningful low-dimensional representations can thus be learned from spatial transcriptomics data that is high-dimensional, sparse, and noisy.

The advantage of contrastive learning lies in self-supervised representation learning that needs no human-labeled annotations. Spatial transcriptomics data usually lack reliable ground-truth labels, but contrastive learning can still learn discriminative feature representations by building positive and negative sample pairs. GraphST fuses GNNs with self-supervised contrastive learning, and it strengthens spatially aware representations by modeling spatial adjacency relationships, which improves performance on tasks such as spatial clustering, multi-slice integration, and cell type identification [3]. In this combined framework, GNNs mainly model spatial topology and neighborhood information, while contrastive learning makes the representations more discriminative and robust. In highly heterogeneous tumors such as breast cancer, this approach identifies tumor cores, invasive margins, stroma-enriched regions, and immune-active zones more accurately, which further supports the analysis of molecular and functional differences across distinct spatial regions.

3. Dataset

This study uses the public 10x Genomics human breast cancer spatial transcriptomics dataset as its experimental data [5]. The dataset was produced with the 10x Genomics Visium Spatial Gene Expression platform, which captures gene expression profiles at different spatial locations (spots) on tissue sections through capture areas carrying spatial barcodes [5]. This breast cancer tissue section contains multiple regions with distinct pathological features, including ductal carcinoma in situ (DCIS), lobular carcinoma in situ (LCIS), and invasive carcinoma, together with non-tumor tissues and areas of different histological appearance, and it shows clear spatial and pathological heterogeneity [5].

According to the quality control information provided by 10x Genomics, this tissue section contains 3,813 tissue-associated spatial locations and 22,968 detected genes in total. The median number of detected genes per spot is 5,394, the median UMI count is 17,531, and the average sequencing read count is about 149,800 [5]. These metrics show that the dataset has high sequencing depth and rich transcriptomic information, which provides a reliable foundation for subsequent spatial representation learning and analysis of breast cancer tissue heterogeneity.

4. Methods




Figure 1. Schematic diagram of the stCL model framework and analysis workflow

The overall analysis workflow of stCL is shown in Figure 1. The original expression matrix is denoted as $X \in \mathbb{R}^{N \times G}$, where N represents the number of spots and G represents the number of genes. Subsequently, high-variance genes are selected as model input to reduce noise while preserving major expression patterns associated with spatial and biological variation [9]. The spatial structure is constructed based on the two-dimensional coordinates of each spot on the tissue section. Let the spatial coordinate of the i th spot be s_i . An edge is added between two spots in the graph if their Euclidean distance is less than a given threshold r :

$$A_{ij} = 1(\|s_i - s_j\|_2 \leq r) \quad (1)$$

A spatial adjacency graph $G=(V, E)$ comes out of this step, with nodes for spots and edges for spatial proximity relationships. The local spatial structure of the tissue section is held in this graph, and drawing on it, the model works with gene expression features and spatial dependencies among neighboring spots together.

The model's main component employs a Graph Attention Network (GAT) to integrate gene expression and spatial neighborhood information. For a node i and its neighbor j , the model first

linearly maps the input feature h_i to W_{h_i} , then computes attention weights based on the features of the target node and its neighboring nodes, and finally updates the node representation:

$$e_{ij} = (\text{LeakyReLU}(a^T [W_{h_i} || W_{h_j}])) \quad (2)$$

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k \in N_i} \exp(e_{ik})}$$

$$h'_i = \sigma(\sum_{j \in N_i} \alpha_{ij} W_{h_j})$$

where N_i denotes the spatial neighbors of node i , α_{ij} represents the contribution weight of neighbor j to node i , and W and a are learnable parameters, while σ is a nonlinear activation function. The multi-head attention mechanism further learns spot representations from different subspaces, concatenating or averaging the outputs of four attention heads to form the final low-dimensional embedding.

The model training jointly optimizes three types of loss functions. The first term is the neighborhood contrastive loss, which constrains representations of the same spot across different attention heads and encourages representations of spatially adjacent spots in the adjacency graph to be closer, while non-adjacent or other global spots serve as negative samples. Cosine similarity is used to measure similarity, and a temperature coefficient τ controls the smoothness of the similarity distribution [3]. The second term is the spatial regularization loss, which penalizes spots that are physically distant but have embeddings that are too close, thereby enhancing consistency between low-dimensional representations and spatial structure [3].

$$L_{reg} = \frac{1}{N^2} \sum_{i,j} (1 - d_z(i, j)) d_s(i, j) \quad (3)$$

where $d_z(i, j)$ is the normalized distance in the embedding space and $d_s(i, j)$ is the normalized distance in the physical space. The third term is the ZINB reconstruction loss, designed to model the common discrete, overdispersed, and zero-inflated expression distributions observed in spatial transcriptomics data. The model predicts the zero-inflation probability π , mean μ , and dispersion parameter θ via the decoder, and minimizes the negative log-likelihood of the ZINB distribution as the reconstruction loss. The overall optimization objective is:

$$L = \lambda_c L_{contrast} + \lambda_r L_{reg} + \lambda_z L_{zinb} \quad (4)$$

where λ_c , λ_r , and λ_z are the weight hyperparameters for the three loss terms.. After training, K-means clustering is applied to the learned low-dimensional representations, followed by label optimization incorporating spatial neighborhood information to make the predicted labels smoother within local spatial regions. The model builds spatial expression maps from gene expression matrices and spatial coordinates. GAT-based multi-view encoding and contrastive objectives produce low-dimensional embeddings, and ZINB decoding, clustering, and label optimization then output the spatial domains.

5. Results




Figure 2. Comparison between stCL-identified spatial domains and pathological annotations (Panels A/B; ARI = 0.6008, NMI = 0.7095)

This study first applied stCL to the breast cancer Visium dataset and compared its identified spatial domains against pathological annotations. As shown in Figure 2, stCL can identify multiple spatially contiguous domains that largely match the pathological regions. Tumor, invasive carcinoma, peritumoral, and relatively healthy areas show clear boundaries, which indicates that stCL captures tissue heterogeneity in breast cancer effectively by combining gene expression with spatial neighborhood information. Compared to clustering based only on expression matrices, stCL gives results with stronger spatial continuity, and this helps cut down local noise and makes region delineation more stable.



Figure 3. Comparison of ARI and NMI performance between stCL and other spatial clustering methods

To evaluate spatial clustering performance quantitatively, this study used the Adjusted Rand Index (ARI) and Normalized Mutual Information (NMI) to compare stCL results with pathological annotations. As shown in Figure 3, stCL achieved an ARI of 0.6008 and an NMI of 0.7095, performing well among the methods compared. Against SpaGCN, stLearn, Leiden, GraphST, SCANPY, STAGATE, SEDR, stGCL, MENDER, and SpaceFlow, stCL posted the highest ARI score [3, 4, 8-14], which points to a high level of consistency between its clustering results and the histopathological annotations. The NMI results also show that stCL keeps pathological region information well, suggesting that the learned low-dimensional representations not only separate spatial regions but also reflect the major structural differences in the histological annotations.

To analyze the molecular characteristics of the different stCL spatial regions, this study ran differential expression analysis with Scanpy's Wilcoxon rank-sum test, grouping by stCL_region and selecting genes with adjusted p-values < 0.05 and log fold changes > 0.25 [9]. Figure 4 shows that each region has its own marker gene profile. The Healthy region is characterized by high expression of MALAT1, TIMP1, C3, and IGFBP7. The Invasive region is enriched for KRT8, S100A16, CCND1, and KRT18, which indicates enhanced epithelial tumor and proliferative features. The Surrounding tumor region is enriched for immunoglobulin-related genes such as IGHG1, IGLC2, IGKC, and IGHG3, along with complement-related genes, pointing to immune responses or components of the tumor microenvironment. The Tumor region displays upregulation of ABHD2, BAMBI, COX6C, S100G, and CSTA, which reflects differences in tumor metabolism, signal transduction, and epithelial status.




Figure 4. Regional expression patterns of representative marker genes across different spatial domains

Based on the differentially expressed gene list, this study went on to run GO and pathway enrichment analysis, with the results shown in Figure 5. The healthy region was mainly enriched in terms such as extracellular matrix structural constituent, extracellular space, and cell periphery, which points to preserved abundant matrix structures and tissue homeostasis features. The invasive region was enriched in cytoplasm and extracellular exosome, suggesting active cellular metabolism and extracellular vesicle processes. The surrounding tumor region was enriched in extracellular space and extracellular region, together with upregulation of immunoglobulin genes, which indicates remodeling of the extracellular environment and immune-related activities. The tumor region was enriched in extracellular exosome and extracellular vesicle, suggesting active intercellular communication and secretion processes within the tumor microenvironment. The associated gene lists can be further reviewed and visualized in ShinyGO.




Figure 5. KEGG pathway enrichment results for differentially expressed genes across spatial domains

Figure 6 summarizes the biological differences among distinct spatial domains from the perspective of the tumor microenvironment [2]. The high-expression spot distribution maps show that regions with elevated expression of different marker genes are not scattered at random, but rather correspond to specific spatial domains defined by stCL. Heatmaps of average regional expression further show that Healthy regions are relatively enriched in TIMP1, C3, and IGFBP7; Invasive regions in KRT8, KRT18, and CCND1; Surrounding tumor regions in BGN, COL1A1, and DCN; and Tumor regions in ABHD2, S100G, and KRT19. These results show that the spatial domains identified by stCL not only line up with pathological annotations but also reflect distinct transcriptional states and tumor microenvironment characteristics. Overall, stCL can capture spatial architecture, molecular expression differences, and potential functional compartmentalization within breast cancer tissues all at once.



Figure 6. Distribution of spots with high expression of representative marker genes in breast cancer spatial transcriptomics slices

6. Discussion and conclusion

On 10x Genomics Visium breast cancer spatial transcriptomics data, this study put stCL to the test for spatial domain identification and biological feature analysis. The identified breast cancer-related regions showed strong spatial continuity, and agreement with the pathological annotations was high overall. Adjusted Rand index (ARI) and normalized mutual information (NMI) scores came out high as well, which means the clustering genuinely reflects histological structural differences. Differential expression and functional enrichment analyses then exposed distinct molecular signatures and functional differences among the spatial domains. Segmentation accuracy and biological meaning arrived together: stCL not only delineates spatial regions precisely but also offers insight into spatial heterogeneity and the tumor microenvironment in breast cancer.

Heterogeneity in breast cancer runs along both spatial and molecular lines. Where traditional transcriptomic analysis loses the tissue location of gene expression, spatial transcriptomics keeps it, so morphology, spatial architecture, and molecular status can be read at once. The regions that stCL carved out here, Tumor, Invasive, Surrounding tumor, and Healthy, lined up well with the pathological morphology. Upregulation of KRT8, KRT18, and CCND1 in the Invasive region marks strong epithelial tumor characteristics and active cell proliferation. Enrichment of immunoglobulin and extracellular matrix-related genes in the Surrounding tumor region points to potentially active immune responses and stromal remodeling. In the Tumor region, ABHD2, S100G, and KRT19 stand out, reflecting specific metabolic and epithelial-associated states. Taken together, molecular heterogeneity in breast cancer tracks tissue spatial structure closely, and cellular populations and microenvironmental components across regions likely act jointly in tumor initiation and progression.

stCL draws its strengths from pairing graph attention networks with a contrastive learning framework. On the graph side, local neighborhood information is aggregated along the spatial adjacency between spots, with attention weights adaptively measuring how much each neighbor contributes to the representation of a target spot. Simple neighborhood aggregation handles the complex, heterogeneous spatial structures of tumor tissues far less flexibly than this attention mechanism does. Contrastive learning contributes the other half, enforcing consistent, discriminative low-dimensional representations across views or attention heads and thereby sharpening the model's grasp of spatial domain boundaries and local expression patterns. A ZINB reconstruction loss rounds it out, fitting the sparsity, over-dispersion, and high share of zero counts typical of spatial transcriptomics data. Gene expression features, spatial neighborhood relationships, and the distributional properties of count data are thus all put to work, which is what makes stCL a fairly complete modeling framework for identifying spatial structures in breast cancer.

In practical terms, the domains stCL identifies carry two kinds of value for breast cancer: supporting pathological region delineation and feeding tumor microenvironment research. Histopathological annotation reads tissue morphology first and foremost, whereas clustering on spatial transcriptomics contributes complementary molecular-level evidence, exposing regions that look alike in morphology yet differ at the molecular level. Differential genes and GO functional enrichment, once combined, point to potential therapeutic targets and regulatory mechanisms inside the microenvironment. Invasive areas with upregulated proliferation and epithelial markers, for instance, may be tracking tumor expansion and invasiveness, and peritumoral enrichment of immune- and stroma-related signaling pathways suggests roles in immune regulation, extracellular matrix remodeling, and cancer cell migration. Validating these region-specific molecular signatures further could open a deeper view of the spatial ecosystem and functional heterogeneity that drive breast cancer progression.

Limitations remain. Visium spots stop short of single-cell resolution, several cell types can share one spot, and expression signals in some spatial regions may carry the imprint of that mixed composition. A single breast cancer tissue section carried the whole analysis, so larger sample sets, other molecular subtypes of breast cancer, and independent cohorts must still test the findings. Spatial transcriptomics data alone drive the current analysis, too; immunohistochemistry, single-cell transcriptomics, or spatial proteomics, folded in later, would help confirm the cellular composition and functional states of the domains stCL identifies. The overall picture, though, is clear. stCL combines gene expression with spatial neighborhood information effectively and identifies biologically meaningful spatial domains within breast cancer tissues. Its results agree well with pathological annotations and lay open the differences across spatial regions in gene expression, functional enrichment, and tumor microenvironment. What this adds up to is a robust computational framework for studying spatial heterogeneity in breast cancer, assisting pathological analysis, and exploring potential therapeutic targets.

References

- [1] Sung, H., et al. (2026). Global cancer statistics 2024: GLOBOCAN estimates of incidence and mortality worldwide for 34 cancers in 186 countries. *CA: A Cancer Journal for Clinicians*, 76(4), e70090. <https://doi.org/10.3322/caac.70090>.
- [2] Ståhl, P. L., et al. (2016). Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science*, 353(6294), 78-82. <https://doi.org/10.1126/science.aaf2403>.
- [3] Long, Y., et al. (2023). Spatially informed clustering, integration, and deconvolution of spatial transcriptomics with GraphST. *Nature Communications*, 14, article 1155. <https://doi.org/10.1038/s41467-023-36796-3>.

- [4] Traag, V. A., et al. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9, article 5233. <https://doi.org/10.1038/s41598-019-41695-z>.
- [5] 10x Genomics. (2026, August 16). Human breast cancer (Block A Section 1). *10x Genomics*. <https://www.10xgenomics.com/datasets/human-breast-cancer-block-a-section-1-1-standard-1-0-0>.
- [6] 10x Genomics. (n.d.). Understanding Space Ranger outputs: Official 10x Genomics support. *10x Genomics*. <https://www.10xgenomics.com/support/software/space-ranger/3.0/analysis/outputs/output-overview>
- [7] Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*. <https://doi.org/10.48550/arXiv.1609.02907>
- [8] Hu, J., et al. (2021). SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network. *Nature Methods*, 18, 1342–1351. <https://doi.org/10.1038/s41592-021-01255-8>
- [9] Wolf, F. A., et al. (2018). SCANPY: Large-scale single-cell gene expression data analysis. *Genome Biology*, 19, Article 15. <https://doi.org/10.1186/s13059-017-1382-0>
- [10] Pham, D., et al. (2023). Robust mapping of spatiotemporal trajectories and cell-cell interactions in healthy and diseased tissues. *Nature Communications*, 14, Article 7739. <https://doi.org/10.1038/s41467-023-43120-6>
- [11] Ren, H., et al. (2022). Identifying multicellular spatiotemporal organization of cells with SpaceFlow. *Nature Communications*, 13, Article 4076. <https://doi.org/10.1038/s41467-022-31739-w>.
- [12] Xu, H., et al. (2024). Unsupervised spatially embedded deep representation of spatial transcriptomics. *Genome Medicine*, 16, Article 12. <https://doi.org/10.1186/s13073-024-01283-x>.
- [13] Yu, N., et al. (2026). stGCL: A versatile cross-modality fusion method based on multi-modal graph contrastive learning for spatial transcriptomics. *Genome Biology*, 27, Article 51. <https://doi.org/10.1186/s13059-025-03896-w>.
- [14] Yuan, Z. (2024). MENDER: Fast and scalable tissue structure identification in spatial omics data. *Nature Communications*, 15, Article 207. <https://doi.org/10.1038/s41467-023-44367-9>.