NACNet: A Histology Context-aware Transformer Graph Convolution Network for Predicting Treatment Response to Neoadjuvant Chemotherapy in Triple Negative Breast Cancer

Qiang Li[1], George Teodoro[2], Yi Jiang[1,3,*], Jun Kong [1,3,*]

[1]Department of Mathematics and Statistics, Georgia State University, Atlanta, 30303, GA, USA

[2]Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte, 31270, Minas Gerais, Brazil

[3]Winship Cancer Institute, Emory University, Atlanta, 30322, GA, US

Abstract

Neoadjuvant chemotherapy (NAC) response prediction for triple negative breast cancer (TNBC) patients is a challenging task clinically as it requires understanding complex histology interactions within the tumor microenvironment (TME). Digital whole slide images (WSIs) capture detailed tissue information, but their giga-pixel size necessitates computational methods based on multiple instance learning, which typically analyze small, isolated image tiles without the spatial context of the TME. To address this limitation and incorporate TME spatial histology interactions in predicting NAC response for TNBC patients, we developed a histology context-aware transformer graph convolution network (NACNet). Our deep learning method identifies the histopathological labels on individual image tiles from WSIs, constructs a spatial TME graph, and represents each node with features derived from tissue texture and social network analysis. It predicts NAC response using a transformer graph convolution network model enhanced with graph isomorphism network layers. We evaluate our method with WSIs of a cohort of TNBC patient (N=105) and compared its performance with multiple state-of-the-art machine learning and deep learning models, including both graph and non-graph approaches. Our NACNet achieves 90.0% accuracy, 96.0% sensitivity, 88.0% specificity, and an AUC of 0.82, through eight-fold cross-validation, outperforming baseline models. These comprehensive experimental results suggest that NACNet holds strong potential for stratifying TNBC patients by NAC response, thereby helping to prevent overtreatment, improve patient quality of life, reduce treatment cost, and enhance clinical outcomes, marking an important advancement toward personalized breast cancer treatment.

Keywords: Whole slide image; Pathological complete response; Graph convolution neural network; Transformer; Triple negative breast cancer; Tumor microenvironment

---

*Correspondence author:

Email:yjiang12@gsu.edu (Y. Jiang); jkong@gsu.edu (J. Kong)

1 Introduction

Breast Cancer is the leading malignant disease and cause of cancer death in women worldwide [1]. Of all breast cancer subtypes, Triple Negative Breast Cancer (TNBC) presents the most aggressive progression, with a dismal prognosis and a high recurrence rate [1]. In current clinical practice, treatment option for this cancer subtype is limited. Although new immunotherapies have recently emerged, chemotherapy remains the primary treatment for patients with TNBC at both early and advanced stages. In particular, because of the absence of estrogen receptors (ER), progesterone receptors (PR), and human epidermal growth factor receptor 2 (HER2), TNBCs do not respond to hormone therapies or HER2-targeted treatments, significantly narrowing the range of therapeutic options. Consequently, chemotherapy, particularly neoadjuvant chemotherapy (NAC), continues to be the standard-of-care and predominant treatment option for this aggressive cancer subtype [2–6].

NAC treatment response is evaluated in resected tissues from the surgery using the residual cancer burden (RCB) metric. Based on RCB scores, two response classes are defined: Pathology Complete Response (pCR) with a post-treatment RCB of zero, and Residual Disease (RD), indicating an incomplete response. Clinically, pCR is often used as an endpoint for reoperative treatment and is strongly correlated with long-term clinical benefits [7, 8].

However, only about 30%~40% of TNBC patients respond well to NAC treatment, while the remaining patients either respond moderately or are refractory to NAC [9]. Despite its clear clinical benefits, for example a reduction in the tumor size or a downgrade in breast cancer before the surgery, NAC can substantially decrease patient life quality with its adverse treatment effects [10]. For non-responders, NAC not only introduces unnecessary toxicity but also delays alternative treatments and surgery, leading to adverse outcomes and overtreatment. Unfortunately, there remains a critical, unmet need to accurately predict NAC treatment response at the time of diagnosis. An accurate treatment response prediction remains a major clinical challenge.

In recent years, machine learning has been successfully applied to histopathology images for predicting survival outcomes and treatment responses in the clinical oncology [11, 12]. Numerous studies have demonstrated the effectiveness of image-based machine learning in improving clinical decision-making. For example, deep learning models have been used to predict early treatment response in metastatic colorectal cancer by identifying subtle morphological changes in tumors, surpassing traditional size-based assessment methods [13]. In breast cancer research, deep learning frameworks have been used to integrate histopathology images with additional data types, such as genomic and radiomic features, to predict NAC response [14–16]. These models provide non-invasive methods for assessing treatment efficacy and guiding personalized therapy strategies. The successful applications of deep learning across various cancer types suggest its potential to revolutionize clinical treatment response prediction, offering more accurate and personalized approaches to support patient care.

Thanks to significant advances in tissue scanning technology, whole slide images (WSIs) of tissue sections can now be routinely produced to capture cell-level tissue details. However, due to their giga-pixel scale, WSIs pose significant computational challenges, making it difficult to directly deploy machine learning and image analysis techniques for large-scale automated analyses. [17–19]. To manage the computational demand, it is a common practice to partition giga-pixel WSIs into numerous small image tiles and analyze them using

a multiple instance learning (MIL) strategy [20–25]. In MIL, tiles from each WSI are treated as instances from a bag that share the same patient-level label [26, 27]. Each tile is assigned a predicted label and these tile-level labels are then aggregated to form a comprehensive patient-level representation [28–30]. However, this MIL strategy may be less effective due to the high tissue heterogeneity within the breast cancer tumor microenvironment (TME), which breaks the MIL assumption that all instances in the negative (RD) bag should have a negative label [31]. Notably, TNBCs exhibit significant TME heterogeneity, where the presence of tumor-infiltrating lymphocytes are associated with a better prognosis and an improved response to chemotherapy [6]. Numerous studies further highlight that TME histological interactions strongly correlated with NAC response in TNBC patients [31–34]. However, MIL-based analyses ignore these spatial TME histology distributions and interactions, remaining insensitive to local tissue context and the global structure organization.

Graph Convolution Network (GCN) has been developed to partially address this problem. It represents each WSI as a graph consisting of nodes and edges, and predicts the WSI label by aggregated features of nodes in local neighborhoods [25, 35–37]. Learning a structured graph requires an effective graph representation [38]. However, constructing a context-aware, effective, and minimal WSI topological representation to hence predictive discriminating power remains an open challenge. Additionally, there is limited research on how to generate biologically meaningful graph node representations and edge connections for cancer grading and subtyping tasks using histopathology WSIs. Most current approaches cutting WSIs into tiles, then develop GCNs at the tile level, and finally aggregate the outcomes from the tile nodes to create overall WSI-level predictions [25, 39]. Although tile-level approaches have advanced research in WSI processing, they fall short in constructing WSI graphs with well-informed node labels and attributes. When the objective is to identify the entire tumor region or capture the TME connectivity in WSIs, where nodes characterize disease stage, it becomes important to incorporate both regional and WSI-level information with accurately labeled nodes and detailed attribute features. Therefore, label-informed WSI graph structure learning methods are needed for these prediction and analysis.

To address the challenges posed by spatial heterogeneity in TNBC WSIs, we developed a multi-step NAC prediction network (NACNet) that leverages a WSI-derived graph representation and a transformer-based graph convolution network (GCN) for NAC response prediction in TNBC patients. Our approach leverages clinically relevant TME information such as cellular (e.g., lymphocyte infiltration) and extracellular matrix (e.g., collagen) heterogeneity for an accurate treatment response prediction. First, we train a deep learning model to recognize local histology labels and produce histology label maps for WSIs. Next, we construct a spatial TME graph using tile-level histology label map, extract context aware graph features, and predict the NAC treatment response using a transformer GCN model that is boosted with Graph Isomorphism Network (GIN) layers [40]. We applied the NACNet to a TNBC patient cohort, and demonstrated that incorporating accurate spatial TME structures enhanced NAC response prediction.

3

2   Dataset

This study used a set of specimens from a TNBC patient cohort receiving NAC treatment at the Decatur Hospital, Georgia, USA. All patient tissues were histologically processed and stained with Hematoxylin and Eosin ($H\&E$). The H&E-stained tissues were scanned using a high-throughput high-resolution digital scanner (Hamamatsu NanoZoomer 2.0-HT C9600-13) at 40x magnification (0.23 μm/pixel). Board-certified pathologists annotated histological features in these slides using an open-source image processing software (QuPath, ver. 0.1.2). The features include adipose tissue, Polyploid Giant Cancer Cells (PGCCs), Microvessel Density (MVD), tumor, stroma, Carcinoma In-Situ (CIS), hemorrhage, mucinous change, apocrine change, immune cells, necrosis, and muscle tissue. The TNBC patient cohort included 105 female patients in total, with 48 patients responding (pCR) and 57 not responding (RD) to the NAC treatment.

3   Methods

We developed a histology context-aware transformer graph convolution network, NACNet, for an enhanced TNBC NAC response prediction with H&E WSIs. NACNet recognizes histology labels of image tiles, constructs spatial TME graphs, characterizes graph nodes with the local node label information including both type and count, tissue texture, and social network analysis (SNA) features. We then used a transformer graph convolution network (GCN) equipped with graph isomorphism network (GIN) layers to analyze these graphs and predict the treatment outcome(Figure 1).
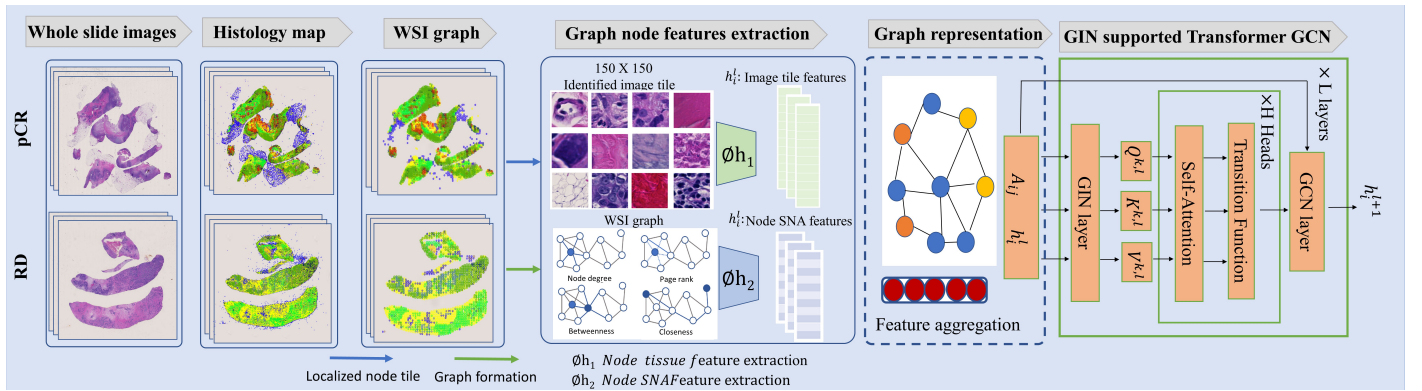

Figure 1: **Overview of the NACNet architecture.** Each WSI is partitioned into non-overlapping image tiles of size 150×150 pixels.A pre-trained convolutional model for image tile classification identifies all image tile classes, resulting in a tile-level histology label map where each pixel represents the histology label of an image tile. A sliding window moves over each histology label map and defines neighbors that share the same histology label as a WSI graph node. Both graph-based SNA and image texture features are used to represent the graph nodes. Nodes within the distance ϵ are connected with an edge. The resulting ϵ-neighborhood graph is provided to a transformer-based GCN with GIN layers. This architecture incorporates the TME spatial information, improving the prediction power.

3.1 Histology map generation

To accommodate computational and memory constraints, we partitioned each WSI into non-overlapping 150 × 150 image tiles and recorded the spatial coordinates of the top left corner of each tile. We trained a convolutional neural network (CNN) classifier to converts each tile into a histology label. Then a WSI became a histology label map (Figure 2a). The histology labels of training image tiles were based on the detailed histology annotations of 12 distinct histology classes by pathologists. Background regions are identified by the cutoff intensity value 220 in all RGB image channels. Tiles with more than 90% background regions are discarded. Thanks to the large number of TNBC related histology components of interest in this study, the histology map can capture diverse and intricate histological characteristics important for TNBC treatment response prediction.

3.2 TME graph construction

Characterizing the spatial histopathological heterogeneity is critical for both understanding cancer progression and for cancer subtyping, biomarker discovery, and treatment response prediction [20, 42]. Our NACNet model creates a graph-based TME representation that takes advantage of the spatial interactions across different tissue histology regions in a WSI (Algorithm 1). The graph, $G = (V, E)$ is undirected, where V is the set of nodes and E is the set of edges connecting adjacent nodes.

As each image tile in our study is 150 × 150 by size, we use a 10×10 sliding window over the histology map to identify context-aware graph nodes by connecting neighboring tiles of the same histology label. As we move the sliding window over the histology map, we spatially cluster the histology map labels within the sliding window. Each resulting cluster is defined as a graph node with its location represented by the cluster centroid. We carefully considered both the physical size differences among histological categories and their clinical significance in breast cancer. Our approach to node formation is uniquely context-aware, aiming to capture the essential pathological features that inform prognosis and treatment strategies. Among our 12 histology categories, several—including tumor, necrosis, immune cells, microvessel density (MVD), stroma, and polyploid giant cancer cells (PGCC)—carry high clinical relevance due to their direct association with tumor progression, prognosis, and treatment response [43–45]. For these critical categories, we assigned a lower threshold $n = 5$ value, ensuring they cluster more readily into nodes and are prominently represented within the TME graph. This approach emphasizes the spatial distribution of these important features within the tissue structure, enhancing the model's ability to focus on key regions that reflect the tumor microenvironment's complexity. Conversely, categories with lower clinical impact, such as hemorrhage, mucinous changes, carcinoma in situ (CIS), adipose tissue, muscle tissue, and apocrine change, were assigned higher thresholds $n = 10$. This prevents over-segmentation of less significant features while still allowing these elements to provide meaningful context within the overall histological landscape. This clinically-weighted clustering strategy not only reduces the impact of misclassifications but also enhances the model's accuracy in constructing a context-aware TME graph. By assigning weights based on clinical relevance, our approach uniquely captures the heterogeneous spatial and clinical characteristics of breast cancer tissues, thus offering a more nuanced and informative
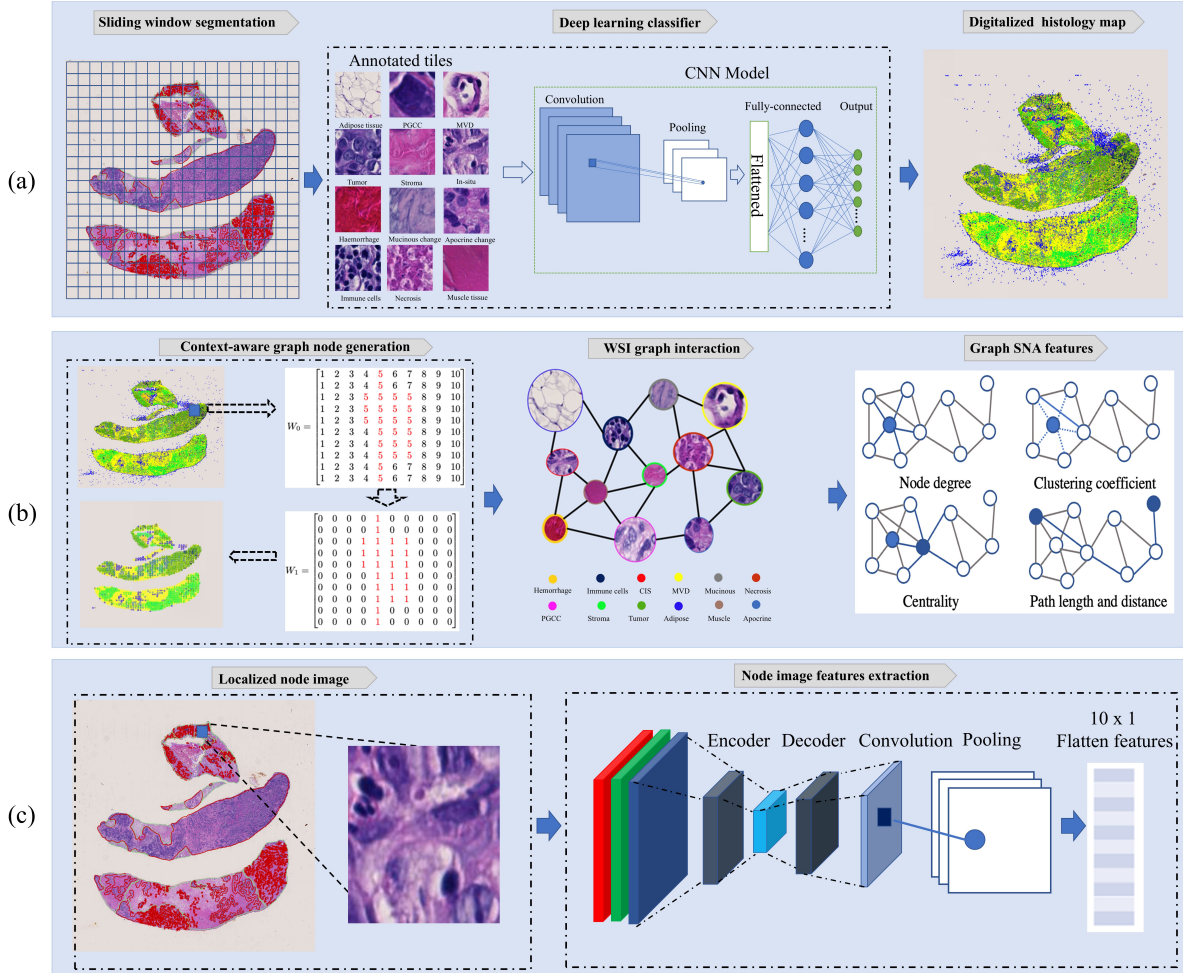
Figure 2: Feature extration from WSI spatial TME graph. (a) Each WSI is partitioned into image tiles. A training set of slides are annotated with one of 12 histology classes. A VGG16 model [41] trained on these annotated slides predict the histology labels for the rest of the titles. All labeled tiles are then combined to produce the histology map for each WSI. (b) Within the histology map, a large cluster of tiles ($n > \eta$) sharing the same histology label is defined as a TME graph node. Any pair of nodes within distance ϵ are connected by an edges. Multiple SNA features are derived from each graph node (e.g., degree, betweenness, page rank, and closeness). (c) An autoencoder connected to a VGG flatten layer is used to find image embeddings of graph nodes. The node size, tissue texture and SNA features are integrated to enable a context-aware graph characterization.

representation for prognostic modeling. Additionally, the number of histology map entries sharing the same histology label in the sliding window is the local node label count (Figure 2b). A pair of nodes with spatial distance less than $\epsilon = 1,500$ pixels (0.25 μm/pixel) [46]. The adjacency matrix $A = [a_{ij}]$, where $a_{ij} = 1$ if an edge exists between node i and j and $a_{ij} = 0$ otherwise, captures the TME edge set E and models histology component spatial interactions in TME.

For each graph node, we use the local node label, node label count, SNA features, and node tissue texture features to capture the spatial histology interactions in the TME graph. We calculate the local node label count by the number of image tiles associated with each node. In addition, we calculate the SNA features [47]

6

Algorithm 1 WSI Graph Representation

    Input: W: WSI; M: Histology map
    Output: G: Graph data; L_c: Subgraph clique distribution; L_e: Edge distribution
    Stage 1: Graph and Node Feature Generation
    for each sliding window m in the M **do**
        $C \leftarrow$ SpatiallyClusterLabels(m)   ▷ Spatially cluster histology map labels within the sliding window
        for each cluster c in C **do**
            $label\_count \leftarrow$ CountCluster(c)   ▷ Count number of labels in the cluster
            if $label\_count > \eta$ **then**
                $V \leftarrow$ CalculateCentroid(c)   ▷ Calculate centroid of the cluster
                $F \leftarrow$ autoencoder(W(centroid(c)))   ▷ Extract image texture features
            end if
        end for
    end for
    for each node v_i in V **do**
        if $d(v_i, v_j) < \epsilon$ **then**
            $E_{ij} \leftarrow (v_i, v_j)$ ▷ Create the edge
        end if
    end for
    return $G(V, E, F)$ ▷ Graph and its representation
    Stage 2: Subgraph Cliques and Edge Distribution
    $L_c(v_i, v_j, v_k) \leftarrow$ allcycles($G, MCL, \lambda = 3$)   ▷ Find all cycles of length three
    $L_e(v_i, v_j) \leftarrow$ allcycles($G, MCL, \lambda = 2$)   ▷ Find edges with all two-node histology label combinations
---

that include node degree, betweenness, pagerank, and closeness (Figure 2b). Specifically, the node degree (or degree centrality) is the number of edges connected to the current node. The betweenness centrality of a node is the sum of the ratio of the number of the shortest paths from node S to T through current node to the number of the shortest paths from S to T over all pairs of S and T. PageRank is a link analysis algorithm originally developed to measuring the relative importance of web documents within the World Wide Web [48]. It assigns a numerical weight to each node in our spatial TME graph that is calculated iteratively as the sum of the ratios of the PageRank of neighboring connected nodes to the number of outbound links from these neighboring connected nodes. Closeness centrality measures the ability of a node to pass information in a graph and is calculated as the reciprocal of the average shortest path distance from the current node to all other nodes in the graph. To include tissue details for the node characterization, we extract node image tile texture features using an autoencoder flatten layer [49] (Figure 2c), resulting in a feature vector of length 12.

For each node, the label, label count, texture features, and SNA features together form the node feature vector h_v of length d. The spatial TME graph feature matrix $H \in \mathbb{R}^{n \times d}$ consists all the features of all node n. The adjacency matrix A captures the edge connections, while the graph feature matrix H represents the graph node attributes. The adjacency matrix A and the feature matrix H together represent the tissue spatial TME graph from each WSI, providing a reliable and standardized method for treatment response prediction.

3.3 NAC treatment response prediction based on the TME graph

Our NACNet aims to improve the accuracy and efficiency of NAC response predictions for TNBC patients using the WSIs. To achieve this goal, NACNet model combines the transformer Graph Convolution Network (GCN) [50] with the Graph Isomorphism Network (GIN) layer [40] to predict the treatment response. The GCN learns the abstract graph features through message passing, where feature vectors in the hidden layers are iteratively updated by aggregation of feature vectors from neighboring nodes [51]. To improve the performance, we combine GCN with a self-attention mechanism introduced by transformer [50].

In our study, we further enhance the model WSI-level representation and discrimination power by a GIN layer with an Multilayer Perceptron (MLP) before the transformer GCN based spatial TME graph prediction module. This integration empowers the model to capture intricate tissue patterns and histology component relationships from the spatial tissue graph. Specifically, the GIN layer is known for data distribution learning in a low-dimensional space. This good merit enables GINs to discriminate different class-conditional graph structures, and map graph structures of same class to similar embeddings with a strong generalizability [52]. In our model architecture, the output of GIN with Multilayer Perceptron (MLP) is further provided to a transformer with a self-attention mechanism. The aggregated features are next provided to a transition function that incorporates residual connections [53] and a normalization layer (LNorm) [54]. Finally, the resulting latent graph embeddings are processed by a GCN layer. In this way, the graph embeddings are iteratively updated by the GIN layer, transformer and the GCN:

$$h_v'^{(k)} = MLP^{(k)} \left((1 + \alpha^{(k)}) \cdot h_v^{(k)} + \sum_{u \in \mathcal{N}(v)} h_u^{(k)} \right) \tag{1}$$

$$\tilde{H}'^{(k)} = Attention \left(H'^{(k)} Q^{(k)}, H'^{(k)} K^{(k)}, H'^{(k)} V^{(k)} \right) \tag{2}$$

$$H^{(k+1)} = GCN(\mathbf{A}, \tilde{H}'^{(k)}) \tag{3}$$

where α is a learnable parameter for the node representation update [40]; $h_v^{(k)}$ is the feature vector of node v at the k-th layer, while $H^{(k)}$ is the graph feature matrix. For each node $v \in \mathcal{V}$, the feature vectors of a set of neighbors (i.e. $\mu \in \mathcal{N}(v)$) by the adjacency matrix $\mathbf{A}$ are uniformly sampled and aggregated with the current node feature vector. Additionally, $V^{(k)}, Q^{(k)}, K^{(k)} \in \mathbb{R}^{d \times d}$ are the value projection weight, query-projection, and key-projection matrix in the transformer, respectively.

3.4 TME histology component spatial interaction characterization

As the complex spatial TME component interactions, e.g. the interactions across tumor, stromal and immune cells, can shed much insight on the underlying disease status, we include multiple TME histology components in the tissue WSIs within a spatial TME graph structure. With this method, we can readily extract TME features of treatment response predictive value and compare the spatial TME component interaction differences between NAC treatment responder (i.e., pCR) and non-responder (i.e., RD) group.

First, we analyze the group differences in the number of subgraph cliques associated with different histology class combinations. To make our analysis amenable, we count the 3-node subgraphs, as outlined in Algorithm 1. We initially identify all subgraphs of cycle length three in the spatial TME graph G by the function *allcycles*, with a Maximum Cycle Length (MCL) set to three. Next, we incrementally count the three-node subgraph clique associated with each of the C_3^l histology class combinations of interest. Next, we identify and count edges connecting two nodes of all histology class label combinations of interest in the same way, but with a MCL set to two. Finally, a statistical t-test is used to assess the discrepancy in subgraph counts between the NAC responder and non-responder group.

In addition, we investigate the predictive value of the TME histology class profile for NAC response prediction. Specifically, we calculate correlation coefficients between histology label count and NAC response (i.e. responders or non-responders) by Pearson correlation and produce the histology label correlation maps to intuitively demonstrate the importance of histology components. Additionally, we compute relevance scores by Minimum-Redundancy-Maximum-Relevance (mRMR) analysis [55] for histology label relevancy importance maps. The histology label correlation and relevancy importance maps by Pearson correlation and mRMR can visually provide insights into the contribution of individual histology labels to the prediction.

3.5   Implementation and evaluation

A VGG16 model for image tile classification: All image tiles are extracted from WSIs with the Python-Openslide package. For image tile classification, the VGG 16 model implemented with TensorFlow Keras 2.3.0 is trained and tested on a workstation equipped with an NVIDIA GTX 3080 GPU. We randomly select 80% and 20% of the labeled image tiles with similar prior histology label distributions for effective training and testing, respectively. For training, we employ a stochastic gradient descent optimizer with a learning rate of 0.001, Mini-batch size of 32, and a dropout rate of 0.2.

Autoencoder for image tile texture extraction: We implement an autoencoder architecture with a VGG16 convolutional layer for image tile texture extraction in our study. The autoencoder consists of an encoder and sequential decoder. The encoder compresses input data into a lower-dimensional latent space using densely connected layers with rectified linear unit activation. Conversely, the decoder reconstructs the original data from this compressed representation, employing a combination of dense layers and a reshape operation to match the input dimensions. The encoded representations are softmax-normalized. This approach enables both an efficient feature extraction and the feature dimensionality reduction.

Model for NAC response: We develop the spatial TME graph-based prediction model with the PyTorch framework. The resulting model is trained with the Adam optimizer [56], a maximum of 250 epochs and a dropout rate of 0.5. Different learning rates are tested, including $e^{-4}, 5e^{-4}, 1e^{-3}, 5e^{-3}$. Two attention heads, two MLP layers, and a hidden size of 256 are included in the model (Supplement TableS1).

Evaluation metrics: We separate patients into two groups: those with pathological complete response (pCR) as the postive group and those with residual disease (RD) as the negative groups. Accuracy (ACC) measures the overall performance of the method, while precision reflects the model's ability to avoid false positives. The area under the ROC curve (AUC-ROC) evaluates the model's discrimination power across various classification

thresholds. The F1 score provides a balanced measure, considering both precision and recall. We use an 8-fold cross-validation for the performance evaluation. In each fold, 87.5% and 12.5% of the patients are used for training and testing, respectively. The mean and standard deviation of testing accuracies from these 8 folds are calculated to assess the overall prediction performance. With the same training and validation schema, we evaluate numerous state-of-the-art methods for a fair comparison.

4 Results

4.1 VGG16 model for histology map generation

With a dataset of 10,000 annotated image tiles, the VGG16 model for image tile histology classification achieves a test accuracy of 90.0%. We visually illustrate representative image tile histology maps and the WSI-derived spatial TME graphs (Figure 3). Through this implementation, we validate the performance of the VGG16 model in generating histology maps, laying a solid foundation and offering valuable insights for further tissue histology interaction analysis. By using the clustering method to identify nodes, we effectively prevent contamination of the resulting TME graph by image noise, such as tissue scratches, ink dots, dust marks, and fingerprints.

4.2 Ablation study and validation

To evaluate the effectiveness of individual NACNet components, we first investigate the impact of individual WSI-derived spatial TME graph node features on the NAC response model performance using 8-fold cross-validation. Next, we examine NACNet's predictve performance by integrating multimodal node features. Both models with single- and multi-modal features are analyzed within the same NACNet pipeline framework. (1) NACNet-I uses only node image tile texture features; (2) NACNet-Lcontains the node label and count; (3) NACNet-S uses node SNA features; (4) NACNet-LI uses the node label, count, and node image tile texture features by autoencoder linked with a VGG flatten layer; (4) NACNet-LS uses node label, its count, and node SNA features; (5) NACNet-IS uses node image tile texture features and node SNA features; (6) NACNet-LIS includes node label, label count, node image tile texture features, and SNA features.

We present the ablation experimental results with evaluation metrics of Accuracy, AUC, Sensitivity, Specificity, Precision, and F1 score in Table 1. The ROCs associated with different ablation models are illustrated in Figure 4. (1) Note the NACNet-LIS achieves the best prediction performance by Accuracy, and Sensitivity. For other metrics, it also presents competitive results. These results suggest the integrative use of multi-modal features help improve NAC response prediction. (2) Recall the image tile feature embeddings are produced by a convolutional auto-encoder and a VGG flatten layer. NACNet-LI works better than NACNet-LS overall, indicating the stronger complementary prediction power from image tile feature embeddings than SNA features. (3) Spatial node histology label and count information plays a critical role for NACNet model performance. Prediction models without node histology label and count information perform poorly. (4) Prediction models with the GIN layer work better than those without this layer overall, suggesting the efficacy of the GIN layer.
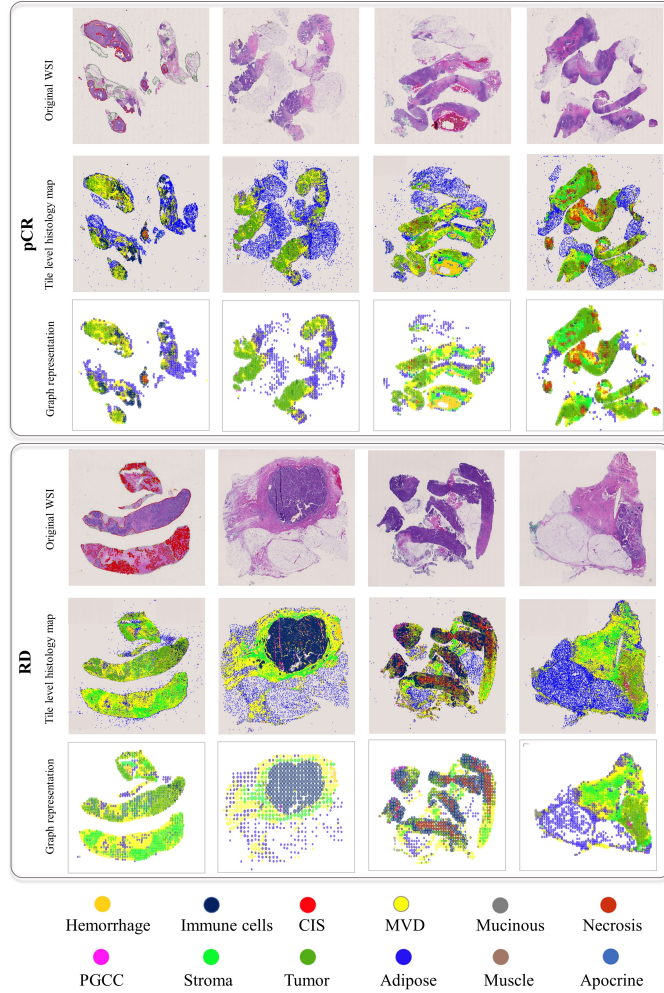
Figure 3: WSI-derived spatial TME graph. We present representative WSIs (top), the corresponding histology maps (middle), and the resulting WSI TME graph node distributions (bottom). The histology labels of image tiles of size 150×150 are classified and assembled to construct the histology map. In total, 12 histology labels are color coded, including hemorrhage, immune cells, carcinoma in situ (CIS), MVD, mucinous changes, necrosis, PGCC, stroma, tumor, adipose tissue, muscle tissue, and apocrine change.

Table 1: Performance of NACNet using different node feature combinations (Mean ± Standard deviation). Models with (∗) do not have a GIN layer. NACNet-I uses texture features by an autoencoder and a VGG flatten layer; NACNet-L uses the node label count; NACNet-LI uses both node label count and image tile texture features; NACNet-LS uses both node label count and SNA features; NACNet-LIS uses node label count, image texture features, and SNA features.

Models	ACC	AUC	Sensitivity	Specificity	Precision	F1
NACNet-L*	0.78 ± 0.12	0.77 ± 0.17	0.57 ± 0.22	0.93 ± 0.09	0.90 ± 0.13	0.67 ± 0.16
NACNet-LI*	0.80 ± 0.19	0.81 ± 0.14	0.83 ± 0.15	0.78 ± 0.25	0.76 ± 0.21	0.79 ± 0.18
NACNet-LS*	0.74 ± 0.11	0.71 ± 0.22	0.54 ± 0.24	0.87 ± 0.13	0.71 ± 0.32	0.60 ± 0.25
NACNet-LIS*	0.81 ± 0.14	0.81 ± 0.14	0.78 ± 0.20	0.83 ± 0.15	0.77 ± 0.19	0.78 ± 0.17
NACNet-IS*	0.48 ± 0.10	0.57 ± 0.10	0.10 ± 0.14	0.11 ± 0.15	0.14 ± 0.19	0.79 ± 0.16
NACNet-I	0.55 ± 0.17	0.56 ± 0.16	0.47 ± 0.23	0.52 ± 0.19	0.45 ± 0.12	0.65 ± 0.22
NACNet-S	0.59 ± 0.18	0.57 ± 0.14	0.48 ± 0.16	0.60 ± 0.25	0.49 ± 0.13	0.66 ± 0.28
NACNet-IS	0.49 ± 0.18	0.48 ± 0.19	0.33 ± 0.23	0.46 ± 0.31	0.35 ± 0.23	0.64 ± 0.21
NACNet-L	0.81 ± 0.11	0.79 ± 0.16	0.76 ± 0.31	0.78 ± 0.11	0.75 ± 0.17	0.80 ± 0.13
NACNet-LI	0.88 ± 0.08	0.83 ± 0.12	**0.96 ± 0.09**	0.83 ± 0.15	0.82 ± 0.15	0.91 ± 0.08
NACNet-LS	0.79 ± 0.09	0.71 ± 0.20	0.62 ± 0.27	0.90 ± 0.08	0.74 ± 0.31	0.66 ± 0.27
NACNet-LIS	**0.90 ± 0.07**	**0.82 ± 0.11**	**0.96 ± 0.08**	**0.88 ± 0.10**	**0.83 ± 0.15**	**0.88 ± 0.09**

4.3 Performance comparison with SOTA

Comparison with GCN methods: Our NACNet model is compared with multiple state-of-the-art (SOTA) graph baseline deep learning methods, including Maximum Entropy Weighted Independent Set Pooling for GNNs (MEWISPool) [57], Deep Graph CNN (DGCNN) [58], SAGPool [59], and SAG Hierarchical pooling [60]. For fair comparisons, we use the same training hyperparameters, loss function, and eight-fold training and testing cross-validation partitions. We present the comparison results in Table 2 (top) and Figure 5(a). Our NACNet model with node information, image texture features and SNA features performs substantially better than the SOTA graph baseline GCN models by all evaluation metrics. Interestingly, most ablated NACNet models in Table 1 outperform the SOTA graph baseline deep learning methods.

Comparison with non-graph deep learning models: We use the multiple instance learning strategy for non-graph deep learning model training. Each image tile is labeled with the same NAC response for the patient. For testing, a majority voting mechanism is used to aggregate tile-level predictions into a final NAC response prediction at the patient level. Specifically, VGG-16, AlexNet, and ResNet [41, 53, 61] are trained and evaluated. The comparisons are presented in Table 2 (bottom) and Figure 5(a). While these methods present performance comparable to graph-based deep learning methods, they are substantially inferior to our NACNet models.

Comparison with traditional machine learning methods: We evaluate multiple traditional machine learning models using histology label proportion, and WSI-graph SNA features with the same 8-fold cross-validation mechanism. The traditional machine learning classifiersinclude Logistic Regression [62], k-Nearest Neighbors (kNN) [63], Support Vector Machines (SVM) [64], Naive Bayes [65], Decision Trees [66], Random Forests [67],
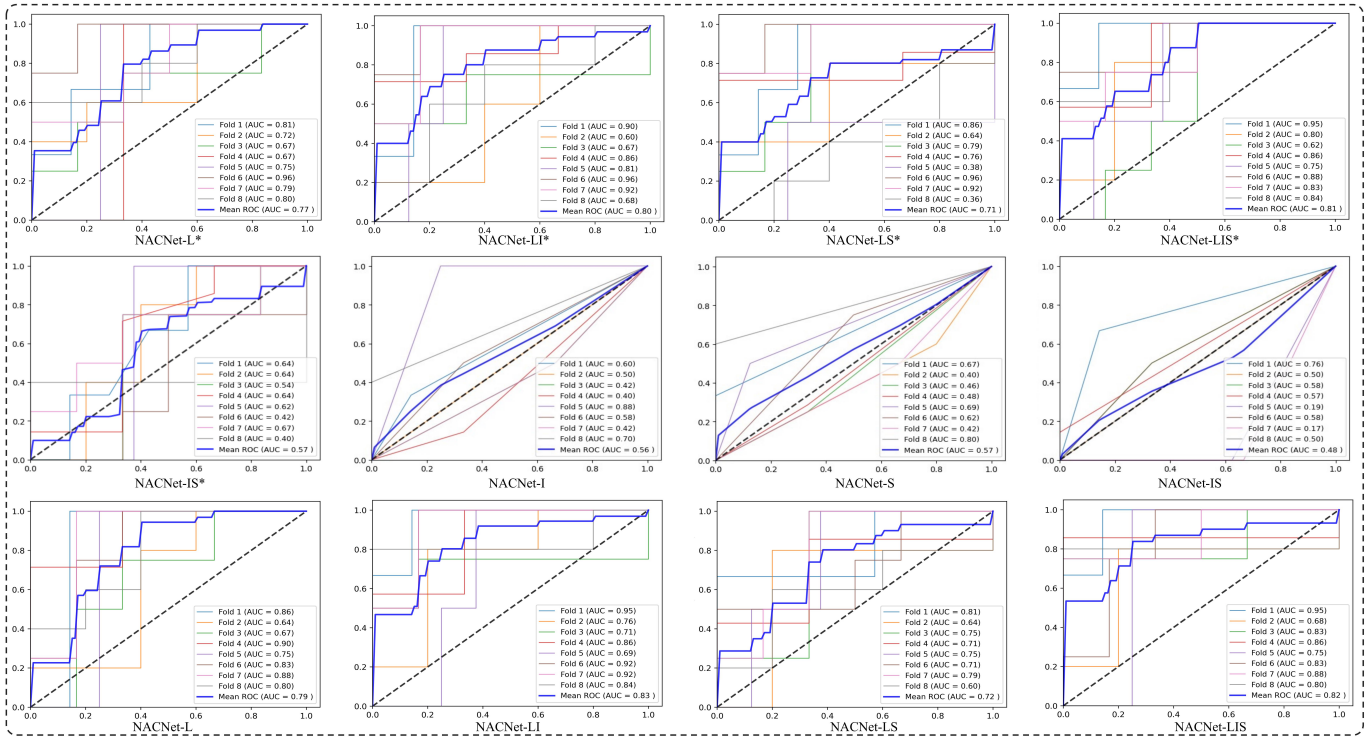
Figure 4: ROC curves of NACNet models in the ablation study. The ROC curves for the NACNet methods are generated for each cross-validation fold. ∗ denotes models without the GIN layer in the network.

Extra Trees [68], and Linear Discriminant Analysis (LDA) [69]. A majority voting mechanism is used to aggregate tile-level predictions into a final NAC response prediction at the patient level. Specifically, these models are trained independently using univariate feature selection [70] and mRMR feature ranking method. The univariate feature selection selects the best features based on univariate ANOVA F-test and is used to identify significant difference. The mRMR feature selection method aims to select features that are highly relevant to the target variable while minimizing redundancy among them. By balancing the trade-off between feature relevance and redundancy, this approach improves the model's predictive performance by focusing on the most informative and distinct features. Effective feature selection reduces data redundancy, minimizes dependency on noisy features in predicting NAC treatment response, mitigates overfitting, ultimately resulting in an improved model accuracy and robustness. The predictions are presented in Table 3 and Figure 5(b), showing that traditional machine learning methods work better than non-graph deep learning methods with the maximum voting strategy. Although feature selection methods improve traditional machine learning model performance, the NACNet model presents a superior performance to traditional machine learning methods overall. This improvement suggests that incorporating node image tile texture features and local histology labels significantly boosts the model's discrimination power for NAC response prediction. Additional results supporting the statistical comparison of treatment prediction accuracy are presented in Supplementary Table S2.

Table 2: Performance of deep learning models for NAC response prediction (Mean ± Standard deviation).

Models	ACC	AUC	Sensitivity	Specificity	Precision	F1
MEWISPool	0.61 ± 0.15	0.68 ± 0.19	0.62 ± 0.17	0.68 ± 0.12	0.70 ± 0.17	0.64 ± 0.12
DGCNN	0.64 ± 0.14	0.78 ± 0.14	0.62 ± 0.16	0.65 ± 0.12	0.74 ± 0.14	0.66 ± 0.12
SAG Global pooling	0.64 ± 0.11	0.77 ± 0.14	0.59 ± 0.10	0.68 ± 0.08	0.71 ± 0.09	0.65 ± 0.08
SAG Hierarchical pooling	0.63 ± 0.11	0.78 ± 0.14	0.64 ± 0.17	0.61 ± 0.10	0.69 ± 0.15	0.64 ± 0.10
GIN	0.64 ± 0.17	0.74 ± 0.15	0.66 ± 0.20	0.63 ± 0.15	0.71 ± 0.17	0.67 ± 0.15
Maximum voting-VGG16	0.48 ± 0.07	0.42 ± 0.20	0.52 ± 0.18	0.36 ± 0.15	0.64 ± 0.16	0.54 ± 0.15
Maximum voting-ResNet 18	0.52 ± 0.17	0.53 ± 0.18	0.60 ± 0.26	0.42 ± 0.20	0.60 ± 0.15	0.58 ± 0.20
Maximum voting - AlexNet	0.54 ± 0.11	0.52 ± 0.15	0.62 ± 0.20	0.39 ± 0.13	0.59 ± 0.12	0.59 ± 0.13
NACNet-LIS	**0.90 ± 0.07**	**0.82 ± 0.11**	**0.96 ± 0.08**	**0.88 ± 0.10**	**0.83 ± 0.15**	**0.88 ± 0.09**

## 4.4 Spatial TME analysis by graph structures

To enable a comprehensive understanding of the underlying biological mechanisms influencing NAC response, we compare the TME histology profiles between NAC treatment responder and non-responder group. The resulting group differences can manifest the impact of specific histology components on the NAC treatment outcome. Additionally, the distributions of the graph edges associated with different pairs of node histology labels are studied between these two patient populations.

TME histology spatial features are well-established as critical factors influencing tumor growth, progression, and metastasis, as cancer cells interact with surrounding stroma and inflammatory cells [71–73]. To investigate their impact on TNBCs, we compare histology profiles between pCR and RD patients, and identify specific histology patterns potentially associations with NAC response. As shown in Figure 6, the proportions of CIS (p=0.000), MVD (p=0.000), stroma (p=0.010), and adipose tissue (p=0.000) differ significantly between the two patient groups. Notably, all WSIs in this study have a similar effective tissue scale, allowing for direct comparison of WSI graph edges and subgraph structures of interest across patient groups.

Figure 7 shows histology pairs with significantly different mean graph edge counts between the two patient groups, based on t-test results. Edges associated with immune-tumor cells (p=0.046), immune-adipose (p=0.031), necrosis-tumor (p=0.013), and necrosis-adipose (p=0.0022) interactions are significantly more enriched in pCR group than in the RD group. Conversely, MVD-stroma (p=0.000) and stroma-adipose (p=0.049) edges are significantly more prevalent in RD patients.

Next, we investigate the impact of 3-node subgraph quantification on NAC response, as shown in Figure 8. Subgraphs involving adipose-stroma-MVD (p=0.007) and MVD-stroma-adipose (p=0.014) are more enriched in RD patients. In contrast, subgraphs containing adipose-tumor-necrosis (p=0.038), adipose-immune cells-tumor (p=0.000), and adipose-MVD-necrosis (p=0.046) are significantly more enriched in the pCR group.

We also generate histology label correlation maps, illustrating the Pearson correlation coefficients between histology label count and NAC response outcome, as well as histology label relevancy importance maps derived from mRMR analysis. These maps offer valuable insights into the decision-making process of the prediction model (Figure 9).

14
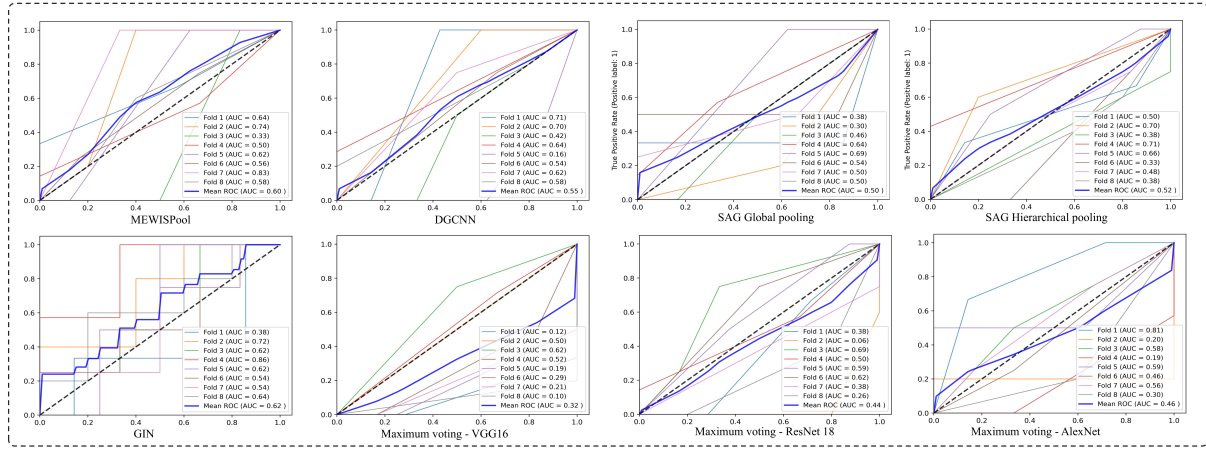
(a)
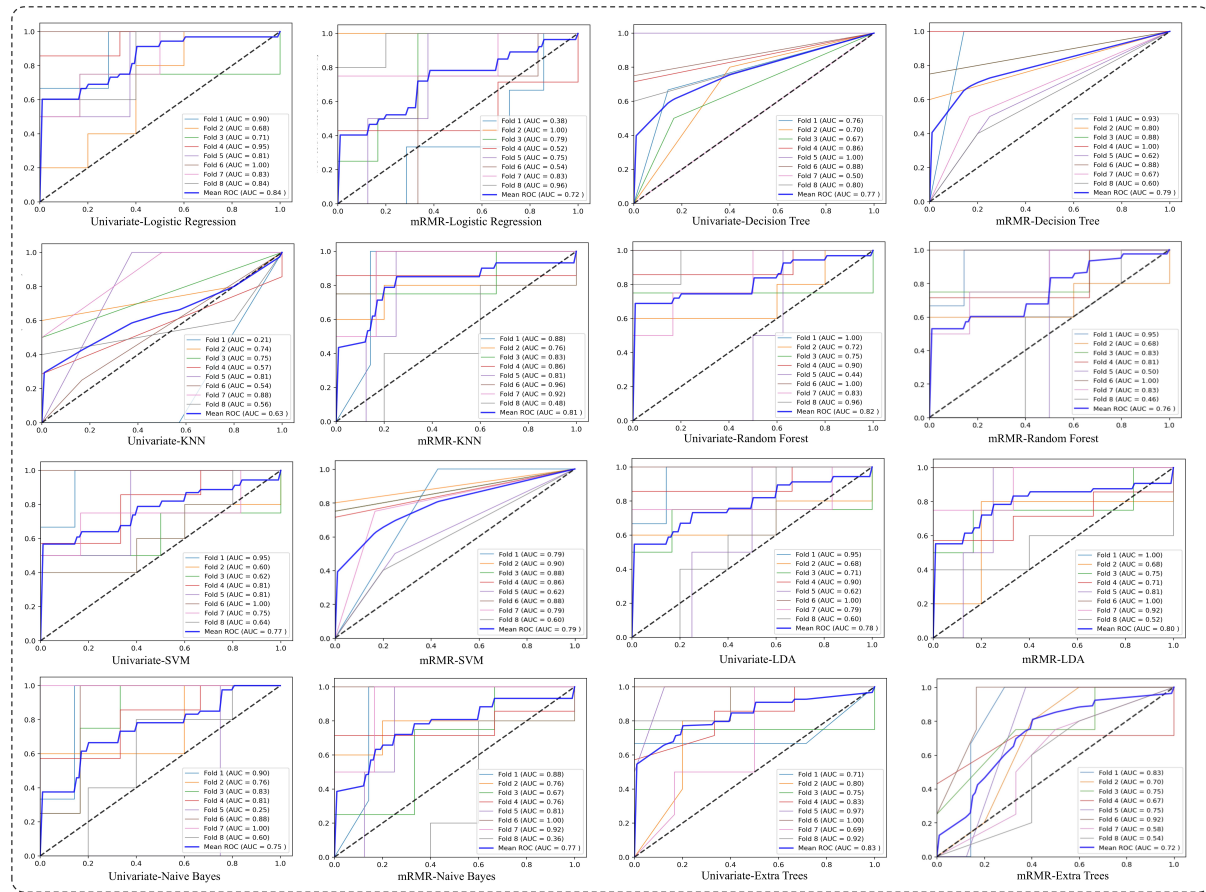
(b)

Figure 5: The ROC curves of methods for comparison. The ROC curves of (a) deep learning methods and (b) traditional machine learning methods.

Furthermore, we visualize in Figure 10(a) spatial attention that targets specific subgraphs used by the prediction model. The subgraph density is represented with iso-contours. In Figure 10(b), the corresponding tissue regions are overlaid with color coded histology labels by our model. Our pipeline unveils the intricate

15

Table 3: NAC treatment response prediction of machine learning models performance (Mean ± Standard deviation).

Models	ACC	AUC	Sensitivity	Specificity	Precision	F1
Univariate-LR	0.75 ± 0.11	0.84 ± 0.10	0.48 ± 0.23	0.96 ± 0.07	0.93 ± 0.12	0.59 ± 0.17
mRMR-LR	0.73 ± 0.15	0.72 ± 0.21	0.70 ± 0.32	0.70 ± 0.15	0.59 ± 0.27	0.63 ± 0.29
Univariate-KNN	0.71 ± 0.11	0.63 ± 0.20	0.29 ± 0.24	**1.00 ± 0.00**	0.63 ± 0.48	0.39 ± 0.31
mRMR-KNN	0.71 ± 0.12	0.81 ± 0.14	0.68 ± 0.18	0.75 ± 0.16	0.65 ± 0.19	0.65 ± 0.15
Univariate-SVM	0.73 ± 0.16	0.77 ± 0.14	0.48 ± 0.25	0.94 ± 0.08	0.85 ± 0.19	0.57 ± 0.21
mRMR-SVM	0.79 ± 0.11	0.79 ± 0.11	0.71 ± 0.17	0.87 ± 0.15	0.78 ± 0.25	0.72 ± 0.17
Univariate-Naive Bayes	0.76 ± 0.13	0.75 ± 0.22	0.81 ± 0.32	0.68 ± 0.17	0.63 ± 0.26	0.70 ± 0.28
mRMR-Naive Bayes	0.73 ± 0.15	0.77 ± 0.18	0.80 ± 0.18	0.66 ± 0.18	0.64 ± 0.17	0.70 ± 0.17
Univariate-Decision Tree	0.78 ± 0.14	0.77 ± 0.14	0.69 ± 0.16	0.85 ± 0.19	0.80 ± 0.22	0.73 ± 0.16
mRMR-Decision Tree	0.81 ± 0.13	0.80 ± 0.14	0.69 ± 0.21	0.91 ± 0.10	**0.90 ± 0.13**	0.72 ± 0.20
Univariate-Random Forest	0.79 ± 0.14	0.83 ± 0.18	0.61 ± 0.32	0.92 ± 0.17	0.78 ± 0.34	0.65 ± 0.30
mRMR-Random Forest	0.74 ± 0.13	0.76 ± 0.18	0.52 ± 0.25	0.90 ± 0.13	0.74 ± 0.33	0.59 ± 0.26
Univariate-LDA	0.74 ± 0.14	0.77 ± 0.14	0.61 ± 0.28	0.82 ± 0.18	0.68 ± 0.33	0.63 ± 0.28
mRMR-LDA	0.80 ± 0.10	0.80 ± 0.16	0.72 ± 0.21	0.88 ± 0.11	0.80 ± 0.18	0.73 ± 0.15
Univariate-Extra Trees	0.61 ± 0.15	0.83 ± 0.11	0.47 ± 0.25	0.95 ± 0.08	0.79 ± 0.34	0.58 ± 0.27
mRMR-Extra Trees	0.76 ± 0.21	0.72 ± 0.12	0.61 ± 0.15	0.84 ± 0.12	0.61 ± 0.15	0.61 ± 0.15
ReliefF-1NN	0.73 ± 0.07	0.74 ± 0.10	0.95 ± 0.09	0.53 ± 0.18	0.60 ± 0.09	0.73 ± 0.07
ReliefF ensembleTree	0.70 ± 0.11	0.69 ± 0.10	0.70 ± 0.20	0.53 ± 0.18	0.58 ± 0.12	0.68 ± 0.14
ReliefF rbf-SVM	0.75 ± 0.11	0.62 ± 0.10	0.48 ± 0.23	0.47 ± 0.23	0.53 ± 0.11	0.62 ± 0.12
ReliefF lin-SVM	0.73 ± 0.15	**0.86 ± 0.07**	0.70 ± 0.32	0.74 ± 0.13	0.71 ± 0.16	0.82 ± 0.11
NACNet-LIS	**0.90 ± 0.07**	0.82 ± 0.11	**0.96 ± 0.08**	0.88 ± 0.10	0.83 ± 0.15	**0.88 ± 0.09**

role of the TNBC TME by exploring spatial TME patterns, extracting local tissue texture information, and identifying potential biomarkers for NAC response prediction. Therefore, it is promising to enable an improved prognosis and treatment planning in practice, enhancing TNBC patient therapy and life quality. While our approach is developed for NAC treatment response prediction in TNBCs, it can be readily extended to other cancer types, contributing to a more comprehensive TME understanding.

5   Discussion

Deep learning has greatly advanced computational pathology over the past decades. However, due to the vast scale of WSIs, much of the work has focused on decision support by processing small WSI tiles independently, overlooking crucial spatial connectivity within tissue. Our approach addresses this limiation by integrating global tissue graph structure signatures with local tissue texture and histology label information to enhance TNBC response prediction.

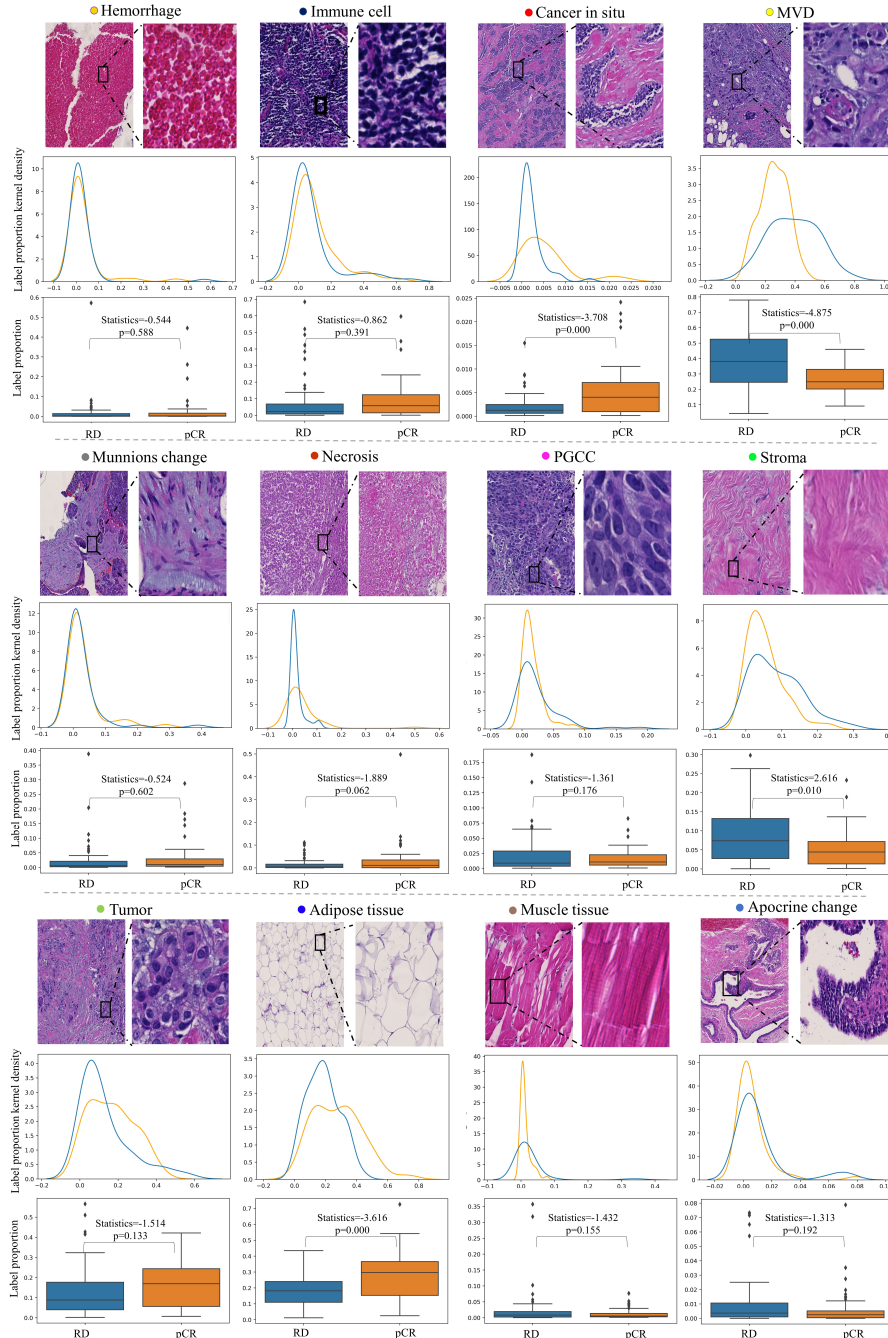Our contribution in this study can be summarized as follows.

Figure 6: Comparison of histology profiles between pCR and RD patients. We plot the individual histology label proportion distributions for both pCR (orange) and RD (blue) TNBC patients. A two-sample t-test is used to assess if individual histology label proportion means are significantly different between these two patient groups. Box plots present quartile values and whiskers extended to data points within the 1.53 interquartile range.

1. We characterize histology features and their spatial interactions using TME graphs within a multi-step framework. First, we assign the histology label to each image tile using a deep learning module trained on $l = 12$ histology classes. Next, we spatially assemble these labeled tiles based on their spatial arrangement,
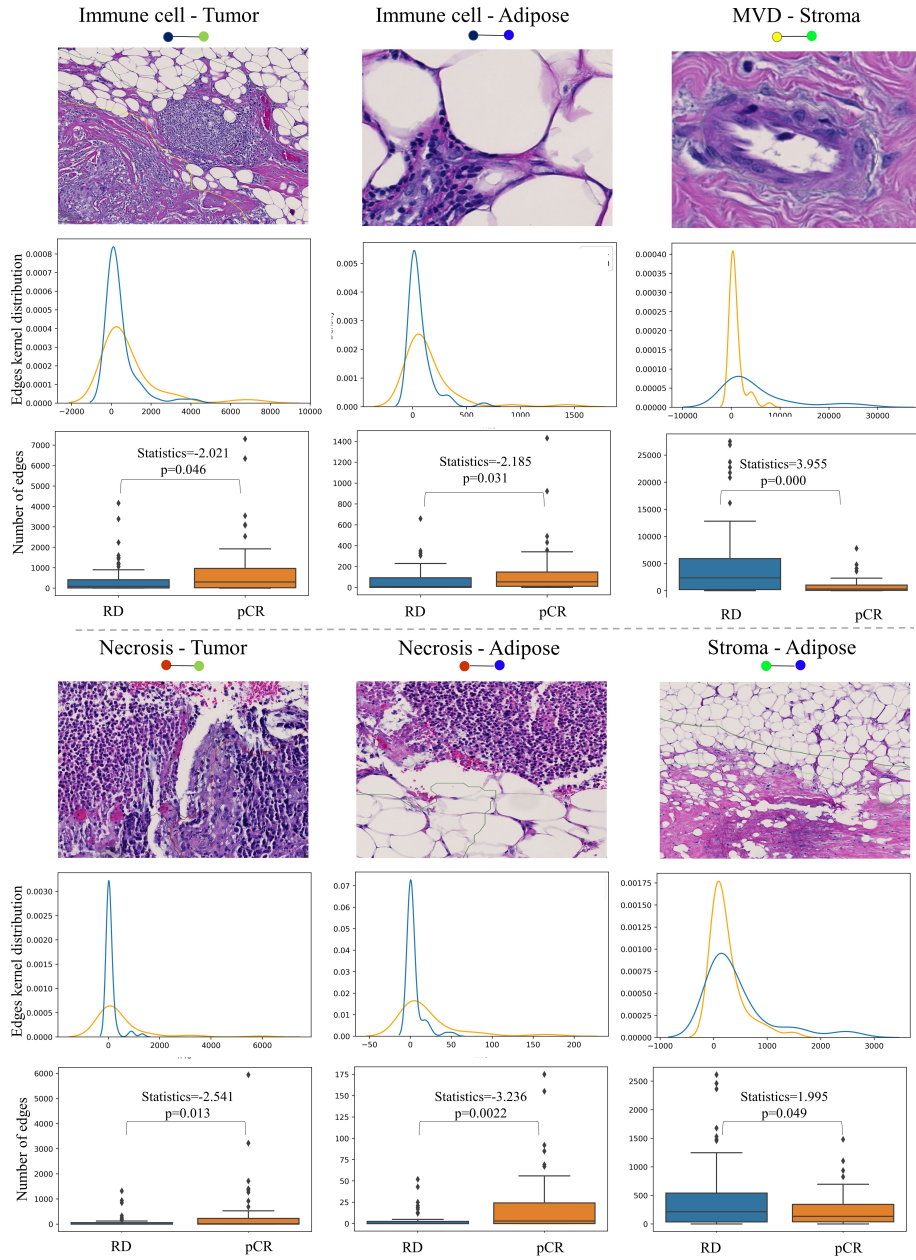
Figure 7: Comparison of WSI-derived graph edge profiles between pCR and RD patients. As we include 12 histology labels in this study, there are 66 edge types associated with different histology label pairs. For each type, the edge counts were compared between pCR and RD patients. For each significant edge type, a typical tissue region associated with the significant edge type is illustrated. Additionally, we plot the individual graph edge distributions for both pCR (orange) and RD (blue) TNBC patients. A two-sample t-test is used to assess if individual graph edge means are significantly different between these two patient groups. Box plots present quartile values and whiskers extended to data points within the 1.53 interquartile range.

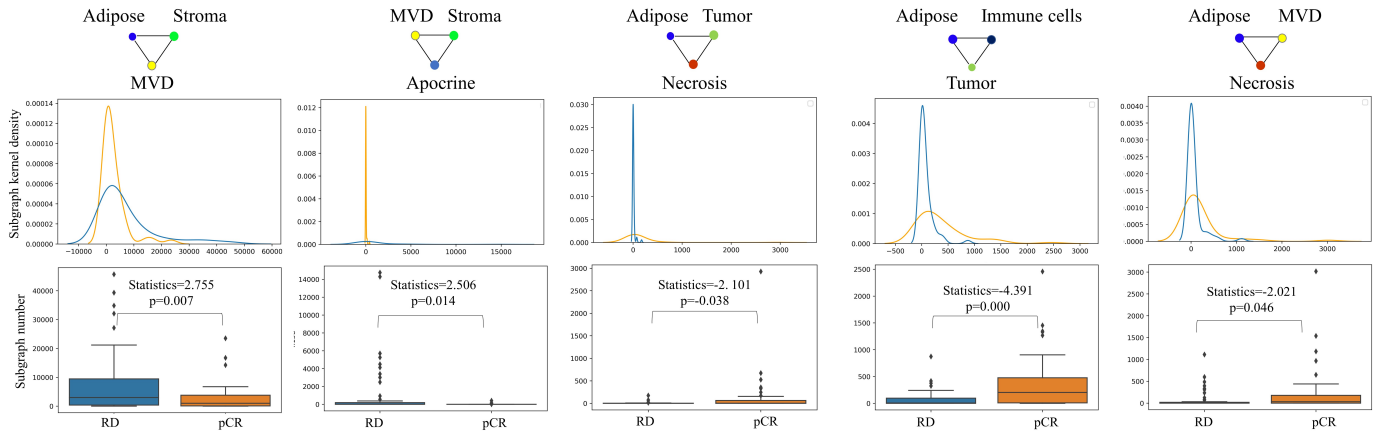generating a tile-level histology label map for each WSI. Finally, we construct a spatial TME graph to

Figure 8: Comparison of WSI-derived subgraph profiles between pCR and RD patients. As we include 12 histology labels in this study, there are 220 3-node subgraph types associated with different histology label triplets. We plot the individual 3-node subgraph distributions for both pCR (orange) and RD (blue) TNBC patients. A two-sample t-test is used to assess if individual 3-node subgraph means are significantly different between these two patient groups. Box plots present quartile values and whiskers extended to data points within the 1.53 interquartile range.

capture the spatial histology organization critical for NAC response prediction. In this graph, each node represents the tile closest to the center of a cluster of tiles sharing the same histology label, while edges connect adjacent nodes based on spatial proximity.

2. To capture predictive features from both local tissue context and global tissue structure organization, we construct two complementary sets of features for graph node representation. First, an autoencoder is used to extract image texture features from tiles associated with each node. Second, social network analysis (SNA) features are computed to enhance node representation. This two-step approach integrates detailed local tissue information with spatial TME arrangement, providing a comprehensive foundation for accurate NAC response prediction.

3. NACNet is a transformer based GCN model designed to enhance the discriminating power of the graph-based prediction. To achieve this, we incorporate GIN layers into the graph transformer network. This layers transform heterogeneous input graphs into useful meta-path graphs and enable end-to-end learning of node representations. By including the GIN layers, NACNet effectively captures the structure of WSI-derived TME graph, leading to improved prediction performance.

NACNet integrates histology label information, local tissue textures, and graph-derived SNA features to differentiate pCR from RD TNBC patients. By multiple evaluation metrics, NACNet achieves a better overall performance than state-of-the-art graph based deep learning architectures, non-graph deep learning models, and traditional machine learning approaches. NACNet efficiently captures local and WSI-level information to predict NAC response for TNBC patients. It processes high-resolution WSIs, predict multiple-class histology labels for local tissue regions, characterize TME spatial histology interactions, predict NAC response outcome, and generate interpretable findings associated with prediction.
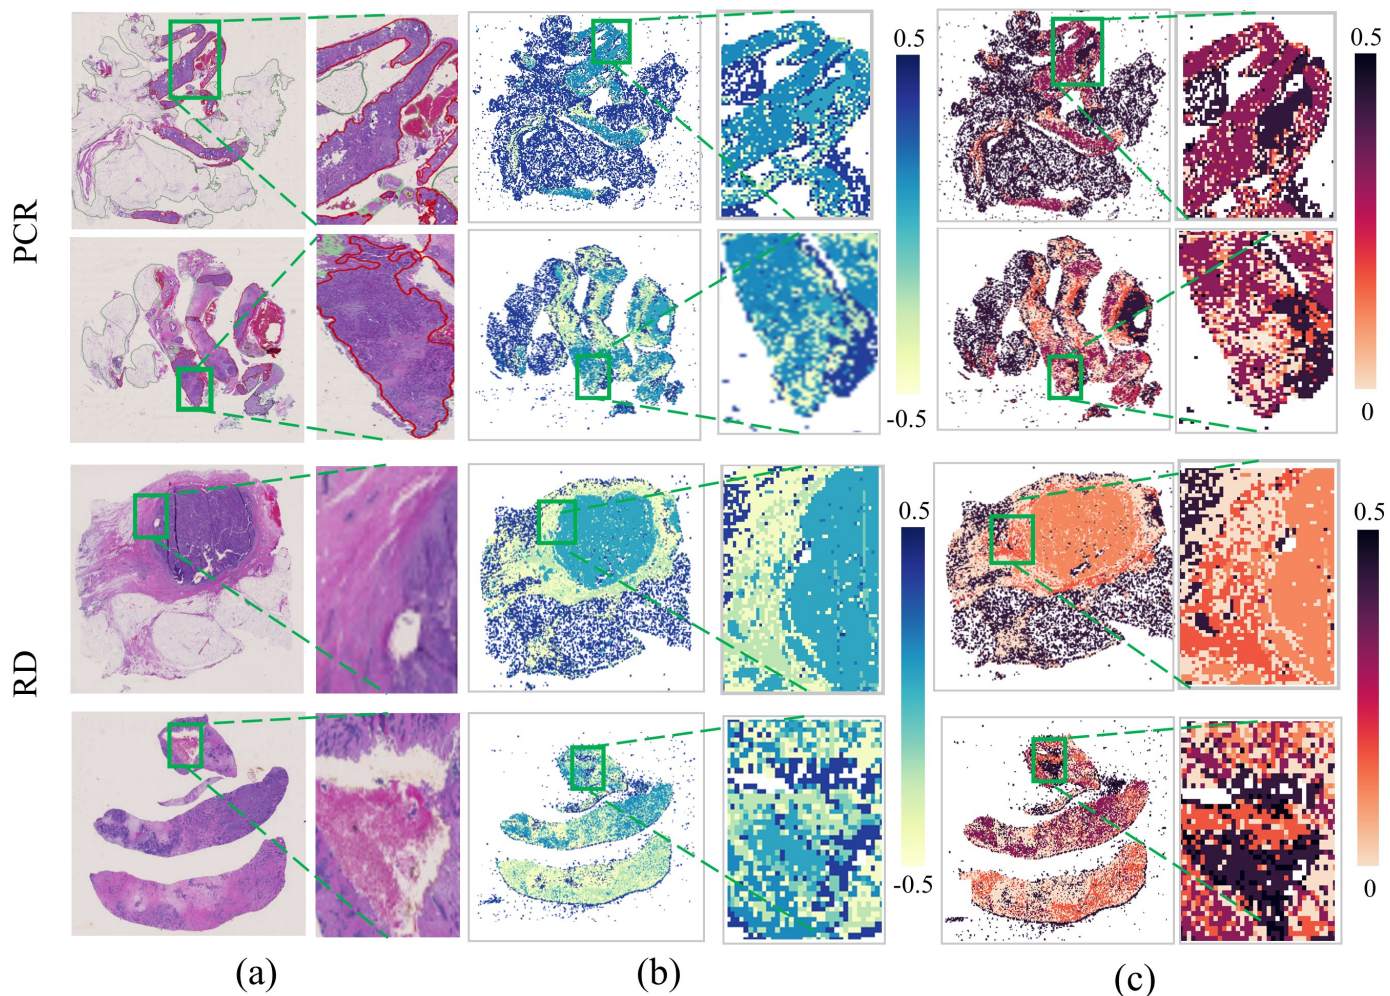
19

Figure 9: Prediction model interpretations with correlation and relevancy importance maps. (a) Typical WSIs with their local tissue regions from pCR and RD group; (b) The histology label correlation map is produced by mapping to the histology map correlation coefficients by the Pearson correlation. The correlation value ranges from -1 to 1 and indicates the strength and direction of the relationship between individual histology labels and prediction outcomes; (c) The importance of graph feature nodes is analyzed using the mRMR (minimum Redundancy Maximum Relevance) method. The histology node importance map is generated by mapping to the histology map the importance values by the mRMR method.

Additionally, our work provides a straightforward way to enable the graph-based TME analysis. It can highlight the spatial histology patterns and the WSI regions that are significantly associated with the NAC response outcome. Unlike other graph-based classification models, NACNet is able to produce the TME attention over WSIs. This is useful as subgraphs and edges connecting different histology regions of interest in the neighborhood can contain biologically interpretable information that is critical for prediction. We perform detailed subgraph analysis and identify edges and three-node cliques of interpretation and prediction value. Such analyses enable biological interpretations especially important in histopathology WSI analysis applications where the detailed characteristics of tissue spatial TME can provide valuable insight on disease progression and
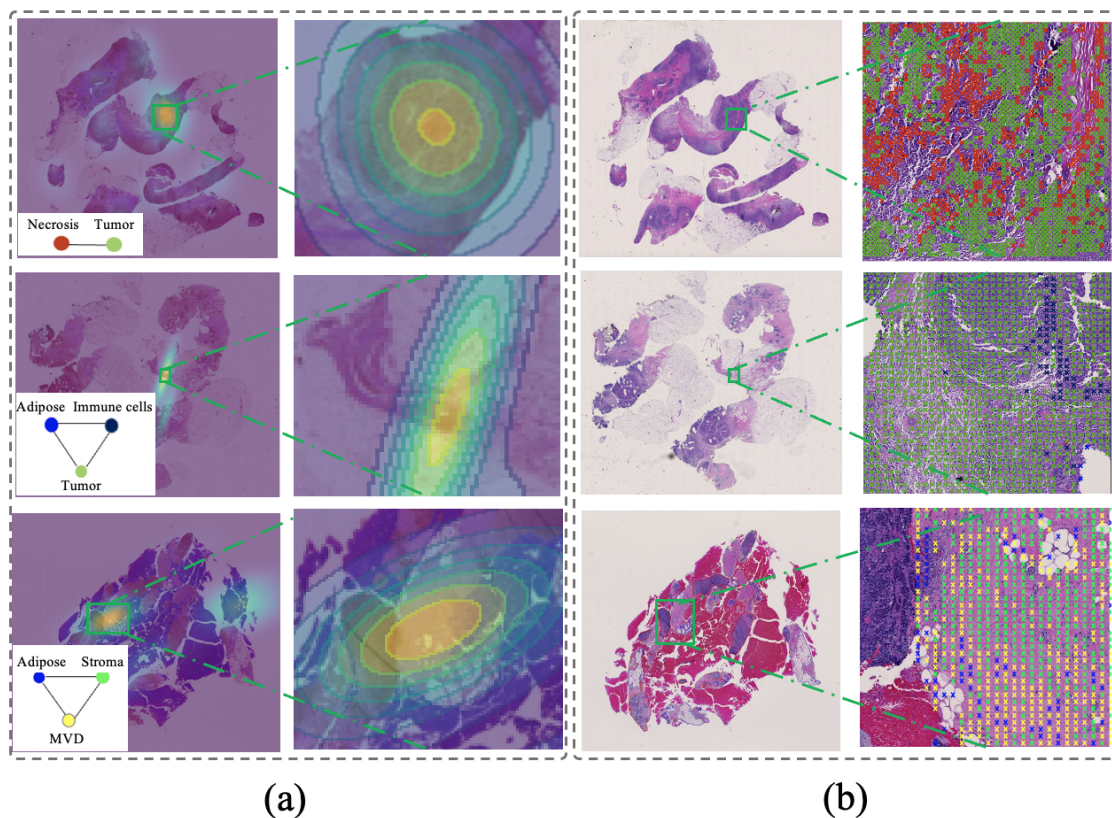
Figure 10: Attention visualization for typical tissue regions. (a) We visualize spatial attention that targets specific subgraphs supporting the model prediction. The subgraph density is represented by iso-contours; (b) The histology labels of corresponding regions of interest are color coded by the same color schema in Figure 3.

treatment response. Moreover, NACNet can be extended to explore the microenvironment reaction analysis with multiple subgraphs in future.

The graph node awareness mechanism can enhance the efficacy of the graph-based WSI analysis. Note the NACNet-LIS model leverages node information, local textures, and SNA graph features, significantly improving the WSI-derived tissue graph representation and contextual understanding. With node features and the contextual node relationships, the NACNet-LIS with intensive graph node awareness allows for a comprehensive understanding of the tissue TME graph structure and properties and achieves a more nuanced understanding of the structure and properties of the graph. Such node awareness mechanism is particularly important in histopathology analysis, where subtle differences in tissue structure can have significant implications for diagnosis and treatment planning. Through ablation experiments, we have demonstrated the necessity of a comprehensive graph node characterization.

Our work presents several avenues for improvement. First, the reliance on detailed image annotations for histology prediction makes the deep learning model labor-intensive. Second, the clustering of discrete points with different labels to form graph nodes presents challenges. To improve node formation, we used a clinically-weighted clustering strategy that considers both the physical size differences and clinical significance of 12 histological categories in breast cancer. By assigning lower thresholds (η) to critical categories like tumor and

necrosis, we ensured these features are prominently represented within the TME graph. For categories of lesser clinical importance, we used a higher threshold (η), reducing unnecessary segmentation while still maintaining relevant context. This approach not only minimizes misclassification impact but also enhances the model's ability to capture breast cancer's pathological characteristics, supporting greater generalizability and clinical interpretability. Future experiments will further validate the effectiveness of this adjusted strategy, ensuring that the resulting TME graph structure holds higher biological significance and diagnostic value in breast cancer applications. Finally, increasing the sample size in future studies will be crucial to further validate the generalizability of our findings.

6 Conclusion

We present a transformer-based GCN model, NACNet, that predicts the NAC treatment response for TNBCs using WSIs of patient biopsy tissues. The model integrates local tissue histology information, tissue textures, and SNA features derived from the spatial TME graph to enhance NAC response prediction. The spatial attention mechanism and SNA features direct transformer GCN attention to tissue regions of prognostic significance, incorporating human prior knowledge. When applied to a TNBC patient dataset, our system achieves 90% prediction accuracy. In addition, we demonstrate the efficacy of an integrative use of node content-awareness features, SNA information, and the GIN spatial attention mechanism on prediction enhancement. Our study also suggests the spatial TME histology patterns carry prognostic values for NAC response prediction for TNBC patients. Our graph-based analysis of histological variables for TME characterization and NAC response prediction is generic and can be tailored to enable treatment response prediction and personalized treatment planning for other cancer types.

Conflict of Interest Statement

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Funding

7 Supplement

Supplementary Tables S1–2.

References

[1] P Boyle. Triple-negative breast cancer: epidemiological considerations and recommendations. Annals of oncology, 23:vi7–vi12, 2012.

[2] Ryuichiro Hirano, Koki Okamoto, Miyu Shinke, Marika Sato, Shigeaki Watanabe, Hitomi Watanabe, Gen Kondoh, Tetsuya Kadonosono, and Shinae Kizaka-Kondoh. Tissue-resident macrophages are major tumor-associated macrophage resources, contributing to early tnbc development, recurrence, and metastases. Communications Biology, 6(1):144, 2023.

[3] Lei-Jie Dai, Ding Ma, Yu-Zheng Xu, Ming Li, Yu-Wei Li, Yi Xiao, Xi Jin, Song-Yang Wu, Ya-Xin Zhao, Han Wang, et al. Molecular features and clinical implications of the heterogeneity in chinese patients with her2-low breast cancer. Nature Communications, 14(1):5112, 2023.

[4] Joe Mehanna, Fady GH Haddad, Roland Eid, Matteo Lambertini, and Hampig Raphael Kourie. Triple-negative breast cancer: current perspective on the evolving therapeutic landscape. International journal of women's health, 11:431, 2019.

[5] Guillermo Arturo Valencia, Patricia Rioja, Zaida Morante, Rossana Ruiz, Hugo Fuentes, Carlos A Castaneda, Tatiana Vidaurre, Silvia Neciosup, and Henry L Gomez. Immunotherapy in triple-negative breast cancer: A literature review and new advances. World Journal of Clinical Oncology, 13(3):219, 2022.

[6] Giampaolo Bianchini, Justin M Balko, Ingrid A Mayer, Melinda E Sanders, and Luca Gianni. Triple-negative breast cancer: challenges and opportunities of a heterogeneous disease. Nature reviews Clinical oncology, 13(11):674–690, 2016.

[7] Patricia Cortazar, Lijun Zhang, Michael Untch, Keyur Mehta, Joseph P Costantino, Norman Wolmark, Hervé Bonnefoi, David Cameron, Luca Gianni, Pinuccia Valagussa, et al. Pathological complete response and long-term clinical benefit in breast cancer: the ctneobc pooled analysis. The Lancet, 384(9938):164–172, 2014.

[8] Angela Pennisi, Thomas Kieber-Emmons, Issam Makhoul, and Laura Hutchins. Relevance of pathological complete response after neoadjuvant therapy for breast cancer. Breast cancer: basic and clinical research, 10:BCBCR–S33163, 2016.

[9] Gianluca Franceschini, Alba Di Leone, Maria Natale, Martin Aleandro Sanchez, and Riccardo Masett. Conservative surgery after neoadjuvant chemotherapy in patients with operable breast cancer. Annali Italiani di Chirurgia, 89:290–290, 2018.

[10] Kaori Sakuma, Masafumi Kurosumi, Hanako Oba, Yasuhito Kobayashi, Hiroyuki Takei, Kenichi Inoue, Toshio Tabei, and Tetsunari Oyama. Pathological tumor response to neoadjuvant chemotherapy using anthracycline and taxanes in patients with triple-negative breast cancer. Experimental and therapeutic medicine, 2(2):257–264, 2011.

[11] Xiyue Wang, Junhan Zhao, Eliana Marostica, Wei Yuan, Jietian Jin, Jiayu Zhang, Ruijiang Li, Hongping Tang, Kanran Wang, Yu Li, et al. A pathology foundation model for cancer diagnosis and prognosis prediction. Nature, pages 1–9, 2024.

[12] Byungsoo Ahn, Damin Moon, Hyun-Soo Kim, Chung Lee, Nam Hoon Cho, Heung-Kook Choi, Dong-min Kim, Jung-Yun Lee, Eun Ji Nam, Dongju Won, et al. Histopathologic image–based deep learning classifier for predicting platinum-based treatment responses in high-grade serous ovarian cancer. Nature Communications, 15(1):4253, 2024.

[13] Lin Lu, Laurent Dercle, Binsheng Zhao, and Lawrence H Schwartz. Deep learning for the prediction of early on-treatment response in metastatic colorectal cancer from serial medical imaging. Nature communications, 12(1):6654, 2021.

[14] Danh-Tai Hoang, Gal Dinstag, Eldad D Shulman, Leandro C Hermida, Doreen S Ben-Zvi, Efrat Elis, Katherine Caley, Stephen-John Sammut, Sanju Sinha, Neelam Sinha, et al. A deep-learning framework to predict cancer treatment response from histopathology images through imputed transcriptomics. Nature Cancer, pages 1–13, 2024.

[15] Hyo-jae Lee, Jeong Hoon Lee, Jong Eun Lee, Yong Min Na, Min Ho Park, Ji Shin Lee, and Hyo Soon Lim. Prediction of early clinical response to neoadjuvant chemotherapy in triple-negative breast cancer: Incorporating radiomics through breast mri. Scientific Reports, 14(1):21691, 2024.

[16] Intae Moon, Jaclyn LoPiccolo, Sylvan C Baca, Lynette M Sholl, Kenneth L Kehl, Michael J Hassett, David Liu, Deborah Schrag, and Alexander Gusev. Machine learning for genetics-based classification and treatment response prediction in cancer of unknown primary. Nature Medicine, 29(8):2057–2067, 2023.

[17] Neeta Kumar, Ruchika Gupta, and Sanjay Gupta. Whole slide imaging (wsi) in pathology: current perspectives and future directions. Journal of digital imaging, 33(4):1034–1040, 2020.

[18] Neofytos Dimitriou, Ognjen Arandjelović, and Peter D Caie. Deep learning for whole slide image analysis: an overview. Frontiers in medicine, 6:264, 2019.

[19] Manuel Stritt, Anna K Stalder, and Enrico Vezzali. Orbit image analysis: an open-source whole slide image analysis tool. PLoS computational biology, 16(2):e1007313, 2020.

[20] Hongyi Duanmu, Pauline Boning Huang, Srinidhi Brahmavar, Stephanie Lin, Thomas Ren, Jun Kong, Fusheng Wang, and Tim Q Duong. Prediction of pathological complete response to neoadjuvant chemotherapy in breast cancer using deep learning with integrative imaging, molecular and demographic data. pages 242–252, 2020.

[21] Han-Byoel Lee, Wonshik Han, Soo-Yeon Kim, Nariya Cho, Kyoung-Eun Kim, Jung Hyun Park, Young Wook Ju, Eun-Shin Lee, Sung-Joon Lim, Jung Ho Kim, et al. Prediction of pathologic complete response using image-guided biopsy after neoadjuvant chemotherapy in breast cancer patients selected based on mri findings: a prospective feasibility trial. Breast Cancer Research and Treatment, 182(1):97–105, 2020.

[22] Hee Jung Shin, Hyeon-Man Baek, Jin-Hee Ahn, Seunghee Baek, Hyunji Kim, Joo Hee Cha, and Hak Hee Kim. Prediction of pathologic response to neoadjuvant chemotherapy in patients with breast cancer using diffusion-weighted imaging and mrs. NMR in Biomedicine, 25(12):1349–1359, 2012.

[23] Yu-Hong Qu, Hai-Tao Zhu, Kun Cao, Xiao-Ting Li, Meng Ye, and Ying-Shi Sun. Prediction of pathological complete response to neoadjuvant chemotherapy in breast cancer using a deep learning (dl) method. Thoracic Cancer, 11(3):651–658, 2020.

[24] Zihao Chen, Maoli Wang, Rudy Leon De Wilde, Ruifa Feng, Mingqiang Su, Luz Angela Torres-de la Roche, and Wenjie Shi. A machine learning model to predict the triple negative breast cancer immune subtype. Frontiers in immunology, 12, 2021.

[25] Yi Zheng, Rushin H Gindra, Emily J Green, Eric J Burks, Margrit Betke, Jennifer E Beane, and Vijaya B Kolachalama. A graph-transformer for whole slide image classification. IEEE transactions on medical imaging, 41(11):3003–3015, 2022.

[26] Stuart Andrews, Thomas Hofmann, and Ioannis Tsochantaridis. Multiple instance learning with generalized support vector machines. In AAAI/IAAI, pages 943–944, 2002.

[27] Richard J Chen, Ming Y Lu, Wei-Hung Weng, Tiffany Y Chen, Drew FK Williamson, Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 4015–4025, 2021.

[28] Angel Cruz-Roa, Ajay Basavanhally, Fabio González, Hannah Gilmore, Michael Feldman, Shridar Ganesan, Natalie Shih, John Tomaszewski, and Anant Madabhushi. Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks. In Medical Imaging 2014: Digital Pathology, volume 9041, page 904103. SPIE, 2014.

[29] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.

[30] Minh Hoai Nguyen, Lorenzo Torresani, Fernando De La Torre, and Carsten Rother. Weakly supervised discriminative localization and classification: a joint learning process. In 2009 IEEE 12th International Conference on Computer Vision, pages 1925–1932. IEEE, 2009.

[31] Giampaolo Bianchini, Carmine De Angelis, Luca Licata, and Luca Gianni. Treatment landscape of triple-negative breast cancer—expanded options, evolving needs. Nature reviews Clinical oncology, 19(2):91–113, 2022.

[32] Leeat Keren, Marc Bosse, Diana Marquez, Roshan Angoshtari, Samir Jain, Sushama Varma, Soo-Ryum Yang, Allison Kurian, David Van Valen, Robert West, et al. A structured tumor-immune microenvironment in triple negative breast cancer revealed by multiplexed ion beam imaging. Cell, 174(6):1373–1387, 2018.

[33] Nils Lycke, Mats Bemark, and Jo Spencer. Mucosal b cell differentiation and regulation. pages 701–719, 2015.

[34] Donald J Coluzzi. An overview of lasers in dentistry. The Alpha omegan, 101(3):125–126, 2008.

[35] Ruoyu Li, Jiawen Yao, Xinliang Zhu, Yeqing Li, and Junzhou Huang. Graph cnn for survival analysis on whole slide pathological images. pages 174–182, 2018.

[36] Joshua Levy, Christian Haudenschild, Clark Barwick, Brock Christensen, and Louis Vaickus. Topological feature extraction and visualization of whole slide images using graph neural networks. pages 285–296, 2020.

[37] Richard J Chen, Ming Y Lu, Muhammad Shaban, Chengkuan Chen, Tiffany Y Chen, Drew FK Williamson, and Faisal Mahmood. Whole slide images are 2d point clouds: Context-aware survival prediction using patch-based graph convolutional networks. pages 339–349, 2021.

[38] William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584, 2017.

[39] Wenqi Lu, Michael Toss, Muhammad Dawood, Emad Rakha, Nasir Rajpoot, and Fayyaz Minhas. Slidegraph+: Whole slide image level graphs to predict her2 status in breast cancer. Medical Image Analysis, 80:102486, 2022.

[40] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? arXiv preprint arXiv:1810.00826, 2018.

[41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.

[42] Heechan Yang, Ji-Ye Kim, Hyongsuk Kim, and Shyam P Adhikari. Guided soft attention network for classification of breast cancer histopathology images. IEEE transactions on medical imaging, 39(5):1306–1315, 2019.

[43] Douglas Hanahan and Robert A Weinberg. Hallmarks of cancer: the next generation. cell, 144(5):646–674, 2011.

[44] Carsten Denkert, Sibylle Loibl, Aurelia Noske, Marc Roller, Berit Maria Müller, Martina Komor, Jan Budczies, Silvia Darb-Esfahani, Ralf Kronenwett, Claus Hanusch, et al. Tumor-associated lymphocytes as an independent predictor of response to neoadjuvant chemotherapy in breast cancer. Journal of clinical oncology, 28(1):105–113, 2010.

[45] Judah Folkman. Role of angiogenesis in tumor growth and metastasis. In Seminars in oncology, volume 29, pages 15–18. Elsevier, 2002.

[46] Callum Christopher Mackenzie, Muhammad Dawood, Simon Graham, Mark Eastwood, et al. Neural graph modelling of whole slide images for survival ranking. In Learning on Graphs Conference, pages 48–1. PMLR, 2022.

[47] Evelien Otte and Ronald Rousseau. Social network analysis: a powerful strategy, also for the information sciences. Journal of information Science, 28(6):441–453, 2002.

[48] Ian Rogers. The google pagerank algorithm and how it works. 2002.

[49] Cheng-Yuan Liou, Wei-Chen Cheng, Jiun-Wei Liou, and Daw-Ran Liou. Autoencoder for words. Neurocomputing, 139:84–96, 2014.

[50] Dai Quoc Nguyen, Tu Dinh Nguyen, and Dinh Phung. Universal graph transformer self-attention networks. In Companion Proceedings of the Web Conference 2022, pages 193–196, 2022.

[51] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907, 2016.

[52] Pinar Yanardag and SVN Vishwanathan. Deep graph kernels. pages 1365–1374, 2015.

[53] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. pages 770–778, 2016.

[54] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. arXiv preprint arXiv:1607.06450, 2016.

[55] Chris Ding and Hanchuan Peng. Minimum redundancy feature selection from microarray gene expression data. Journal of bioinformatics and computational biology, 3(02):185–205, 2005.

[56] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.

[57] Amirhossein Nouranizadeh, Mohammadjavad Matinkia, Mohammad Rahmati, and Reza Safabakhsh. Maximum entropy weighted independent set pooling for graph neural networks. arXiv preprint arXiv:2107.01410, 2021.

[58] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. 32(1), 2018.

[59] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. In International conference on machine learning, pages 3734–3743. PMLR, 2019.

[60] Zhen Zhang, Jiajun Bu, Martin Ester, Jianfeng Zhang, Chengwei Yao, Zhi Yu, and Can Wang. Hierarchical graph pooling with structure learning. arXiv preprint arXiv:1911.05954, 2019.

[61] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25, 2012.

[62] David W Hosmer. Applied logistic regression, 2. 2000.

[63] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. IEEE transactions on information theory, 13(1):21–27, 1967.

[64] Corinna Cortes and Vladimir Vapnik. Support-vector networks. Machine learning, 20:273–297, 1995.

[65] Harry Zhang. The optimality of naive bayes. Aa, 1(2):3, 2004.

[66] J. Ross Quinlan. Induction of decision trees. Machine learning, 1:81–106, 1986.

[67] Leo Breiman. Rd. Machine learning, 45:5–32, 2001.

[68] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. Machine learning, 63:3–42, 2006.

[69] Suresh Balakrishnama and Aravind Ganapathiraju. Linear discriminant analysis-a brief tutorial. Institute for Signal and information Processing, 18(1998):1–8, 1998.

[70] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. the Journal of machine Learning research, 12:2825–2830, 2011.

[71] Florian R Greten and Sergei I Grivennikov. Inflammation and cancer: triggers, mechanisms, and consequences. Immunity, 51(1):27–41, 2019.

[72] Yiping Jiao, Junhong Li, Chenqi Qian, and Shumin Fei. Deep learning-based tumor microenvironment analysis in colon adenocarcinoma histopathological whole-slide images. Computer Methods and Programs in Biomedicine, 204:106047, 2021.

[73] Mohamed Amgad, James Hodge, Maha Elsebaie, Clara Bodelon, Samantha Puvanesarajah, David Gutman, Kalliopi Siziopikou, Jeffery Goldstein, Mia Gaudet, Lauren Teras, et al. A population-level computational histologic signature for invasive breast cancer prognosis. 2023.