

SRNet: A spatial-relationship aware point-set classification method for multiplexed pathology images

Yan Li

lix4266@umn.edu

Dept. of Computer Sci. & Eng.
University of Minnesota
Minneapolis, Minnesota, USA

Majid Farhadloo

farha043@umn.edu

Dept. of Computer Sci. & Eng.
University of Minnesota
Minneapolis, Minnesota, USA

Santhoshi Krishnan

snk2@rice.edu

Dept. of Electrical & Computer Eng.
Rice University
Houston, Texas, USA

Timothy L Frankel

timofran@med.umich.edu

Dept. of Surgery
University of Michigan
Ann Arbor, Michigan, USA

Shashi Shekhar

shekhar@umn.edu

Dept. of Computer Sci. & Eng.
University of Minnesota
Minneapolis, Minnesota, USA

Arvind Rao

ukarvind@med.umich.edu

Dept. of Computational Medicine &
Bioinformatics
University of Michigan
Ann Arbor, Michigan, USA

ABSTRACT

Point-set classification for multiplexed pathology images aims to distinguish between the spatial configurations of cells within multiplexed immuno-fluorescence (mIF) images of different diseases. This problem is important towards aiding pathologists in diagnosing diseases (e.g., chronic pancreatitis and pancreatic ductal adenocarcinoma). This problem is challenging because crucial spatial relationships are implicit in point sets and the non-uniform distribution of points makes the relationships complex. Manual morphologic or cell-count based methods, the conventional clinical approach for studying spatial patterns within mIF images, is limited by inter-observer variability. The current deep neural network methods for point sets (e.g., PointNet) are limited in learning the representation of implicit spatial relationships between categorical points. To overcome the limitation, we propose a new deep neural network (DNN) architecture, namely spatial-relationship aware neural networks (SRNet), with a novel design of representation learning layers. Experimental results with a University of Michigan mIF dataset show that the proposed method significantly outperforms the competing DNN methods, by 80%, reaching 95% accuracy.

CCS CONCEPTS

• **Applied computing** → **Bioinformatics**; • **Information systems** → *Geographic information systems*.

KEYWORDS

point set, classification, spatial relationship, deep learning, co-location pattern, cross-K function, Bioinformatics, pathology diagnosis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DeepSpatial '21, August, 2021, KDD-organized Virtual Conference

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Yan Li, Majid Farhadloo, Santhoshi Krishnan, Timothy L Frankel, Shashi Shekhar, and Arvind Rao. 2021. SRNet: A spatial-relationship aware point-set classification method for multiplexed pathology images. In *Proceedings of DeepSpatial '21: 2nd ACM SIGKDD Workshop on Deep Learning for Spatiotemporal Data, Applications, and Systems (DeepSpatial '21)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Point-set classification for multiplexed pathology images aims to distinguish between the spatial configurations of cells within multiplexed immuno-fluorescence (mIF) images of different diseases. Advances in the field of multiplexed and anti-body based imaging methods have promoted the development of mIF images, which facilitates bio marker-specific cell species and sub species identification [32]. An example of a multiplexed immunofluorescence image is shown in Figure 1. A point set from multiplexed pathology images records the location and the attributes (e.g., surface phenotype markers) of the cells in a mIF image. For example, Figure 2 shows a sample point set from a mIF image. The location of each cell is represented by its pixel coordinates whose origin is at the top left corner of the image. The cell attributes are the existence of surface phenotype markers (e.g., Epithelial), where "pos" means the presence of a phenotype marker and "neg" otherwise. Figure 3 illustrates the spatial distribution of "pos" phenotype markers in a mIF image of chronic pancreatitis.

Classifying point sets from mIF images is important because it provides a novel way for pathologists to diagnose diseases. For example, in the context of chronic pancreatitis and pancreatic ductal adenocarcinoma, the point sets from mIF images describe the spatial relationships between the diseases' cells, which reveals information about how the interactions between these cells vary.

This problem is challenging due to the following three reasons. First, the points are distributed non-uniformly in the space, which results in complex spatial relationships. Second, the contributions of different spatial relationships vary between different classification tasks, which requires that the representation of the relationships be adjusted to meet the need of specific tasks. Third, spatial relationships between cells of different types are both crucial and implicit




Figure 1: A sample multiplexed immunofluorescence (mIF) image, with the different colours signifying the fluorescence corresponding to different surface bio-markers on the cells imaged. Image courtesy Dr. Timothy L. Frankel.

Cell.X.Position	Cell.Y.Position	Treg	APC	Epithelial	HelperT	PDL1_CD3	PDL1_CD8	PDL1_FoxP3	CD4	CTLs
867	12	neg	neg	neg	neg	neg	neg	neg	pos	neg
27	17	neg	neg	neg	neg	neg	neg	neg	pos	neg
1104	20	neg	neg	neg	neg	neg	neg	neg	pos	neg
1114	24	neg	neg	neg	neg	neg	neg	neg	pos	neg

Location Attributes (e.g., surface phenotype markers)

Figure 2: A point set from a multiplexed pathology image.




Figure 3: A map of a point set from a sample Chronic Pancreatitis mIF image.

in point sets, and the small number of available learning samples makes it difficult for deep neural networks (DNNs) to learn these spatial relationships without appropriate neural network architectures.

Manual morphologic or cell-count based methods, which are the conventional clinical approaches for studying spatial patterns within mIF images, are limited by inter-observer variability. Substantial efforts have been made to apply machine learning techniques to automate the pathology diagnosis process alongside the expansion of digital imaging techniques. In particular, deep neural networks (DNNs) have been extensively studied in a large number of pathology diagnosis applications, including pixel/patch-level region-of-interest detection [4, 7, 9] as well as image-level decision

[15, 38] for various diseases, which have shown state-of-the-art results. However, most of the existing DNN-based applications take images as the input and are inapplicable for our problem. The disadvantage of working with raw images is that the variation in staining and artifacts present across all images in a given cohort may influence analysis. In contrast, point sets offer a simplified representation of cell locations and neighbourhoods, invariant of cell borders and cellular morphology. Recently, as point cloud data from LiDAR scanners have become increasingly popular, the representation of point sets has attracted more attention [30]. However, current methods mainly focus on point sets with few numerical attributes, such as signal strength, and they do not handle categorical attributes specifically. Hence, they do not take full advantage of the spatial relationships between points of different categories.

To overcome these limitations, we propose a new DNN architecture, namely SRNet, with novel design of spatial-configuration based representation learning layers. Experiments show that the proposed methods yield much higher accuracy than the competing DNN methods. Our contributions can be summarized as follows:

- We introduce a deep neural network architecture, SRNet, to learn a representation of the spatial relationships between points of different categories that are not captured by the commonly used statistics such as the cross-K function and the participation ratio.
- We conduct rich experimental studies to evaluate the accuracy of the proposed methods. The discovered crucial patterns are verified by domain scientists, confirming the method's potential to help pathologists identify novel spatial relationships between different cell types (e.g., immune cells and tumor cells) in the micro-environment.

Scope: The scope of this study is limited to analyzing point datasets representing the location and types of cells derived from multiplexed immuno-fluorescence (mIF) images to distinguish between diseases. Analyzing mIF images without converting them to point sets is outside the scope of this paper. In addition, we do not evaluate the proposed method with larger datasets due to a lack of public benchmarks. Field trials to evaluate the clinical value of the proposed method also fall outside the scope of this study.

Outline: The rest of the paper is organized as follows. Section 2 describes the application domain of the study. Section 3 introduces the formal definition of the problem and provides a short description of the dataset. Related work is reviewed in Section 4. Our proposed methods are described in Section 5. Section 6 presents the evaluation of the proposed methods. Section 7 concludes the paper and outlines future work.

2 APPLICATION DOMAIN

Biopsies are the standard procedure in practice for disease diagnosis, including cancers. In this procedure, a sample of tissue is removed from the body, chemically treated, sliced into thin sections, placed on a glass slide, and stained with specific chemicals to enhance contrast for visual inspection [2]. A pathologist then performs a macroscopic examination of the specimen and describes various features such as type of cells present, their distribution, and other important diagnostic features.

With developments in whole slide digital imaging and antigen-based staining technology (e.g., Time-Of-Flight mass cytometry (CyToF)[8], and Co-detection by Indexing (CODEX)[10]), it is possible to not only isolate cell nuclei in an image, but also to determine types and sub-types of each cell type present in the image based on the cell's surface chemistry with high throughput, and the potential to scale up to more than 30 markers [16, 24]. These novel technologies have played an important role in the era of cancer immunotherapeutic treatment regimens [5, 25], which involves the treatment of diseases by inducing, enhancing, or suppressing an immune response in the patient. This treatment regimen has been gaining increasing attention due to its potential in the treatment of cancers which are non-responsive to conventional methods such as radiotherapy and chemotherapy [26, 28]. As this treatment regimen utilizes the immunoregulatory cells of the patient in eliminating tumorous cells, there is a growing interest in understanding the interplay between various cells in a spatially informed manner in the tumor microenvironment [36, 37]. As an example, for tumor infiltrating lymphocytes (TILs) to be able to induce cell death, these cells must have direct or proximal contact with tumor cells [1]. Thus, the distance between tumor and immune cells is an important indicator for determining disease progression and treatment effect.

Emerging research in this area has begun to highlight the importance of spatial organization among cell phenotypes for cancer diagnosis and prognosis [36]. Currently, visual inspection and cell-counting by a pathologist are the methods used to describe the spatial organization of cells, which is fraught with inter-observer variability and inconsistency between studies. Also, although we have some information regarding which immune cell features predict a positive response, there is a lack of reliable methods to identify which patients will benefit from immunotherapeutic measures based on their individual immune cell make-up. The development and adoption of spatially informed methods both for tumor and disease micro-environment quantification generally would help in developing optimal treatment plans tailored to each patient. Additionally, it would be prudent to leverage the power of algorithmic intelligence in the pathology domain, as it can provide insights which cannot be captured visually by a pathologist.

3 PROBLEM DEFINITION & DATA DESCRIPTION

Given a collection of categorical point sets (e.g., cells with different surface phenotype markers) from multiplexed immuno-fluorescence (mIF) images and the class labels of the point sets (e.g., different diseases), the goal of this study is to train a machine learning model that distinguishes between the point sets of different classes. The primary objective is to achieve a high classification accuracy.

We define a categorical point set as a collection of points, where each individual point belongs to a single category and is located in 2-D Euclidean space. This study was conducted on 199 anonymized point sets derived from mIF images belonging to two disease groups, namely chronic pancreatitis (i.e., class 1) and pancreatic ductal adenocarcinoma (PDAC) (i.e., class 2), which had 56 and 143 point sets, respectively. In the original dataset, cell surface markers indicate nine phenotypes. Each cell might be associated with one or more

phenotype. To transform the original point sets into categorical point sets, we considered any point that had a single phenotype marker as belonging to the category corresponding to that phenotype, and we replaced every point that had multiple phenotype markers with a group of points, having one marker each and then assigned the points to multiple categories corresponding to the phenotype of each point's marker. Generating point sets from multiplexed pathology images is beyond the scope of this paper, and we treat point sets as given inputs.

4 RELATED WORK

The history of deep neural network (DNN) methods that directly take point sets as the input dates back to PointNet [30], which learns point features independently through multiple fully connected neural network layers and aggregates them into a shape feature using a max pooling layer. These methods have been widely used for 3D shape classification and semantic segmentation as the point clouds collected from LiDAR scanners have become increasingly popular. PointNet++ [31] defines multi-scale regions and uses PointNet to learn their features. It then hierarchically aggregates the regions' features, so it can capture local configurations and learn fine-grained patterns. Similar to PointNet++, the idea of spatially partitioning points and then recursively aggregating them has been extensively explored. For example, KD-trees are used in [13, 21] to spatially partition points based on point density.

Meanwhile, much effort has been made to introduce DNN architectures that were originally designed for other data formats (e.g., imagery and time series). For example, convolutional neural network (CNN) models are studied in the spectral domain (e.g., RGCNN [33]) and the spatial domain (e.g., Pointwise convolution [18]). Recursive neural network (RNN) models are applied with the assumption that "order matters" [34], and there are autoencoders that learn the representation of point sets [17]. However, these models are not specifically designed to handle multi-categorical point sets and do not take full advantage of the spatial relationships between different categories of points.

5 PROPOSED APPROACHES: SRNET

The cross-category spatial neighborhood relationships is an important component in the spatial configuration of points. In pathology diagnosis, the spatial correlations between different types of immune cells may vary with diseases, which inspires us to introduce a deep neural network (DNN) method, namely spatial-relationship aware neural network (SRNet), with novel representation layers to represent point sets with the spatial relationships between different categories of points in them.

5.1 Spatial-Relationship Quantification

An intuitive way of representing the spatial relationships of point sets consisting of various categories is to utilize measures quantifying the relationships. In this subsection, we present two measures for spatial relationships widely used in spatial data mining and spatial statistics, and how they can be used in classification tasks.

5.1.1 Participation ratio. The participation ratio quantifies the degree to which a category tends to be involved in a co-location pattern. Co-location patterns [19, 35] refer to set of categorical

point sets that tend to be located in close proximity, such as a point set of Nile crocodiles and Egyptian plovers [22].

A co-location pattern [19] has three defining concepts. First, a co-location pattern is in the form of a set of categories. Second, a neighborhood clique is a set of points within which every pairwise distance is smaller than a threshold. Third, an instance of a spatial co-location pattern is a neighborhood clique composed of one point from every category in the pattern. The participation ratio (PR) of a category in a co-location pattern is then defined as the ratio of the points in the category that are within the instance of the pattern, which is calculated as:

$$PR(c_i, p) = \frac{|c_i \text{ points in the instances of } p|}{|c_i \text{ points}|}, \quad (1)$$

where c_i is a category and p is a spatial co-location pattern, and $|\cdot|$ yields the cardinality of a set. The value of a participation ratio is between 0 and 1. The greater the value, the more likely c_i points are located nearby the points of other categories in the pattern p .

For the sake of computational efficiency, in this study we only consider the spatial co-location patterns composed of two categories, so Equation 1 can be transformed as:

$$PR(c_i, c_j, d) = \frac{|c_i \text{ points with } c_j \text{ in } SN(c_i, d)|}{|c_i \text{ points}|}, \quad (2)$$

where $SN(c_i, d)$ yields a circular spatial neighborhood with a radius of d around a c_i point. Given a point set containing points belonging to k categories and a neighborhood distance threshold, there will be $k(k-1)$ participation ratios. An important hyperparameter that affects the value of the participation ratio is the neighborhood distance threshold. Participation ratios with different neighborhood distance thresholds imply the relationships between points in different spatial scales, so we compute the participation ratios with a collection of l neighborhood distance thresholds. Therefore, we can use a vector of $k(k-1)l$ participation ratios as the representation of a point set with k categories.

To validate that the spatial relationships quantified by participation ratios may be useful for distinguishing between the point sets of different diseases, we plot the probability density distribution of four participation ratios in the dataset we described in Section 3 using histograms in Figure 4. Each histogram has ten equal-width bins that represent the range of participation ratio values, and the area of each bin is the probability density of the bin. As can be seen, the probability distribution of a participation ratio varies with category pairs as well as with neighborhood distance thresholds, and in Figure 4a and 4c, the probability distributions for the two diseases are notably different. Therefore, $PR(APC, Treg, 100)$ and $PR(APC, Treg, 200)$ may be used to distinguish the point sets of the two diseases.

5.1.2 Ripley's cross-K function. The participation ratio, $PR(c_i, c_j, d)$, can be thought of as the expectation that c_j points exist in the spatial neighborhood c_i point. However, the existence of c_j points does not tell the whole story about the distribution of c_j points in a c_i points' spatial neighborhood. Ripley's cross-K function, instead,




Figure 4: Examples of the probability distribution of participation ratios.

focuses on the number of c_j points in c_i points' spatial neighborhood. It is defined in the following form:

$$\text{cross-K}(c_i, c_j, d) = \frac{E(|c_j \text{ in } SN(c_i, d)|)}{E(|c_j \text{ in entire study area}|)}, \quad (3)$$

where c_i and c_j are two categories, d is a neighborhood distance threshold, $SN(c_i, d)$ yields the circular spatial neighborhood of a c_i point with a radius of d , and $E(\cdot)$ returns the expectation. The value of a cross-K function is non-negative. The greater the value, the more c_j points are located nearby the c_i points. Similar to how we represent a point set using its participation ratios, given l neighborhood distance thresholds, we can also represent a point set with k categories using a vector contains $k(k-1)l$ cross-K function values.

5.2 Proposed SRNet Architecture

In the definitions of the participation ratio and the cross-K function, a core component is the representation of the spatial neighborhood of points. Given an ordered category pair (c_i, c_j) , and a spatial neighborhood distance threshold d , the participation ratio uses the existence of c_j points and the cross-k function uses the count of c_j points to represent the distribution of c_j points in the spatial neighborhood of c_i points. However, in addition to existence and count, there may be other patterns that describe the spatial relationships between c_i and c_j points. Hence, we design a DNN model that uses a spatial-relationship aware neural network (SRNet) that learns the spatial distribution of c_j points in c_i points' spatial neighborhood for every ordered category pair (c_i, c_j) , and then to generate a representation of point sets. The point-set representation can then be fed into a fully connected neural network for classification.

Figure 5 shows the overall architecture of SRNet. The input of the approach is a categorical point set denoted as $X \in \mathbb{R}^{N \times (D+1)}$, where N is the number of points and $D = 2$ is the spatial dimensions. Each point has one categorical attribute, and there are k categories in total. Similar to using the participation ratios or the cross-K




Figure 5: Overview of the SRNet architecture.




Figure 6: The architecture of the spatial relationship layer.

function values to represent point sets, the SRNet uses a DNN layer (spatial relationship layer) to learn the spatial relationship measures of all $k(k-1)$ ordered category pairs. This architecture facilitates the integration of human expert knowledge by concatenating the learned spatial relationship measures with the measures provided by human experts (e.g., the participation ratio, the cross-K function). The architecture of the spatial relationship layer, shown in Figure 6, has three main components: a spatial neighborhood layer (Section 5.2.1), a spatial distribution attention layer (Section 5.2.2), and a weighted average pooling layer. For every ordered category pair (c_i, c_j) , the spatial neighborhood layer generates a representation of the spatial distribution of c_j points in every c_i point's spatial neighborhood, and the spatial distribution attention layer learns the attention to be paid to each c_i point according to the spatial distribution of c_i points. Then, the weighted average pooling layer aggregates the spatial neighborhood representation of every c_i point with different weights to calculate the spatial relationship measures of pair (c_i, c_j) . Finally, the spatial relationship measures of all ordered category pairs are concatenated to generate the overall representation of the point set, denoted as $Y \in \mathbb{R}^{k \times (k-1) \times W}$, where W is the feature dimension of the spatial relationship measures of a category pair.




Figure 7: The architecture of the spatial neighborhood layer.

5.2.1 Spatial neighborhood layer. Given an ordered category pair (c_i, c_j) , a spatial neighborhood layer is applied to represent the spatial distribution of c_j points within every individual c_i point's spatial neighborhood independently. The input of this layer is a c_i point and the c_j points in its spatial neighborhood, and its output is a vector representing the spatial distribution of the c_j points. There are two main steps in this layer, namely, spatial location representation and spatial distribution summarization (Figure 7).

Spatial location representation focuses on representing the relative location of a c_j point in the spatial neighborhood of a c_i point. The most commonly used representation of a relative location is the difference of coordinates. However, it was reported in [27] that the difference of coordinates failed to convey the information of various spatial distributions. Recently, Gao et al. proposed a representational model that uses the hexagon patterns of the grid cells to form a high-dimensional vector representation of 2D locations (x) , based on the following theorem whose proof is given in [14].

THEOREM 5.1. Let $\Psi(x) = (e^{i\langle a_j, x \rangle}, j = 1, 2, 3)^T \in \mathbb{C}^3$ where $e^{i\theta} = \cos \theta + i \sin \theta$ and $\langle a_j, x \rangle$ is the inner product of a_j and x . $a_1, a_2, a_3 \in \mathbb{R}^2$ are 2D vectors such that the angle between each pair is $2\pi/3, \forall j, \|a_j\| = 2\sqrt{\alpha}$. Let $C \in \mathbb{C}^{3 \times 3}$ be a random complex matrix such as $C * C = I$. Then $\phi(x) = C\Psi(x)$, $M(\Delta x) = C \text{diag}(\Psi(\Delta x))C^*$ satisfies

$$\phi(x + \Delta x) = M(\Delta x)\phi(x) \quad (4)$$

and

$$\langle \phi(x + \Delta x), \phi(x) \rangle = d(1 - \alpha \|\Delta x\|^2) \quad (5)$$

where $\phi(x)$ is the representation of location x , $d = 3$ is the dimension of $\phi(x)$, and Δx is a small displacement from x .

Based on Theorem 5.1, Mai et al. [27] introduced a multi-scale location representation model by using sine and cosine functions of different frequencies in $\Psi(x)$, inspired by the multi-scale periodic representation of grid cells in mammals [3]. In this model, $\Psi(x)$ is represented as a concatenation of the position embedding (PE) at S

scales, $PE(x) = [PE_1(x); \dots; PE_s(x); \dots; PE_S(x)]$,

$$PE_s(x) = [PE_{s,1}(x); PE_{s,2}(x); PE_{s,3}(x)], \quad (6)$$

$$PE_{s,j}(x) = [\cos(\frac{\langle x, a_j \rangle}{\lambda_{min} \cdot g^{s/(S-1)}}); \sin(\frac{\langle x, a_j \rangle}{\lambda_{min} \cdot g^{s/(S-1)}})], \quad (7)$$

$\forall j = 1, 2, 3,$

where $a_1 = [1, 0]^T$, $a_2 = [-1/2, \sqrt{3}/2]^T$, $a_3 = [-1/2, -\sqrt{3}/2]^T$ are unit vectors, the angles between every pair of vectors is $2\pi/3$, $\lambda_{min}, \lambda_{max}$ are the minimum and maximum grid scales, and $g = \frac{\lambda_{max}}{\lambda_{min}}$. The matrix multiplication $C\Psi(x)$ is represented as $NN(PE(x))$, where $NN(\cdot)$ are fully connected ReLU layers. Therefore, the location of a c_j point relative to a c_i point can be represented as $NN(PE(\Delta x))$, where Δx is the difference of their coordinates.

Given a collection of relative location representations of c_j points in a c_i point's spatial neighborhood, a max pooling layer is applied to summarize the relative locations to get the representation of the c_i point's spatial neighborhood. Pointnet[30] has theoretically and experimentally demonstrated that with enough neurons, a max pooling layer is able to learn to summarize a point distribution [30].

5.2.2 Spatial distribution attention layer. To get the representation of the spatial relationship measures of pair (c_i, c_j) , an average pooling layer is used to aggregate the representation of c_j points' distribution in all the spatial neighborhoods of c_i points. However, it is questionable whether all c_i points should contribute equally to the spatial relationship measures. In their study of the spatial co-location patterns, Barua and Sander discovered that the spatial distribution of the points belonging to a category within a co-location pattern affected the statistical significance of the pattern's participation ratio where all points contributed equally [6]. A potential reason is the existence of spatial auto-correlation. In other words, the spatial neighborhoods of nearby points are similar. If all points contribute equally, the spatial neighborhood of a point away from other points may be overwhelmed by the spatial neighborhoods of the points in clusters. Therefore, we introduce a spatial distribution attention layer to determine the attention paid to each c_i point when generating the spatial relationship measures of (c_i, c_j) . The layer first generates the representation of the spatial distribution of c_i points in each c_i point's spatial neighborhood independently using the proposed spatial neighborhood layer. Then it estimates the attention paid to each c_i point according to the representations using multiple fully connected ReLU layers. This method is similar to the application of farthest point sampling (FPS) in PointNet++ [31], which selects subsets of representative points to learn local features. Instead of using a greedy heuristic as in FPS, the proposed spatial distribution attention layer uses neural network layers to adjust the attention to points.

6 EXPERIMENT

Our experimental evaluation has two components: (1) a comparison of the proposed methods with the state-of-the-art deep neural network (DNN) point set classification methods; and (2) an analysis of the importance of the spatial relationship measures.

6.1 Classification Accuracy Comparison

We have conducted two sets of experiments: (1) comparing our proposed methods: handcrafted features using classic spatial measure (i.e., participation ratio or cross-k function) and learned features using SRNet, each combined with a simple neural network classifier, with the state-of-the-art (SOTA) DNN point set classification methods (i.e., PointNet and PointNet++), (2) comparing handcrafted features combined with simple classification models with the SOTA DNN point set classification methods. The experiments are designed to answer the following questions: 1) did the proposed method yield more accurate classification results than the competing DNN methods? 2) how do the spatial relationship measures used to represent point sets affect classification accuracy? 3) how does the choice of classification method affect accuracy? Classification accuracy is measured by AUC-ROC, precision, recall, F1 score, and accuracy. The candidate methods compared were as follows.

- **PointNet[30]:** PointNet is a neural network architecture that directly consumes point sets for applications ranging from object classification to part segmentation.
- **PointNet++[31]:** PointNet++ is a hierarchical neural network architecture that applies PointNet recursively to capture local structure and recognize fine-grained patterns and complex scenes.
- **PR + DT / RF / NN:** The point set representation composed of the participation ratios (Section 5.1.1) is fed into a decision tree / random forest / fully connected neural network model for classification.
- **cross-K + DT / RF / NN:** The point set representation composed of the cross-K function values (Section 5.1.2) is fed into a decision tree / random forest / fully connected neural network model for classification.
- **SRNet / +PR / +cross-K:** The point set representation learned by the SRNet model proposed in Section 5.2 without human expert knowledge / with the participation ratio measures / with the cross-K function measures is fed into a fully connected neural network model for classification.

The implementation of both PointNet and PointNet++ are available on GitHub¹. The decision tree, the random forest, and the fully connected neural network methods were implemented using the Python scikit-learn package [29]. The maximal depth of the decision tree methods was set to 4, and the maximal depth and the number of estimators of the random forest methods were set to 3 and 1000. Other hyperparameters were kept as the default values. The fully connected neural network classifier had two hidden ReLU layers with 4096 neurons and a sigmoid layer as the output layer.

The SRNet method was implemented using PyTorch, where the spatial neighborhood of each point was set as a circle with a radius of 200, and the minimal grid cell size, the maximal grid cell size, and the number of scales of the multi-scale location representation layers were set at 1, 100, and 10 respectively. All the spatial neighborhood layers shared the same architecture and parameters. The fully connected ReLU layers in the spatial neighborhood layers had four hidden layers, and the hidden layer dimension was set at 256. The feature dimension of the learned spatial relationship measures

¹Link to PointNet repository: <https://github.com/charlesq34/pointnet>. Link to PointNet++ repository: <https://github.com/charlesq34/pointnet2>.




Figure 8: The classification accuracy of the methods using neural network classifiers.

of each ordered category pair was 32. The SRNet and the neural network classifier were trained using the Adam optimization algorithm with the learning rate of 10^{-4} to minimize the cross entropy loss of the classification results and the ground truth.

We used the dataset described in Section 3. Since the original dataset only had 199 point sets, we used 5-fold cross-validation and augmented the number of point sets by partitioning, flipping, and rotating the original point sets. To get subsets of a point set and keep spatial relationship information in each subset, instead of randomly sampling points, we partitioned the minimum bounding rectangle (MBR) of the point set horizontally by 20% and 80% and then 80% and 20%, and used the 80% subsets. The subsets were then flipped both horizontally and vertically. Finally, the flipped subsets were rotated by 90 degrees three times. Thus, after data augmentation, there were $199 \times 2 \times 4 \times 4 = 6368$ point sets in total.

Table 1 shows the mean and standard deviation (in parentheses) of classification accuracy measures of the candidate methods. The highest accuracy is highlighted in bold. It is evident that the proposed methods, even a very simple model (e.g., the decision tree model) with a well-defined spatial relationship measures (e.g., the participation ratio), were much more accurate than the competing DNN point set classification methods (i.e., PointNet and PointNet++).

A comparison of the classification accuracy of the methods using neural network classifiers (Figure 8), shows that the methods using classic spatial relationship measures (PR+NN and cross-K+NN) and those using measures learned by the proposed SRNet (SRNet, SRNet+PR, SRNet+cross-K) had much higher accuracy than the competing DNN methods. This indicates that the proposed SRNet was able to learn spatial relationship measures that were missed by the competing DNN methods. Moreover, the accuracy of the SRNet+PR and SRNet+cross-K methods was higher than that of the PR+NN and cross-K+NN methods, respectively, which means the proposed SRNet is able to learn features that are not captured by the participation ratio and the cross-K function but that were useful for the classification task.

Finally, the classification accuracy of methods using the same point set representation (e.g., PR+DT v.s. PR+NN) indicates that complex models yielded more accurate results. However, the effect of choosing different classification methods on classification accuracy was not as significant as the effect of point set representation.




Figure 9: First two layers of the decision trees trained using the entire dataset in Section 3

6.2 Analysis of Spatial Relationship Measures

The goal of the second set of experiments was to analyze the category pairs whose spatial relationship measures are important for classifying the point sets of the two diseases, as this provided a way to discover the interactions between cells that varied with diseases.

In the PR+DT and cross-K+DT methods, the feature vectors composed of the participation ratios and cross-K function values were fed into decision tree models. Since in every node of the decision tree model, a feature is selected greedily to divide samples into two groups according to a heuristic (e.g., the information gain), the selected features indicate which category pairs contain high variation in their spatial relationships. Figure 9 shows the first two layers of the decision trees trained using the entire dataset described in Section 3. As can be seen, the spatial relationships between HelperT cells and CD4 cells and between Treg cells and HelperT cells were significantly different under the micro environment of the two diseases.

In the PR+RF and cross-K+RF methods, the feature vectors composed of the participation ratios and cross-K function values were fed into random forest models. Feature importance in the random forest models can be measured by the mean impurity decrease, which also implies the spatial relationships between the category pairs vary a lot in the point sets of the two diseases. Table 2 lists the top ten important features in the PR+RF and cross-K+RF models trained using the entire dataset. As can be seen, both the participation ratio features and cross-K function features indicate that the spatial relationships between the HelperT and Treg cells are most useful for distinguishing the point sets of the two diseases.

For the PR+NN, cross-K+NN, and the SRNet methods, we evaluated the importance of the spatial relationship measures, namely, the participation ratio, the cross-K function value, and the representation learned by SRNet, through permutation feature importance. Permutation feature importance measures the increase in the prediction error of the model after we permute the feature's values. In this experiment, the importance of the spatial relationship measures of an ordered category pair was measured by the classification accuracy after exchanging the corresponding elements in the representation vectors. The lower the accuracy, the more important the measures of ordered category pair. In the dataset described in Section 3 the most important ordered category pairs were (HelperT, Treg), (HelperT, CD4), (CTLs, Treg), and (APC, Treg).

Table 1: Classification accuracy results.

Method	AUC-ROC	Precision	Recall	F1 score	Accuracy
PointNet	0.518 (0.026)	0.352 (0.079)	0.518 (0.026)	0.421 (0.120)	0.508 (0.160)
PointNet++	0.529 (0.089)	0.412 (0.138)	0.529 (0.089)	0.421 (0.138)	0.529 (0.089)
PR+DT	0.903 (0.027)	0.955 (0.028)	0.911 (0.036)	0.932 (0.016)	0.905 (0.021)
PR+RF	0.979 (0.011)	0.936 (0.025)	0.949 (0.027)	0.942 (0.022)	0.917 (0.031)
PR+NN	0.980 (0.016)	0.948 (0.035)	0.954 (0.041)	0.950 (0.025)	0.929 (0.035)
cross-K+DT	0.852 (0.011)	0.911 (0.027)	0.914 (0.058)	0.911 (0.027)	0.874 (0.031)
cross-K+RF	0.955 (0.028)	0.852 (0.019)	0.967 (0.017)	0.906 (0.015)	0.856 (0.023)
cross-K+NN	0.938 (0.027)	0.908 (0.037)	0.933 (0.046)	0.919 (0.025)	0.883 (0.036)
SRNet	0.939 (0.030)	0.951 (0.038)	0.884 (0.066)	0.914 (0.031)	0.884 (0.039)
SRNet+PR	0.985 (0.015)	0.967 (0.002)	0.962 (0.040)	0.964 (0.020)	0.950 (0.014)
SRNet+cross-K	0.964 (0.022)	0.953 (0.028)	0.909 (0.047)	0.930 (0.028)	0.904 (0.037)

Table 2: Top 10 important features obtained in the PR+RF and cross-K+RF methods.

Rank	PR+RF feature	cross-K+RF feature
1	PR(HelperT, Treg, 1)	cross-K(Treg, HelperT, 100)
2	PR(HelperT, CD4, 1)	cross-K(HelperT, Treg, 200)
3	PR(HelperT, Treg, 50)	cross-K(HelperT, Treg, 50)
4	PR(HelperT, Treg, 200)	cross-K(HelperT, Treg, 100)
5	PR(HelperT, Treg, 100)	cross-K(HelperT, Treg, 1)
6	PR(HelperT, Treg, 150)	cross-K(Treg, HelperT, 50)
7	PR(CD4, Treg, 150)	cross-K(Treg, HelperT, 1)
8	PR(CD4, Treg, 200)	cross-K(HelperT, Treg, 150)
9	PR(CD4, Treg, 100)	cross-K(Treg, HelperT, 200)
10	PR(APC, Treg, 100)	cross-K(Treg, HelperT, 150)

6.3 Clinical Implications

From a clinical perspective, the results highlight some key cell phenotype relationships that may directly or indirectly play a role in the disease micro-environment. Specifically, the relationship between CTLs and T-regs, and Helper T-cells and T-regs are of particular interest from an immunological standpoint. Cytotoxic Lymphocytes(CTLs) are the cells that actively seek out and kill cancer cells in the environment on activation of the immune system[11]. On the other hand, under normal conditions, the T-regulatory cells have a regulating effect on the immune response of the locale [23]. It has been observed that T-regulatory cells play a more functional role in the cancer micro-environment, and there is potential for some interplay between the two cell phenotypes from a functional standpoint. Due to this, there is a tendency for them to co-localize at a higher frequency with CTLs, and potentially inhibit their function [12]. This may be due to physiologic suppression of activated CTLs, or pathological polarization of CD4 positive cells by tumor secreted factors in the tumor micro-environment[20]. Further investigation on a larger cohort to confirm the potential discriminatory power of the pairwise interactions observed in this experiment would be warranted.

The identification of the cell-pairs opens up a potential for a novel method to capture the difference in cellular arrangements

across different diseases. This also alludes to the influence of cell-cell distances and their relative placement in the state of the micro-environment [36]. Along with reinforcing known relationships, these features would also serve to offer new insight into potential cell-cell relationships that were either unknown or little explored in previous studies. In the age of increasing focus on personalized treatment paradigms, the utilization of a spatially-aware approach would assist physicians in making more informed treatment plans.

7 CONCLUSION & FUTURE WORKS

In this paper, we proposed a deep learning point-set classification method, namely SRNet, for multiplexed pathology images. SRNet provides a novel way for pathologists to diagnose diseases. Instead of classifying multiplexed immuno-fluorescence (mIF) images directly, we first converted mIF images to point sets representing the cells on mIF images, and then classified the point sets. An experimental evaluation showed that the proposed SRNet can learn spatial relationship measures that are not captured by classic measures, and the classification accuracy of using the learned measures significantly outperformed the SOTA deep learning point-set classification methods, reaching 95% accuracy (about 80% more accurate). In addition, the proposed methods helped to discover pairs of cell types that might inspire new pathology findings.

In the future, we will compare the proposed method on point sets with the methods directly analyzing mIF images without converting them to point sets. We also plan to identify larger mIF images and other spatial pathology datasets for larger and broader evaluation of the proposed method. In addition, the proposed SRNet focuses on the spatial relationships between two cell types, and we plan to extend its capability by taking the relationships between multiple cell types into consideration.

ACKNOWLEDGMENTS

This material is based upon work supported by the NSF under Grants No. 2040459, 1737633, 1901099, 1218168, and 1916518, the USDOD under Grants No. HM0476-20-1-0009, the NIH under Grant No. UL1 TR002494, KL2TR002492, and TL1 TR002493, the USDA under Grant No. 2017-51181-27222, and the OVPR Infrastructure Investment Initiative, Minnesota Supercomputing Institute (MSI), and Provost's Grand Challenges Exploratory Research and International Enhancements Grants at the University of Minnesota. A.R.

and S.K. were supported by CCSG Bioinformatics Shared Resource 5 P30 CA046592, NCI R37CA214955-01A1, a gift from Agilent technologies, a Research Scholar Grant from the American Cancer Society (RSG-16-005-01), Institutional support from the University of Michigan, the Propelling Original Data Science (PODS) MIDAS grant from the University of Michigan, and a Precision Health investigator award to A.R. along with L.Rozek and M.Sartor.

REFERENCES

- [1] [n.d.]. Immunological Synapse. *SpringerReference* ([n.d.]). https://doi.org/10.1007/springerreference_32407
- [2] 2020. Biopsy. <https://www.cancer.net/navigating-cancer-care/diagnosing-cancer/tests-and-procedures/biopsy>
- [3] Alison Abbott and Ewen Callaway. 2014. Nobel prize for decoding brain's sense of place. *Nature News* 514, 7521 (2014), 153.
- [4] Eirini Arvaniti, Kim S Fricker, Michael Moret, Niels Rupp, Thomas Hermanns, Christian Fankhauser, Norbert Wey, Peter J Wild, Jan H Rueschoff, and Manfred Claassen. 2018. Automated Gleason grading of prostate cancer tissue microarrays via deep learning. *Scientific reports* 8, 1 (2018), 1–11.
- [5] Souptik Barua, Penny Fang, Amrith Sharma, Junya Fujimoto, Ignacio Wistuba, Arvind U.K. Rao, and Steven H. Lin. 2018. Spatial interaction of tumor cells and regulatory T cells correlates with survival in non-small cell lung cancer. *Lung Cancer* 117 (2018), 73–79. <https://doi.org/10.1016/j.lungcan.2018.01.022>
- [6] Sajib Barua and Jörg Sander. 2014. Mining Statistically Significant Co-location and Segregation Patterns. *IEEE Trans. Knowl. Data Eng.* 26, 5 (2014), 1185–1199. <https://doi.org/10.1109/TKDE.2013.88>
- [7] Babak Ehteshami Bejnordi, Maeve Mullooly, Ruth M Pfeiffer, Shaoqi Fan, Pamela M Vacek, Donald L Weaver, Sally Herschorn, Louise A Brinton, Bram van Ginneken, Nico Karssemeijer, et al. 2018. Using deep convolutional neural networks to identify and classify tumor-associated stroma in diagnostic breast biopsies. *Modern Pathology* 31, 10 (2018), 1502–1512.
- [8] S. C. Bendall, E. F. Simonds, P. Qiu, E.-A. D. Amir, P. O. Krutzik, R. Finck, R. V. Bruggner, R. Melamed, A. Trejo, O. I. Ornatsky, and et al. 2011. Single-Cell Mass Cytometry of Differential Immune and Drug Responses Across a Human Hematopoietic Continuum. *Science* 332, 6030 (2011), 687–696. <https://doi.org/10.1126/science.1198704>
- [9] Nicolas Coudray, Paolo Santiago Ocampo, Theodore Sakellariopoulos, Navneet Narula, Matija Snuderl, David Fenyö, Andre L Moreira, Narges Razavian, and Aristotelis Tsigaris. 2018. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature medicine* 24, 10 (2018), 1559–1567.
- [10] Gajalakshmi Dakshinamoorthy, Jaskirat Singh, Joseph Kim, Nadya Nikulina, Roya Bashier, Sejal Mistry, Maria E. Gallina, Atri Choksi, Meenu Perera, Ashley Wilson, and et al. 2019. Abstract 490: Highly multiplexed single-cell spatial analysis of tissue specimens using CODEX. *Immunology* (2019). <https://doi.org/10.1158/1538-7445.sabcs18-490>
- [11] Bagher Farhood, Masoud Najafi, and Keywan Mortezaee. 2018. CD8 cytotoxic T lymphocytes in cancer immunotherapy: A review. *Journal of Cellular Physiology* 234, 6 (2018), 8509–8521. <https://doi.org/10.1002/jcp.27782>
- [12] Anita Feichtenbeiner, Matthias Haas, Maik Böttner, Gerhard G. Grabenbauer, Rainer Fietkau, and Luitpold V. Distel. 2013. Critical role of spatial interaction between CD8 and Foxp3 cells in human gastric cancer: the distance matters. *Cancer Immunology, Immunotherapy* 63, 2 (2013), 111–119. <https://doi.org/10.1007/s00262-013-1491-x>
- [13] Matheus Gadelha, Rui Wang, and Subhransu Maji. 2018. Multiresolution Tree Networks for 3D Point Cloud Processing. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VII (Lecture Notes in Computer Science, Vol. 11211)*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer, 105–122. https://doi.org/10.1007/978-3-030-01234-2_7
- [14] Ruiqi Gao, Jianwen Xie, Song-Chun Zhu, and Ying Nian Wu. 2019. Learning Grid Cells as Vector Representation of Self-Position Coupled with Matrix Representation of Self-Motion. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- [15] Hrushikesh Garud, Sri Phani Krishna Karri, Debodoot Sheet, Jyotirmoy Chatterjee, Manjunatha Mahadevappa, Ajoy K Ray, Arindam Ghosh, and Ashok K Maity. 2017. High-magnification multi-views based classification of breast fine needle aspiration cytology cell samples using fusion of decisions from deep convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 76–81.
- [16] Yury Goltsev, Nikolay Samusik, Julia Kennedy-Darling, Salil Bhatte, Matthew Hale, Gustavo Vazquez, Sarah Black, and Garry P. Nolan. 2017. Deep profiling of mouse splenic architecture with CODEX multiplexed imaging. (2017). <https://doi.org/10.1101/203166>
- [17] Kaveh Hassani and Mike Haley. 2019. Unsupervised Multi-Task Feature Learning on Point Clouds. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. IEEE, 8159–8170. <https://doi.org/10.1109/ICCV.2019.00825>
- [18] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. 2018. Pointwise Convolutional Neural Networks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. IEEE Computer Society, 984–993. <https://doi.org/10.1109/CVPR.2018.00109>
- [19] Yan Huang, Shashi Shekhar, and Hui Xiong. 2004. Discovering colocation patterns from spatial data sets: a general approach. *IEEE Transactions on Knowledge and data engineering* 16, 12 (2004), 1472–1485.
- [20] Magdalena Huber, Corinna U. Brehm, Thomas M. Gress, Malte Buchholz, Bilal Alashkar Alhamwe, Elke Von Strandmann, Emily P. Slater, Jörg W. Bartsch, Christian Bauer, Matthias Lauth, and et al. 2020. The Immune Microenvironment in Pancreatic Cancer. *International Journal of Molecular Sciences* 21, 19 (2020), 7307. <https://doi.org/10.3390/ijms21197307>
- [21] Roman Klokov and Victor Lempitsky. 2017. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *Proceedings of the IEEE International Conference on Computer Vision*. 863–872.
- [22] Timothy F Leslie, Cara L Frankenfeld, and Matthew A Makara. 2012. The spatial food environment of the DC metropolitan area: Clustering, co-location, and categorical differentiation. *Applied Geography* 35, 1-2 (2012), 300–307.
- [23] Zhiyuan Li, Dan Li, Andy Tsun, and Bin Li. 2015. FOXP3 regulatory T cells and their functional regulation. *Cellular & Molecular Immunology* 12, 5 (2015), 558–565. <https://doi.org/10.1038/cmi.2015.10>
- [24] Jia-Ren Lin, Benjamin Izar, Shu Wang, Clarence Yapp, Shaolin Mei, Parin M Shah, Sandro Santagata, and Peter K Sorger. 2018. Highly multiplexed immunofluorescence imaging of human tissues and tumors using t-CyCIF and conventional optical microscopes. *eLife* 7 (2018). <https://doi.org/10.7554/elife.31657>
- [25] Steve Lu, Julie E. Stein, David L. Rimm, Daphne W. Wang, J. Michael Bell, Douglas B. Johnson, Jeffrey A. Sosman, Kurt A. Schalper, Robert A. Anders, Hao Wang, and et al. 2019. Comparison of Biomarker Modalities for Predicting Response to PD-1/PD-L1 Checkpoint Blockade. *JAMA Oncology* 5, 8 (2019), 1195. <https://doi.org/10.1001/jamaoncol.2019.1549>
- [26] Iwona Lugowska, Paweł Teterycz, and Piotr Rutkowski. 2018. Immunotherapy of melanoma. *Współczesna Onkologia* 2018, 1 (2018), 61–67. <https://doi.org/10.5114/wo.2018.73889>
- [27] Gengchen Mai, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, and Ni Lao. 2020. Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- [28] K Noel Maslhi. 2001. Fighting infection using immunomodulatory agents. *Expert Opinion on Biological Therapy* 1, 4 (2001), 641–653. <https://doi.org/10.1517/14712598.1.4.641>
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [30] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 652–660.
- [31] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 5099–5108.
- [32] Wei Chang Colin Tan, Sanjna Nilesh Nerurkar, Hai Yun Cai, Harry Ho Man Ng, Duoduo Wu, Yu Ting Felicia Wee, Jeffrey Chun Tatt Lim, Joe Yeong, and Tony Kiat Hon Lim. 2020. Overview of multiplex immunohistochemistry/immunofluorescence techniques in the era of cancer immunotherapy. *Cancer Communications* 40, 4 (2020), 135–153. <https://doi.org/10.1002/cac2.12023>
- [33] Gusi Te, Wei Hu, Amin Zheng, and Zongming Guo. 2018. RGCNN: Regularized Graph CNN for Point Cloud Segmentation. In *2018 ACM Multimedia Conference on Multimedia Conference, MM 2018, Seoul, Republic of Korea, October 22-26, 2018*, Susanne Boll, Kyoung Mu Lee, Jiebo Luo, Wenwu Zhu, Hyeran Byun, Chang Wen Chen, Rainer Lienhart, and Tao Mei (Eds.). ACM, 746–754. <https://doi.org/10.1145/3240508.3240621>
- [34] Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. 2016. Order Matters: Sequence to sequence for sets. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).
- [35] Yiqun Xie, Emre Eftelioglu, Reem Y Ali, Xun Tang, Yan Li, Ruhi Doshi, and Shashi Shekhar. 2017. Transdisciplinary foundations of geospatial data science. *ISPRS International Journal of Geo-Information* 6, 12 (2017), 395.
- [36] Yinyin Yuan. 2016. Spatial Heterogeneity in the Tumor Microenvironment. *Cold Spring Harbor Perspectives in Medicine* 6, 8 (2016). <https://doi.org/10.1101/cshperspect.a026583>

- [37] Hongming Zhang and Jibei Chen. 2018. Current status and future directions of cancer immunotherapy. *Journal of Cancer* 9, 10 (2018), 1773–1781. <https://doi.org/10.7150/jca.24577>
- [38] Ling Zhang, Le Lu, Isabella Nogues, Ronald M Summers, Shaoxiong Liu, and Jianhua Yao. 2017. DeepPap: deep convolutional networks for cervical cell classification. *IEEE journal of biomedical and health informatics* 21, 6 (2017), 1633–1643.