Visinity: Visual Spatial Neighborhood Analysis for Multiplexed Tissue Imaging Data

Simon Warchol[1,2], Robert Krueger[1,2,3], Ajit Johnson Nirmal[3], Giorgio Gaglia[4], Jared Jessup[2], Cecily C. Ritch[4], John Hoffer[3], Jeremy Muhlich[3], Megan L. Burger[5], Tyler Jacks[5], Sandro Santagata[3,4], Peter K. Sorger[3], and Hanspeter Pfister[2]

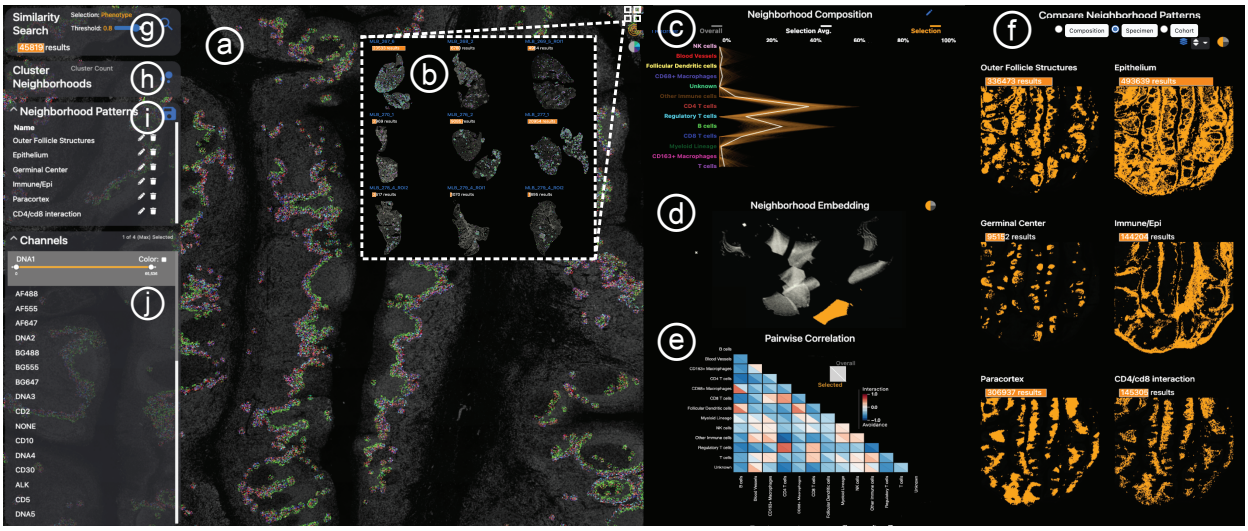

Fig. 1: Visinity interface. a) Image viewer: multiplex whole-slide tissue images highlighting spatial cell arrangement. b) Cohort view: search, apply, compare spatial patterns across different specimens. c) Neighborhood composition view: visualizes cell types that make up cell neighborhoods; d) UMAP embedding view: encodes cells with similar neighborhood as dots close to each other; e) Correlation matrix: pairwise interactions between cells; f) Comparison & summary view: different small multiple encodings of extracted patterns. g) Neighborhood search: finds cells with similar neighborhood; h) Interactive clustering: automated detection of neighborhood patterns; i) Annotation panel: save and name patterns; j) Channel selection: color and combine image channels.

Abstract— New multiplexed tissue imaging technologies have enabled the study of normal and diseased tissues in unprecedented detail. These methods are increasingly being applied to understand how cancer cells and immune response change during tumor development, progression, and metastasis as well as following treatment. Yet, existing analysis approaches focus on investigating small tissue samples on a per-cell basis, not taking into account the spatial proximity of cells, which indicates cell-cell interaction and specific biological processes in the larger cancer microenvironment. We present Visinity, a scalable visual analytics system to analyze cell interaction patterns across cohorts of whole-slide multiplexed tissue images. Our approach is based on a fast regional neighborhood computation, leveraging unsupervised learning to quantify, compare, and group cells by their surrounding cellular neighborhood. These neighborhoods can be visually analyzed in an exploratory and confirmatory workflow. Users can explore spatial patterns present across tissues through a scalable image viewer and coordinated views highlighting the neighborhood composition and spatial arrangements of cells. To verify or refine existing hypotheses, users can query for specific patterns to determine their presence and statistical significance. Findings can be interactively annotated, ranked, and compared in the form of small multiples. In two case studies with biomedical experts, we demonstrate that Visinity can identify common biological processes within a human tonsil and uncover novel white-blood networks and immune-tumor interactions.

Index Terms—Visual analytics, spatial analysis, tissue imaging

1 Introduction

Tobler's first law of geography, *"Everything is related to everything else, but near things are more related than distant things"* [72] emphasizes the importance of spatial proximity. This is not limited to geographical phenomena. It is also applicable to many biological systems [51]. Biological tissues comprise numerous cell types that function together in multi-cellular units that are crucial in development, normal physiology, and disease. In cancer, the interactions between tumor cells and immune cells are of particular interest because these contacts dictate whether tumor growth is controlled or proceeds unrestrained [17]. Recent tissue imaging methods permit the identification and quantification of tumor and immune cell types within cancer tissue. Important functional interactions between cells can be inferred by identifying cells that are next to each other. In addition, higher-order arrangements of cells that may represent a structural or functional component of tissue can be quantified by determining which cells tend to neighbor each other (spatial neighborhood analysis). These 'recurrent neighborhoods' can assemble to compose more extensive spatial patterns.

Detecting such patterns poses substantial challenges. Recent research found that spatial patterns must be investigated across large regions of tissues to yield biologically and statistically meaningful results [41], necessitating methods capable of spatial analysis at

- [1] *Simon Warchol and Robert Krueger contributed equally to this work.*
- [2] *Harvard John A. Paulson School of Engineering and Applied Sciences*
- [3] *Laboratory of Systems Pharmacology, Harvard Medical School*
- [4] *Brigham and Women's Hospital*
- [5] *David H. Koch Institute for Integrative Cancer Research, MIT*

large scales. Experts in the fields of pathology, cancer research, and systems pharmacology thus have acquired whole-slide images from tissue sections using fluorescence microscopy techniques, such as CyCIF [39], with an overall size of up to 60k×60k. This results in highly multiplexed tissue images, often larger than 100GB in size. Moreover, a single experiment can involve imaging ten or more specimens as part of a larger cohort. Identifying spatial neighborhood patterns in such data requires scalable computational methods. However, such techniques alone cannot fully replace the human mind; experts have expansive domain knowledge of cell and tissue morphology formed through years of visually investigating tissues. There is thus a need to facilitate visual human-in-the-loop data exploration, permitting these experts to guide pattern identification and verification. Yet existing visual approaches [63, 70, 71] for spatial neighborhood analysis, by design, only scale to single images representing small regions of tissue and are limited in their interactive capabilities.

We addressed these challenges as a team of visualization researchers, cancer researchers, pathologists, and cell biologists via a process of goal specification, iterative design, and tool deployment in a biomedical research laboratory. We make the following contributions: **(1)** A domain-specific human-in-the-loop workflow to visually analyze, extract, and summarize spatial interaction patterns within and across datasets. This workflow enables both exploratory analysis and confirmatory analysis through semi-automatic pattern detection and visual querying. Identified patterns can be annotated with information about their biological context, compared, and saved for continued study. **(2)** A scalable and flexible computational pipeline to quantify cellular neighborhoods and their spatial arrangements into larger spatial microenvironments (patterns). This pipeline quantifies the spatial neighborhood of each cell as a vector of surrounding cell types in a defined query range. We group similar neighborhoods and verify their significance through permutation testing, allowing for the identification of meaningful spatial patterns within and across tissues. **(3)** A scalable visual analytics system named *Visinity* to interactively analyze the computed neighborhood patterns in and across the large whole-slide tissue image data. Visinity consists of a web-based multiplex image viewer with different rendering modes and superimposed neighborhood encodings. Image exploration is linked to projections and parallel coordinates highlighting frequent neighborhoods and their composition. Small multiple arrangements of these views allow side-by-side comparison and summarize findings.

We evaluate the applicability of our approach in two hands-on case studies with biomedical experts. We first demonstrate that our system can detect well-established spatial patterns of immune cells in a healthy human tonsil specimen. Second, we analyze a cohort of specimens from a genetically engineered mouse model of lung cancer, revealing immune cell interactions that are an area of cutting-edge research in oncology. We report on user feedback on Visinity's functionalities and demonstrate the tool's computational scalability.

2 Related Work

2.1 Visual Spatial Analysis of Biomedical Imaging Data

A wide variety of bioimaging data viewers (e.g., OMERO Pathviewer [27], ViV [43]), Napari [68], Cytomine [61], Minerva [56]) and visual analysis tools (e.g., ParaGlyder [47], Vitessce [13], Facetto [31], Scope2Screen [24]) are used to study multiplexed tissue images and derived feature data. Generally, these systems focus on visual exploration and cell-type identification but they are not intended to analyze interactions between cell types and larger spatial neighborhoods that tissue microenvironments are composed of.

A small subset of tools goes beyond single-cell analysis to investigate neighborhood patterns. CytoMAP [71] is a computational toolbox designed to analyze spatial patterns in highly-multiplexed imaging data. Similar to our approach, it uses radial queries to compute local neighborhoods and visualizes their composition and arrangements. However, the static plots that CytoMAP provides do not allow for interactive exploration and search. ImaCytE [70] and HistoCat [63], and Halo [3] offer interactive spatial analysis capabilities through linked views. These approaches visualize cell-cell interaction as matrices, superimposed links in the image space (Halo), interaction networks (HistoCAT), and as aggregated glyph-based representations of frequent neighborhoods (ImaCytE). While the proposed visual encodings were evaluated as effective for exploring cell interactions, their confirmatory analysis capabilities are limited. Visinity, by contrast, offers various visual querying capabilities to search for specific interactions.

Most existing systems do not scale to the large datasets our users work with. ImaCytE supports computation and rendering of tens of thousands of cells, whereas whole-slide tissue images often contain upwards of a million cells. Visinity enables this through scalable multi-resolution WebGL rendering, spatial indexing, and algorithms that operate at interactive rates. Additionally, like CytoMAP, most approaches offer isolated analysis of one tissue at a time, whereas Visinity enables users to analyze and compare across specimens. Somarakis et al. [69] support such cohort analysis of pairwise cell-cell interactions through explicit visual encodings (raincloud plots and heatmaps), but data sizes are limited to 10^5 cells per dataset (Visinity scales to 10^7 cells). Finally, to test the statistical significance of identified spatial patterns, HistoCAT and ImaCYtE rely on permutation testing. We extend these methods with efficient parallelization, precomputation, and visualization to make them scalable, understandable, and interactively adjustable (see **Sec 6**).

2.2 Visualization of Spatial Interaction

Our approach also draws more broadly on work visualizing spatial interactions outside of the biomedical domain, such as movement and communication between geographic areas. A straightforward approach is to display such interactions in their spatial dimensions, e.g., on top of a map or image. Advantages of this approach are the familiarity of reading maps as well as emphasizing the spatial auto-correlation in the interaction data. ArcGIS [1], among other geographical information systems, offers statistical methods and visual encodings [4, 65] to compute and display spatial dependencies, involving spatial correlation, clustering, and alignment of spatial objects (shape, center, orientation). Results are usually superimposed on the map. To show spatial interaction, flow maps [82] are a common practice. Varying opacity [78], spatial aggregation [54, 74], and edge bundling [54] are common methods to resolve clutter in these views. However, while flow maps are well suited for tracking interactions across large distances, cell interactions in tissue usually form a more planar graph with local connectedness. Glyph overlays [60] can indicate a direction (i.e., tensor fields) and additional features without obscuring the underlying data. In our data, interaction is not explicitly defined but indicated through spatial proximity. We thus decided not to emphasize interaction by visual edges or glyphs. Instead, we use minimalistic color-coding to highlight cell types in user-based selections and contours (concave-hulls) to emphasize the unity of detected neighborhood clusters, while keeping the underlying image data visible.

Interaction patterns have also been displayed in abstract (non-spatial) views where they can be visualized in aggregation, such as node-link diagrams [5] and matrices [82]. Other systems [18, 33], similar to Visinity, apply a combination of coordinated spatial and abstract views [59], offering different perspectives on the data.

Visinity also draws on visual querying to search spatial interaction patterns. PEAX [35] introduces visual querying for pattern search in sequential data where users query by example and interactively train a classifier to find similar patterns. Krueger et al. [55] propose a visual interface to sketch, query, display, and refine spatial interactions between moving objects. Our approach adapts and task-tailors this workflow for neighborhood analysis in multiplexed tissue data. Search can be triggered by selecting existing patterns in the image or by explicitly sketching a neighborhood composition.

2.3 Computational Spatial Analysis Methods

Relevant spatial analysis methods can be categorized into: (A) statistical methods measuring spatial distribution/ correlation of single or pairwise features, (B) approaches detecting higher-order feature interaction such as topics and motifs from tabular data, and (C) image-based approaches to find reoccurring spatial features.

A) Spatial (Auto-)Correlation Methods: Moran's I [46] and Geary's C [12] are commonly used methods to quantify spatial autocorrela-

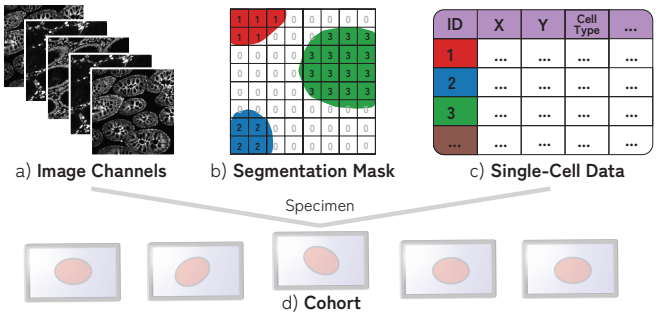

Fig. 2: A specimen consists of **(a)** multi-channel image data, **(b)** segmentation mask of cells (often $> 10^6$ cells), and **(c)** single-cell data containing information about the position, cell type, and marker intensity values for each cell. **(d)** Specimens are often part of cohorts.

tion, which determines how a variable is distributed spatially. Such correlation can be computed for every data feature (e.g., gene/protein expression) in isolation to describe their spatial organization. Ripley's K function [58], among others, extends this by computing random, dispersed, or clustered distribution patterns at varying scale, for one (univariate) or between two features (bivariate).

B) Higher-order Interactions between Multiple Features: To identify higher-order patterns, the spatial relatedness of objects can also be modeled as a network. Ribeiro et al. summarize the field of subgraph counting and motif discovery [57]. However, the number of cells (nodes in the graph) in combination with a variety of cell types (node attributes) renders current motif discovery algorithms computationally infeasible for interactive setups, especially when considering interaction at multiple scales. The complexity of most algorithms grows exponentially with motif sizes [77]. Other approaches compute groups (topics or clusters) based on probabilities and distances. Zhu et al. [83] use a Hidden-Markov random field to model spatial dependency of gene expression, and Spatial-LDA [75] is a probabilistic topic modeling approach, which is also applied in the biomedical field [48]. Most similar to our approach, stLearn [53] and CytoMAP [71] rely on traditional distance-based clustering. They extract per-cell features from histology images, collect neighborhood information for each cell (represented as vectors), and cluster these vectors to reveal neighborhoods patterns. We formalize and combine these distance-based concepts into a computational neighborhood quantification pipeline. We improve existing methods by inverse distance weighting [2], adding similarity search, and creating a scalable implementation allowing to change neighborhood sizes on the fly and analyze spatial patterns at different scales.

C) Image-based Approaches: Other methods directly operate on the image. Common deep learning approaches [62] include representation learning for comparing and finding similar image(-regions) [10, 80] and (region-based) convolutional neural networks [15], for object classification and localization. Bengio et al. [8] surveys representation learning for multi-view data (e.g., multiplexed image data). However, supervised approaches are difficult in a biomedical context because of a lack of labeled data and their explainability. By building our approach on single-cell data (see **Sec 3**) derived from multiplexed images we were able to develop a more scalable and quantifiable approach.

3 Whole-slide Multiplexed Tissue Imaging Data

The data created by our biomedical experts consists of multiplexed images of tissue generated by iterative staining with antibodies that recognize specific proteins followed by imaging with a high-resolution optical microscope in successive cycles. Our collaborators commonly use cyclic immunofluorescence (CyCIF) [40] to generate this data, though other imaging technologies [14, 16] are capable of producing similar image data. Individual cells in the image are then segmented. Based on the relative protein expression levels that are present within each cell, most cells can be assigned to a specific cell type [36, 50].

This process, therefore, yields the following data for a single specimen (Fig. 2): **(a)** *a multi-channel tissue image*, where each channel corresponds to a different 'marker' (often proteins are recognized by an antibody), **(b)** *a cell segmentation mask* that conveys the location of each individual cell in the image space, and **(c)** *single-cell data*: a feature table that includes position and cell type for each cell. **(d)** Specimens can be parts of 'cohorts' that provide biological replicates and are then investigated in conjunction with one another. This data is sizeable; single slides range from $2cm^2$ to $6cm^2$ in size and can contain up to 10^7 cells. The tissue image, given the number of channels, contains up to 10^9 pixels, resulting in an image file as large as 200GB. All analysis in this paper uses the OME standard for microscopy image data [37] and is generated with the MCMICRO [64] image processing pipeline.

4 Goal and Task Analysis

To understand the needs of domain experts in the field, we surveyed a group of 6 biologists and 3 pathologists, all affiliated with either Harvard Medical School, Dana-Farber Cancer Institute, or Brigham and Women's Hospital. From the questionnaire and monthly meetings over a period of one year, we identified and refined a set of high-level domain **goals** from which we derived specific **tasks** as guidelines for an effective visual analytics system. In essence, we fulfilled the translator role put forth in the design study methodology by Sedlmair et al. [66], requiring a comfort level with task abstractions in computer science.

4.1 Goals

G1. Experts are interested in identifying how specific cell types attract or repel each other (cell-cell interaction). When immune and cancer cells are frequently observed close to one another (in each other's spatial neighborhood) they are likely to interact. For instance, the interaction of two types of immune cells (B and T-cells) is a central tenet of protective immunity. These interactions can occur at various scales (cells directly proximate to each other or in the same large region of tissue).

G2. With these neighborhoods as building blocks, experts seek to understand their spatial arrangement within the tissue image. These neighborhood patterns can be equally distributed throughout the image or appear in proximity, forming biologically meaningful spatial structures (groups). Germinal centers, which are regions within lymph nodes where B cells proliferate, are an example of such a micro-structure.

G3. Experts seek to validate the statistical significance of identified neighborhood patterns and larger spatial structures **(G1, G2)** within a specimen and between specimens. By understanding how often they appear and how properties such as composition, size, etc., vary between patterns, one can determine their clinical relevance and whether they are subject to closer investigation.

G4. Finally, experts want to connect patterns present in an image back to biological and clinical information. They hope to determine how the presence of specific patterns correlates to specific cancer therapies, the growth of tumors, and immune response to those tumors, with an overall goal of improving cancer diagnosis and treatment.

Our survey showed that the scientists are interested in performing both exploratory and confirmatory analysis to achieve these goals. In their highly experimental settings, discovering novel spatial neighborhood patterns and thereby generating new hypotheses is of similar importance as the ability to express and verify existing hypotheses.

4.2 Tasks

T1: Visually explore spatial neighborhoods (G1, G2) from different perspectives. This includes navigation, visual identification, and selection of regions of interest in the tissue as a means for exploring the spatial neighborhoods present in tissue or within many tissues.

T2: Group similar cell interactions, through which experts can identify the larger structures formed by these cell-cell interaction patterns (**G2**). These grouping strategies must scale to the large data, even grouping the patterns present in multiple images at once, while also being interactively configurable to incorporate users' domain knowledge.

T3: Express and search for hypotheses (**G3**). This includes the ability to query for specific user-defined cell-cell interactions as well as search by example, i.e., find additional occurrences of the pattern present in a ROI. It also includes querying across cohorts.

T4: Compare the contents and spatial expression of patterns identified within a single tissue dataset as well as across a cohort (**G3**),

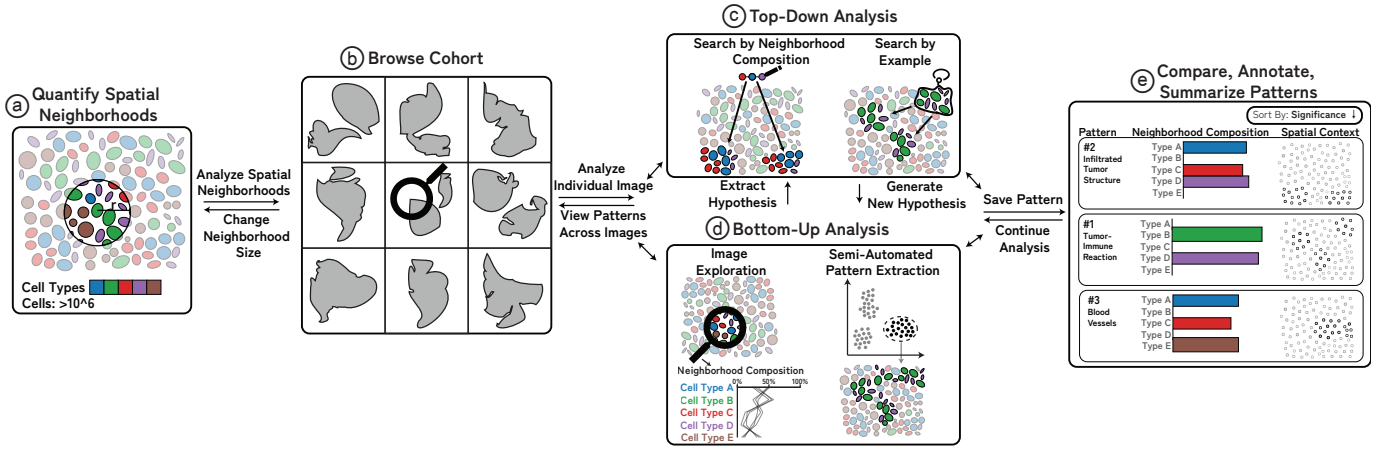

Fig. 3: **Visinity Workflow**: (a) Neighborhood quantification: users choose a spatial range triggering neighborhood vector computation; (b) Browse cohort: small multiples of specimen to gain overview of neighborhood patterns; (c) Bottom-up analysis: explore spatial arrangements and cell-type composition of neighborhoods, generate hypotheses, cluster, and extract patterns; (d) Top-down analysis: two visual querying capabilities allow hypothesis generation and search for similar patterns. (e) Pattern annotation and comparison within and across datasets.

while taking into account the biological and clinical context of regions or tissues (**G4**). In this context, spatial expression refers to the presence of a pattern within one or many tissues.

T5: Rank the presence and statistical significance of pattern within a tissue, building on the previous task. This augments comparison and helps experts better understand these patterns (**G1, G2**) and how they differ within and across specimens (**G3**).

T6: Extract, annotate, and save found patterns along with biological and clinical information (**G4**), as analysis is ongoing and not limited to isolated sessions. This information may relate to the source of the tissue, specific treatments involved, or patient outcomes. This allows for continued analysis, where gained knowledge is applied to new tissue samples (**G3**) and insights can be shared with other experts.

5 Workflow

From the identified goals and tasks (Sec. 4) and in following bi-weekly sessions with the experts, we extracted an iterative human-in-the-loop workflow (Fig. 3) for spatial neighborhood analysis that guided the design of our visual analytics system Visinity (see Sec. 7 and Fig. 1).

After importing the data, users can specify a neighborhood size, depending on if they are interested in more local or global interaction patterns (Fig. 3, a). This size can be changed throughout the analysis, triggering a neighborhood quantification process, described in Sec. 6. Once neighborhoods are quantified, users can gain an overview of the tissue environments across all specimens in their cohort **(T0)** in the form of small multiples (Fig. 3, b). Users can start to explore the image (bottom-up analysis) (d) and select regions of interest in the tissue to visualize spatial neighborhoods present in the tissue **(T1)**. This is covered in detail in Sec. 7. To aid exploration, they can interactively cluster neighborhoods in an entire tissue **(T2)**, extracting patterns in a semi-automated manner (Sec. 8.1). Top-down analysis (d,**T3**) to test and refine existing hypotheses can either be done by specifying the cell types involved in a neighborhood pattern or by selecting a region of interest in the tissue as an example of a spatial pattern (Sec. 8.2). Both trigger a search within and across tissue images for matching neighborhoods. Based on the contents, spatial context within the tissue, and statistical significance of results (quantified through permutation testing), users refine their hypotheses and can save and annotate identified patterns **(T6)**. Finally, a user can **(e)** compare these saved patterns within a dataset and across datasets. When comparing across tissues, users can determine if a pattern in one tissue is present and statistically significant in another **(T4, T5)**, as described in Sec. 9.

6 Quantifying Spatial Neighborhoods

According to our biomedical experts, a cell's neighborhood is defined by the cells in spatial proximity within a defined spatial distance. Groups of cells with similar neighborhoods form patterns that assemble to compose more extensive spatial arrangements. To explore **(T1)**, group **(T2)**, search for **(T3)**, compare **(T4)**, rank **(T5)**, and save **(T6)** cell neighborhoods and the patterns they form, we propose a neighborhood quantification pipeline, building on existing work in the field [70, 71]. We extend these methods with a more scalable implementation to interactively investigate patterns of different length-scales in whole-slide imaging data, with spatial distance weighting to reflect the neighborhood influence of cells, and with higher-order permutation testing to determine the significance of found proximity patterns.

Our computational pipeline works in 5 steps (1-5):

Step 1: We build a ball tree with the coordinates of each cell in the image. This takes $O(n+k)$ to perform spatial range queries, where n is the number of points and k the number of points returned in that range [28] (Fig. 6,a).

Step 2: We create a feature vector representing each cell's neighborhood of size 1 x c (where c is the number of cell types in the dataset). Each column corresponds to the fraction of a cell's overall neighborhood occupied by a specific cell type. These values are linearly weighted such that cells closer to the center of the neighborhood radius contribute more to the overall neighborhood. The resulting vector is normalized. This representation builds on existing approaches [70, 71] and was reaffirmed by feedback from our collaborators, who said that it was highly interpretable and fit with the hypotheses they had regarding the cell-cell interactions present in a dataset **(T3)** (Fig. 6, b).

Step 3: Repeat this process for every cell in a dataset, resulting in one vector for each cell **(c)**. We generate a matrix representing the neighborhoods in a dataset, where each row is the neighborhood of a specific cell (Fig. 6, c).

Step 4: Cells with similar spatial neighborhoods are represented by similar neighborhood vectors. To find cells with similar neighborhoods we run a distance-based nearest neighbor search **(T3)** with a threshold distance defining a cutoff of what's considered similar. To find groups of similar neighborhood **(T2)**, we utilize partition-based clustering (see Sec. 8.1 for more details). Together with our experts and based on literature [23], we evaluated vector comparison based on Euclidean distance to achieve the most satisfying results. Further advantages are fast computation, simplicity, and interpretability.

Step 5: Inspired by similar approaches [63,70], we use permutation testing [22] to determine the patterns' statistical significance **(T5)** within a specimen and across a cohort. We count individual neighborhoods that match a given neighborhood pattern based on a user-defined similarity threshold (**Sec. 8.2**). We then randomly shuffle the assigned cell types in the data, recomputing neighborhood vectors (steps **a** - **c**), and calculating the number of matching neighborhoods for each permutation. Eq. 1 describes our P value calculation.

$$P = \frac{\Sigma(\text{Permuted Data Matches} \geq \text{Actual Matches})}{\text{\# of Permutations}} \quad (1)$$

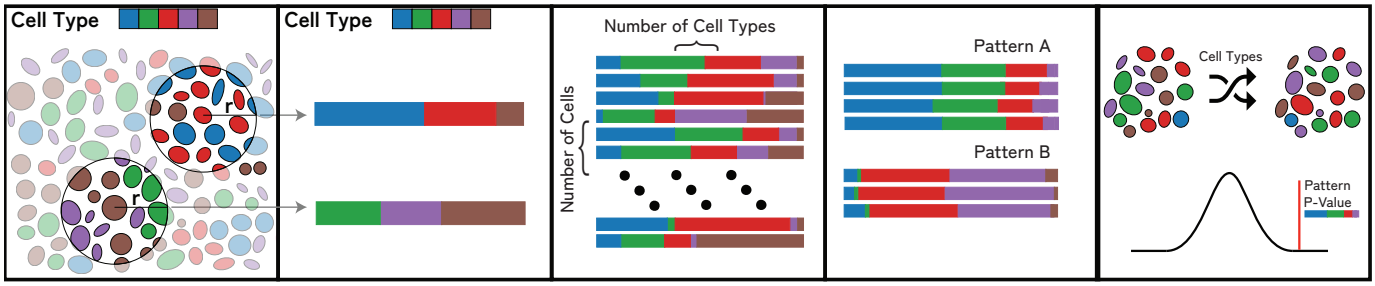

Fig. 4: **Neighborhood Quantification**: (**Step 1:**) For a cell in an example microenvironment, find all proximate cells within a specified radius. (**2**) Each cell's neighborhood is a feature vector that represents the weighted presence of each cell type in the neighborhood. (**3**) Repeat this process for each cell, resulting in a neighborhood vector for each cell in an image. (**4**) Groups of similar neighborhood vectors correspond to spatial patterns. (**5**) Randomly permute cell types in an image to determine patterns' significance.

7 Visualizing Spatial Neighborhoods

To realize the identified workflow **Sec. 5** and task **(T1-T6)** we developed Visinity - a visual analytics system for spatial neighborhood analysis. Visinity's interface (Fig. 1) consists of coordinated views offering different perspectives on detected neighborhoods, including their composition and spatial occurrence in the tissue.

Cohort Overview. After data import and neighborhood quantification (Sec. 6), users can start the analysis with an overview of all specimens in a cohort. Together with the biomedical experts, we chose small multiples of image thumbnails as a sufficient and compact way to comparatively summarize **(T4)** the images in a cohort and their morphology (Fig 1, b). Users can zoom and pan into these thumbnails to begin exploration and select a specific dataset for thorough investigation **(T1)** in the image viewer (a). As analysis progresses, spatial neighborhood patterns identified in a single specimen are detected and visualized in other members of the cohort by highlighting their spatial presence in linked images. Specimens can also be sorted by the number of matching neighborhoods to the pattern currently being investigated or by the statistical significance of that pattern within a specimen (see Sec. 9).

Spatial Exploration of Tissue Morphology. To support visual exploration of the spatial neighborhoods in tissue images (**T1**, Fig. 3, d), we offer a scalable image viewer (Fig 1, a), allowing navigation via zooming+panning. The viewer builds on our previous work Facetto [31] and Scope2Screen [24]. Through multi-channel rendering, pseudo-colored channels can be blended together into a single view, enabling users to analyze the expression level of multiple markers at once.

To mark and filter for cell neighborhoods in an ROI, users can employ an interactive lasso tool. Fig. 5, a shows a selected spatial neighborhood in a lymphoid nodule from a healthy tonsil tissue. We render selected cells with superimposed outlines so that the underlying tissue image and morphology is still visible. These outlines are optionally colored by cell type using a categorical color scale (Fig. 5, a). In this example, the region is composed primarily of two cell types, B cells and T cells, that are part of the immune system. The encoding is effective for a few image channels and cell types (Fig. 6, a), but becomes increasingly difficult to comprehend when combining multiple pseudocolored channels with categorical cell coloring. Thus, users can switch to an alternative mode that visualizes detected neighborhood patterns with a concave hull, emphasizing their unity while maintaining a non-occlusive view of the underlying image-channels (Fig. 6, b).

Composition of Spatial Neighborhoods. To provide more information about what cell types spatial neighborhoods are composed of **(T1)**, we include a parallel coordinate plot (Fig. 5, b). We chose PC plots over bar charts and box plots to emphasize the occurrence of each cell type in a spatial neighborhood while also encoding the distribution and correlation between the features. Fig. 5,b shows the composition of a selected tonsil region (a). Here, each poly-line represents the neighborhood of a cell. Each individual axis is defined by the influence of a specific cell type in the neighborhood. Two distinct neighborhoods are represented, one containing more B cells and one containing more CD4 T cells. To emphasize correlations between cell types on adjacent axes, we employ an axis reordering strategy [6, 25]. Fig. 8, a shows a negative correlation between CD4 T cells and B cells. Our experts were also interested in investigating specific pairwise interactions between cell types **(T3)**. We address this by allowing users to reorder axes by drag&drop. To compare **(T4)** the current selection to the overall composition of a specimen or cohort, the PC plot optionally encodes the neighborhoods in the entire specimen or cohort in gray, behind the current selection in orange (Fig. 1, c).

Embedding of Neighborhood Vectors. While a PC plot emphasizes neighborhood composition and can reveal correlation, it can be hard to distinguish patterns from one another due to increasing occlusion and visual clutter. To make neighborhood patterns more distinct, we perform dimensional reduction on the neighborhood vectors and visualize the reduced 2D data in a scatterplot (Fig. 5 c), where each point reflects a cell's spatial neighborhood. Cells that are close to another share a similar neighborhood and form spatial groups, while distant points (cells) indicate different neighborhood compositions. For cohort data, we create a shared embedding of the neighborhoods across the individual datasets allowing users to identify and compare similar spatial neighborhoods and discrepancies across specimens.

Dimensional reduction is a conventional and familiar practice for our intended users. There is significant biological precedent for the use of t-SNE [30, 70, 71, 73] and UMAP [7, 44, 71] to investigate spatial features in tissue images. The popularity of these approaches stems from the emphasis on local neighborhood preservation in both methods [30, 79]. T-SNE has limitations [76], namely the impact of hyperparameters, the potential lack of preservation of global structures, and higher runtime. We find UMAP advantageous because of its better preservation of global structures, especially when visualizing multiple specimens. Compared to variational autoencoders (VAEs) [29], we found that UMAP produced better spatial separation in the 2D layout, allowing users to better distinguish patterns from another and for easier selection. To perform this embedding in a scalable manner, we use the RAPIDS GPU UMAP implementation [49], which is capable of embedding million-cell neighborhood matrices in a few minutes (see Sec. 10).

Users can navigate in the scatterplot (embedding view) and make selections as they would with the image view which highlights the selection in the coordinated views. E.g., a user may notice a region of high density in the embedding Fig 5, c) and select it in order to understand the cell types that compose that visual group (b) and the locations (a) of those neighborhoods within a tissue image (image viewer) or across multiple images (cohort view). Likewise, users can select a spatial region in the cohort view and image viewer and review the greater neighborhood pattern it is part of in the embedding.

Pairwise Correlation Between Cell Types. To address the needs to explore pairwise interaction between two specific cell types **(G1)**, **(T1)** we offer a correlation matrix visualization. We chose a matrix over a node-link diagram to avoid clutter when encoding relationships between every pair of cell types. Matrix visualizations provide a consistent, and compact layout, making it easy for users to compare interactions in a specific pattern to those in the overall specimen and to other patterns **(T4)**. Inspired by existing geographic [42] and biomedical [71] approaches, we compute the Pearson correlation coefficient between each pair of cell types within the computed neighborhood vectors to quantify these pairwise relationships. Two cell types with a strong positive correlation tend to be found in similar neighborhoods, whereas two cell types with a strong negative correlation tend to avoid each

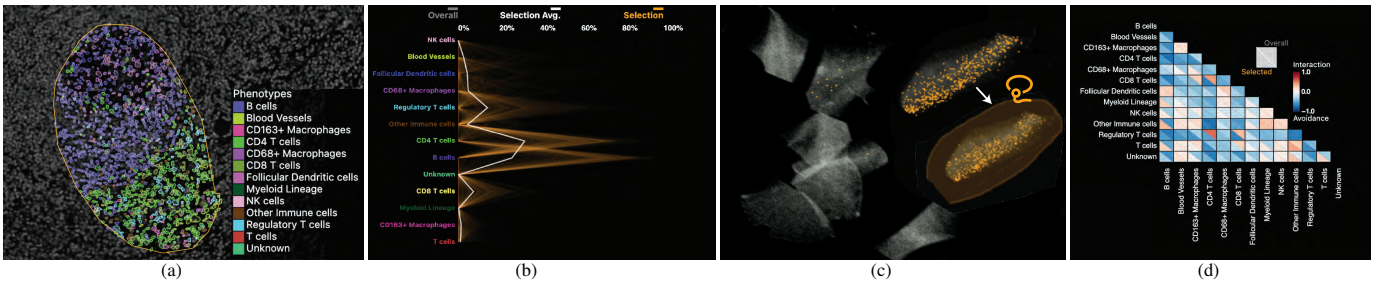

(a) (b) (c) (d)

Fig. 5: **Visual Exploration Through Linked Views**: (**a**) Selected ROI to investigate the spatial neighborhoods. Cell types are displayed with color-coded segmentation outlines. (**b**) Neighborhood composition in a PC plot - orange lines represent neighborhoods selected, exhibiting two discrete patterns. (**c**) Interactive 2D UMAP embedding of all neighborhood vectors in a specimen in grey; current selection visualized in orange. Users can select a region to explore similar neighborhoods. (**d**) Pairwise cell-cell interactions visualized as correlation matrix.

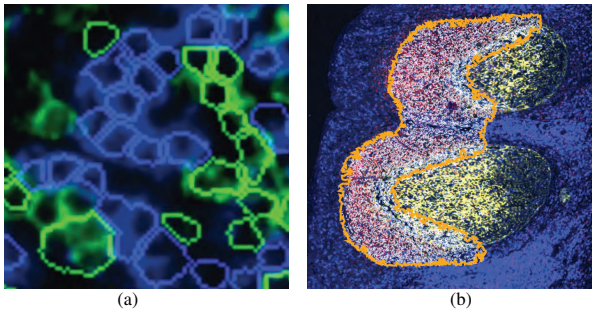

(a) (b)

Fig. 6: **Cell Outlines or Concave Hull**: Two view modes: (a) coloring cell outlines by cell type; (b) outlining patterns with a concave hull;

other. We use a diverging red-white-blue color palette to visualize these correlation values. To support comparison, we split each field into two triangles (Fig. 1, e), one representing the correlation of cells types in the overall image or cohort and one representing correlations within currently selected neighborhoods. Likewise, selecting a triangle in the matrix filters the other views to highlight neighborhoods with respective pairwise cell type correlation above a user-defined threshold. In Fig. 5, d, B cells have a negative correlation to all other cell types (all correlations are blue), indicating that the region of B cells is very homogeneous. CD4 T cells, meanwhile, have a positive (red) correlation with CD8 T cells, blood vessels, and regulatory T cells, indicating these different cell types are interacting.

8 Semi-Automated Analysis

Beyond exploration through linked views (**T1**), Visinity offers semi-automated methods to cope with the large and high-dimensional ($\geq$ million cells per specimen) data. This allows to automatically group spatial neighborhoods (**T2**) through interactive clustering and search for neighborhood patterns to test and refine hypotheses (**T3**) at scale.

8.1 Detecting Spatial Patterns Through Clustering

To automatically cluster the cells based on similar neighborhood vectors we chose to use EM (expectation-maximization) clustering for Gaussian mixture models [81]. Depending on the specific biomarkers, derived cell types, and cellular neighborhoods, we can only make assumptions about the underlying model with many latent unknown variables. The EM algorithm finds (local) maximum likelihood parameters of that statistical model given our sample data [11]. Furthermore, the EM algorithm can detect clusters that vary in shape and density, compared to, e.g., k-means, and its runtime still scales to cluster our datasets interactively without precomputation.

The interface enables users to either cluster the neighborhoods in a specific dataset and apply this clustering to the rest of the cohort, or to run the clustering on all specimens in a cohort and then drill down into individual specimens to explore results in more detail. Computed clusters are listed with other saved patterns in a list view (see Fig 1, i) and users can click on them to visualize them in each of the linked views. The clustering can be used in an iterative manner; based on the spatial context in the image, the cluster's neighborhood composition, and where the cluster lies in the embedding, a user may choose the number of clusters they desire, allowing the detection of sub-structures that may exist within a given neighborhood pattern as well as the macro-structures that contain these neighborhoods are a part of. We visualize these clusters by coloring their cells by cell type and emphasize unity with a concave hull, as we do for regional selections. Fig. 10 shows identified clusters in the tonsil data, including the B cell follicles investigated earlier and the *'Paracortex'*, containing many different types of T cells.

8.2 Hypotheses Testing With Visual Querying

Visual Querying By Region of Interest. For any region of interest found during data exploration, there might be similar regions in the specimen or throughout the cohort. Visinity enables users to execute such queries for any detected or manually selected neighborhood pattern (query-by-example), revealing similar spatial neighborhoods both within the image and in all images in the cohort (Fig. 5, g).

The computational process is outlined in Fig. 7: For any selected region (a), we first compute the mean of the neighborhood vectors in that region (b). We then compare this representative vector v_1 with each other neighborhood vector v_2 in the image using the Euclidean similarity score [67]: $\frac{1}{1+d(v_1,v_2)}$, with distance d between vectors. Neighborhoods above a specified similarity threshold are then highlighted in each of the linked views (Fig. 7, d) and the statistical significance of this hypothesis is visualized as determined through permutation testing (Sec. 6). We visualize the significance of the query and number of matching neighborhood (Sec. 9).

Representing an ROI by its mean vector is an approximation of the neighborhoods in the region. As the region increases in size, particularly when the region contains multiple patterns of interest, this gets more and more inaccurate. However, when such a region is selected, the composition and embedding views provide context to the homogeneity of this selection, allowing a user to investigate potentially discrete patterns in isolation. Moreover, to accommodate the analysis of regions of tissue of any size, users can modify the neighborhood radius to fit their needs. Through iterative design with frequent feedback from our expert collaborators, we found this approach was effective at finding similar neighborhoods, as demonstrated in Sec. 11.1 and Sec. 11.2.

Visual Querying By Neighborhood Composition. Experts indicated that they often were interested in investigating interactions between specific cell types (**T3**). To support expression and search for their hypotheses, which is a key part of the Visinity workflow (Fig. 5, f), users can look for specific interactions by defining a custom query vector in the neighborhood composition view by clicking and dragging to create a polyline representing a query vector. The current neighborhood composition and overall composition are both optionally displayed behind the query vector, allowing the user to match trends reflected in these views while building a query (Fig 8, a). The results, which represent the interaction of the specified cells, are visualized in the linked views (Fig 8). A user can perform query refinement or expansion by adjusting the similarity threshold; increasing this threshold yields neighborhoods that more precisely match the query (Fig 8, b) while decreasing the threshold casts a wider net (c). The same refinement can be applied to ROI-based queries.

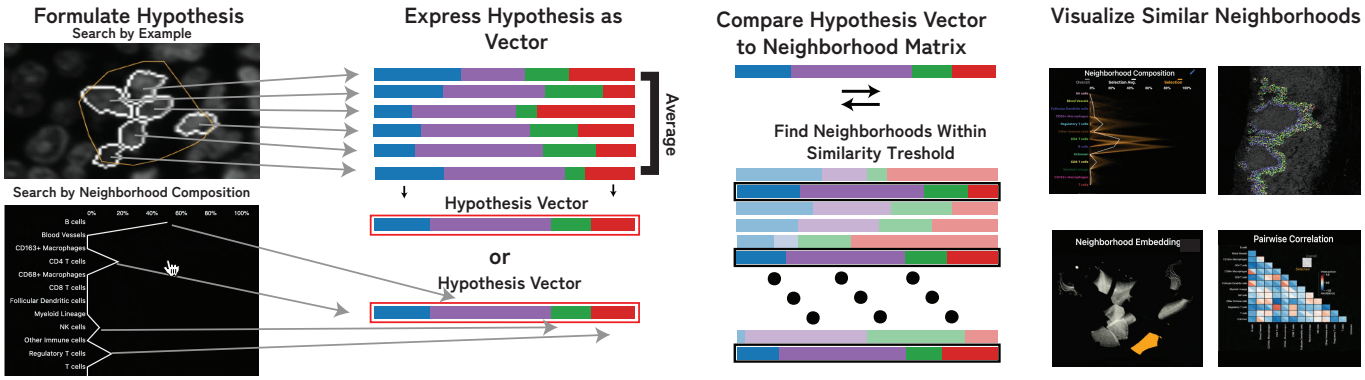

Fig. 7: **Hypothesis Testing through Visual Querying**: **(a)** Users can test spatial pattern hypotheses in the form of regions of interest or approximate neighborhood composition. **(b)** We formulate this hypothesis as a neighborhood vector and **(c)** compare that neighborhood to every neighborhood in the image. **(d)** Results above a set similarity threshold are visualized in each of the linked spaces.

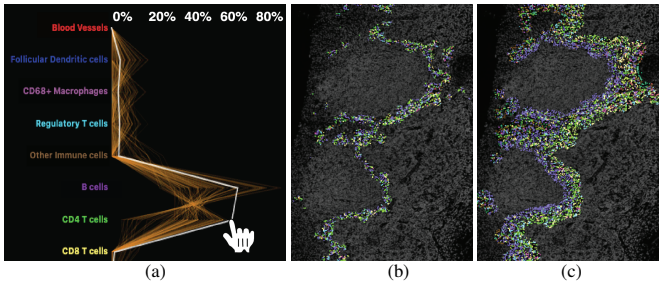

(a) (b) (c)

Fig. 8: **Search By Neighborhood Composition**: **(a)** A user can sketch a custom neighborhood and increase/decrease the threshold to find more/less similar neighborhood patterns **(b, c)**. Here, interactions between B cells (left) and T cells (right), outlining B cell follicles.

8.3 Visual Querying Across Specimen

When a user queries a specimen for a pattern, we search across the entire cohort for matching neighborhoods, which are displayed in the cohort section of the comparison view (Fig. 1, f). We additionally display the number of results and significance of the query within each specimen above the spatial summary, as discussed in Sec. 9. Users sort related specimens within the list by significance and number of results **(T5)**.

9 Visual Comparison and Result Summarization

While using Visinity, users can save and retrieve patterns and label them with clinical or biological context **(T6)**, which are displayed in a pattern list in the interface, as reflected in the proposed workflow (Fig. 3, e). Once saved, experts can compare the contents (neighborhood composition) and spatial expression (location in the image) of these patterns **(T4)**. Comparison can happen on two levels: between identified neighborhood patterns, and between specimens.

Comparison Between Neighborhood Patterns. To summarize and compare identified neighborhood patterns, users can choose between spatial and compositional comparison (Fig. 1, f). Small multiples of bar charts (Fig. 10) represent the composition of each saved pattern, such that users can compare the presence of each cell type between patterns. We chose bar charts over the initial PC plot (Sec 7) as they are a concise and simple visual encoding which allows users to easily compare the neighborhood composition of many patterns **(T4)**. We also provide spatial context regarding the patterns in a concise manner with scatterplots (Fig. 10), which encode the position of each pattern in the image space. This provides a stripped-down visual summary of the tissue imaging data. A user can select these small multiples to view them in full detail in each of the linked views and, in the case of comparison across images, to search for that pattern within a new image. On top of each plot, a bar represents the number of cells with that given neighborhood above each specimen thumbnail (Figure 9). The color of the bar encodes the statistical significance (see Sec 6, step 5) on a single hue white-orange color scale.

Comparison Between Specimens. It is also of interest to understand where and how frequently a specific neighborhood pattern appears in the different specimens **(T4)**. After a pattern is selected from the pattern list (Fig 1, i), a bar shows the number of matching neighborhoods above each specimen thumbnail of the Cohort View (Fig. 10). Again, we encode the number of cells contained in the pattern with its computed significance as bar above the respective plot.

10 Scalable Implementation

Throughout the design process, we emphasized scalability, leveraging methods and interfaces that support simultaneous analysis of gigapixel images containing more than a million cells. We feature a JavaScript client / Python server architecture with web-based frontend visualization and efficient backend computation. We built on our previous work (Scope2Screen [24], Minerva [21, 56]), storing images in the Zarr [45] format and rendering them using WebGL. The embedding view, linked image thumbnails, and spatial comparison views use the regl-scatterplot [34] library, which allows for panning, zooming, and selection in datasets containing as many as 20 million points. We combined efficient methods with strategic pre-computation to quantify and analyze spatial neighborhoods. Neighboring cells are queried using scikit-learn's [52] Ball-tree index structure. The neighborhood vector computation is compiled with Numba [32], and thus translated into efficient machine-code. When the neighborhood radius changes, we recompute all neighborhoods and save them as Zarr arrays.

To evaluate the runtime performance of this approach (Fig. 9) we generated random test data ranging in size from 10^5 cells to 10^7 cells. This synthetic data maintains the same cell density as the Tonsil dataset investigated in Case Study 1 (Sec. 11.1). We randomly give each cell one of the 13 cell types assigned in the same tonsil dataset, maintaining the same incidence of each cell type as was present in the tonsil. The greatest computational bottleneck is neighborhood quantification at increasingly large neighborhood sizes, which is primarily a result of the time needed to identify neighboring cells. We found the ball tree used to search neighbors outperformed various other similarity search implementations [19, 26] on 2D data. However, as these libraries add GPU-accelerated range search, this may no longer be the case. When computing neighborhoods, we also randomly reassign the cell types in a dataset to create new permuted versions of this neighborhood matrix, which is similarly compiled and saved as a compressed Zarr array. Similarity search across the true neighborhood vectors, as well as the permuted versions, is conducted in parallel. Neighborhood clustering is also performed using scikit-learn [52]. To interactively cluster large datasets, we train the Gaussian mixture model on a 10% random sub-sample of the cells in an image before assigning clusters across the entire specimen. As demonstrated in Fig. 9, our permutation testing, clustering, and search implementations scale linearly as the number of cells increases. Our source code will be made available upon acceptance of this paper.

11 Case Study Evaluation

We present two case studies that demonstrate the utility of our system. In each 90-minute in-person session, the experts were given a brief walk-through of the system's features and then steered the system themselves throughout the study. While analyzing the tissues, the experts were instructed to think-aloud regarding the biological context

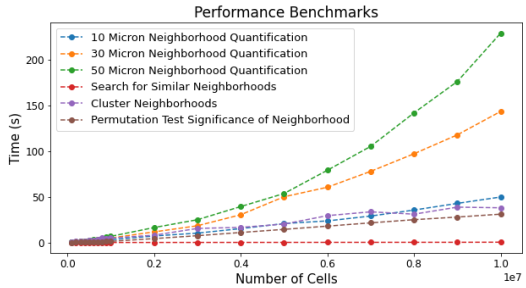

Fig. 9: Runtime evaluation for steps in the neighborhood computation pipeline. Datasize is increased gradually.

and motivation for their analysis and to provide usability feedback for the system. After completing the session, each expert filled out a brief survey quantifying the usefulness and intuitively of Visinity's features.

11.1 Case Study 1: Human Tonsil

Tonsils, which are lymphoid organs at the back of the throat, are a part of the immune system and help defend against foreign organisms. They have been extensively studied as they are dense tissues with distinct morphology and contain many immune structures and interactions. In this case study, a senior pathologist investigated a single tissue specimen of a normal palatine human tonsil. We demonstrate that Visinity can be used to identify such known spatial arrangements of immune cells.

Data: The tonsil was scanned to generate a 17,231 x 12,312 micron whole-slide image of 1.3 million cells. The imaging data is 62 GB and contains 45 channels, each 26,482 pixels × 19,065 pixels in size. Each cell in the specimen was assigned by a computational biologist to one of 13 different cell types based on patterns of protein expression.

Analysis: We began the analysis by turning on the DNA channel in the image viewer, outlining the nuclei of all cells in the specimen in the tissue image. Zooming and panning in the image revealed the expected tonsil morphology and architecture. As an initial exploratory step, we clustered the spatial neighborhoods in the tonsil into ten distinct groups. The pathologist immediately noticed three clusters occupying adjacent regions in the image that corresponded to known biological structures and interactions. One cluster formed distinct circular and oval-shaped groups. The pathologist identified this cluster as representing germinal centers which are areas of B cells, a type of white blood cell that proliferates before migrating out of the tonsil to secrete antigen-binding antibodies (Fig 10, a). Another cluster outside of the germinal centers contained numerous T cells of various types (helper, cytotoxic, and regulatory T cells) and specialized blood vessels that shuttle B and T cells in and out of the tonsil . The pathologist identified this as the T-cell zone of the tonsil (Fig 10, b). Between these clusters was a band of Helper T cells binding to B cells, facilitating the development of the B cells into antibody-secreting plasma cells or memory B cells, which remember information about specific antigens so the body can defend against them in the future (Fig 10, e).

The pathologist then investigated the distribution of macrophages in the tonsil, which are cells that engulf and consume other cells and attract immune cells. In this tissue, two types of macrophages were identified based on the proteins they express, CD68 and CD163. The CD68 macrophages were primarily in the germinal centers, while CD163 macrophages were located in the zone around the lymphoid follicles. We selected each type of macrophage in the cell-type legend and found that these neighborhoods occupied two distinct regions of high density in the embedding. We saved both patterns, labeling each by the type of macrophage it contained. In the spatial comparison view, the pathologist noted that the CD68 macrophage neighborhoods were located in the interior of the tonsil, near and within the B cell follicles, whereas CD163 macrophages were closer to the epithelium that covers much of the surface of the tonsil (Fig 10, c,d). The comparison bar charts demonstrated that while both macrophage neighborhoods contained roughly the same fraction of Helper T cells and Regulatory T cells, the CD68 macrophage neighborhoods were far richer in B cells, whereas the CD163 macrophage neighborhoods had more T cells and fewer B cells suggesting differential roles in B cell development.

In addition to identifying common immune interactions and distributions, the analysis revealed less well-appreciated patterns. The correlation matrix showed that Cytotoxic T cells and Regulatory T cells had the highest correlation in the specimen. We selected this index in the matrix to find neighborhoods containing both cell types. Regulatory T cells help prevent the immune system from attacking healthy cells and the pathologist noted that the co-localization with Cytotoxic T cells is consistent with a modulatory role for the regulatory cells in controlling cytotoxic T cells, which kill other cells. In the composition view, most of the neighborhoods were shown to contain a roughly equal proportion of the two cell types as well as some helper T cells, but some neighborhoods were visible with a higher incidence of Regulatory T cells. Sketching such a pattern and searching for similar neighborhoods revealed small groups of cells underneath the tonsil membrane and outside of the B cell follicles (Fig. 10, f). The pathologist indicated that it was interesting to find Regulatory T cells clustered together since they are generally more evenly dispersed throughout tissues.

11.2 Case Study 2: Lung Cancer in Mouse Tissues

Mice share many biological similarities with humans and are commonly used as a model organism to study cancer biology. During this case study, a senior biologist investigated 10 lung tissues belonging to a cohort of mice that were genetically modified to develop multiple small tumor nodules. Within this cohort, half of the specimens were from mice that develop tumors which elicit a very poor immune response (immune-poor) while the other half were engineered to express an antigen that strongly activates Cytotoxic T cells, which drives a robust infiltration of immune cells into the tumors (immune-rich). The objective of this study is to identify spatial organization and molecular features of immune cells that prevent tumor growth and to characterize the features of the immune response between the two specimen types.

Data: This cohort was composed of 10 specimens, representing 740 GB of imaging data. Each image contains 31 channels and is more than 600 million pixels in size. In total, these specimens contain 2.6 million cells, each assigned to one of 17 cell types by the senior biologist.

Analysis: As lungs contain many regions with low cell density, we chose a relatively wide 50 micron neighborhood radius. The biologist began in the cohort view to get an overview for each specimen and its morphology and then looked at the cell types with the highest pairwise correlation across the cohort in the matrix visualization, which were Dendritic cells and T cells. These cells are known to interact; Dendritic cells are messengers, presenting antigens to the surface of T cells, activating the T cells (helper or killer T cells depending on the antigen), which then combat tumors. The biologist indicated that identifying this interaction confirmed the accuracy of the cell-typing across the cohort.

The biologist next investigated the neighborhood embedding and changed the plot to color by cell type; doing this revealed a region containing many different immune cells, which was confirmed by selecting this region and visualizing the neighborhoods it contained in the composition view. Adjoining this region in the embedding, the biologist identified a large cluster composed of a high percentage of epithelial cells. The biologist selected one half of the epithelial region of the embedding; when visualized spatially across all specimens, this region corresponded to many of the epithelial cells that make up the airways in the lung. The other half was identified as containing many of the areas of tumor. The biologist clicked on a specimen in the cohort view to investigate it individually. Turning on the channel for TTF1, a biomarker often used to recognize lung tumors, allowed the biologist to identify many separate tumor nodules throughout the tissue, two of them in close proximity to each other. After selecting one of the tumors with the lasso tool, the outlined cells and composition view showed mostly epithelial cells, as well as a small number of macrophages and Cytotoxic T cells on its periphery. The biologist identified this tumor as not infiltrated by immune cells (immune-poor). We used neighborhood similarity search to identify this pattern across both this single specimen and the entire cohort of specimens, identifying several other structures within the specimen that the biologist determined to be immune-poor tumor. When assessed across the cohort and sorted by the number of matching neighborhoods, the biologist found that far fewer of these

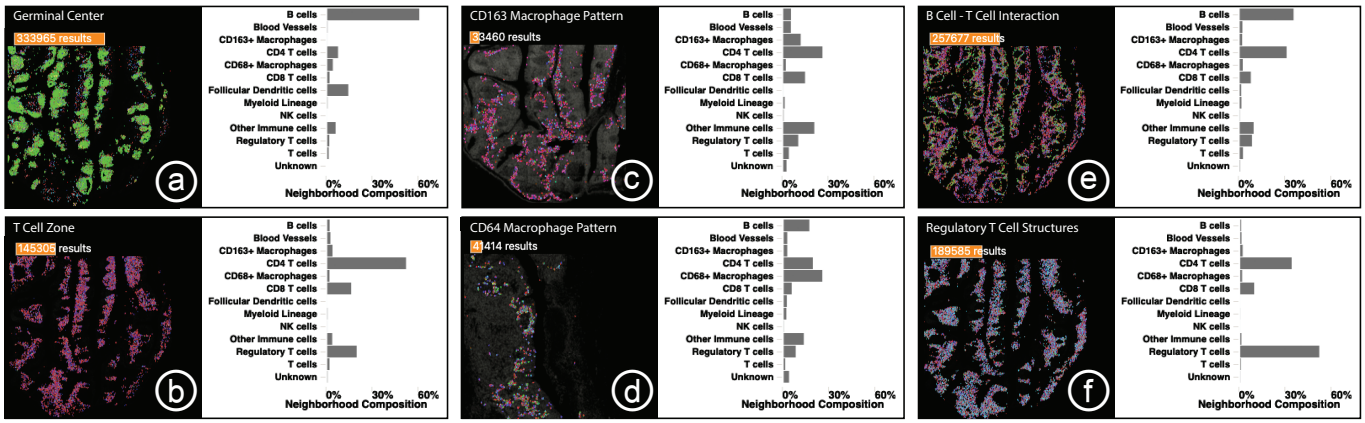

Fig. 10: **Case Study 1**: Small multiple comparison views summarize the biological structures and processes found in a human tonsil, including key interactions between B cells, T cells, and macrophages as well as unexpected regulatory T cell structures.

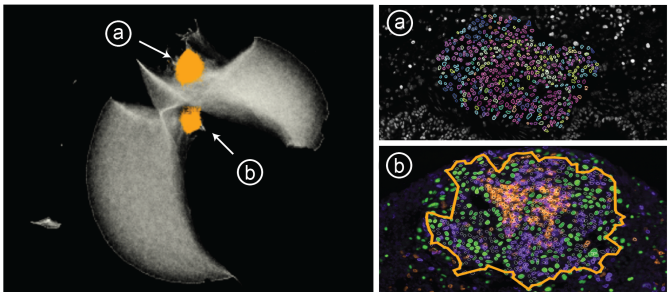

Fig. 11: **Case Study 2**: Distinct regions in the embedding represent **(a)** immune structures far from tumors and **(b)** tumors infiltrated with immune cells. Tumor cells (green), B cells (orange) and T cells (purple).

structures were identified in the immune-rich cohort than in the immune-poor cohort, consistent with the lack of tumor antigen expression and immune activation in the immune-poor cohort. The biologist then selected an adjacent tumor in the image which did not belong to the noninfiltrated tumor pattern. The linked views demonstrated that this tumor contained many B cells as well as Cytotoxic and Helper T cells. Despite arising in an animal that was not engineered to express a tumor antigen, this tumor was infiltrated by immune cells, highlighting the natural variability present in animal models of disease.

11.3 Feedback and Survey

We collected think-aloud feedback [9] from the pathologist (P1) and biologist (B1) during the case studies. We also held additional 30-minute hands-on sessions with two biologists not involved in the project. The users then filled out a survey to rate Visinity's features on a 5 point Likert scale ranging from strongly disagree to strongly agree.

All users rated the interface design as intuitive and accessible (5x strongly agree). They particularly liked the search by region of interest (4x strongly agree) and based on composition (4x agree). Both P1 and C1 emphasized that searching for known cell-cell interactions is a powerful way to verify that cell types and begin analysis of a tissue. While all users strongly agreed that the matrix helps to identify co-localization of cell types and can highlight pairwise interaction patterns, P1 indicated that the lack of stronger correlations and, thus light shades of blue or red, in the Tonsil dataset made it more difficult to interpret. The composition view was rated helpful, and intuitive (3x strongly agree, 1 agree). B1 suggested allowing users to switch between displaying the relative and normalized presence of a cell type in neighborhoods on the x axis of the PC plot; some types occurred infrequently in the mouse cohort, making it hard to compare their incidence in neighborhoods. In addition to looking at cell types, P1 and the two additional biologists noted that while the relationships between certain markers and cell types are known, incorporating marker intensity values more directly might lead to discovery of new relationships.

P1 particularly liked switching between user exploration and cross-sample testing. As an extension, he suggested to add additional functionality to facet specimen into individual sets, in order to facilitate regional comparison in a similar way. When comparing patterns across specimens, users found that patterns were either wholly statistically significant or insignificant. B1 noted that while permutation testing is the standard way of determining significance, other statistical approaches that did not assume that the axis of variability was similar to the axis of information might yield more nuanced results. B1 stated that moving from a single sample to groups of samples made Visinity uniquely suited to robust and reliable biological discovery and was not included in any other data visualization tool he had been exposed to.

12 Conclusion and Future Work

We present Visnity, a visual analytics system for investigating spatial neighborhood patterns within and across tissues in a flexible human-in-the-loop workflow. We demonstrate the applicability of our system to identify biologically meaningful spatial neighborhood patterns in cancer tissues. In addition to this ongoing work, we identify three key avenues of future research, ranging from the short to the long term.

Extracting Image-Based Features. Our methods for quantifying and identifying spatial patterns are built for single-cell information, namely position in the image and cell type. However, the resolution of novel imaging methods is increasing. Reducing high-detailed images of cells to single intensity values or type classification causes a huge loss of valuable information. Computer vision models could help to capture biological structures based on their shape, as well as marker polarization within cells, indicating if cells attract or repel each other. While many deep learning models for image-based feature extraction are a black box, visual analytics could aid interpretability and steerability.

Moving Beyond 2D Imaging. Recent developments in tissue imaging have begun to produce multiplexed high-resolution 3D volumes. While our scalable method for quantifying neighborhoods adapts well to 3D data, such multi-volumetric data poses new challenges in designing suitable visual encodings, scalable volume+surface rendering strategies, constraint and guided navigation, and interaction.

Identifying Spatial Signatures. As spatial analysis of tissue imaging data becomes more prevalent, it would be helpful to extend Visinity's existing ability to save significant spatial patterns into a knowledge database similar to the Molecular Signature Database [38] that serves as a vital reference in genomics or even assemble the relations into a knowledge graph representation as provided in the INDRA project [20]. Accompanying visual interfaces to query and explore attributes and relationships of biological relevant spatial interaction patterns could greatly benefit the digital pathology community, allowing the discovery of new patterns and easing communication of cancer research.

13 Acknowledgements

This work is supported by the Ludwig Center at Harvard Medical School, NIH/NCI grant NCI U54-CA225088 and NSF grant IIS-1901030.

REFERENCES

[1] About ArcGIS | Mapping & Analytics Software and Services, https://www.esri.com/en-us/arcgis/about-arcgis/overview, last accessed 3/30/2022. last accessed 3/30/2022.

[2] How inverse distance weighted interpolation works—ArcGIS Pro | Documentation, https://pro.arcgis.com/en/pro-app/latest/help/analysis/geostatistical-analyst/how-inverse-distance-weighted-interpolation-works.htm, last accessed: 3/30/2022.

[3] Indica Labs. HALO, https://indicalab.com/halo/, last accessed: 3/30/2022.

[4] An overview of the Spatial Statistics toolbox—Help | ArcGIS Desktop.

[5] N. Andrienko, G. Andrienko, G. Fuchs, and P. Jankowski. Scalable and privacy-respectful interactive discovery of place semantics from human mobility traces. *Information Visualization*, 15(2):117–153, Apr. 2016.

[6] M. Ankerst, S. Berchtold, and D. Keim. Similarity clustering of dimensions for an enhanced visualization of multidimensional data. In *Proceedings IEEE Symposium on Information Visualization*, pages 52–60, 1998.

[7] E. Becht, C.-A. Dutertre, I. W. H. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell. Evaluation of UMAP as an alternative to t-SNE for single-cell data. *bioRxiv*, page 298430, Jan. 2018.

[8] Y. Bengio, A. Courville, and P. Vincent. Representation learning: a review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, Aug. 2013.

[9] S. Carpendale. Evaluating Information Visualizations. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization: Human-Centered Issues and Perspectives*, pages 19–45. Springer, Berlin, Heidelberg, 2008.

[10] G. Chechik, V. Sharma, U. Shalit, and S. Bengio. Large Scale Online Learning of Image Similarity Through Ranking. *J. Mach. Learn. Res.*, 11:1109–1135, Mar. 2010.

[11] X. Fan, Y. Yuan, and J. S. Liu. The EM Algorithm and the Rise of Computational Biology. *Statistical Science*, 25(4), Nov. 2010.

[12] R. C. Geary. The Contiguity Ratio and Statistical Mapping. *The Incorporated Statistician*, 5(3):115–146, 1954.

[13] Gehlenborg Lab. Vitessce.

[14] M. J. Gerdes, C. J. Sevinsky, A. Sood, S. Adak, M. O. Bello, A. Bordwell, A. Can, A. Corwin, S. Dinn, R. J. Filkins, D. Hollman, V. Kamath, S. Kaanumalle, K. Kenny, M. Larsen, M. Lazare, Q. Li, C. Lowes, C. C. McCulloch, E. McDonough, M. C. Montalto, Z. Pang, J. Rittscher, A. Santamaria-Pang, B. D. Sarachan, M. L. Seel, A. Seppo, K. Shaikh, Y. Sui, J. Zhang, and F. Ginty. Highly multiplexed single-cell analysis of formalin-fixed, paraffin-embedded cancer tissue. *Proceedings of the National Academy of Sciences*, 110(29):11982, July 2013.

[15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Region-Based Convolutional Networks for Accurate Object Detection and Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):142–158, Jan. 2016.

[16] Y. Goltsev, N. Samusik, J. Kennedy-Darling, S. Bhate, M. Hale, G. Vazquez, S. Black, and G. P. Nolan. Deep Profiling of Mouse Splenic Architecture with CODEX Multiplexed Imaging. *Cell*, 174(4):968–981.e15, Aug. 2018.

[17] H. Gonzalez, C. Hagerling, and Z. Werb. Roles of the immune system in cancer: from tumor initiation to metastatic progression. *Genes & development*, 32(19-20):1267–1284, Oct. 2018.

[18] D. Guo. Flow Mapping and Multivariate Visualization of Large Spatial Interaction Data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1041–1048, Nov. 2009.

[19] R. Guo, P. Sun, E. Lindgren, Q. Geng, D. Simcha, F. Chern, and S. Kumar. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. In *International Conference on Machine Learning*, 2020.

[20] B. M. Gyori, J. A. Bachman, K. Subramanian, J. L. Muhlich, L. Galescu, and P. K. Sorger. From word models to executable models of signaling networks using automated assembly. *Molecular Systems Biology*, 13(11):954, Nov. 2017. Publisher: John Wiley & Sons, Ltd.

[21] J. Hoffer, R. Rashid, J. L. Muhlich, Y.-A. Chen, D. P. W. Russell, J. Ruokonen, R. Krueger, H. Pfister, S. Santagata, and P. K. Sorger. Minerva: a light-weight, narrative image browser for multiplexed tissue images. *Journal of open source software*, 5(54):2579, 2020.

[22] J. W. L. Hooton. Randomization tests: statistics for experimenters. *Computer Methods and Programs in Biomedicine*, 35(1):43–51, May 1991.

[23] L.-Y. Hu, M.-W. Huang, S.-W. Ke, and C.-F. Tsai. The distance function effect on k-nearest neighbor classification for medical datasets. *SpringerPlus*, 5(1):1304–1304, Aug. 2016.

[24] J. Jessup, R. Krueger, S. Warchol, J. Hoffer, J. Muhlich, C. C. Ritch, G. Gaglia, S. Coy, Y. -A. Chen, J. -R. Lin, S. Santagata, P. K. Sorger, and H. Pfister. Scope2Screen: Focus+Context Techniques for Pathology Tumor Assessment in Multivariate Image Data. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2021.

[25] S. Johansson and J. Johansson. Interactive Dimensionality Reduction Through User-defined Combinations of Quality Metrics. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):993–1000, 2009.

[26] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734*, 2017.

[27] Josh Moore, Melissa Linkert, Colin Blackburn, Mark Carroll, Richard K. Ferguson, Helen Flynn, Kenneth Gillen, Roger Leigh, Simon Li, Dominik Lindner, William J. Moore, Andrew J. Patterson, Blazej Pindelski, Balaji Ramalingam, Emil Rozbicki, Aleksandra Tarkowska, Petr Walczysko, Chris Allan, Jean-Marie Burel, and Jason Swedlow. OMERO and Bio-Formats 5: flexible access to large bioimaging datasets at scale. volume 9413, Mar. 2015.

[28] H. M. Kakde. Range searching using kd tree. *Florida State University*, 2005.

[29] D. P. Kingma and M. Welling. Auto-encoding variational bayes. 2014. arXiv: 1312.6114 [stat.ML].

[30] D. Kobak and P. Berens. The art of using t-SNE for single-cell transcriptomics. *Nature Communications*, 10(1):5416, Nov. 2019.

[31] R. Krueger, J. Beyer, W.-D. Jang, N. W. Kim, A. Sokolov, P. K. Sorger, and H. Pfister. Facetto: Combining Unsupervised and Supervised Learning for Hierarchical Phenotype Analysis in Multi-Channel Image Data. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):227–237, Jan. 2020.

[32] S. K. Lam, A. Pitrou, and S. Seibert. Numba: A LLVM-Based Python JIT Compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, LLVM '15, New York, NY, USA, 2015. Association for Computing Machinery.

[33] T. v. Landesberger, M. Görner, R. Rehner, and T. Schreck. A System for Interactive Visual Analysis of Large Graphs Using Motifs in Graph Editing and Aggregation, 2009.

[34] F. Lekschas. flekschas/regl-scatterplot, Mar. 2021.

[35] F. Lekschas, B. Peterson, D. Haehn, E. Ma, N. Gehlenborg, and H. Pfister. PEAX: Interactive Visual Pattern Search in Sequential Data Using Unsupervised Deep Representation Learning. *Computer Graphics Forum*, 2020.

[36] J. H. Levine, E. F. Simonds, S. C. Bendall, K. L. Davis, E.-a. D. Amir, M. D. Tadmor, O. Litvin, H. G. Fienberg, A. Jager, E. R. Zunder, R. Finck, A. L. Gedman, I. Radtke, J. R. Downing, D. Pe'er, and G. P. Nolan. Data-Driven Phenotypic Dissection of AML Reveals Progenitor-like Cells that Correlate with Prognosis. *Cell*, 162(1):184–197, July 2015.

[37] S. Li, S. Besson, C. Blackburn, M. Carroll, R. K. Ferguson, H. Flynn, K. Gillen, R. Leigh, D. Lindner, M. Linkert, W. J. Moore, B. Ramalingam, E. Rozbicki, G. Rustici, A. Tarkowska, P. Walczysko, E. Williams, C. Allan, J.-M. Burel, J. Moore, and J. R. Swedlow. Metadata management for high content screening in OMERO. *High-throughput Imaging*, 96:27–32, Mar. 2016.

[38] A. Liberzon, A. Subramanian, R. Pinchback, H. Thorvaldsdóttir, P. Tamayo, and J. P. Mesirov. Molecular signatures database (MSigDB) 3.0. *Bioinformatics (Oxford, England)*, 27(12):1739–1740, June 2011.

[39] J.-R. Lin, M. Fallahi-Sichani, J.-Y. Chen, and P. K. Sorger. Cyclic Immunofluorescence (CycIF), A Highly Multiplexed Method for Single-cell Imaging. *Current protocols in chemical biology*, 8(4):251–264, Dec. 2016.

[40] J.-R. Lin, B. Izar, S. Wang, C. Yapp, S. Mei, P. M. Shah, S. Santagata, and P. K. Sorger. Highly multiplexed immunofluorescence imaging of human tissues and tumors using t-CyCIF and conventional optical microscopes. *eLife*, 7:e31657, July 2018.

[41] J.-R. Lin, S. Wang, S. Coy, M. Tyler, C. Yapp, Y.-A. Chen, C. N. Heiser, K. S. Lau, S. Santagata, and P. K. Sorger. Multiplexed 3D atlas of state transitions and immune interactions in colorectal cancer. *bioRxiv*, 2021.

[42] A. Malik, R. Maciejewski, N. Elmqvist, Y. Jang, D. S. Ebert, and W. Huang. A correlative analysis process in a visual analytics environment. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 33–42, Oct. 2012.

[43] T. Manz, I. Gold, N. H. Patterson, C. McCallum, M. S. Keller, I. Herr, Bruce W, K. Börner, J. M. Spraggins, and N. Gehlenborg. *Viv: Multiscale Visualization of High-Resolution Multiplexed Bioimaging Data on the Web*. OSF Preprints, Aug. 2020.

[44] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. 2020. arXiv: 1802.03426 [stat.ML].

[45] A. Miles, jakirkham, M. Durant, M. Bussonnier, J. Bourbeau, T. Onalan,

J. Hamman, Z. Patel, M. Rocklin, shikharsg, R. Abernathey, J. Moore, V. Schut, r. dussin, E. S. de Andrade, C. Noyes, A. Jelenak, A. Banihirwe, C. Barnes, G. Sakkis, J. Funke, J. Kelleher, J. Jevnik, J. Swaney, P. S. Rahul, S. Saalfeld, john, T. Tran, p. bot, and sbalmer. zarr-developers/zarr-python: v2.5.0, Oct. 2020.

[46] P. A. P. Moran. Notes on Continuous Stochastic Phenomena. *Biometrika*, 37(1/2):17–23, 1950.

[47] E. Mörth, I. S. Haldorsen, S. Bruckner, and N. N. Smit. ParaGlyder: Probe-driven Interactive Visual Analysis for Multiparametric Medical Imaging Data. In N. Magnenat-Thalmann, C. Stephanidis, E. Wu, D. Thalmann, B. Sheng, J. Kim, G. Papagiannakis, and M. Gavrilova, editors, *Advances in Computer Graphics*, Lecture Notes in Computer Science, pages 351–363, Cham, 2020. Springer International Publishing.

[48] A. J. Nirmal, Z. Maliga, T. Vallius, B. Quattrochi, A. A. Chen, C. A. Jacobson, R. J. Pelletier, C. Yapp, R. Arias-Camison, Y.-A. Chen, C. G. Lian, G. F. Murphy, S. Santagata, and P. K. Sorger. The spatial landscape of progression and immunoediting in primary melanoma at single cell resolution. Technical report, May 2021.

[49] C. J. Nolet, V. Lafargue, E. Raff, T. Nanditale, T. Oates, J. Zedlewski, and J. Patterson. Bringing UMAP Closer to the Speed of Light with GPU Acceleration, 2021. _eprint: 2008.00325.

[50] J. A. Overton, R. Vita, P. Dunn, J. G. Burel, S. A. C. Bukhari, K.-H. Cheung, S. H. Kleinstein, A. D. Diehl, and B. Peters. Reporting and connecting cell type names and gating definitions through ontologies. *BMC Bioinformatics*, 20(5):182, Apr. 2019.

[51] G. Palla, D. S. Fischer, A. Regev, and F. J. Theis. Spatial components of molecular tissue biology. *Nature Biotechnology*, 40(3):308–318, Mar. 2022.

[52] 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. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[53] D. Pham, X. Tan, J. Xu, L. F. Grice, P. Y. Lam, A. Raghubar, J. Vukovic, M. J. Ruitenberg, and Q. Nguyen. stLearn: integrating spatial location, tissue morphology and gene expression to find cell types, cell-cell interactions and spatial trajectories within undissociated tissues. Technical report, May 2020.

[54] D. Phan, L. Xiao, R. Yeh, and P. Hanrahan. Flow map layout. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pages 219–224. IEEE, 2005.

[55] R. Krueger, T. Tremel, and D. Thom. VESPa 2.0: Data-Driven Behavior Models for Visual Analytics of Movement Sequences. In *2017 International Symposium on Big Data Visual Analytics (BDVA)*, pages 1–8, Nov. 2017.

[56] R. Rashid, Y.-A. Chen, J. Hoffer, J. L. Muhlich, J.-R. Lin, R. Krueger, H. Pfister, R. Mitchell, S. Santagata, and P. K. Sorger. Online narrative guides for illuminating tissue atlas data and digital pathology images. *bioRxiv*, page 2020.03.27.001834, Jan. 2020.

[57] P. Ribeiro, P. Paredes, M. E. P. Silva, D. Aparicio, and F. Silva. *A Survey on Subgraph Counting: Concepts, Algorithms and Applications to Network Motifs and Graphlets*. 2019.

[58] B. D. Ripley. The second-order analysis of stationary point processes. *Journal of Applied Probability*, 13(2):255–266, June 1976.

[59] J. C. Roberts. State of the Art: Coordinated amp; Multiple Views in Exploratory Visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)*, pages 61–71, July 2007.

[60] T. Ropinski, S. Oeltze, and B. Preim. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics*, 35(2):392–401, Apr. 2011.

[61] U. Rubens, R. Hoyoux, L. Vanosmael, M. Ouras, M. Tasset, C. Hamilton, R. Longuespée, and R. Marée. Cytomine: Toward an Open and Collaborative Software Platform for Digital Pathology Bridged to Molecular Investigations. *Proteomics. Clinical applications*, 13(1):e1800057, Jan. 2019.

[62] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *arXiv:1409.0575 [cs]*, Jan. 2015.

[63] D. Schapiro, H. W. Jackson, S. Raghuraman, J. R. Fischer, V. R. T. Zanotelli, D. Schulz, C. Giesen, R. Catena, Z. Varga, and B. Bodenmiller. histoCAT: analysis of cell phenotypes and interactions in multiplex image cytometry data. *Nature methods*, 14(9):873–876, Sept. 2017.

[64] D. Schapiro, A. Sokolov, C. Yapp, J. L. Muhlich, J. Hess, J.-R. Lin, Y.-A. Chen, M. K. Nariya, G. J. Baker, J. Ruokonen, Z. Maliga, C. A. Jacobson, S. L. Farhi, D. Abbondanza, E. T. McKinley, C. Betts, A. Regev, R. J. Coffey, L. M. Coussens, S. Santagata, and P. K. Sorger. MCMICRO: A scalable, modular image-processing pipeline for multiplexed tissue imaging. *bioRxiv*, 2021.

[65] L. M. Scott and M. V. Janikas. Spatial Statistics in ArcGIS. In M. M. Fischer and A. Getis, editors, *Handbook of Applied Spatial Analysis: Software Tools, Methods and Applications*, pages 27–41. Springer, Berlin, Heidelberg, 2010.

[66] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012.

[67] T. Segaran. *Programming collective intelligence : building smart web 2.0 applications*. O'Reilly, Sebastapol, California, 1st edition edition, 2007.

[68] N. Sofroniew, T. Lambert, K. Evans, J. Nunez-Iglesias, G. Bokota, M. Bussonnier, G. Peña-Castellanos, P. Winston, K. Yamauchi, D. D. Pop, alisterburt, Pam, Z. Liu, A. C. Solak, L. Gaifas, G. Buckley, A. Sweet, G. Lee, J. Rodríguez-Guerra, J. Bragantini, N. Clack, L. Migas, M. W. Mendonça, V. Hilsenstein, R. Haase, Hector, J. Freeman, P. Boone, A. R. Lowe, and C. Gohlke. napari/napari: 0.4.15rc1, Mar. 2022.

[69] A. Somarakis, M. E. Ijsselsteijn, S. J. Luk, B. Kenkhuis, N. F. C. C. de Miranda, B. P. F. Lelieveldt, and T. Höllt. Visual cohort comparison for spatial single-cell omics-data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):733–743, Feb. 2021.

[70] A. Somarakis, V. Van Unen, F. Koning, B. P. Lelieveldt, and T. Hollt. ImaCytE: Visual Exploration of Cellular Microenvironments for Imaging Mass Cytometry Data. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019.

[71] C. R. Stoltzfus, J. Filipek, B. H. Gern, B. E. Olin, J. M. Leal, Y. Wu, M. R. Lyons-Cohen, J. Y. Huang, C. L. Paz-Stoltzfus, C. R. Plumlee, and others. CytoMAP: a spatial analysis toolbox reveals features of myeloid cell organization in lymphoid tissues. *Cell reports*, 31(3):107523, 2020.

[72] W. R. Tobler. A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic Geography*, 46(sup1):234–240, June 1970.

[73] L. van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.

[74] T. von Landesberger, F. Brodkorb, P. Roskosch, N. Andrienko, G. Andrienko, and A. Kerren. MobilityGraphs: Visual Analysis of Mass Mobility Dynamics via Spatio-Temporal Graphs and Clustering. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):11–20, Jan. 2016.

[75] X. Wang and E. Grimson. Spatial Latent Dirichlet Allocation. page 8.

[76] M. Wattenberg, F. Viégas, and I. Johnson. How to Use t-SNE Effectively. *Distill*, 2016.

[77] E. Wong, B. Baur, S. Quader, and C.-H. Huang. Biological network motif detection: principles and practice. *Briefings in Bioinformatics*, 13(2):202–215, June 2011.

[78] J. Wood, A. Slingsby, and J. Dykes. Visualizing the Dynamics of London's Bicycle-Hire Scheme. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 46(4):239–251, Dec. 2011.

[79] J. Xia, Y. Zhang, J. Song, Y. Chen, Y. Wang, and S. Liu. Revisiting Dimensionality Reduction Techniques for Visual Cluster Analysis: An Empirical Study. *IEEE Transactions on Visualization and Computer Graphics*, PP:1–1, 2021.

[80] T. Xiao, C. J. Reed, X. Wang, K. Keutzer, and T. Darrell. Region Similarity Representation Learning. *arXiv preprint arXiv:2103.12902*, 2021.

[81] M.-S. Yang, C.-Y. Lai, and C.-Y. Lin. A robust EM clustering algorithm for Gaussian mixture models. *Pattern Recognition*, 45(11):3950–3961, Nov. 2012.

[82] Y. Yang, T. Dwyer, S. Goodwin, and K. Marriott. Many-to-Many Geographically-Embedded Flow Visualisation: An Evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):411–420, Jan. 2017.

[83] Q. Zhu, S. Shah, R. Dries, L. Cai, and G.-C. Yuan. Identification of spatially associated subpopulations by combining scRNAseq and sequential fluorescence in situ hybridization data. *Nature Biotechnology*, Oct. 2018.