

Barrio: Customizable Spatial Neighborhood Analysis and Comparison for Nanoscale Brain Structures

Jakob Troidl^{1,3}, Corrado Cali², Eduard Gröller³, Hanspeter Pfister¹, Markus Hadwiger⁴ and Johanna Beyer¹

¹Harvard University, ²University of Turin, ³TU Wien, ⁴KAUST

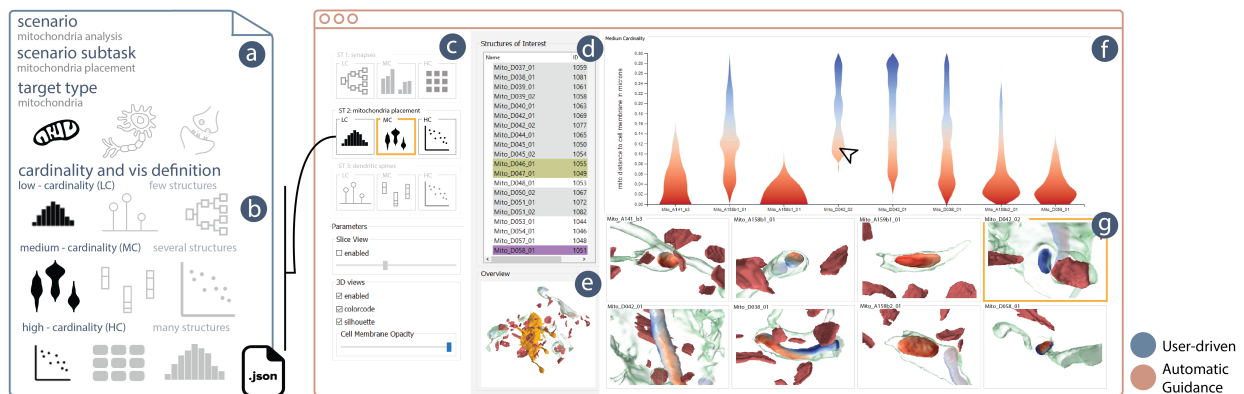


Figure 1: Customized scalable comparison of spatial neighborhoods. Barrio supports user-defined analysis scenarios (a,b) to automatically adjust visualizations, view parameters, and the user interface for the comparative analysis of nanoscale brain structures. Barrio suggests the comparison method that is best suited for the number of user-selected neighborhoods that are being compared in the analysis view (f). Users first specify the biological target of their analysis (a), and define preferred visualization methods for comparisons of different cardinalities (i.e., for two to n structures) (b). Barrio automatically adjusts the visualizations and UI, including the settings panel (c), a list view of instances of a target type (d), the 3D overview (e), the quantitative analysis views (f), and the detailed 3D neighborhood views (g).

Abstract

High-resolution electron microscopy imaging allows neuroscientists to reconstruct not just entire cells but individual cell substructures (i.e., cell organelles) as well. Based on these data, scientists hope to get a better understanding of brain function and development through detailed analysis of local organelle neighborhoods. In-depth analyses require efficient and scalable comparison of a varying number of cell organelles, ranging from two to hundreds of local spatial neighborhoods. Scientists need to be able to analyze the 3D morphologies of organelles, their spatial distributions and distances, and their spatial correlations. We have designed Barrio as a configurable framework that scientists can adjust to their preferred workflow, visualizations, and supported user interactions for their specific tasks and domain questions. Furthermore, Barrio provides a scalable comparative visualization approach for spatial neighborhoods that automatically adjusts visualizations based on the number of structures to be compared. Barrio supports small multiples of spatial 3D views as well as abstract quantitative views, and arranges them in linked and juxtaposed views. To adapt to new domain-specific analysis scenarios, we allow the definition of individualized visualizations and their parameters for each analysis session. We present an in-depth case study for mitochondria analysis in neuronal tissue and demonstrate the usefulness of Barrio in a qualitative user study with neuroscientists.

CCS Concepts

• **Human-Centered Computing** → Spatial neighborhood analysis, Visual comparisons, Neuroscience, Scientific visualization;

1. Introduction

Neurobiologists are acquiring ever-increasing amounts of high-resolution image data of the mammalian brain, aiming to tackle questions regarding brain physiology, the onset of diseases, and the

emergence of consciousness. Connectomics researchers are interested in the detailed connectivity between neurons at nanoscale resolution [LD11], as well as in the analysis and comparison of local spatial neighborhoods around neuronal structures and cell organelles. However, mammalian brains are staggeringly complex,

with tens of millions of interconnected neurons, dozens of cell substructures (i.e., cell organelles) per neuron, and billions of synapses. Thus, the number of interesting neighborhoods and comparisons of neighborhoods in such a dataset is even higher, ranging from detailed pair-wise comparisons to the comparison and clustering of dozens, hundreds, or thousands of structures. For example, to evaluate how spatial variations in neuronal tissue might influence brain development and neuronal connectivity, scientists need to compare different biological structures (such as synapses), analyze distances and spatial correlations (e.g., between synapses and cell organelles), and look at how certain biological features are spatially distributed. The specific details of this analysis (e.g., the biologist's goals, comparison target types, measurement types) vary for each scientist. While one expert might be interested in the distribution of mitochondria within nerve cells, others might only focus on the spatial neighborhoods of synapses. Therefore, in this work, we have developed a framework that supports a highly customizable spatial neighborhood analysis workflow of scientists and a flexible comparative analysis, which scales to their data sizes.

Most visualization tools for connectomics provide either 3D views for exploring large, segmented electron microscopy (EM) volumes [BAAK*13, BHAA*13, HBJP12], or focus on the analysis of synaptic connectivity [AABS*14, SBS*13]. Little work has focused on the scalable visual comparative analysis of well-defined cell neighborhoods in EM data. Furthermore, while all of these tools target neuroscientists, none of them can adapt to the highly-individualized and specific domain goals that neuroscientists often have. Most tools designed to be general enough to support several types of analyses result in complex user interfaces and unintuitive workflows. On the other hand, tools designed for a specific analysis task typically do not generalize well to other tasks, often summarized as the expressivity vs. usability trade-off [FCOO12].

In this design study, we present Barrio, a novel tool for the interactive analysis of spatial neighborhoods in high-resolution 3D electron microscopy (3DEM) data. In particular, we focus on neighborhoods around cell organelles. The main contribution of this work is a *customizable framework for spatial neighborhood analysis* of 3DEM data, facilitating the analysis of spatial distributions, distances, and 3D morphology. Scientists can specify new, individualized analysis scenarios and appropriate visualizations that support their analysis tasks. Second, we contribute a *scalable visual comparison* method for spatial structures that allows the comparison of a wide range of element cardinalities, e.g., ranging from two to hundreds of structures within a single comparison. Third, we report on a *detailed case study of using Barrio to analyze the spatial neighborhood of mitochondria*, the novel visual encodings we designed for it, and the insights gained by the scientist. Finally, we also report on a qualitative user study with eight domain experts.

2. Related Work

Visualization for Connectomics. Visualization approaches have been applied to the entire neuroscience and connectomics data acquisition and analysis pipeline [HHM*17], ranging from segmentation [BSL18, GG90] and visual proofreading of segmentations [AABH*16, HKT*18, GWB*21] to data exploration and analysis [BAAK*13, AABS*14, MAAB*18, CBB*16]. Data ex-

ploration typically focuses either on the original large-scale microscopy data and its segmentation, or on higher-level connectivity information [PKB*12] and uses tools that are geared towards the neuroscience workflow and data rather than general visualization systems. Structural connectivity analysis of neuronal cells often relies on abstract visual metaphors such as subway maps [AABS*14] or circuit diagrams [SBS*13] that highlight synaptic connections. On the other hand, exploring the general morphology of the data typically relies on scalable volume visualization of the original or segmented electron microscopy data [HBJP12, BHAA*13]. Few works have focused on the visual analysis of local energy consumption [ACA*19] and interactions between cell organelles [MAAB*18]. However, none of these approaches analyze neighborhoods of selected neuronal structures and also do not support comparative visualizations of them.

Visualization Specification and Generation Frameworks.

Frameworks for the specification and generation of visualizations can be classified into low level [BOH11, BH09] and high level [SWH14, SH14, SRHH16, SMWH17] approaches, based on the level of abstraction they provide in their interface. Many techniques use declarative methods (i.e., defining *what* not *how*) to specify data visualizations. Low-level visualization specifications like D3 [BOH11] require users to compose a visualization programmatically by specifying its components, while high-level approaches may support interactive data binding and visualization specification using drag and drop [SH14]. Barrio is inspired by these approaches, but offers an interface for *domain-specific visualization specification*. Users can define domain-specific analysis scenarios and map different visualizations to them. Our system parses this input and automatically creates customized views.

Spatial Neighborhood Analysis. Identifying and searching for patterns in cell neighborhoods is an active area of research. For highly multiplex imaging data, tools such as CytoMap [SFG*20] and Facetto [KBJ*20] have been proposed, which support the search for visual patterns in cell neighborhoods and hierarchical clustering of cell types for phenotype analysis, respectively. Other approaches focus on cell clustering based on Hidden Markov random fields concerning gene expression and localization data [Pet15], phenotype analysis for cellular screens [DSG*17], analysis of cell interactions [NDBB17], or the evaluation of patterns in cell populations [FGBHOdS05]. All of these approaches use multiplex imaging or microscopy data at a much lower resolution than EM data and, therefore, focus on larger cell clusters and areas. Jorstad et al. [JNC*15] allow users to manually extract quantitative measures from reconstructed meshes of cell organelles in high-resolution EM data. However, they do not focus on neighborhood analysis and comparisons between organelles.

Comparative Visualization. Comparative visualization allows users to assess similarities and differences between multiple data points or data sets. Visual comparisons can be classified into juxtaposition, superposition, and explicit encoding of differences [GAW*11], or into data-level and image-level approaches [PP95]. Comparing many structures simultaneously requires scalable visualization methods and encodings, such as for large genome matrices [LBK*18] or trees [MGT*03], or the combination of spatial and non-spatial features [MMH*13]. We use




Figure 2: Structure of neurons and cell organelles. Nerve signals are transmitted from axons to dendrites via synapses. Synapses often form between a presynaptic bouton and a dendritic spine. Neurons contain cell organelles, such as mitochondria or ribosomes.

juxtaposition and superposition to compare neuronal structures and spatial neighborhoods in a scalable way, both in spatial and non-spatial views. For comparing scalar and flow fields, contour trees and graph matching have been used [HSKK01, SWC*08], as well as graph comparison for analyzing brain connectivity [ABHR*13]. For 3D structures, Schmidt et al. [SPA*14] compare surface meshes against a reference mesh. For our data, reference meshes do not exist due to the variability in neurite shapes and branching patterns. Thus, we compare structures in both, 3D and data analysis views.

3. Background in Neuroscience

Neuroscience Fundamentals. The human brain consists of billions of interconnected nerve cells, or *neurons*. Neurons receive input on tree-like structures called *dendrites* and transmit signals via long tubular structures called *axons* (see Fig. 2). Axons and dendrites (collectively called *neurites*) connect via *synapses* to neighboring cells [Stu08]. The presynaptic areas of axons are called *boutons*, while the postsynaptic protrusions of dendrites are called *spines*. In addition to synaptic connectivity, subcellular structures of neurons, so-called *cell organelles*, are also of high interest to neuroscientists. *Mitochondria*, for example, are cell organelles responsible for providing energy to the cell and its synapses [SW06]. They vary significantly in their shape, size, and location within a cell. Our approach studies the spatial neighborhoods of these structures, to gain insights into how they mutually influence each other.

Connectomics Workflow and Data. Our collaborators start data acquisition by dissecting solid blocks of brain tissue stained for electron microscopy. Next, they image individual slices with an electron microscope (EM) at a pixel resolution of 5×5 nanometers [KMWL08], with a slice thickness of 15 nanometers, and register them into a 3D volume. Scientists segment neurites, synapses, mitochondria, spines, and boutons using semi-automatic tools like Ilastik [SSKH11], TrackEM [CSS*12], or VAST [BSL18]. Additionally, they measure distances between those structures (details in Sec. 7). As with any imaging method, EM provides a snapshot of a current physiological state that might change over time. However, it still provides a valuable glimpse into the neuroanatomy of a specimen. Our collaborators have collected and reconstructed six data sets of layer one of the somatosensory cortex with a focused ion beam scanning electron microscope (FIB-SEM), which are all five cubic microns in size. Three data sets show brain tissue of four-month-old mice, while the other three show tissue of 24-month-old mice [CWB*18]. Each dataset contains roughly 900 segmented structures (400 neurites, 300 cell organelles, and 200 synapses).

4. Goal & Task Analysis

Following a problem-driven design study approach [SMM12], we have identified domain goals, associated analysis tasks, and general system requirements in semi-structured interviews with four experienced neuroscientists at the University of Turin and the Harvard Center of Brain Science. All scientists work with 3DEM data and have more than ten years of experience in the field. Over the course of several months, we regularly met with our main collaborator (who is also a co-author) to iterate on design prototypes.

4.1. Domain Goals

Our collaborators' main objective is the detailed analysis of local neighborhoods in their neuronal 3DEM data sets. In particular, they want to analyze, compare, and correlate morphological features of interest in selected cells, such as spines, boutons, or mitochondria, to discover previously unknown biological patterns. Interactive exploration is crucial, as our collaborators need to investigate the detailed 3D morphology to recognize subtle patterns or areas of interest that cannot be easily detected with automatic methods. The particular domain-specific goals of the neuroscientists are:

G1 - Proximity Analysis for Cell Organelles. The collaborators are concerned with the spatial aspects of how cell organelles relate to each other. On the one hand, they are interested in analyzing **distances** between a varying number of structures (e.g., *What are the distances of the closest three synapses to a mitochondrion?*). On the other hand, they are interested in **distributions** (e.g., *What percentage of a mitochondrion is closer than 30 nm to the cell boundary?*). Combined with 3D spatial exploration of the dataset, our collaborators want to investigate whether there are any patterns that indicate connectivity strength based on the underlying spatial arrangement of cell organelles. Our initial interviews showed that distances and distance distributions are important as a metric for neuroscientists, since the influence of cell organelles and synapses is spatially restricted to their immediate neighborhood.

G2 - Comparative Analysis of Neuronal Structures. In addition to looking at a single spatial neighborhood, our collaborators need to compare neighborhoods and structures of interest. Their goal is to perform comparisons with **different biological targets**, to compare between: a) different data sets, b) different instances of a structure of interest (e.g., dendrites), or c) different attributes of a single structure of interest (e.g., synapse distances to a single mitochondrion). Additionally, they need to be able to perform comparisons of **different cardinalities**, ranging from a detailed comparison of two structures to comparing hundreds of structures at the same time. This feature allows them to test whether local findings also apply to a larger data set and, conversely, allows them to drill down from a larger perspective to highly-detailed local instances.

G3 - Customizable Analysis of User-Specific Hypotheses. Our scientists want to fine-tune their visual analysis to their current scientific objective. In addition to general comparisons and proximity analysis in their data, they want to customize analyses steps and define which visualizations to use for the domain-specific tasks. Therefore, their system should be easily customizable to allow for different views, UI elements, and interactions, and guide the users through their analysis based on their scientific objectives.

4.2. Analysis Tasks

Based on the above goals, we have derived a set of analysis tasks scientists need to perform:

T1 - Explore original EM image data and segmented structures. Before diving into a detailed analysis, scientists first need to get an overview and verify the correctness of the segmentation (G1).

T2 - Explore the spatial neighborhood of a single structure. Once a structure of interest has been identified, domain scientists explore its spatial features and 3D morphology (G1, G3).

T3 - Quantify single distances and distance distributions in a spatial neighborhood. After their initial exploration, scientists need to analyze proximity and distances more accurately (G1, G3).

T4 - Compare the spatial neighborhoods of several structures. Analyzing the variability of certain features in different local spatial neighborhoods is essential for discovering trends and patterns when studying biological features (G2, G3).

T5 - Compare neighborhoods of different data sets. Scientists want to compare different specimens to analyze the stability of certain neighborhood features and differences between them (G2, G3).

4.3. Higher-Level Requirements

To support the domain goals and analysis tasks for many different users and their individual domain questions, we have extracted some higher-level design requirements that Barrio must fulfill:

R1 - Flexibility and Expandability. While all of our collaborators want to analyze and compare spatial neighborhoods, the details vary for each scientist. For example, one scientist is interested in mitochondria, another one in endoplasmic reticula. To support both (and other) usage scenarios, Barrio has to be customizable and expandable to allow scientists or technicians to adjust the provided visualization capabilities for specific domain questions (G3).

R2 - Scalable Comparisons. To support the visual comparison of different biological targets and at different cardinalities, our system needs to provide comparisons at different scales and allow users to switch between them easily (G2).

R3 - Ease of Use and Automatic Guidance. In addition to being flexible and expandable, the system must also support neuroscientists with little or no programming experience. Biological data is inherently complex. Therefore, we need to limit the cognitive load for users by employing methods that automatically guide and support them in their data exploration. Hence, Barrio should provide a human-readable interface and file format for customization and automatically guide users in their specific analysis tasks (G1-3).

R4 - Shareable Analysis Configuration. After discovering an interesting observation in a data set, experts often want to recreate their results later or share them with colleagues. Results often rely on complex user settings that are hard to reproduce in other environments. Hence, Barrio should support saving and sharing of user-specific analysis configurations.

5. Barrio Toolkit Design

The design of Barrio tackles the requirements R1 and R3 by supporting user-defined analysis scenarios and subtle automatic guid-

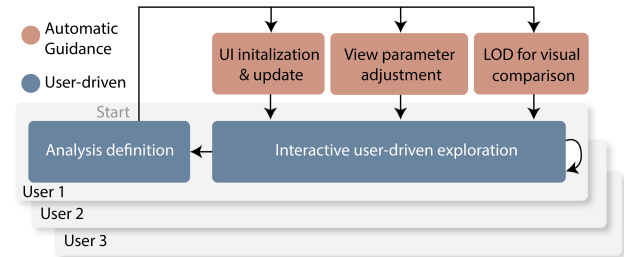


Figure 3: Barrio workflow and automatic user guidance. Barrio automatically initializes and updates the user interface, adjusts views and view parameters, and adapts the level of detail (LOD) for visual neighborhood comparison based on the user-specified analysis definition, while the user interactively explores the data.

ance throughout the analysis. Based on the custom user-defined analysis scenario, Barrio suggests fitting visualizations, adjusts view parameters, and offers different comparison modes. Sec. 6 describes our visual neighborhood comparison approach (R2).

5.1. User Workflow

The user starts by loading a custom analysis scenario and its data. The specification of the scenario is done prior to the interactive analysis (see Sec. 5.2). Users perform an initial exploration in 3D and select structures of interest for further investigation. Next, users can perform a detailed analysis, such as comparing user-selected structures in detail to extract distances, detect distribution patterns, or test hypotheses. In all these steps, users are supported and guided by Barrio to streamline the analysis process and minimize the mental load. For instance, Barrio selects appropriate visualizations based on the current data cardinality or displays 3D views of the specified neighborhood automatically. Fig. 3 shows the user workflow and which parts of the framework are adjusted automatically.

5.2. Custom Analysis Scenarios

Barrio allows users to define custom analysis scenarios that fine-tune the general neighborhood analysis (R1).

Scenario Specification. Conceptually, a *scenario* defines the analysis a neuroscientist wants to perform and breaks it down into individual and concrete scenario subtasks (see Fig. 1a). Each scenario subtask is defined in terms of data mappings, default visualizations, comparison modes, and supported user interactions. We will refer to the tasks in a user-defined analysis scenario as *scenario subtasks* to distinguish them from the general analysis tasks identified in our goal and task analysis in Sec. 4. Below, we describe the concrete elements comprising a scenario.

Scenario: An analysis scenario represents a higher-level question a specific domain scientist might have, such as “How do mitochondria influence their spatial neighborhood?”.

Scenario Subtasks: Each analysis scenario comprises one or several scenario subtasks, which are specific analysis steps a user needs to perform to answer the higher-level scenario question. For example, a scenario subtask in a mitochondria analysis scenario would be to *analyze the synaptic neighborhood of a single mitochondrion*, i.e., synapses in close proximity to the mitochondrion.

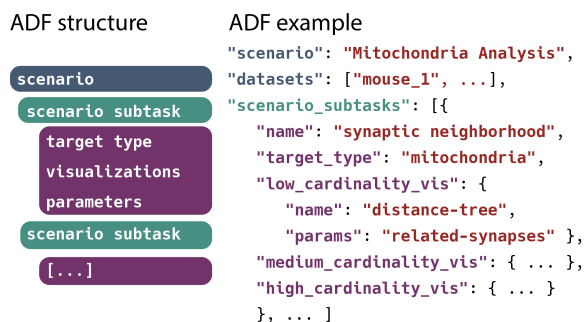


Figure 4: Structure and example of an analysis definition file (ADF). The ADF defines analysis scenarios, scenario subtasks, target types, visualizations and UI parameters.

To support these steps, each scenario subtask can define visualizations, comparison modes, and interaction parameters that help the user in the analysis. A high-level view of a scenario, its subtasks, and visualization specifications is shown in Fig. 4.

Target type: Each scenario subtask can define a target type, which refers to the type of biological structure that is under investigation for the task. Target types can be any biological structures that have been segmented or labeled in the original EM dataset. Note that the target type refers to a biological class (e.g., all synapses) and not to a specific instance of that class (e.g., synapse x). During the interactive analysis, users can then select specific instances of the target type for further investigation. Defining a target type ensures that visualizations and user interface elements automatically focus on that type (e.g., show a list of all instances of that type, or automatically center 3D views on those instances).

Visualizations and data mappings: Each scenario subtask can specify customized visualizations that are best suited to complete the task. Barrio supports a set of standard information visualization views, such as bar charts, scatterplots, or violin plots, as well as several custom information visualization views for neuroscience tasks. Each view can be customized by changing the specific data mapping, such as which attributes to show on the axes of a scatterplot. For example, to see information about the proximity of a structure to the nearest cell membrane, we can color code the distance onto the surface of the structure, or show a violin plot of the same data. Furthermore, Barrio is extensible and allows to implement and plug in additional views that are not part of the original framework (see Sec. 7). Finally, we can define comparative views for different data cardinalities (i.e., different numbers of objects in a comparison). This means that we can specify a high-detailed view for low-cardinality comparisons, and overview visualizations for high-cardinality comparisons (R2). We describe our scalable comparison approach in Sec. 6.

Interaction parameters: In addition to specifying visualizations and tasks, a scenario can also define interaction parameters, such as user-adjustable sliders for filtering data prior to visualization.

Analysis Definition File. Custom analysis scenarios are defined by users in the form of *analysis definition files* (ADFs), as shown in Fig. 4. ADFs are written in a human-readable JSON format and do not require programming skills (R3). Barrio reads the ADF upon startup and configures the user-interface accordingly. The

JSON based format allows domain experts to easily store, edit, and share analysis configurations, which supports the reproducibility of their analysis results (R4). We provide sample ADFs that users can further fine-tune. Barrio's documentation includes detailed information about supported visualizations and their parameters [VCG22b].

5.3. Automatic User Guidance

Based on the analysis scenario defined in the ADF, Barrio automatically adjusts the user interface, views, and interactions to help experts in their analysis (R3). Our goal is to minimize the users' mental load during the interactive analysis, so that they can focus on their domain science rather than on operating the analysis interface. Barrio includes the following guidance features (see Fig. 3):

User Interface Initialization: Barrio automatically adjusts the user interface based on the user-specified analysis scenario in the ADF. This includes updating the view arrangement and adding necessary user input elements such as sliders or checkboxes. For example, for analyzing mitochondria, scientists might want a slider to interactively change the distance threshold of what they consider a "close" synapse. Furthermore, we automatically show preview icons of the three different comparison levels for each scenario subtask (see Fig. 1c). The icons allow users to choose their current scenario subtask quickly and to switch comparison mode manually, if desired.

View and View Parameter Adjustment: Barrio automatically updates the list of available visualizations and adjusts views and view parameters whenever a new scenario or scenario subtask is selected by the user. For example, whenever a scenario defines biological target types (e.g., mitochondria), we automatically highlight all instances of that type throughout the system. We show all instances of that type in the list view and blend out other structures (see Fig. 1d). Additionally, spatial neighborhood views and every spatial neighborhood analysis will automatically be centered around the specified target type. When users switch between scenario subtasks with different target types, the views and UI adjust accordingly.

Level of Detail for Visual Comparisons: Barrio can automatically adjust the level of detail in our scalable comparison approach, depending on the number of structures selected for comparison (see Sec. 6). This allows us to adjust the level of detail and abstraction level in a comparative visualization based on data cardinality.

6. Visual Neighborhood Analysis

Analyzing and comparing spatial neighborhoods is challenging, especially in biological tissue that is highly variable and exhibits a complex three-dimensional structure. Looking at not just one but multiple neighborhoods solely in 3D, easily leads to cognitive overload and does not provide quantitative comparisons. Therefore, Barrio offers two complementary visualization methods. First, we can show juxtaposed high-detail 3D renderings of the selected local neighborhoods. This allows neuroscientists to explore the original 3D data and get an intuition about the morphology and spatial arrangement of interesting structures. In addition, we have developed a comparison approach for the linked *analysis views* that scales from small to large comparisons and that supports a more quantitative analysis of neighborhoods, described below.

We define a spatial neighborhood as a sub-volume centered around a chosen structure of interest (e.g., a mitochondrion). By default, the neighborhood contains the structure of interest and its surrounding cell (see Fig. 6). Alternatively, the user can determine the size of a neighborhood by interactively adjusting a distance threshold with a maximum of 6 microns around a structure of interest.

6.1. Scalable Comparison Approach

Our comparison method supports two levels of comparison, orthogonal to each other. On the one hand, we support different biological target types. On the other hand, we support comparisons of different cardinalities (i.e., a varying number of objects). Both approaches together allow comparisons to range from overview to detail in a variety of different usage scenarios (R2). The scalable comparisons are integrated into the custom analysis scenarios, where Barrio suggests and adjusts comparative views based on the current scenario subtask and the number of selected structures.

Comparison Targets. Barrio supports comparing different instances of the same biological target type (e.g., comparing two dendrites). It also supports the comparison of cellular substructures of the same structure of interest (e.g., synapses of the same dendrite). In addition, Barrio realizes the comparison of different data sets (e.g., data sets of neuronal tissue in different developmental stages) (see Fig. 9). Users can define comparison targets in the ADF by setting the *target type* and specifying the data mapping for comparisons. Barrio then automatically adjusts the UI and views based on the specified comparison target.

Cardinality of Comparisons. In their analysis, scientists can move from a single interesting structure to many, to verify whether identified features correspond to trends or patterns. Alternatively, they can start with an overview of many structures and then drill down to see more details. A highly detailed comparison takes up more screen space and only supports a low number of objects. Comparing a high number of objects would not be ideal with a high-detail comparison view. We have designed a scalable comparison approach to support comparisons of sets with different cardinalities, and have integrated our method into Barrio. In particular, Barrio provides comparisons at three different scales, as shown in Fig. 5. In our visual design, we use either a single visualization, a single visual element, or a single visual mark per neighborhood for the low-, medium-, and high-cardinality views, respectively. This follows the intuitive notion of adjusting the amount of data we show per neighborhood, depending on the number of neighborhoods we compare. The *low-cardinality* comparison mode is designed to compare a handful of objects (two to five) and offers detailed juxtaposed views on the objects being compared, arranged as small multiples. The *medium-cardinality* comparison mode supports up to a dozen structures, while the *high-cardinality* comparison view scales up to hundreds of objects. The latter view is an overview visualization and makes trends, patterns, and outliers visible.

Barrio can either automatically switch between comparison modes, depending on the number of selected objects, or let users choose the desired comparison mode. The actual visualization for each comparison mode is based on the scenario subtasks specified in the ADF. Fig. 7 and Fig. 8 show the different cardinality views for the analysis scenario described in Sec. 8.

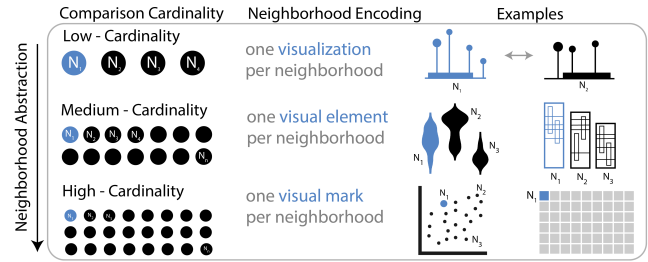


Figure 5: Scalable comparisons at three different cardinalities. Low-cardinality comparisons (top row) analyze a small number of spatial neighborhoods and use juxtaposed small multiples. Medium- and high-cardinality comparisons support increasingly larger numbers of neighborhoods and encode a single neighborhood as a single visual element (i.e., a violin in a violin plot) or a single visual mark (i.e., a dot in a scatterplot), respectively.

6.2. Visual Elements

Next, we describe the visual elements of our research prototype (see Fig. 1) and how they support visual neighborhood analysis.

3D Neighborhood View. The neighborhood view (see Fig. 1g) is the starting point of many analyses and examines the local spatial neighborhood of a user-selected structure of interest in detail (T2). Neuroscientists can analyze the morphology and three-dimensional anatomy of the selected structure while at the same time seeing nearby cellular substructures. Detailed three-dimensional visualizations of neighborhoods are essential for understanding the spatial arrangements and also provide context for the abstract analysis views, which quantify neighborhood features. For example, in the neighborhood view, we can show a selected mitochondrion, its surrounding dendrite, as well as nearby synapses. We display the selected structure of interest and the surrounding region in a high-resolution 3D surface rendering and support different rendering styles (see Fig. 6). We support focus and context rendering by displaying surrounding neuronal structures as silhouettes to highlight cell organelles located inside a cell. In addition, users can enable a slice view of the original EM data and move it through the neighborhood view (T1). With the slice view users proofread segmentation accuracy on-the-fly. They can analyze features in the EM data that have not been segmented and reconstructed yet, such as the distribution of individual glycogen granules.

We have implemented the neighborhood view as small multiples (see Fig. 1g). We juxtapose views of multiple local neighborhoods next to each other (T4, T5). Small multiples can be linked to each other for a synchronized navigation between all views.

Spatial Context View. This view provides an initial overview of currently selected neighborhoods and all its segmented structures (see Fig. 1e). Scientists can see the spatial location of selected structures in the context of the entire volume and they can examine how different local neighborhoods are related to each other, and whether they share synapses or exhibit clusters of certain structures (T1). The spatial overview is linked to all other views so that hovering over a structure in the neighborhood view highlights the respective structure in the spatial context view and vice versa.

Analysis View. While the 3D neighborhood views are great for




Figure 6: Detailed 3D neighborhood view. The dendrite (green) and close synapses (red). (a) 2D slice overlay of the original EM data to show biological context. (b) Silhouette enhancing, and semi-transparent surface rendering to highlight how enclosed mitochondria (blue) are positioned within the dendrite. (c) Color-coding of a mitochondrion showing the distance between the mitochondrion and its surrounding cell membrane.

getting a spatial understanding of a neighborhood, it is still hard to compare multiple neighborhoods in a more quantitative way. Therefore, Barrio offers *analysis views* (see Fig. 1f) that abstract certain neighborhood characteristics and support direct comparisons between neighborhoods (T3, T4, T5). Fig. 5 shows examples for the different chart types in Barrio, such as violin plots, or scatter plots. A complete list is given in the documentation [VCG22b]. The actual visualization in the analysis view is based on the number of selected neighborhoods (i.e., the cardinality of the current comparison) and on the specific analysis scenario. Analysis views and 3D neighborhood views are connected by brushing and linking.

Segmentation Tree List View. To easily select structures of interest, we provide a hierarchical tree list view of all segmented objects in the data set. Scientists can quickly select structures by navigating the data hierarchically (e.g., from a neuron to its cell organelles). Alternatively, users can choose to only see the list of objects of the current biological target type (see Fig. 1d), to quickly select structures of interest for their comparative analysis.

6.3. Interaction

Brushing & Linking: Barrio supports brushing and linking between analysis views, neighborhood views, and spatial context views. Hovering over a region in the analysis view highlights the corresponding neighborhood view, which is helpful when comparing many structures. We also highlight structures like synapses or dendritic spines in the 3D views when selected in the analysis view so that users can connect the insight they gained from the analysis view to the detailed 3D representation.

Linked 3D Navigation: To facilitate fast and intuitive visual analysis of multiple neighborhoods, users can link zooming and navigation for all 3D neighborhood views.

Neighborhood Selections: Based on initial user feedback, Barrio supports two main approaches to select neighborhoods for further investigation. Users can either start with a high-level comparison of all neighborhoods (using a high-cardinality comparison view). Subsequently they can select interesting neighborhoods based on outliers or clusters in the analysis view. This interaction mode is pri-

marily used by scientists to look at new datasets that they have not analyzed before. Alternatively, users can select individual neighborhoods in the tree list view and look at them in small multiple 3D neighborhood views before starting a more quantitative analysis. This mode is particularly useful if scientists already have a neighborhood or region of interest in mind that they want to investigate. To further aid the user in finding interesting neighborhoods, when selecting one structure in the tree list view, we automatically highlight spatially close structures in the list view as well.

7. Data Processing and Implementation

Barrio Architecture. Barrio is a modular application that separates the *Barrio-core* from the *user interface* and the *visualization stack*. The *visualization stack* provides a general API for all visualizations in our system. Whenever a new analysis scenario is loaded, the core module re-initializes and updates the user interface and different views of the visualization stack, as specified in the ADF. For flexibility, Barrio's visualization stack is extensible. Additional D3-based data visualization modules have to implement our visualization interface and can then be further customized via the ADF. Since Barrio already supports standard visualizations (e.g., scatterplots, bar charts, etc.), this feature is primarily meant to support novel custom visualizations specific to a particular usage scenario.

Data Preprocessing. We display surface meshes of the reconstructed structures as well as the original EM image volume. Therefore, we use a standardized file format [VCG22b] to store surface meshes, imaging volumes, any manually labeled meta data, and automatically computed measures such as distances. We extract meshes from human proofread segmentation volume with Neuronorm [JNC*15] and store them in wavefront format [Wav21]. Synapses are stored as meshes of the segmented postsynaptic densities. Additionally, labeled mesh attributes provided by domain scientists, such as spine lengths, spine volumes, or surface areas are stored in JSON format and can be used as visualization parameters in the ADF. Finally, we pre-compute and store the closest distances between neuronal structures as well as distance distributions to neighboring structures based on evenly sampled reconstructed surface meshes using the CGAL library [FGK*00]. To compute the distance between two structures, we use the distance between their two closest surface points (i.e., minimum distance), as this is of biological significance to our collaborators. To compute a distance distribution for structure *A* in relation to structure *B*, we compute the minimum distance of each surface point of structure *A* to *B*. CGAL uses an axis-aligned bounding box (AABB) tree to efficiently compute distances. Since we use meshes generated from human proofread segmentation masks for the distance computations, the distances can be considered accurate. We used a Windows 10 machine with a AMD Ryzen 9 3900X CPU, a GeForce RTX 2080 Super GPU, and 128 GB RAM for preprocessing and for the user study. Preprocessing 1GB of data takes about one hour.

Implementation. Barrio is implemented in C++ 17, OpenGL 4.3, and Qt5. The information visualizations are using D3.js [BOH11] and Phylotree.js [SWKP18]. We have implemented a state-of-the-art real-time transparency rendering algorithm using per-pixel linked lists [Wol11] for the rendering of semi-transparent surfaces, which requires a current GPU. Barrio is open-source [VCG22a].




Figure 7: Analysis views generated for ST1: Synaptic neighborhood of mitochondria. (a) to (c): low-, medium-, and high-cardinality comparison views. (a) distance tree for mitochondrion 'Mito 1' showing its surrounding synapses (red) and their closest mitochondria (blue), (b) bar chart showing three mitochondria and the spatial distance of their surrounding synapses, (c) distance matrix of all mitochondria and their neighboring synapses.

Scalability. We precompute all distances and store them in a hash table. Our current datasets contain roughly 1,100 segmented structures, where each structure is stored as a mesh. For rendering local neighborhoods, we only need to load the meshes of the structures inside that neighborhood onto the GPU. Our local neighborhoods, in practice, contain between 1-30 structures whose meshes need to be rendered. We currently load the entire hash table into memory and perform rendering in-core. When comparing multiple datasets, we load one hash table per dataset. For significantly larger datasets, we envision an initial filtering step prior to Barrio, to limit the number of structures being compared. Filtering could be done in our high-cardinality views (e.g., selecting elements in a scatterplot) or by querying systems such as ConnectomeExplorer [BAAK*13]. Finally, the number of side-by-side comparison views is limited by the available screen space. When comparing many neighborhoods, users should either switch to a higher cardinality comparison mode, or perform filtering prior to looking at juxtaposed views.

8. Case Study: Mitochondria Analysis

We demonstrate the utility of Barrio based on a case study with one of our collaborating domain experts (also a co-author) who performed a detailed spatial neighborhood analysis of mitochondria. The visual encodings used in this scenario were developed over the course of several months in collaboration with our collaborators, to create custom views for the comparison of neuronal structures. After a one-hour training session, the expert operated Barrio independently and also updated the default ADF based on his analysis interests. The corresponding ADF is included in the supplemental material. Due to the COVID-19 pandemic, the case study was performed remotely, with our collaborator steering the user interface via the TeamViewer software. The case study consists of the user-defined analysis scenario 'Mitochondria Analysis' and examines different aspects of mitochondria neighborhoods. The first two subtasks ST1 and ST2 are described in Sec. 8.1 and Sec. 8.2. Details on the third subtask are given in the supplemental material.

8.1. Synaptic Neighborhood of Mitochondria (ST1)

Domain Goal: Here, the scientist wanted to analyze the spatial relationships of mitochondria and their surrounding synapses (i.e., synapses connecting to the mitochondrion's cell) to understand the role of mitochondria in synapse formation and strengthening.

Visualizations: Fig. 7 shows the visualizations defined in the ADF by the domain expert for comparing mitochondria (and their synaptic neighborhood) at low-, medium-, and high-cardinality. For *low-cardinality* comparisons, our collaborator used juxtaposed synaptic distance trees (Fig 7a), one tree for each mitochondrion under investigation. Inspired by phylogenetic trees [SWKP18], the tree encodes the mitochondrion (root, blue), its surrounding synapses (leaves), as well as the distance between structures (edge length). This view gives a quick insight into the synapse arrangement of the immediate neighborhood, while abstracting complex 3D morphology. The *medium-cardinality* comparison uses a grouped bar chart and gives a compact view on several mitochondria and their distances to synapses (Fig 7b). The *high-cardinality* view uses a distance matrix to compare dozens to hundreds of mitochondria simultaneously (Fig 7c). Rows correspond to mitochondria and columns to synapses. Distances between mitochondria and synapses are color-coded, darker colors representing closer distances.

Interactive analysis: The expert started his analysis by using the list view to select an axonal mitochondrion, as well as a spatially close dendritic mitochondrion. Next, he compared both in the distance tree, looked at their synapses in the 3D neighborhood views, and found a distant synapse of the dendritic mitochondrion that was particularly large. He enabled the EM slice to study the surroundings of this synapse and spotted glycogen granules, cell organelles responsible for providing energy, in its vicinity. He hypothesized that these glycogen granules might be responsible for the genesis of a large synapse otherwise distant from its related mitochondrion (i.e., energy provider). After this initial discovery, our collaborator extended his analysis to four axonal and four dendritic mitochondria, which led Barrio to show the medium-cardinality visualization. Using the grouped bar chart, he could confirm his hypothesis that axonal mitochondria have much closer synapses than dendritic mitochondria. However, he discovered a close synapse to a dendritic mitochondrion in the bar chart (see highlighted bar in Fig. 7b), which he confirmed to be a shaft synapse by using the 3D view. Shaft synapses are not formed at dendritic spines but are directly located on the cell membrane of the dendrite, and their functionality is not fully understood yet.

8.2. Mitochondria Placement (ST2)

Domain Goal: Here, neuroscientist wanted to investigate how close mitochondria come to the cell membrane of their neurite and if cell regions with close proximity to mitochondria differ from other areas. Since mitochondria provide energy to the cell, such morphological influences are of biological interest but are difficult to analyze with previous tools that give no quantitative feedback.

Visualizations: For this task, the expert needed to look at the detailed 3D position of a mitochondrion, but also at quantitative visualizations of the distance distribution (Fig. 8), which he defined in the ADF. Barrio supports color-coding structures in the 3D views, based on user-defined attributes. For this task, the scientist defined that each point on a mitochondrion's surface encoded its distance to the closest point on the cell membrane. Red regions indicate close proximity, while blue areas reveal regions more distant to the cell membrane (Fig. 6c). For *low-cardinality* comparisons, the expert




Figure 8: Analysis views generated for ST2: Mitochondria placement. (a) to (c): low-, medium-, high-cardinality comparison views. (a) cumulative histogram showing the percentage of a single mitochondrion's surface that is within a certain distance to the surrounding cell membrane (red means close, blue means far), (b) violin plots for four mitochondria, each showing the distance distribution of the mitochondrion's surface to the surrounding cell membrane (color and vertical position encode the distance, the width of the plot encodes the number of surface points with that distance), (c) scatter plot showing many mitochondria (dots) and their volume correlated to their distance to the cell membrane.

used small multiples of cumulative histograms for each mitochondrion, encoding the distances from each surface point of the mitochondrion to the cell membrane (Fig. 8a). This view allows users to judge whether a mitochondrion follows closely the cell membrane, or whether it only comes close in a small area. The *medium-cardinality* comparison uses juxtaposed violin plots, where each violin shape corresponds to a single mitochondrion (Fig. 8b), allowing a more scalable comparison of several structures. The *high-cardinality* comparison is based on a scatterplot where each dot represents a single mitochondrion (Fig. 8c). The horizontal axis represents the percentage of a mitochondrion surface that is a user-defined distance (here: 50 nm) or closer to the cell membrane. On the vertical axis, we encode the mitochondrion's volume.

Interactive analysis: The expert started his investigation by comparing two axonal and two dendritic mitochondria. First, he studied the immediate proximity of the mitochondria using the EM slice view (see Fig. 6a). Next, he observed a steeper curve in the cumulative histogram for the axonal mitochondria, indicating that axonal mitochondria are in general closer to the cell membrane. The scientist used the violin plots to identify mitochondria with a wider variety in their distance distributions, which all turned out to be dendritic mitochondria. He studied these mitochondria in the 3D view to find other neuronal structures close to them. He discovered some synapses, which he marked for further investigation. In a last step, the scientist searched for mitochondria neighborhoods by examining outliers in the scatterplot. While analyzing these outliers in the 3D views, he discovered that boutons often develop synapses in regions where the mitochondrion is distant from the cell membrane. This pattern has not been observed before, and the domain scientist is now studying this behavior in a follow-up analysis.

Case Study Findings: During this case study, the scientist made new discoveries in the dataset that he had worked with for months. He is now performing in-depth follow-up analyses on bouton development close to mitochondria (ST2) and mitochondrial spine coverage (ST3). He was also able to adjust the ADF file without prior programming experience, and was comfortable steering and conducting the analysis session in Barrio after a one-hour training session.

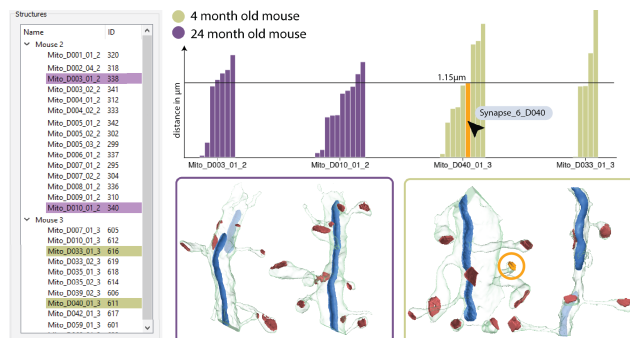


Figure 9: Comparing synapse locations in two different mouse data sets (purple and green). The 3D views show the four selected mitochondria neighborhoods. Grouped bars show the distances between mitochondria and their neighboring synapses. Hovering over a bar highlights the corresponding synapse in the 3D neighborhood view (orange). The young mouse (green, right) exhibits a wider variety of mitochondria-synapse distances than the adult mouse (purple, left), indicating synaptic pruning over time.

9. Qualitative User Study

In addition to the case study, we conducted a qualitative user study with eight experts to assess the effectiveness and usefulness of Barrio. To demonstrate the ease of sharing analysis scenarios using our ADF format, and to be able to compare the analysis sessions of different users, we asked each expert to use the 'Mitochondria Analysis' ADF file as a starting point. We describe the related subtasks in Sec. 8 and the supplementary material. After their initial analysis, each expert fine-tuned Barrio to their individual analysis needs. For example, one scientist loaded two datasets of brain tissue at different ages to analyze the change of *synaptic neighborhoods over time* (see Fig. 9). Others focused on synapses as target types, or analyzed synapses that were spatially close to, but *unconnected* to selected neurites, which required minor modifications to the ADF.

Participants. We evaluated Barrio with eight experts (P1-P8, 1 male, 7 female) from the Neuroscience Institute at the University of Turin and the Center for Brain Science at Harvard University. All experts (three professors, one postdoctoral researcher, two senior Ph.D. students, and two undergraduates) are experienced in analyzing 3DEM data of mice brains.

Setup. We met with each participant for a two-hour session remotely over Zoom video conferencing. After a one hour introduction to Barrio and its user interface, we started a data exploration session. In the study, we used a data set provided by the domain experts showing layer one of the somatosensory cortex of a four-month-old mouse with 1,088 segmented structures.

Feedback. Fig. 10 shows user ratings for general real-world applicability, tool effectiveness, and relevance for each of the three scenario subtasks. Seven out of eight participants rate questions Q1-Q4 positively, with the remaining participant being neutral. All participants agreed that our tool improved the analysis of 3DEM data. This is not surprising; other than Barrio, no study participant currently has visual tools available to analyze spatial neighborhoods of brain tissue. Prior to Barrio, P2 used Blender for visual exploration and manually computed distances between structures. P2 did not

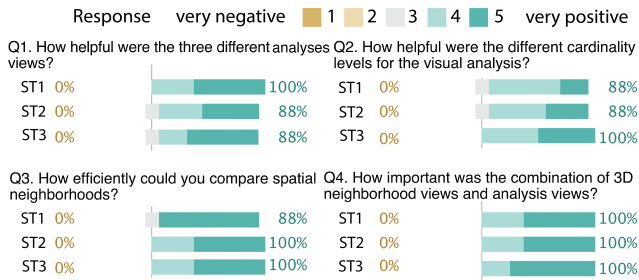


Figure 10: Barrio user ratings on the three case study subtasks (ST1 - ST3). We show neutral responses (3) in gray, positive responses (4, 5) on the right side in teal and the percentage of participants who responded positively.

have any support for visual comparisons other than taking screenshots. P6, P7, and P8 previously wrote small Python scripts to analyze brain tissue visually.

9.1. Findings

Combining EM data with reconstructed data is essential. Domain experts highly appreciated overlaying the 3D meshes with the EM image data. It allowed them to verify segmentation accuracy and enabled them to spot unsegmented properties of brain tissue.

Experts use high-cardinality comparisons to spot trends. Most participants looked at the high-cardinality visualization first to gain an overview and to spot outliers or neighborhoods of interest, which they then examined more closely in the neighborhood view.

Larger data sets could improve neighborhood analysis. Due to the high resolution of EM data, currently, our data sets do not include full reconstructions of nerve cells, but only small parts of them. Neurites and cell organelles can get cut off, which limits the quantitative analysis of those incomplete structures.

Custom analysis scenarios simplified the workflow. By being able to fine-tune the analysis in the ADF file, scientists were able to focus on their specific analysis needs and interests. Therefore, participants were immediately engaged in the tool, as they saw potential benefits, even for their highly individualized scientific goals.

Barrio improves the state of the art. All participants indicated that Barrio is useful (four rated it useful, and four rated it very useful) and that the software has a high potential as a visual analysis tool for 3DEM data. P1, P2, P3, P6, and P7 plan to specify additional analysis scenarios for their data sets.

10. Discussion

Additional features. During the user study, participants suggested a couple of additional features for Barrio. For instance, it would be interesting to extend Barrio with visualizations for analyzing the texture of mitochondria in greater detail. Domain experts P6 and P7 have classified mitochondria based on their texture in EM images into bright, dark, and combined mitochondria, and they are interested in relating these categorizations with our spatial neighborhood analysis. Additionally, P6 and P7 pointed out that knowledge about spatial neighborhoods of cell organelles such as mitochondria could help to classify cell types in the future.

Limitations. First, some preprocessing steps, such as segmentation, currently still require manual work. For instance, the dendritic spines and the cell body of all axons and dendrites were manually labeled. However, manual data preprocessing makes the analysis of larger data sets difficult. Second, for specifying ADFs, we offer a human-readable format that does not require programming. However, it still requires manual editing and a technical description of analysis scenarios. In the future, we could further simplify the ADF specification for users by offering a visual interface to it.

Expressivity vs. Usability Trade-off. An important lesson learned were the nuances in the expressivity/usability tradeoff. On the one hand, our tool is fine-tuned to specific data and the 3DEM workflow (i.e., high in usability), but at the same time, Barrio is flexible enough to reflect the needs of individual researchers (i.e., high in expressivity). This flexibility is crucial to accommodate various neuroscientists working on 3DEM data. However, it comes with the additional burden of modifying or defining custom scenarios and ADFs. To ease this burden, we provide template ADFs, a tutorial, and guidelines on writing custom scenarios [VCG22b]. Furthermore, ADFs are human-readable, allowing scientists to share their analysis scenarios or start from a provided template.

Flexibility is key. In addition to the flexibility afforded by custom scenarios, we found that flexibility in terms of data size and the size of visual comparisons is crucial. Barrio adjusts its views based on the user-defined scenario and the number of neighborhoods in a comparison. These automatic adjustments based on data and user goals simplify and streamline a users' workflow and can be applied far beyond the domain of neuroscience.

11. Conclusions and Future Work

We plan to extend Barrio to more usage scenarios involving other cell organelles, such as ribosomes and neuromuscular junctions. So far, our collaborators are not segmenting and reconstructing those structures, which limits their usage in Barrio. We further plan to use the neighborhood characteristics discovered in Barrio to investigate the extent to which they allow us to automatically label neuronal structures (e.g., to differentiate excitatory and inhibitory synapses). We also plan to incorporate topology-centered methods for the classification of neurites.

We believe that our neighborhood analysis approach is generalizable to other areas, such as single-cell analysis or the exploration of neighborhoods in city models. Supporting customized analysis scenarios can provide helpful functionality in many domain-specific visual analysis tools where users have varying analysis goals. Finally, we believe that Barrio can advance neuroscience research, not just for the presented usage scenario of mitochondria analysis, but also for future novel domain questions that have yet to be specified by neuroscientists.

12. Acknowledgements

We thank C. Nobre, K. Yin, M. Boido, X. Wang, M. Santos, N. Marino, M. Lorenzati and S. Stanga for their feedback. This research was supported by the National Science Foundation (NSF) Award Number IIS-1901030, King Abdullah University of Science and Technology (KAUST), and by the University of Torino.

References

- [AABH*16] AL-AWAMI A., BEYER J., HAEHN D., KASTHURI N., LICHTMAN J. W., PFISTER H., HADWIGER M.: NeuroBlocks – Visual Tracking of Segmentation and Proofreading for Large Connectomics Projects. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 738–746. doi:10.1109/TVCG.2015.2467441. 2
- [AABS*14] AL-AWAMI A., BEYER J., STROBELT H., KASTHURI N., LICHTMAN J. W., PFISTER H., HADWIGER M.: NeuroLines: A Subway Map Metaphor for Visualizing Nanoscale Neuronal Connectivity. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2369–2378. doi:10.1109/TVCG.2014.2346312. 2
- [ABHR*13] ALPER B., BACH B., HENRY RICHE N., ISENBERG T., FEKETE J.: Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), CHI '13, pp. 483–492. URL: <https://doi.org/10.1145/2470654.2470724>, doi:10.1145/2470654.2470724. 3
- [ACA*19] AGUS M., CALÌ C., AL-AWAMI A., GOBBETTI E., MAGISTRETTI P., HADWIGER M.: Interactive Volumetric Visual Analysis of Glycogen-derived Energy Absorption in Nanometric Brain Structures. *Computer Graphics Forum* 38, 3 (2019), 427–439. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13700>, doi:https://doi.org/10.1111/cgf.13700. 2
- [BAAK*13] BEYER J., AL-AWAMI A., KASTHURI N., LICHTMAN J. W., PFISTER H., HADWIGER M.: ConnectomeExplorer: Query-Guided Visual Analysis of Large Volumetric Neuroscience Data. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2868–2877. doi:10.1109/TVCG.2013.142. 2, 8
- [BH09] BOSTOCK M., HEER J.: Protovis: A Graphical Toolkit for Visualization. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1121–1128. doi:10.1109/TVCG.2009.174. 2
- [BHAA*13] BEYER J., HADWIGER M., AL-AWAMI A., WON-KI J., KASTHURI N., LICHTMAN J. W., PFISTER H.: Exploring the Connectome: Petascale Volume Visualization of Microscopy Data Streams. *IEEE Computer Graphics and Applications* 33, 4 (2013), 50–61. URL: <http://ieeexplore.ieee.org/document/6562713/>, doi:10.1109/MCG.2013.55. 2
- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D3 Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. doi:10.1109/TVCG.2011.185. 2, 7
- [BSL18] BERGER D., SEUNG H. S., LICHTMAN J. W.: VAST (Volume Annotation and Segmentation Tool): Efficient Manual and Semi-Automatic Labeling of Large 3D Image Stacks. *Frontiers in Neural Circuits* 12 (2018). URL: <https://www.frontiersin.org/articles/10.3389/fncir.2018.00088/full>, doi:10.3389/fncir.2018.00088. 2, 3
- [CBB*16] CALÌ C., BAGHABRA J., BOGES D., HOLST G., KRESHUK A., HAMPRECHT F., SRINIVASAN M., LEHVÄSLAIHO H., MAGISTRETTI P.: Three-dimensional immersive virtual reality for studying cellular compartments in 3D models from EM preparations of neural tissues. *Journal of Comparative Neurology* 524, 1 (2016), 23–38. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.23852>, doi:https://doi.org/10.1002/cne.23852. 2
- [CSS*12] CARDONA A., SAALFELD S., SCHINDELIN J., ARGANDA-CARRERAS I., PREIBISCH S., LONGAIR M., TOMANCAK P., HARTENSTEIN V., DOUGLAS R.: TrakEM2 Software for Neural Circuit Reconstruction. *PLoS One* 7, 6 (2012), e38011. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0038011>, doi:10.1371/journal.pone.0038011. 3
- [CWB*18] CALÌ C., WAWRZYNIAC M., BECKER C., MACO B., CANTONI M., JORSTAD A., NIGRO B., GRILLO F., DE PAOLA V., FUA P., OTHERS: The effects of aging on neuropil structure in mouse somatosensory cortex—A 3D electron microscopy analysis of layer 1. *PLoS One* 13, 7 (2018), e0198131. doi:https://doi.org/10.1371/journal.pone.0198131. 3
- [DSG*17] DINKLA K., STROBELT H., GENEST B., REILING S., BOROWSKY M., PFISTER H.: Screenit: Visual Analysis of Cellular Screens. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 591–600. doi:10.1109/TVCG.2016.2598587. 2
- [FCOO12] FREITAS A., CURRY E., OLIVEIRA J. G., O'RIAIN S.: Querying Heterogeneous Datasets on the Linked Data Web: Challenges, Approaches, and Trends. *IEEE Internet Computing* 16, 1 (Jan. 2012), 24–33. URL: <http://ieeexplore.ieee.org/document/6051410/>, doi:10.1109/MIC.2011.141. 2
- [FGBH0dS05] FERNANDEZ-GONZALEZ R., BARCELLOS-HOFF M., ORTIZ-DE SOLORZANO C.: A tool for the quantitative spatial analysis of complex cellular systems. *IEEE Transactions on Image Processing* 14, 9 (2005), 1300–1313. doi:10.1109/TIP.2005.852466. 2
- [FGK*00] FABRI A., GIEZEMAN G.-J., KETTNER L., SCHIRRA S., SCHÖNHERR S.: On the design of CGAL a computational geometry algorithms library. *Software: Practice and Experience* 30, 11 (2000), 1167–1202. doi:https://www.cgal.org/. 7
- [GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C., ROBERTS J.: Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309. URL: <https://doi.org/10.1177/1473871611416549>, doi:10.1177/1473871611416549. 2
- [GG90] GLASER J., GLASER E.: Neuron imaging with NeuroLucida—a PC-based system for image combining microscopy. *Computerized Medical Imaging and Graphics* 14, 5 (1990), 307–317. doi:https://doi.org/10.1016/0895-6111(90)90105-K. 2
- [GWB*21] GONDA F., WANG X., BEYER J., HADWIGER M., LICHTMAN J. W., PFISTER H.: VICE: Visual Identification and Correction of Neural Circuit Errors. *Computer Graphics Forum* 40, 3 (2021), 447–458. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14320>, doi:10.1111/cgf.14320. 2
- [HBJP12] HADWIGER M., BEYER J., JEONG W.-K., PFISTER H.: Interactive Volume Exploration of Petascale Microscopy Data Streams Using a Visualization-Driven Virtual Memory Approach. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2285–2294. URL: <http://ieeexplore.ieee.org/document/6327233/>, doi:10.1109/TVCG.2012.240. 2
- [HHM*17] HAEHN D., HOFFER J., MATEJEK B., SUISSA-PELEG A., AL-AWAMI A., KAMENSKY L., GONDA F., MENG E., ZHANG W., SCHALEK R., WILSON A., PARAG T., BEYER J., KAYNIG V., JONES T., TOMPKIN J., HADWIGER M., LICHTMAN J., PFISTER H.: Scalable Interactive Visualization for Connectomics. *Informatics* 4, 3 (2017), 29. URL: <https://www.mdpi.com/2227-9709/4/3/29>, doi:10.3390/informatics4030029. 2
- [HKT*18] HAEHN D., KAYNIG V., TOMPKIN J., LICHTMAN J., PFISTER H.: Guided proofreading of automatic segmentations for connectomics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 9319–9328. doi:10.1109/CVPR.2018.00971. 2
- [HSKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T.: Topology matching for fully automatic similarity estimation of 3D shapes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (Aug. 2001), SIGGRAPH '01, pp. 203–212. URL: <https://doi.org/10.1145/383259.383282>, doi:10.1145/383259.383282. 3
- [JNC*15] JORSTAD A., NIGRO B., CALÌ C., WAWRZYNIAC M., FUA P., KNOTT G.: NeuroMorph: A Toolset for the Morphometric Analysis and Visualization of 3D Models Derived from Electron Microscopy Image Stacks. *Neuroinformatics* 13, 1 (2015), 83–92. URL: <https://doi.org/10.1007/s12021-014-9242-5>, doi:10.1007/s12021-014-9242-5. 2, 7

- [KBJ*20] KRUEGER R., BEYER J., JANG W.-D., KIM N. W., SOKOLOV A., SORGER P. K., PFISTER H.: Facetto: Combining Unsupervised and Supervised Learning for Hierarchical Phenotype Analysis in Multi-Channel Image Data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 227–237. doi:10.1109/TVCG.2019.2934547. 2
- [KMWL08] KNOTT G., MARCHMAN H., WALL D., LICH B.: Serial Section Scanning Electron Microscopy of Adult Brain Tissue Using Focused Ion Beam Milling. *Journal of Neuroscience* 28, 12 (2008), 2959–2964. URL: <https://www.jneurosci.org/content/28/12/2959>, doi:10.1523/JNEUROSCI.3189-07.2008. 3
- [LBK*18] LEKSCHAS F., BACH B., KERPEDEJIEV P., GEHLENBORG N., PFISTER H.: HiPiler: Visual Exploration of Large Genome Interaction Matrices with Interactive Small Multiples. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 522–531. doi:10.1109/TVCG.2017.2745978. 2
- [LD11] LICHTMAN J. W., DENK W.: The Big and the Small: Challenges of Imaging the Brain's Circuits. *Science* 334, 6056 (Nov. 2011), 618–623. URL: <https://www.sciencemag.org/lookup/doi/10.1126/science.1209168>, doi:10.1126/science.1209168. 1
- [MAAB*18] MOHAMMED H., AL-AWAMI A., BEYER J., CALÌ C., MAGISTRETTI P., PFISTER H., HADWIGER M.: Abstractocyte: A Visual Tool for Exploring Nanoscale Astroglial Cells. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 853–861. doi:10.1109/TVCG.2017.2744278. 2
- [MGT*03] MUNZNER T., GUIMBRETIÈRE F., TASIRAN S., ZHANG L., ZHOU Y.: TreeJuxtaposer: scalable tree comparison using Focus+Context with guaranteed visibility. *ACM Transactions on Graphics* 22, 3 (2003), 453–462. URL: <https://doi.org/10.1145/882262.882291>, doi:10.1145/882262.882291. 2
- [MMH*13] MARIÉS A., MAYS N., HUNT M., WONG K. F., LAYTON W., BOUDREAU R., ROSANO C., MARAI G. E.: GRACE: A Visual Comparison Framework for Integrated Spatial and Non-Spatial Geriatric Data. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2916–2925. doi:10.1109/TVCG.2013.161. 2
- [NDBB17] NHU H., DRIGO R. A. E., BERGGREN P.-O., BOUDIER T.: A novel toolbox to investigate tissue spatial organization applied to the study of the islets of Langerhans. *Scientific reports* 7, 1 (2017), 1–12. doi:10.1038/srep44261. 2
- [Pet15] PETTIT J.: *Spatial analysis of complex biological tissues from single cell gene expression data*. PhD Thesis, University of Cambridge, 2015. URL: <https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.708750>. 2
- [PKB*12] PFISTER H., KAYNIG V., BOTHA C., BRUCKNER S., DERCKSEN V., HEGE H., ROERDINK J.: Visualization in Connectomics. *arXiv:1206.1428 [cs, q-bio]* (2012). URL: <http://arxiv.org/abs/1206.1428>. 2
- [PP95] PAGENDARM H., POST F.: Comparative visualization: approaches and examples. *Visualization in Scientific Computing*, pp. 95–108. URL: <https://elib.dlr.de/37368/>. 2
- [SBS*13] SORGER J., BÜHLER K., SCHULZE F., LIU T., DICKSON B.: neuroMAP — Interactive graph-visualization of the fruit fly's neural circuit. In *IEEE Symposium on Biological Data Visualization (BioVis)* (2013), pp. 73–80. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6664349>, doi:10.1109/BioVis.2013.6664349. 2
- [SFG*20] STOLTZFUS C., FILIPEK J., GERN B., OLIN B., LEAL J., WU Y., LYONS-COHEN M., HUANG J., PAZ-STOLTZFUS C., PLUMLEE C., OTHERS: CytoMAP: a spatial analysis toolbox reveals features of myeloid cell organization in lymphoid tissues. *Cell reports* 31, 3 (2020), 107523. doi:10.1016/j.celrep.2020.107523. 2
- [SH14] SATYANARAYAN A., HEER J.: Lyra: An Interactive Visualization Design Environment. *Computer Graphics Forum* 33, 3 (2014), 351–360. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12391>, doi:10.1111/cgf.12391. 2
- [SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2431–2440. doi:10.1109/TVCG.2012.213. 3
- [SMWH17] SATYANARAYAN A., MORITZ D., WONGSUPHASAWAT K., HEER J.: Vega-Lite: A Grammar of Interactive Graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 341–350. doi:10.1109/TVCG.2016.2599030. 2
- [SPA*14] SCHMIDT J., PREINER R., AUZINGER T., WIMMER M., GRÖLLER E., BRUCKNER S.: YMCA — Your mesh comparison application. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), pp. 153–162. doi:10.1109/VAST.2014.7042491. 3
- [SRHH16] SATYANARAYAN A., RUSSELL R., HOFFSWELL J., HEER J.: Reactive Vega: A Streaming Dataflow Architecture for Declarative Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 659–668. URL: <http://ieeexplore.ieee.org/document/7192704/>, doi:10.1109/TVCG.2015.2467091. 2
- [SSKH11] SOMMER C., STRAEHLE C., KÄTHE U., HAMPRECHT F. A.: Ilastik: Interactive learning and segmentation toolkit. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro* (2011), pp. 230–233. doi:10.1109/ISBI.2011.5872394. 3
- [Stu08] STUFFLEBEAM R.: Neurons, synapses, action potentials, and neurotransmission. *Consortium on Cognitive Science Instruction* (2008). 3
- [SW06] SWEENEY D., WILLIAMSON B.: *Biology: Exploring Life: Laboratory Manual*. Pearson Education, Incorporated, 2006. 3
- [SWC*08] SCHNEIDER D., WIEBEL A., CARR H., HLAWITSCHKA M., SCHEUERMANN G.: Interactive Comparison of Scalar Fields Based on Largest Contours with Applications to Flow Visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1475–1482. doi:10.1109/TVCG.2008.143. 3
- [SWH14] SATYANARAYAN A., WONGSUPHASAWAT K., HEER J.: Declarative interaction design for data visualization. In *Proceedings of the 27th annual ACM symposium on User interface software and technology* (2014), UIST '14, pp. 669–678. URL: <https://doi.org/10.1145/2642918.2647360>, doi:10.1145/2642918.2647360. 2
- [SWKP18] SHANK S., WEAVER S., KOSAKOVSKY POND S.: phyloTree.js - a JavaScript library for application development and interactive data visualization in phylogenetics. *BMC Bioinformatics* 19, 1 (July 2018), 276. URL: <https://doi.org/10.1186/s12859-018-2283-2>, doi:10.1186/s12859-018-2283-2. 7, 8
- [VCG22a] VCG: Barrio codebase, 2022. last accessed: 04/17/2022. URL: <https://github.com/VCG/Barrio>. 7
- [VCG22b] VCG: Barrio Documentation, 2022. last accessed: 04/17/2022. URL: <https://github.com/VCG/Barrio/wiki>. 5, 7, 10
- [Wav21] WAVEFRONT: wavefront format, 2021. URL: https://en.wikipedia.org/wiki/Wavefront_.obj_file. 7
- [Wol11] WOLFF D.: *OpenGL 4.0 shading language cookbook*. Packt Publishing Ltd, 2011. 7