

Improving the Usability of Virtual Reality Neuron Tracing with Topological Elements

Torin McDonald, Will Usher, Nate Morrical, Attila Gyulassy, Steve Petruzza, Frederick Federer, Alessandra Angelucci, and Valerio Pascucci




Fig. 1: Left to right: A connected graph of ridge-like structures is extracted from the Morse-Smale complex (MSC), containing a superset of the possible neuron segments in the data. Our MSC-guided semi-automatic tracing tool enables users to rapidly trace paths and view a live preview as they do so (orange line). When satisfied with the trace, they can add it to the reconstruction (white line).

Abstract—Researchers in the field of connectomics are working to reconstruct a map of neural connections in the brain in order to understand at a fundamental level how the brain processes information. Constructing this wiring diagram is done by tracing neurons through high-resolution image stacks acquired with fluorescence microscopy imaging techniques. While a large number of automatic tracing algorithms have been proposed, these frequently rely on local features in the data and fail on noisy data or ambiguous cases, requiring time-consuming manual correction. As a result, manual and semi-automatic tracing methods remain the state-of-the-art for creating accurate neuron reconstructions. We propose a new semi-automatic method that uses topological features to guide users in tracing neurons and integrate this method within a virtual reality (VR) framework previously used for manual tracing. Our approach augments both visualization and interaction with topological elements, allowing rapid understanding and tracing of complex morphologies. In our pilot study, neuroscientists demonstrated a strong preference for using our tool over prior approaches, reported less fatigue during tracing, and commended the ability to better understand possible paths and alternatives. Quantitative evaluation of the traces reveals that users' tracing speed increased, while retaining similar accuracy compared to a fully manual approach.

Index Terms—Virtual Reality, Morse-Smale Complex, Semi-automatic Neuron Tracing

1 INTRODUCTION

A central goal within the field of neuroscience is to understand how the dense, interconnected neural circuits in the brain communicate and process information, and how this processing relates to behavior. The field of connectomics was founded to understand the fundamental wiring map of the brain in order to comprehend these neural circuits at a mechanistic level. Through analyzing neuron structure and connectivity, neuroanatomists can gain a deeper understanding of fundamental brain functions and new insights about brain diseases and treatments.

However, obtaining a comprehensive wiring diagram for even relatively small and simple mammalian brains, such as that of a mouse, is a massive undertaking [8, 13, 44, 61]. Projects focusing on species with larger brains more similar to humans, such as non-human primates (NHP), are even more challenging. Although recent advancements in high-resolution tissue labeling [35], optical tissue clearing [15, 42, 70] and imaging [44, 61] have made it possible to image NHP brains at large scales and high resolutions, the technology for extracting the imaged

neuron morphologies has struggled to keep up.

Current efforts to improve the speed of neuron morphology extraction have largely focused on fully automatic techniques. Automatic techniques take a stack of images and attempt to extract the imaged neuron structures, without user input. The DIADEM (DIGital reconstructions of Axonal and DEndritic Morphology) Challenge [21] was proposed in 2009 to motivate improvement of these techniques. The ultimate goal of this community effort was to increase the speed that neurons could be traced by 20 \times . However, at the end of the challenge no algorithm had achieved this goal due to the laborious post-processing required to correct errors [34]. Peng et al. [49] reported that this post-processing step can take longer than a manual tracing. Although additional efforts to improve automatic reconstruction are ongoing [50], in practice the bulk of neuron tracing is done manually [40] or with a semi-automatic method.

Manually tracing neurons is a difficult and time-consuming process. Tracing is typically done on a desktop, using standard software (e.g., NeuroLucida [39], Vaa3D [52]). The data is displayed as a 2D set of images or 3D volume, and the user clicks along the neuron to draw a path. The lack of ability to directly make selections in 3D or navigate the data in 3D introduces additional usability challenges on top of the already difficult task of tracing. To address this issue, Usher et al. [67] proposed a Virtual Reality (VR) based tool for manual neuron tracing and found that neuroscientists using the tool performed similar quality traces in less time.

Semi-automatic neuron tracing methods have been proposed to provide a compelling alternative to both manual and fully automatic neuron tracing [38, 43, 51, 52]. When using a semi-automatic method, the user provides coarse guidance to the algorithm, e.g., through a set of start and end points or clicks. The algorithm then extracts the neuron structure

- Torin McDonald, Will Usher, Nate Morrical, Attila Gyulassy and Valerio Pascucci are with the SCI Institute, University of Utah. torin@sci.utah.edu.
- Steve Petruzza is with the SCI Institute, University of Utah and Utah State University.
- Frederick Federer and Alessandra Angelucci are with the Moran Eye Institute, University of Utah.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

between these guide points. Semi-automatic methods can significantly reduce the amount of time taken to trace a neuron by integrating the neuroscientist's guidance into the algorithm, reducing the amount of post-processing manual cleanup required.

In this work we propose a new semi-automatic neuron tracing framework that builds on topological analysis methods [22], developed through direct collaboration with expert neuroanatomists. Our approach uses the Morse-Smale complex (MSC) to precompute a superset of potential paths that follow neurons. Having access to this superset of traces allows neuroanatomists to quickly trace along the neuron of interest by selecting subsets of these paths. We implement our semi-automatic method within a virtual reality neuron tracing system to provide an intuitive environment to work with the 3D data. In a pilot study with neuroanatomists, we find that our approach provides significant benefits, retaining trace accuracy while improving speed and reducing fatigue. Moreover, the additional guides provided by our tool assist the user in interpreting the data. Our contributions are:

- A novel topologically guided framework for real-time semi-automatic neuron tracing
- An intuitive interaction design for using this framework in VR
- A comparison of our approach against widely used semi-automatic methods as well as previous manual tracing methods through a user study with domain experts.

2 BACKGROUND AND RELATED WORK

The neuron morphology reconstruction workflow has a number of components, with one of the most time consuming being the physical tracing of neurons. To provide context for neuron tracing, we describe the typical reconstruction workflow in practice (Section 2.1). We then review current automatic and semi-automatic neuron tracing methods and their limitations (Section 2.2) and the state of the art in immersive environments (Section 2.3). Finally, we review the Morse-Smale complex and examples of its application to analysis tasks in other scientific domains (Section 2.4).

2.1 Neuron Tracing Workflow

Modern methods for acquiring neuron microscopy data use viral vectors carrying genes for fluorescent proteins [35]. When injected into the tissue these vectors induce fluorescence within the structures to be imaged, labelling them at high resolution. The brain tissue is then rendered optically transparent using a clearing technique such as CLARITY [15], PACT [70], or SWITCH [42], and imaged in blocks with a confocal or two-photon microscope. These methods allow for imaging large blocks of tissue or entire brains, and can produce terabytes of high-resolution image stacks.

To reconstruct the labeled neurons from these image stacks, neuroanatomists use commercial tools like NeuroLucida [39] or open-source tools like Vaa3D [52]. These tools display the collected image stacks as either a set of 2D slices or as a 3D volume, where the user can trace manually by drawing lines along the structures of interest, or guide a semi-automatic algorithm along the structures to extract them. Once the desired neurons have been reconstructed, they can be used in brain function simulations or overlaid on functional maps of the brain, to understand the connectivity between brain regions. Although fully automatic algorithms are supported by standard tools, they are less widely used in practice due to issues with image quality and ambiguity. Our coauthor's lab employs several trained undergraduates responsible for the bulk of the neuron tracing work, with additional tracing done by graduate students and research scientists.

2.2 Automatic and Semi-Automatic Neuron Tracing

Today, neuron tracing remains a crucial bottleneck in the field of connectomics [40]. A large body of work has been devoted to developing new methods to accelerate this process, either through fully automatic algorithms or semi-automatic user-guided algorithms.

A significant ongoing effort in the community has been to develop and evaluate fully automatic algorithms for neuron reconstruction. Two

community efforts, the DIADEM Challenge [21] and the ongoing BigNeuron Project [47, 50], seek to provide a test bed for evaluating new reconstruction algorithms. Results from the DIADEM challenge suggest that the current state-of-the-art automatic tracing algorithms are not suitable for widespread use in practice, because significant manual post-processing is required [49]. Recent work has begun applying Machine Learning techniques to the neuron segmentation problem [5, 62]; however, as with other ML-based approaches, a large amount of training data must be provided as input to the algorithm, that must be produced using existing techniques. In contrast to work segmenting space-filling neurons imaged using electron microscopy [29], our approach works on sparsely labeled neurons as linear structures. For a full review of recent advances we refer to the recent survey by Magliaro et al. [36] and that by Acciai et al. [3].

Due to the challenges in using fully automatic methods in practice, semi-automatic algorithms have found a growing interest in the community. When using a semi-automatic reconstruction algorithm, the user guides the algorithm along the neuron by tracing roughly along the neuron or clicking to mark start, branch, and end points to connect. By integrating more guidance from the neuroscientist into the algorithm, the amount of additional post-processing cleanup required can be reduced, while still decreasing the time spent tracing compared to a fully manual trace. For example, Vaa3D's semi-automatic approach uses a pixel based shortest path algorithm [51] to connect the start point and one or more markers placed by the user. NeuroLucida 360's [38] semi-automatic tracing works similarly, where the user traces along the feature to guide the algorithm to important features. Neuromantic [43] uses a 3D extension of Meijering et al.'s 2D steerable Gaussian filter algorithm [41] for semi-automatic reconstruction.

However, these methods all work in the context of traditional desktop software, taking 2D inputs from a mouse and providing 2D imagery through a monitor. For example, Vaa3D's *Virtual Finger* [53] casts rays through the volume to find the potentially selected objects as the user draws a line with the mouse. Thus, users may need to perform multiple interactions and camera rotations to find and select the desired feature, to work around occluders or ambiguous hits in the ray casting process. Furthermore, such methods typically operate on the underlying image data and are computationally intensive, thereby impacting interactivity.

2.3 Immersive Environments

There has been a growing interest in using virtual reality or immersive environments for neuron tracing and visualization in general to overcome the limitations of traditional 2D desktop interaction and visualization modalities. Existing tools such as Vaa3D have announced early VR system support, and other new VR-specific tools have been released [2, 67]. In contrast to desktop software, VR and immersive systems allow users to visualize and interact with their data directly in 3D, providing a more intuitive interface and allowing for better understanding of 3D structures [30, 31, 55].

Usher et al. [67] proposed a virtual reality system for room scale or seated manual neuron tracing. In their evaluation, they found that domain experts could perform similar quality traces to standard desktop software in less time, achieving a roughly 2× speedup. Moreover, they found that experts reported the VR tool to be more intuitive and less fatiguing, with the immersive visualization aiding their understanding of the data. However, their tool supports only manual tracing and thus, while faster than working on a desktop, would still require a significant amount of time and effort to trace large data sets.

Immersive systems have also been proposed for visualization of electron and wide-field microscopy data. Agus et al. [4] proposed a model for simulating lactate absorption to compute lactate absorption maps on 3D segmented electron microscopy data sets. The absorption maps can then be visualized in a VR or CAVE environment, by rendering the segmented neuron meshes colored by absorption rate. Boges et al. [9, 10] proposed a virtual reality tool for creating, editing, and visualizing skeletonizations of brain cells in electron microscopy data, along with their segmented surface mesh. The neuron skeleton is created by the user with the assistance of a guidance system that moves their points placed inside the mesh to its center. In the use case we




Fig. 2: From left to right: the saddle-maximum arcs of the MSC are simplified with increasing persistence thresholds. Red lines represent saddle-maximum arcs, and red and green points correspond to maxima and saddles, respectively. A subset of the arcs cover neurons (left, middle), however, over-simplification discards faint features (right).

target, a segmentation of the data is not available. Instead, we leverage the Morse-Smale complex to guide users when tracing through a volumetric representation. Boorboor et al. [11] proposed a data processing and feature extraction pipeline, the output of which could be visualized in an immersive display wall visualization system implemented with Unity. Sicut et al. [60] presented DXR, a Unity based toolkit for easily developing immersive visualization applications. Fulmer et al. [20] presented a web-based immersive neuron visualization system using Unity to explore online databases of neuron data in a HoloLens.

2.4 Topological Analysis

Topological methods have been shown to be highly effective in extracting application-specific features of interest (e.g., [6, 7, 12, 24, 25, 32, 56, 63, 69]). In many cases, the features of interest can be defined directly in terms of topological structures such as the Contour Tree or the Morse-Smale complex (MSC). Recent advances in the computation of the MSC [57, 59] have made it readily available to the broader scientific community through open source libraries such as TTK [66] and MSCEER [22]. In this work, we build our framework for neuron tracing on top of the MSC.

Given a compact d -manifold M , a scalar function $f : M \rightarrow R$ is a Morse function if its critical points are non-degenerate and have distinct values. A critical point occurs where the gradient vanishes, $\nabla f = 0$, and is non-degenerate if its Hessian is non-singular.

For three-dimensional domains, a critical point is either a *minimum*, *1-saddle*, *2-saddle*, or *maximum*. An integral line in f is a path in M whose tangent vector is parallel to the gradient of f at each point along the path. The lower limit of the integral line is called the *origin*, and the upper limit the *destination*. These lower and upper limits occur at critical points of f . *Ascending* and *descending* manifolds are obtained as clusters of integral lines having common origin and destination, respectively. The descending manifolds of f partition M into a cell complex called the *Morse complex*. Symmetrically, the ascending manifolds also partition M into a cell complex. A Morse function f is a *Morse-Smale function* if the ascending and descending manifolds of its critical points intersect only transversally. The 0- and 1-dimensional cells of the intersection of ascending and descending manifolds form the *1-skeleton* of the Morse-Smale complex. Practically, the MSC *1-skeleton* is formed by *nodes* and *arcs*. *Nodes* are the critical points of the MSC, and *arcs* are the integral lines connecting critical points which differ in index by one. For a complete visual overview of the components of the MSC we refer to Gyulassy et al. [26].

One of the greatest advantages of using a topological approach is the multiscale analysis made possible by persistent homology [17], which allows for the simplification of noisy topological features. Persistence simplification repeatedly removes critical point pairs that form the birth and death of a topological feature, based on the lifespan of the feature in a sweep of the range of the function. Low-intensity noise often appears as low-persistence features that can be removed while maintaining an overall connected structure. The MSC 1-skeleton has well understood rules that govern its persistence simplification through repeatedly removing pairs of critical points connected by arcs, and reconnecting their neighborhoods [27]. Our use of the MSC is motivated by the observation that the ridge-like structures formed by the 1-skeleton of the simplified MSC, composed of the arcs between 2-saddles and maxima, correspond to the center lines of the vast majority of neurons in the data, as shown in Figure 2.

The components of the MSC have been used in practice to extract fea-

tures of interest in a range of application domains. For example, these components define features in the electron density field in the quantum theory of atoms in molecules: maxima occur at atom locations; 2-saddle-maximum arcs are covalent bonds; and descending 3-manifolds are atomic basins [6]. In other domains, specially selected subsets of the MSC can be used to extract features: descending 2-manifolds represent bubbles in mixing fluids [32]; 2-saddle-maximum arcs can be used to extract the core of a porous solid [24] as well as the filamentary structure of galaxies [63] or structural materials [54]; descending 2- and 1-manifolds identify lithium diffusion pathways [25]; and ascending 2-manifolds define burning regions in combustion simulations [12]. In each application, the features of interest were computed by identifying the appropriate topological abstraction, and then selecting a subset of the topological features that correspond to the quantities under study.

While our work is inspired by these approaches, the images generated from fluorescence microscopy pose a massive challenge for automated analysis. In addition to high-intensity noise, images of neurons have uneven staining, shadows, alignment artifacts from stitching image tiles, and other unexplained gaps in the signal which require manual intervention. This poses a challenge to topological methods that report what is in the scalar function itself, faithfully representing artifacts and noise along with the desired signal.

3 TOPOLOGY GUIDED NEURON TRACING

A semi-automatic method for neuron tracing in a VR environment faces two main constraints. First, maintaining interactivity with the high framerate requirements of VR places a hard time constraint on point queries, neighborhood queries, and path computation. As image volumes can reach tens to hundreds of gigabytes in size, on-the-fly computation on the raw voxel data is unable to meet these constraints. Second, the image quality varies dramatically even within a single volume. Traditional skeletonization and ridge-extraction techniques that take a background or foreground threshold will either over- or under-connect features of interest within the volume. Our topology-based ridges meet both requirements: (i) a fast and scalable precomputation produces a sparse data structure that is fast to query when tracing and (ii) the MSC-graph produced adapts to local image quality without relying on a global value threshold.

3.1 Computing MSC-Graphs

In the images produced through the imaging process described in Section 2.1, high-intensity values correspond to the labeled cell bodies (soma), and the structures connecting them to each other (dendrites and axons), which form each neuron. When tracing these structures, the user aims to produce a path that follows the center line of these ridge-like structures. Our approach in this work is to generate *every possible ridge-like path* first, turning the neuron reconstruction task into an efficient and interactive subset selection on a sparse data structure. This is in sharp contrast to existing semi-automatic and automatic methods, that attempt to mimic the manual extraction process by computing the single most-likely path for the user [43, 48, 71].

Our first task is to extract the set of all possible ridge lines from the scalar field. Historically, ridge lines have been defined with techniques relating to the alignment of the principle directions of curvature and the gradient, Eberly et al. [16] provide an excellent overview. However, locally defined ridge lines have major limitations for the task of acting as an acceleration structure for neuron reconstruction. Height ridges do not necessarily form an interconnected network, with segments ending where the local image no longer looks like a ridge. Pruning ridge lines by intensity further disconnects the network, exacerbating the problem. Instead, we use a topological approach based on the MSC that identifies ridge-like structures that are close enough to true ridge lines, but also form an interconnected network ideal for path computations. Figure 3 compares ridge-lines computed with iterative thinning, second-derivative ridges, and the MSC.

Although integral lines do not merge for continuous functions, the computational methods utilize discrete Morse theory [19, 23], where 2-saddle-maximum paths may merge. Our data structure inserts a new node at each merger to maintain the property that any vertex




Fig. 3: Comparison of ridge structure computation. Both (b) skeletonization through iterative thinning [33] and (c) second-derivative ridge detection [64] have gaps in the structure, whereas (d) the MSC-graph [26] produces connected paths. Parameters were adjusted in (b) and (c) to maximally connect the network; however, neither approach was able to recover from gaps in the signal or reproduce faint axons.




Fig. 4: The image preprocessing step reduces the effect of noise, creating a sparser initial MSC-graph with a better geometric embedding.

along a path can be mapped to an arc of the data structure. In the remainder of this text, we use *MSC-graph* to refer to this modified 1-skeleton. Figure 2 illustrates the full MSC and its successive coarsening through persistence simplification. The steps for computing our data structure are image preprocessing, MSC computation and persistence simplification, dividing arcs to form the MSC-graph, and geometric smoothing of the MSC-graph.

Image Preprocessing. Noise in microscopy images leads to poor geometric embedding of the MSC-graph, since it is sensitive to discontinuities in the image gradient (Figure 4a). The microscopy images used in this study show characteristic *speckle noise*, high-frequency, high-amplitude intensity spikes. A common practice is to first remove noise before further analysis [3]. While many sophisticated approaches exist for noise removal, a simple combination of median and Gaussian blur filters (with radius 2) led to MSC-graphs with both fewer arcs and better embedding (Figure 4b). This filter combination produced excellent results for all images in the study coming from different staining techniques, tissue samples, magnifications, and microscopes. However, different acquisitions may require alternative denoising approaches.

MSC Computation and Simplification. We use a standard approach based on discrete Morse theory to construct a discrete gradient field [23, 57], available in the open-source MSCEER [22] library. The library computes a discrete representation of the gradient using a parallel local filter, after which it traces integral paths in the gradient field to construct the 1-skeleton of the MSC. MSCEER also supports computing the MSC at a user-specified persistence simplification threshold [27], which we use to simplify extraneous features created by noise. Higher thresholds will produce coarser complexes and sparser MSC-graphs; however, selecting too high of a threshold may remove faint but desirable features. We empirically found that selecting a low threshold—as low as 1% of the function range—is sufficient to remove a large portion of the noise while keeping the majority of faint neurons (Figure 5). Finally, the 2-saddle-maximum arcs of the MSC are split as needed to remove geometric overlap, forming the MSC-graph.

MSC-Graph Post-Processing. The discrete gradient used in computing the MSC and MSC-graph produces arcs whose segments follow the staircase-like structure of voxels aligned with the underlying grid axes. These arcs are smoothed using a simple averaging of neighbor positions to produce more natural appearing results and improve arc length calculations (Figure 6).




Fig. 5: Low-threshold persistence simplification removes many extraneous features due to noise while preserving faint but desirable ones.




Fig. 6: The arcs and critical points from the original MSC-graph in its discretized form (left), and after applying the smoothing process (right).

3.2 A Fast and Efficient Querying Framework

Our semi-automatic tracing method is based on adapting the MSC-graph into a data structure that can quickly answer nearest neighbor and shortest path queries based on user input and navigation. In VR, we require that all queries must return and update the rendering in well under 11ms to maintain 90 FPS, otherwise they can cause the visualization to skip frames, which is disorienting to users.

The key interactions in our system are picking the closest point in the MSC-graph, picking all points within some radius, and computing shortest paths between points in the MSC-graph. The geometry of each arc of the MSC-graph is stored as an ordered set of 3D points. To accelerate nearest point and radius queries on the MSC-graph, we insert these points into a k -d tree, which is built at start up when loading the data. Each point contains a reference to the arc in the MSC-graph from which it originated. To find the nearest point to the controller, the k -d tree is queried and the corresponding arc looked up in the MSC-graph.

To find the best path reconstructing a neuron between some start and end point, we query the MSC-graph using a weighted shortest path algorithm (Figure 7c). Paths corresponding to neurons have higher intensity values in the image, thus we bias path selection towards paths through voxels with higher data values. The weight of each arc is the integral of $w(p) = \epsilon + 1.0 - I(p)$ over the arc, where I is the *normalized image intensity*, p is a point in the domain, and ϵ is a small constant to avoid zero-cost arcs. The arc weights are computed at startup to improve runtime performance. When a selection query returns a point inside an existing arc, we symbolically split the arc at the point and recompute the weights of each of the new arc segments.

To extract the shortest path in the graph, we use A^* [28]. A^* is similar to Dijkstra’s algorithm, but prioritizes paths in the graph likely to be the optimal path based on a heuristic, which reduces the number of nodes in the graph that are processed in the search. We select a standard heuristic that uses the L2 distance to the point being queried.

4 TRACING TOOL DESIGN

Tracing neurons in complex, noisy, and poorly imaged regions often encountered in real-world data is not a straightforward task. When tracing, the neuroscientist must make careful decisions to determine connectivity in such regions. Similarly, determining whether or not a neuron is branching is difficult in such regions, and those where many neurons intersect or pass in close proximity to each other. The design of our tool is motivated by two primary goals: to aid the user in better understanding the data to facilitate the decision making process, and to increase the efficiency and ease of tracing neurons.

Drawing from common best practices for design studies [58], we worked closely with domain scientists, our co-authors, to develop the tool through multiple iterations of testing and feedback. Starting from an existing tool for manual neuron tracing [67] (Section 4.1), we

integrated a prototype of our semi-automatic tracing method, which we refined based on feedback from domain scientists (Section 4.2).

4.1 Virtual Reality Neuron Tracing Framework

Our VR software framework is built using OpenVR, supporting room scale and seated modes using the variety of VR headsets available today, and is implemented in C++ and OpenGL. In this work, we use an HTC Vive Pro. The framework supports intuitive interaction modes for navigating volumetric data and manual neuron tracing, along with streaming and rendering large volume data through the IDX format [46] and an integrated data caching system. We provide a brief overview of this framework, which served as the starting point for implementing our MSC-guided tracing method, in the following two sections.

4.1.1 Tracing and Navigation

Tracing neurons and navigating the data are the two primary 3D interactions performed when working on a neuron reconstruction, and thus must be intuitive and quick to perform. In the VR framework, one of the Vive controllers is mapped to tracing and the other to navigation. Both interactions are initiated by holding the trigger on the respective controller and moving it, to either trace along a neuron, or directly grab the volume and translate it. Releasing the trigger then ends the interaction, either stopping the trace or releasing the volume. The user can also navigate by walking in the virtual space. To trace a neuron, the user moves the controller through space following the structure as a line is drawn from the tip of the controller. Navigating with the controller is used to either stream new data from disk when working with large volumes, or to reduce the amount of motion the user must perform, e.g., if working in a chair.

Capturing the branching structure of a neuron is critical to reconstructing its connectivity, for use in subsequent analysis tasks. Thus, it is important to support an intuitive workflow for tracing the branches of a neuron. The user can start a trace off any point on an existing one to begin a new branch, or follow a branch back to its parent trace to connect it. If mistakes are made during the tracing process, a quick undo operation can be performed by pressing the trackpad. Corrections can be made when reviewing a trace by selecting portions of the trace with the controller and deleting them by pressing the trackpad. The user can then re-trace the removed section to correct it. The tool uses haptic feedback to improve the user's perception of physically selecting points on existing traces when creating branches or making edits by playing a haptic click when hovering a point.

4.1.2 Rendering

As scientists will potentially use the tool for hours on end to perform their work, providing a comfortable experience is critical to avoid motion sickness or discomfort. To meet the high-resolution and frame rate demands of VR, the framework follows best practices from VR game development [68]. The tool renders a 256^3 subregion of the volume to keep the rendering cost within an 11ms time budget to achieve 90FPS, and limits the amount of data paged onto the GPU each frame. The data paging system allows exploration of arbitrary sized volumes using the IDX format [46], combined with a CPU and GPU on-demand data paging system that loads data onto the GPU as needed. The volume is stored in a sparse 3D texture, and is rendered by a standard single-pass GLSL volume ray caster supporting both volume and implicit isosurface raycasting (e.g., [65]). Traces are rendered as thick lines using OpenGL line primitives. To further reduce the number of pixels (and thus rays) that must be shaded each frame, the renderer uses the `NV_clip_space_w_scaling` extension (available on NVIDIA GTX 10 series and newer GPUs) to reduce the rendering resolution at the edges of the eye, approximating a foveated rendering approach. As users are typically focused at the center of the image on the data being analyzed, the reduced resolution around the edges of the image is not disruptive.

4.2 Semi-Automatic Tracing in Virtual Reality

A key focus in the design of our tool is on enabling users to work with the MSC-graph in an intuitive manner, as neuroscientists are unlikely

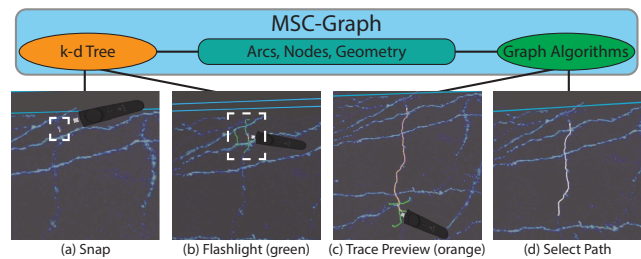


Fig. 7: The MSC-graph is leveraged to aid tracing in our VR system. A trace interaction proceeds as follows: (a) When interacting, user selections are snapped to the center of the selected ridge. (b) The flashlight is used to preview local potential connections for guidance. (c) After picking a start point, the trace selected from the MSC-graph is displayed as a live preview while the user navigates to the end point. (d) Users can then accept the selection and add it to their reconstruction.

to be familiar with the underlying topological framework. Displaying the entire set of arcs computed in the MSC-graph is distracting and overwhelming (Figure 5a), and may lead to following arcs in the MSC-graph that do not correspond to the desired neurons. In our design, we put the data and the neuroscientist's interpretation of it first, and display the arcs on-demand as the user hovers the tracing controller over regions of the data. Through design feedback from the domain experts, we developed a "flashlight" that can be used to preview the potential neuron traces in a region, leveraging the query system described in Section 3.2 (Section 4.2.1). When tracing, we leverage the same query system to compute and display a preview on the fly to the user, integrating the tracing and proofediting tasks into a single efficient workflow (Section 4.2.2). To allow quickly tracing through larger regions of the data we also integrated a zoom feature, which switches to a lower resolution representation of the data to avoid frame drops. An overview of how the MSC-graph is leveraged in our tool to provide guidance to users is shown in Figure 7.

4.2.1 Augmenting Visualization to Aid Users

Our initial prototype only displayed the user's current selection (e.g., Figure 7c) when tracing. Although this is valuable to display the selection so far, it does not provide guidance on where to go next. To assist the user in choosing where to continue the trace, especially at difficult decision points, we provide an additional visual aid, the "flashlight" (Figure 7b).

The flashlight displays a preview of the nearby arcs in the MSC-graph as green lines. The arcs passing within a small ball around the controller are queried from the *k-d* tree and line fragments outside the ball are discarded to avoid clutter. In noisy or poorly imaged regions and ambiguous crossings the flashlight allows users to peek into the underlying MSC-graph to gain additional information about the potential connections to aid the decision making process. For example, in regions with imaging gaps, the flashlight can be used to check if the underlying ridge line continues across the gap or not, and if it reconnects later to a region with better imaging quality. Similarly, when deciding on crossings or branchings in regions with many nearby neurons the flashlight previews the set of possible paths which are most likely given the topology of the data. By providing additional information that is less reliant on imaging quality and visual representation, the flashlight is able to supplement the neuroscientist's domain knowledge to make more informed decisions when tracing.

4.2.2 Tracing

When using the MSC-guided tracing tool, both the flashlight and the closest arc in the MSC-graph are highlighted, giving a small live preview of the arcs in the underlying complex. The closest arc is shown as an orange line (Figure 7a,b), displaying the arc that will be selected when starting a trace. To begin an MSC-guided trace, the user places the controller next to the arc they want to start at and presses the trigger to pick a starting point within the arc, "snapping" the selection to the arc. As they move the controller along the neuron being traced, we recompute the candidate trace using the fast weighted shortest path

algorithm on the MSC-graph discussed previously to update the preview in real time. To end the trace and add the displayed preview to the reconstruction, the user presses the trigger at the desired end point (Figure 7c,d). The live preview allows users to check that the selected path accurately follows the desired neuron before selecting the trace.

When tracing manually, the user must hold the trigger for the duration of the trace, and trace each neuron segment individually and precisely. However, when using our MSC-guided tool, they simply pick the start point and navigate to the desired end point, while checking that the preview follows the desired structure. The MSC-guided tool snaps the trace to the neuron’s center line, producing an accurate trace without requiring significant physical effort or precision from the user. If the preview jumps away from the desired neuron, either because a shorter path is found or the MSC-graph lacks a needed connection, the user places an end point before the problematic section and continues a new MSC-guided or manual trace off the end point.

To trace the branches of a neuron, the user can choose between a manual or MSC-guided trace to start from some point on the existing tree, or start a new trace and reconnect it to the parent. When connecting a new MSC-guided branch to an existing parent tree, it is possible that no arc exists in the MSC-graph to connect the two. In this case, we create a connection to join the new branch with the nearest point on the parent tree, if the gap is less than a few voxels.

When evaluating our first prototype, the neuroscientists found the MSC-guided tracing mode especially useful for tracing long axons through the data. The ability to simply place the start point and navigate to the end point reduced the amount of physical effort required, allowing them to focus on interpreting the data rather than precise interaction. However, to navigate to the end point users would spend a large amount of time translating the volume to find the end of the axon, limiting the speed at which they could trace.

To allow navigating at larger scales or getting an overview or close up view of the data, we added a zoom interaction, performed by holding the grip buttons on both controllers and moving them further apart or closer together. Zooming out increases the number of voxels visible in the focus region, potentially degrading performance below 90 FPS. To maintain an acceptable framerate, we leverage the support for multiresolution queries provided by the IDX format [46]. In addition to the full resolution volume cache, we run a second lower resolution one in parallel. When zoomed out, we switch to the lower resolution data to reduce rendering load; when the user returns to the original zoom level we switch back to the full resolution data. The caches are run simultaneously to ensure that the required data is available to provide a smooth transition between levels when zooming in or out.

5 EVALUATION

To evaluate the effectiveness of our MSC-guided semi-automatic neuron tracing method we study both the effectiveness of the underlying topological framework, and the design of our semi-automatic tracing tool in VR. First, to demonstrate that the MSC-graph provides an effective framework for neuron tracing we perform an offline comparison against semi-automatic methods available in current desktop software, solely comparing the path-finding capabilities of each approach (Section 5.2). We then evaluate our MSC-guided tracing tool in virtual reality through a pilot study with trained neuroanatomists and undergraduate students (Section 5.3).

5.1 Experimental Setup

We evaluate our approach using two data sets. The first data set is the *Neocortical Layer 1 Axons* data set [45], made publicly available for the DIADEM challenge [14]. The data set is a $1464 \times 1033 \times 76$ volume made from six aligned subvolumes containing 34 axons imaged from a mouse brain. The resolution of the data is $\approx 0.08 \mu\text{m}/\text{pixel}$ in X and Y and $\approx 1 \mu\text{m}/\text{pixel}$ along Z. The data set includes a reference trace for each neuron, which we use as one point of comparison in our evaluation. The reference traces were used in the DIADEM challenge and produced manually using NeuroLucida. Throughout the text we will refer to these traces as the “DIADEM reference traces”. We also compare against traces created manually by an expert in VR during




Fig. 8: The DIADEM reference (blue) was made using standard desktop software, and consists of coarse line segments. The VR reference (orange) consists of finer segments, and follows the neuron more closely.



Fig. 9: Examples of incorrect traces (b,c) compared to a reference trace (a). Traces (b) and (c) score 0.875 and 0.5 with the DIADEM metric, respectively. The error in (c) misses a subtree containing more important features than are missed in (b), impacting subsequent analysis significantly, and is scored lower as a result.

a previous study [67], referred to as the “VR reference traces”. As neuron traces are produced by hand by experts, there is an inherent subjectivity in each trace informed by the expert’s knowledge of the data and imaging process, and as such no true “ground truth” is available to compare against. During a review of the DIADEM traces we observed that they would often drift from the neuron, following a straighter path than the underlying data, while the VR reference traces followed the structure more closely (Figure 8).

The second data set, *Cell Bodies*, was provided by A. A.’s laboratory [18] and is a $1024 \times 1024 \times 314$ volume with a resolution of $0.331 \mu\text{m}/\text{pixel}$ in X and Y and $1.5 \mu\text{m}/\text{pixel}$ in Z. The volume was imaged from a Marmoset visual cortex and contains multiple cell bodies with axons that have a complex branching structure. The data set contains significantly more noise, overlapping neurons, and poorly imaged regions than the *Neocortical Layer 1 Axons* data set, making it more difficult to trace. The data set does not include a set of reference traces, in our comparisons we used the most experienced tracer’s (subject 1) manual trace as the reference.

Scoring Trace Quality. We use the DIADEM metric [21] to compare traces against the references and expert judgment to assess their accuracy and quality. The DIADEM metric accounts for the length and topology of a trace, scoring how well a trace captures the branching structure of a neuron on a scale of 0 (dissimilar) to 1 (identical). The score is penalized for missing branches, excess branches, incorrectly placed branches, and differences in branch length (see Figure 9). Branches and endpoints are matched to the reference by checking within a fixed radius of the reference points. The DIADEM score correlates reasonably well with expert judgment of trace quality and remains widely used in the field [36]. However, the score focuses on the topology of the trace and does not account for geometric differences. Furthermore, the traces produced in our tool are more finely sampled along the trace, which we have found to pose difficulties for the DIADEM scoring method as the length of the trace will be longer than a similar trace composed of coarser line segments (see Figure 8). Developing metrics for robustly comparing reconstructions of the same neuron remains an open problem [37,47]. Thus, we also evaluate traces using expert judgment, and have one of the expert neuroanatomists rate the quality of an anonymized subset of traces. The expert rates the trace on how well it follows the center-line of a neuron and how accurately branches and end points have been traced.

Table 1: DIADEM scores for traces extracted with Vaa3D and our MSC-guided method compared to the reference traces. We find that the MSC-guided method computes better and more consistent traces.

| Reference Trace | Vaa3D | MSC-guided |
|-----------------|-------------|-------------|
| DIADEM | 0.45 ± 0.37 | 0.74 ± 0.26 |
| VR | 0.53 ± 0.38 | 0.81 ± 0.22 |

Table 2: DIADEM scores (average and standard deviation) comparing users' manual and MSC-guided traces to DIADEM reference traces (top), and an expert's manual trace in VR (bottom).

| User | Manual Score | MSC-guided Score |
|-------------------------------------|--------------|------------------|
| Compared to DIADEM Reference Traces | | |
| 1 | 0.51 ± 0.35 | 0.43 ± 0.36 |
| 2 | 0.45 ± 0.37 | 0.35 ± 0.35 |
| 3 | 0.35 ± 0.35 | 0.25 ± 0.31 |
| 4 | 0.26 ± 0.35 | 0.23 ± 0.28 |
| 5 | 0.42 ± 0.34 | 0.32 ± 0.31 |
| Compared to Expert VR Traces | | |
| 1 | 0.46 ± 0.46 | 0.40 ± 0.41 |
| 2 | 0.56 ± 0.45 | 0.28 ± 0.39 |
| 3 | 0.36 ± 0.42 | 0.24 ± 0.38 |
| 4 | 0.34 ± 0.42 | 0.30 ± 0.37 |
| 5 | 0.44 ± 0.41 | 0.33 ± 0.38 |

5.2 Algorithm Comparison

We first perform an offline study to assess the quality of the neurons computed by our MSC-graph based algorithm compared to Vaa3D's semi-automatic tracing method [51, 52]. Vaa3D is a widely used open-source software suite for neuron reconstruction, and is the designated platform for testing algorithms in the Big Neuron Project [47]. Vaa3D's semi-automatic tracing works similar to our MSC-guided tool: given a start and end point it will attempt to trace the neuron between them.

The guide points are generated from the existing DIADEM and expert VR traces by extracting the start, branch, and end points to simulate user clicks. The points are extracted automatically through a depth-first traversal of the neuron tree, and supplied in that order. This approximates how a user would interact with a semi-automatic method, clicking along the structure to mark key points on the neuron and letting the algorithm extract the structure. Our MSC-guided method works between a single start and end point, thus both start and end points are provided for each segment of the neuron. Along with point to point extraction, Vaa3D can trace from a start point to connect to a set of points, which is the mode we use in our evaluation. In both cases, the paths between the points are automatically generated.

We found that the traces computed by our MSC-guided method follow the neuron more accurately and achieved a higher DIADEM score than those computed using Vaa3D's semi-automatic method (Table 1). This can be partly attributed to the MSC-guided method placing greater weight on the user's guidance, extracting the shortest path between the given end points as they are clicked in order. Vaa3D's semi-automatic method treats the set of points as hints and is not guaranteed to connect them in the same order.

Our MSC-guided approach is able to provide faster neuron extraction times, as required for VR. Vaa3D took 28.64s per reconstruction on average, while our MSC-guided method took just 0.029s, a speedup of 986×. Our MSC-guided method computes shortest paths on the relatively sparse MSC-graph (~10k nodes) compared to the implied adjacency graph of the original voxel data (~100M voxels). The pre-computation to build the MSC-graph is fast and scalable [22, 23]; on a laptop with an i7-7700HQ CPU the image filtering and blurring takes 26s using ImageJ, after which the MSC-graph is computed in 134s using MSCEER [22]. Although the total time is longer for a single neuron, the MSC-graph can be reused for all neurons in the volume, allowing real-time path computation in VR.





Fig. 10: An MSC-guided trace (green) has lower DIADEM score than a manual (blue) VR trace (0.529 vs. 0.833, respectively), when compared against the corresponding DIADEM (red) reference trace, even though the manual trace contains an extra branch. Both the MSC-guided and manual traces were rated acceptable by an expert. It is possible that differences in the length of segments, or failure to match branch points between MSC-guided and DIADEM traces contribute to a low score.



(a) Manual traces. (b) MSC-guided traces.

Fig. 11: A comparison of the same trace from all users. The left shows how manual traces can be inconsistent, while the right shows the high consistency using the MSC-guided method.

5.3 User Study

We tested our MSC-guided software infrastructure through a pilot study, where neuroscientists performed real traces on a varied set of tasks. We evaluate our tool using quantitative measures of trace quality and speed, and extensive qualitative feedback from the user group on the overall impact on workflow and user experience.

User group. We conducted our pilot study with five users of varying levels of experience from A. A.'s laboratory. Subjects 1 and 5 are senior neuroanatomists, subjects 2 and 3 are undergraduates with 2–3 years of experience tracing neurons using NeuroLucida, and subject 4 is a graduate student with little prior tracing experience but with advanced knowledge of neuron morphology. This range of experience levels provides a representative sample of a typical connectomics lab.

Description of tasks. Each subject traces a set of 34 neurons twice over two separate sessions, spaced at least three days apart to ensure users did not learn specific traces from a previous session. The first 28 neurons come from the *Neocortical Layer 1 Axons* data set, the last six are from the *Cell Bodies* data set. For each neuron the start point is marked in space and the user instructed to trace the neuron to its perceived end points. In each session the users traced the first half of the neurons manually and the second half using any desired combination of the MSC-guided tool and manual tracing. In the second session the set of neurons which were traced manually and with the MSC tool is flipped. At the start of their first session, each user completed a training session to learn the manual and MSC-guided tools.

User training. The goal of the training was to reduce the effect of learning on evaluating the accuracy and speed of manual and MSC-guided tracing. The training process acquainted users with the VR system, and demonstrated the capabilities and limitations of our MSC-guided tracing, and how to address them. The user was given four starting points in one of the test data sets and asked to trace the neurons to completion. When they finished tracing a neuron, they were shown their trace compared to a reference. To emphasize how their trace may have deviated from the reference, their trace was colormapped by the shortest distance to a corresponding point on the reference. By the end of the training, users demonstrated an intuitive feel for how the MSC-guided tool operated. In particular, the senior neuroanatomist

Table 3: Speedups for total task time of the MSC-tool vs. manual tracing. Time is in seconds, distance is in voxels, and rate is the distance traced over the time, the total across each task set is reported. Overall, the improvement in tracing times was higher using the MSC-guided method.

| User | Session 1 | | | | | | Session 2 | | | | | | Total | | |
|------|---------------|-------|-------|-------------------|-------|-------|---------------|-------|-------|-------------------|-------|-------|--------|------------|---------------|
| | Manual: Set A | | | MSC_guided: Set B | | | Manual: Set B | | | MSC_guided: Set A | | | Manual | MSC_guided | |
| | Time | Dist | Rate | Time | Dist | Rate | Time | Dist | Rate | Time | Dist | Rate | Rate | Rate | Speedup |
| 1 | 3275 | 34453 | 10.5 | 2532 | 23809 | 9.40 | 2337 | 23963 | 10.25 | 2401 | 29804 | 12.41 | 11.06 | 10.48 | 0.94 × |
| 2 | 1403 | 19911 | 14.19 | 982 | 19751 | 20.11 | 1267 | 18883 | 14.90 | 1251 | 25257 | 20.18 | 14.12 | 19.91 | 1.41 × |
| 3 | 1979 | 39555 | 19.98 | 1525 | 41329 | 27.10 | 2369 | 38820 | 16.38 | 1957 | 35237 | 18.00 | 17.63 | 22.5 | 1.27 × |
| 4 | 1823 | 9844 | 5.39 | 2142 | 17391 | 8.11 | 2345 | 11539 | 4.92 | 1524 | 10397 | 6.82 | 5.14 | 7.57 | 1.47 × |
| 5 | 1334 | 17443 | 13.07 | 1350 | 37825 | 28.01 | 2366 | 20532 | 8.67 | 1406 | 41906 | 29.80 | 9.95 | 27.3 | 2.74 × |




Fig. 12: Users tracing manually (pink) can significantly deviate from the neuron. The MSC-guided trace (orange), automatically snaps to the center of the neuron, producing a more accurate trace. These traces were produced by subject 3.

recommended the training process not only for familiarizing users with the tool, but also to teach novices the reasoning behind selecting paths, branch points, and start and end points.

Interactive session. Each of the sessions was held in a room with enough space to walk around and explore the data, using a 2.5m×2m VR area, under supervision of the authors. While the users completed each task, their interactions, i.e., tracing, panning, zooming, head and hand movement, and controller button clicks, were recorded by the tool and exported to a JSON file for analysis. On average, each session took one to one-and-a-half hours. Users were able to walk freely or sit as desired for any portion of the session.

5.3.1 Quantitative Evaluation of Traces

Tracing neurons is a challenging task and essentially relies on the expertise and accuracy of the user. This makes the task subjective, and can lead to disagreement on so-called “reference traces” produced by other experts have produced. Hence in this scenario, we have to consider that scoring the user traces against a fixed reference does not necessarily mean the reference is more accurate.

When using the DIADEM metric, we find that the quality of the traces produced using fully manual tracing and our MSC-guided semi-automatic tracing are similar, with traces made using the MSC tool receiving slightly lower DIADEM scores (Table 2). Note that *even compared to a reference manual VR trace*, DIADEM scores for traces also completed manually in VR were on average below 0.4. This indicates that either the DIADEM metric is very sensitive, or that the tracing task is so challenging that even experts disagree and produce different traces. Finally, the DIADEM score, while a reasonable metric, was found to produce unexpectedly low scores on traces rated by our two experts as acceptable.

To perform an additional validation of the traces we asked one of the senior neuroanatomists in A.A.’s lab to review a subset of them. The expert reviewed a series of traces, each containing three anonymized traces of the same neuron: one traced using the manual method, one using the semi-automatic method, and one being the DIADEM reference. The expert noted that, due to the ambiguity in the data, most of the traces were potentially correct. However, for almost every set reviewed, the expert rated the MSC-guided traces as more accurate than those performed manually (see Figure 10).

In reviewing the user traces, we found a higher degree of geometric inconsistencies in the manual traces compared to the MSC-guided

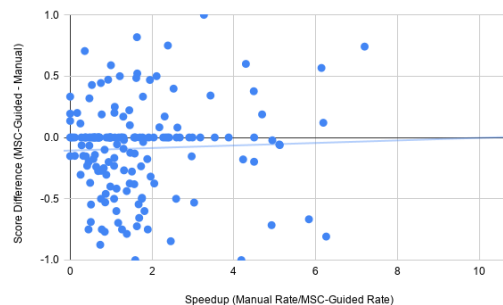


Fig. 13: The correlation between score difference (MSC-guided score minus manual score), and the speedup (Manual Rate/MSC-guided rate) of all individual tracing tasks. There is no correlation between the score difference and the speedup.

traces (see Figure 11). Figure 11a shows how the hand motions of tracing neurons manually can significantly vary, while MSC-guided tracing (Figure 11b) snaps the traces to the centerline of the neuron. The MSC-guided traces have more consistent and higher geometric accuracy as a result. Figure 12 highlights how geometrically inaccurate manual traces can be. If a user is not precise in their hand movements, they can produce traces that clearly deviate from the intended path. We found this to frequently occur in places where the neurons bend either slightly or significantly.

Table 3 shows the total session time, distance traced, and the rate of tracing in voxels per second for each set of tasks. We found that for the majority of users, the MSC-guided approach reduced the time it took to complete each tracing task compared to manual tracing. We explored whether there was a correlation between the speedup of a tracing task and the score of that task (Figure 13). The Pearson correlation coefficient between the two was 0.044, indicating no association between speedup and score, i.e., equivalent quality traces were produced in less time. The users did not show consistent speedup or slowdown between the first and second sessions, indicating that memory of traces that have already been seen in the first session did not impact tracing in the second session. We note that user 1 was slower when using the MSC-guided method, as he spent additional time exploring the data with the flashlight to find connections in sparse regions below the visible range of the volume or isosurface visualization.

We remark on the significant variability in the total distance traced between users, both within and between manual and MSC-guided tasks. Empirically, this variability highlights the subjectivity in neuron tracing, where some users decide that neurons in certain complex or poorly imaged areas either branch or continue, while others decide they terminate. We did not find a consistent increase or decrease in the total distance of neurons traced between MSC-guided and manual methods.

The median speedup across all users was 1.41×, with subject 5, an expert in their field, achieving the highest speedup of 2.74×. Neuron tracing is a time consuming process, with typical real world traces taking days, weeks, or months. Reducing the time it takes to trace neurons by these factors could save hours or days of work, significantly accelerating the neuron reconstruction pipeline.

5.3.2 Qualitative Expert Evaluation and Feedback

In this section we detail our users’ qualitative feedback regarding the design and usability of the MSC-guided tracing tool, discussing both benefits and limitations. During the study we collected feedback from

users through a survey completed after each session and open-ended discussions. The survey focused on the usability and usefulness of the MSC-guided tool for tracing neurons, with questions rated on a 5-point Likert scale. Open-ended discussions were used to solicit general feedback on the design of the tool and general comments or issues regarding its use in practice.

User Experience. Overall feedback from users on the MSC-guided tool was positive. In the survey all subjects reported preferring the MSC tool over manual tracing, finding it less fatiguing and more comfortable to use for long periods. In the sessions when the MSC-guided tool was made available to them, 84% of the total tracing time was spent using the tool on average across all users. The ability to quickly switch between the MSC-guided and manual tools was reported to be valuable when resolving topological issues.

Four of five subjects reported the MSC-guided tool was more comfortable to use and allowed them to focus more on the data, as the tool requires less precise physical interaction than manual tracing. Subject 4 did not report a significant perceived difference in terms of interaction effort. All users found the MSC tool's live preview trace valuable to review the trace that would be selected before selecting it, allowing for proofediting the trace on the fly during tracing. The interactive aspect of the preview was reported to be particularly valuable when visualizing possible paths at intersecting neurons.

Four of five subjects reported that the flashlight feature was valuable for navigating the data, subject 2 did not find feature useful. Subjects 1 and 5, both experts in their field, commented that the local preview provided by the flashlight was especially helpful when tracing in poorly imaged regions and determining the endpoints of branches. The zoom functionality was also found to be useful when navigating through the data, with users zoomed out for the majority of time spent navigating.

Manual vs. MSC-guided Tracing. In our initial evaluation, the majority of users found the MSC-tool more challenging to learn. A frequent comment we received was that it took a few traces to become comfortable with the tool and to learn the cases where the tool would follow the desired path and where it would not. As a result of these comments, we extended our training process to the one described in Section 5.3 to help users familiarize themselves with the new technique. With this new training in place, we no longer received comments from users about the tool being difficult to use. Although some users still reported becoming more comfortable with the tool over time, the difference in tracing performance was not as significant.

Users reported the MSC-guided tool to be especially useful when tracing long axons through large portions of the volume. When using the MSC tool they would let it follow the neuron for them, and focus on navigating to the end point of the axon to finish the trace. The ability to zoom out and cover larger portions of the volume when tracing with the MSC-tool accelerated this process significantly over manual tracing. When tracing manually this task is more difficult, as users must frequently swap between tracing and navigating to create an accurate trace. During the manual portion of the second session, many subjects lamented being unable to use the MSC-guided method in these sections.

There are a number of regions in the data where a neuron may appear to end or fade due to issues with the tissue labeling or imaging process. All subjects reported that the MSC-guided tool was particularly helpful in resolving these portions of the data. One of the senior neuroanatomists, subject 1, reported that using the MSC-guided tool helped him analyze these cases more carefully. On one trace this led to him ultimately determining that a neuron did not end where he initially thought it did. Subjects 4 and 5 made similar comments, noting that the MSC-guided tool and the flashlight preview helped them make decisions at potential branch and termination points.

Finally, during a visual inspection of the traces we found that in many cases the MSC-guided trace followed the neuron more closely than the DIADEM and VR reference traces (e.g., Figures 10 to 12). The DIADEM trace clearly shows the limitations of tracing manually on the desktop, where users click and place points to construct straight line segments between them. The manual traces performed in VR, while producing a more refined line, show the difficulties of maintaining the precise hand motion required to follow the neuron center accurately.

The MSC-guided tool alleviates the need for such precision by automatically following the ridge line of the neuron, requiring the user to only provide a coarse set of inputs.

Comparison to Desktop Software. Expert 5 has extensive experience using Imaris [1] in his usual workflow. Imaris is a proprietary software package for microscopy image analysis that provides a semi-automatic tracing feature with similar functionality. When asked to compare the two methods, he reported that our MSC-guided tracing algorithm performed similarly to Imaris. Moreover, he reported that the flashlight feature, which Imaris does not have, was valuable in resolving low resolution and poorly imaged regions of the data. He also noted that the 3D navigation and interaction capabilities used in VR made the tracing experience easier and more intuitive, compared to the large number of view point adjustments that must be made when tracing on a desktop.

6 SUMMARY AND FUTURE WORK

We have presented a novel semi-automatic neuron tracing method based on the topological features extracted from the Morse-Smale complex. We implemented our MSC-guided tracing tool within an existing VR environment, and demonstrated that it improves neuron tracing performance over manual tracing in VR and semi-automatic methods on a desktop. When using our MSC-guided tracing tool, experts were able to produce acceptable quality traces with less fatigue and in far less time. By leveraging the fast online computation time of our method, we are able to show a live preview of the trace to the user, removing the need for extensive post-process proofediting. The neuroanatomists' qualitative feedback indicates that, although more work remains to be done, our MSC-guided tool is a promising approach to accelerate neuron tracing, especially when tracing long range connections or in poorly imaged regions and ambiguous regions.

Through iterating on the system design and training process with experts, we have improved both the usability and interpretability of our MSC-guided tool. The flashlight and preview features work together to guide the neuroanatomist through the data and aid their decision making process. At the same time, the zoom feature enables them to trace through the data faster. Finally, our improved training process helps users get up to speed with the system faster to use the tool effectively when tracing.

A major challenge faced in our work were the structural and geometric differences in traces considered "reference". The MSC-guided tracing method significantly reduces the geometric variation between traces, enabling the design of better tools for evaluating where expert traces disagree. The high quality and consistent traces produced could be used to build a training set for labeling portions of an MSC-graph or image volume in a machine learning approach to advance automated neuron computation. Although the results of our pilot study are promising, we have also found areas for improvement. Users expressed interest in seeing additional information beyond the flashlight, which may make the tool more informative for new and experienced users. In addition, the trace previews and flashlight could be augmented by using a certainty measure to colormap the preview.

ACKNOWLEDGEMENTS

The authors would like to thank Pavol Klacansky for assistance in the initial integration of the tool. This work was funded in part by NSF OAC awards 1842042, 1941085, NSF CMMI awards 1629660, LLNL LDRD project SI-20-001. This material is based in part upon work supported by the Department of Energy (DoE), National Nuclear Security Administration (NNSA), under award DE-NA0002375. This research was supported in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the DoE Office of Science and the NNSA. This work was performed in part under the auspices of the DoE by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. This work is supported by in part by grants from the NIH (R01 EY026812, R01 EY019743, BRAIN U01 NS099702), the NSF (IOS 1755431, EAGER 1649923), and The University of Utah Neuroscience Initiative, to A.A., a grant from Research to Prevent Blindness, Inc. and a core grant from the NIH (EY014800) to the Department of Ophthalmology, University of Utah.

REFERENCES

- [1] Imaris. <https://imaris.oxinst.com/>.
- [2] syGlass. <https://www.syglass.io/>.
- [3] L. Acciai, P. Soda, and G. Iannello. Automated Neuron Tracing Methods: An Updated Account. *Neuroinformatics*, 2016.
- [4] M. Agus, D. Boges, N. Gagnon, P. J. Magistretti, M. Hadwiger, and C. Cali. GLAM: Glycogen-derived Lactate Absorption Map for visual analysis of dense and sparse surface reconstructions of rodent brain structures on desktop systems and virtual environments. *Computers & Graphics*, 2018.
- [5] N. Apthorpe, A. Riordan, R. Aguilar, J. Homann, Y. Gu, D. Tank, and H. S. Seung. Automatic neuron detection in calcium imaging data using convolutional networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 29, 2016.
- [6] H. Bhatia, A. Gyulassy, V. Lordi, J. E. Pask, V. Pascucci, and P.-T. Bremer. TopoMS: Comprehensive topological exploration for molecular and condensed-matter systems. *Journal of Computational Chemistry*, 2018.
- [7] A. Bock, H. Doraiswamy, A. Summers, and C. Silva. TopoAngler: Interactive Topology-Based Extraction of Fishes. *IEEE Transactions on Visualization and Computer Graphics*, 2018.
- [8] D. D. Bock, W.-C. A. Lee, A. M. Kerlin, M. L. Andermann, G. Hood, A. W. Wetzel, S. Yurgenson, E. R. Soucy, H. S. Kim, and R. C. Reid. Network anatomy and in vivo physiology of visual cortical neurons. *Nature*, 2011.
- [9] D. Boges, M. Agus, R. Sicut, P. J. Magistretti, M. Hadwiger, and C. Cali. Virtual reality framework for editing and exploring medial axis representations of nanometric scale neural structures. *Computers & Graphics*, 2020.
- [10] D. Boges, C. Cali, P. J. Magistretti, M. Hadwiger, R. Sicut, and M. Agus. Immersive Environment for Creating, Proofreading, and Exploring Skeletons of Nanometric Scale Neural Structures. *Smart Tools and Apps for Graphics - Eurographics Italian Chapter Conference*, 2019.
- [11] S. Boorboor, S. Jadhav, M. Ananth, D. Talmage, L. Role, and A. Kaufman. Visualization of Neuronal Structures in Wide-Field Microscopy Brain Images. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [12] P.-T. Bremer, G. Weber, V. Pascucci, M. Day, and J. Bell. Analyzing and Tracking Burning Structures in Lean Premixed Hydrogen Flames. *IEEE Transactions on Visualization and Computer Graphics*, 2010.
- [13] K. L. Briggman, M. Helmstaedter, and W. Denk. Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, 2011.
- [14] K. M. Brown, G. Barrionuevo, A. J. Canty, V. De Paola, J. A. Hirsch, G. S. Jefferis, J. Lu, M. Snippe, I. Sugihara, and G. A. Ascoli. The DIADEM data sets: Representative light microscopy images of neuronal morphology to advance automation of digital reconstructions. *Neuroinformatics*, 2011.
- [15] K. Chung, J. Wallace, S.-Y. Kim, S. Kalyanasundaram, A. S. Andalman, T. J. Davidson, J. J. Mirzabekov, K. A. Zalocusky, J. Mattis, A. K. Denisin, et al. Structural and molecular interrogation of intact biological systems. *Nature*, 2013.
- [16] D. Eberly, R. Gardner, B. Morse, S. Pizer, and C. Scharlach. Ridges for image analysis. *J. Math. Imaging Vis.*, 1994.
- [17] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, 2000.
- [18] F. Federer. Neurons in Marmoset Visual Cortex. <https://klacansky.com/open-scivis-datasets/>.
- [19] R. Forman. A user's guide to discrete morse theory. *Se'minaire Lotharingien de Combinatoire*, 2002.
- [20] W. Fulmer, T. Mahmood, Z. Li, S. Zhang, J. Huang, and A. Lu. ImWeb: Cross-platform immersive web browsing for online 3D neuron database exploration. In *Proceedings of the 24th International Conference on Intelligent User Interfaces - IUI '19*, 2019.
- [21] T. A. Gillette, K. M. Brown, and G. A. Ascoli. The DIADEM Metric: Comparing Multiple Reconstructions of the Same Neuron. *Neuroinformatics*, 2011.
- [22] A. Gyulassy. MSCEER: Morse-Smale Complex Extraction, Exploration, Reasoning. <https://github.com/sci-vi-sus/MSCEER>, 2018.
- [23] A. Gyulassy, P.-T. Bremer, and V. Pascucci. Shared-Memory Parallel Computation of Morse-Smale Complexes with Improved Accuracy. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [24] A. Gyulassy, M. Duchaineau, V. Natarajan, V. Pascucci, E. Bringa, A. Higginbotham, and B. Hamann. Topologically clean distance fields. *IEEE Transactions on Computer Graphics and Visualization*, 2007.
- [25] A. Gyulassy, A. Knoll, K. C. Lau, B. Wang, P.-T. Bremer, M. E. Papka, L. A. Curtiss, and V. Pascucci. Interstitial and Interlayer Ion Diffusion Geometry Extraction in Graphitic Nanosphere Battery Materials. *IEEE Transactions on Visualization and Computer Graphics*, 2016.
- [26] A. Gyulassy, N. Kotava, M. Kim, C. D. Hansen, H. Hagen, and V. Pascucci. Direct Feature Visualization using Morse-Smale Complexes. *IEEE Transactions on Visualization and Computer Graphics*, 2012.
- [27] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. A topological approach to simplification of three-dimensional scalar functions. *IEEE Transactions on Visualization and Computer Graphics*, 2006.
- [28] P. E. Hart, N. J. Nilsson, and B. Raphael. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 1968.
- [29] W.-K. Jeong, J. Beyer, M. Hadwiger, A. Vazquez, H. Pfister, and R. T. Whitaker. Scalable and interactive segmentation and visualization of neural processes in em datasets. *IEEE Transactions on Visualization and Computer Graphics*, 2009.
- [30] B. Laha, D. A. Bowman, and J. J. Socha. Effects of VR system fidelity on analyzing isosurface visualization of volume datasets. *IEEE Transactions on Visualization and Computer Graphics*, 2014.
- [31] B. Laha, K. Sensharma, J. D. Schiffbauer, and D. A. Bowman. Effects of immersion on visual analysis of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 2012.
- [32] D. Laney, P.-T. Bremer, A. Mascarenhas, P. Miller, and V. Pascucci. Understanding the structure of the turbulent mixing layer in hydrodynamic instabilities. *IEEE Transactions Visualization and Computer Graphics*, 2006.
- [33] T.-C. Lee, R. L. Kashyap, and C.-N. Chu. Building skeleton models via 3-D medial surface/axis thinning algorithms. *CVGIP: Graph. Models Image Process.*, 56(6):462–478, Nov. 1994.
- [34] Y. Liu. The DIADEM and Beyond. *Neuroinformatics*, 2011.
- [35] L. Luo, E. M. Callaway, and K. Svoboda. Genetic Dissection of Neural Circuits. *Neuron*, 2008.
- [36] C. Magliaro, A. L. Callara, N. Vanello, and A. Ahluwalia. Gotta Trace 'em All: A Mini-Review on Tools and Procedures for Segmenting Single Neurons Toward Deciphering the Structural Connectome. *Frontiers in Bioengineering and Biotechnology*, 2019.
- [37] D. Mayerich, C. Bjornsson, J. Taylor, and B. Roysam. NetMets: Software for quantifying and visualizing errors in biological network segmentation. *BMC bioinformatics*, 2012.
- [38] MBF Bioscience. NeuroLucida 360. <https://www.mbfbioscience.com/neuroLucida360>.
- [39] MBF Bioscience. NeuroLucida version 2019. <https://www.mbfbioscience.com/neuroLucida>.
- [40] E. Meijering. Neuron tracing in perspective. *Cytometry Part A*, 2010.
- [41] E. Meijering, M. Jacob, J.-C. Sarria, P. Steiner, H. Hirling, and M. Unser. Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images. *Cytometry Part A: The Journal of the International Society for Analytical Cytology*, 2004.
- [42] E. Murray, J. Cho, D. Goodwin, T. Ku, J. Swaney, S.-Y. Kim, H. Choi, Y.-G. Park, J.-Y. Park, A. Hubbert, M. McCue, S. Vassallo, N. Bakh, M. Frosch, V. Wedeen, H. Seung, and K. Chung. Simple, Scalable Proteomic Imaging for High-Dimensional Profiling of Intact Systems. *Cell*, 2015.
- [43] D. R. Myatt, T. Hadlington, G. A. Ascoli, and S. J. Nasuto. Neuroman-tic – from Semi-Manual to Semi-Automatic Reconstruction of Neuron Morphology. *Frontiers in Neuroinformatics*, 2012.
- [44] S. W. Oh, J. A. Harris, L. Ng, B. Winslow, N. Cain, S. Mihalas, Q. Wang, C. Lau, L. Kuan, A. M. Henry, M. T. Mortrud, B. Ouellette, T. N. Nguyen, S. A. Sorensen, C. R. Slaughterbeck, W. Wakeman, Y. Li, D. Feng, A. Ho, E. Nicholas, K. E. Hirokawa, P. Bohn, K. M. Joines, H. Peng, M. J. Hawrylycz, J. W. Phillips, J. G. Hohmann, P. Wahnoutka, C. R. Gerfen, C. Koch, A. Bernard, C. Dang, A. R. Jones, and H. Zeng. A mesoscale connectome of the mouse brain. *Nature*, 2014.
- [45] V. D. Paola, A. Holtmaat, G. Knott, S. Song, L. Wilbrecht, P. Caroni, and K. Svoboda. Cell Type-Specific Structural Plasticity of Axonal Branches and Boutons in the Adult Neocortex. *Neuron*, 2006.
- [46] V. Pascucci and R. J. Frank. Global static indexing for real-time exploration of very large regular grids. In *Supercomputing, ACM/IEEE 2001 Conference*, 2001.
- [47] H. Peng, M. Hawrylycz, J. Roskams, S. Hill, N. Spruston, E. Meijering, and G. A. Ascoli. BigNeuron: Large-Scale 3D Neuron Reconstruction from Optical Microscopy Images. *Neuron*, 2015.
- [48] H. Peng, F. Long, and G. Myers. Automatic 3D neuron tracing using

- all-path pruning. *Bioinformatics*, 2011.
- [49] H. Peng, F. Long, T. Zhao, and E. Myers. Proof-editing is the Bottleneck Of 3D Neuron Reconstruction: The Problem and Solutions. *Neuroinformatics*, 2011.
- [50] H. Peng, E. Meijering, and G. A. Ascoli. From DIADEM to BigNeuron. *Neuroinformatics*, 2015.
- [51] H. Peng, Z. Ruan, D. Atasoy, and S. Sternson. Automatic reconstruction of 3D neuron structures using a graph-augmented deformable model. *Bioinformatics*, 2010.
- [52] H. Peng, Z. Ruan, F. Long, J. H. Simpson, and E. W. Myers. V3D enables real-time 3D visualization and quantitative analysis of large-scale biological image data sets. *Nature Biotechnology*, 2010.
- [53] H. Peng, J. Tang, H. Xiao, A. Bria, J. Zhou, V. Butler, Z. Zhou, P. T. Gonzalez-Bellido, S. W. Oh, J. Chen, A. Mitra, R. W. Tsien, H. Zeng, G. A. Ascoli, G. Iannello, M. Hawrylycz, E. Myers, and F. Long. Virtual finger boosts three-dimensional imaging and microsurgery as well as terabyte volume image visualization and analysis. *Nature Communications*, 2014.
- [54] S. Petruzza, A. Gyulassy, S. Leventhal, J. J. Baglino, M. Czabaj, A. D. Spear, and V. Pascucci. High-throughput feature extraction for measuring attributes of deforming open-cell foams. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [55] Prabhat, A. Forsberg, M. Katzourin, K. Wharton, and M. Slater. A Comparative Study of Desktop, Fishtank, and Cave Systems for the Exploration of Volume Rendered Confocal Data Sets. *IEEE Transactions on Visualization and Computer Graphics*, 2008.
- [56] B. Rieck, U. Fugacci, J. Lukasczyk, and H. Leitte. Clique Community Persistence: A Topological Visual Analysis Approach for Complex Networks. *IEEE Transactions on Visualization and Computer Graphics*, 2018.
- [57] V. Robins, P. Wood, and A. Sheppard. Theory and Algorithms for Constructing Discrete Morse Complexes from Grayscale Digital Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011.
- [58] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 2012.
- [59] N. Shivashankar and V. Natarajan. Parallel Computation of 3D Morse-Smale Complexes. *Computer Graphics Forum*, 2012.
- [60] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister. DXR: A Toolkit for Building Immersive Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [61] L. Silvestri, A. Bria, L. Sacconi, G. Iannello, and F. S. Pavone. Confocal light sheet microscopy: Micron-scale neuroanatomy of the entire mouse brain. *Optics Express*, 2012.
- [62] S. Soltanian-Zadeh, K. Sahingur, S. Blau, Y. Gong, and S. Farsiu. Fast and robust active neuron segmentation in two-photon calcium imaging using spatiotemporal deep learning. *Proceedings of the National Academy of Sciences*, 2019.
- [63] T. Sousbie. The persistent cosmic web and its filamentary structure - I. theory and implementation. *Monthly Notices of the Royal Astronomical Society*, 2011.
- [64] C. Steger. An unbiased detector of curvilinear structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1998.
- [65] S. Stegmaier, M. Strengert, T. Klein, and T. Ertl. A Simple and Flexible Volume Rendering Framework for Graphics-Hardware-based Raycasting. In *Volume Graphics, 2005. Fourth International Workshop On*, 2005.
- [66] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, and M. Michaux. The Topology ToolKit. *IEEE Transactions on Visualization and Computer Graphics*, 2018. <https://topology-tool-kit.github.io/>.
- [67] W. Usher, P. Klacansky, F. Federer, P.-T. Bremer, A. Knoll, J. Yarch, A. Angelucci, and V. Pascucci. A Virtual Reality Visualization Tool for Neuron Tracing. *IEEE Transactions on Visualization and Computer Graphics*, 2018.
- [68] A. Vlachos. Advanced VR Rendering. GDC, 2015.
- [69] G. H. Weber, S. E. Dillard, H. Carr, V. Pascucci, and B. Hamann. Topology-controlled volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 2007.
- [70] B. Yang, J. B. Treweek, R. P. Kulkarni, B. E. Deverman, C.-K. Chen, E. Lubeck, S. Shah, L. Cai, and V. Gradinaru. Single-cell phenotyping within transparent intact tissue through whole-body clearing. *Cell*, 2014.
- [71] J. Yang, M. Hao, X. Liu, Z. Wan, N. Zhong, and H. Peng. FMST: An Automatic Neuron Tracing Method Based on Fast Marching and Minimum Spanning Tree. *Neuroinformatics*, 2018.