

Recut: a Concurrent Framework for Sparse Reconstruction of Neuronal Morphology

Karl Marrett, Muye Zhu, Yuze Chi, Chris Choi, Zhe Chen, Hong-Wei Dong, Chang Sin Park, X. William Yang, Jason Cong

Abstract— Interpreting the influx of microscopy and neuroimaging data is bottlenecked by neuronal reconstruction’s long-standing issues in accuracy, automation, and scalability. Rapidly increasing data size is particularly concerning for modern computing infrastructure due to the wall in memory bandwidth which historically has witnessed the slowest rate of technological advancement. Recut is an end to end reconstruction pipeline that takes raw large-volume light microscopy images and yields filtered or tuned automated neuronal reconstructions that require minimal proofreading and no other manual intervention. By leveraging adaptive grids and other methods, Recut also has a unified data representation with up to a 509× reduction in memory footprint resulting in an 89.5× throughput increase and enabling an effective 64× increase in the scale of volumes that can be skeletonized on servers or resource limited devices. Recut also employs coarse and fine-grained parallelism to achieve speedup factors beyond CPU core count in sparse settings when compared to the current fastest reconstruction method. By leveraging the sparsity in light microscopy datasets, this can allow full brains to be processed in-memory, a property which may significantly shift the compute needs of the neuroimaging community. The scale and speed of Recut fundamentally changes the reconstruction process, allowing an interactive yet deeply automated workflow.

Introduction

Morphology, the 3D shape of single neurons, and topology, the arrangement, coverage and connectivity of such cells, may be a contributing facet of neuronal function and pathology. Morphology in particular is a critical component of cell type taxonomy: the process of discriminating and classifying different neuron types by genetic, proteomic, connectivity or pathological characteristics(1). A common method to studying morphology is reconstruction, the process where fluorescently labeled cells in 3D images are segmented and compacted into ball and stick models(2). Recent reviews have named a scalable(3), end-to-end(4) pipeline tool for single-cell reconstruction as the greatest need of the community. The Recut pipeline comprehensively addresses these needs.

Reconstruction establishes a *coverage topology* of the segmented neuronal regions and a set of vertices with connections—the connectome—which is generally described in compacted terms via a graph. This graph $G = (V, E)$ is composed of vertices $v_{i,j,k}$ and edges e , where each vertex has a fixed and unique location in 3D space. The neuroscience community and reconstruction tooling has aligned on the the SWC standard(2), a text file format where each line specifies a vertex in the reconstructed graph with its 3D coordinates, radii and parent vertex. Since vertices only carry a directed edge to their parent, the SWC format is termed a *tree* as opposed to

a general graph. Once the tree of all neurons is constructed, it is partitioned into subtrees that represent individual neurons. With individual neurons established, their morphology is analyzed in aggregate.

Challenges and Related Work.

Accuracy. Dozens of computer-assisted implementations exist for reconstruction as documented by comprehensive literature surveys and reviews(4, 5). Graph-based reconstruction methods are particularly suited in modeling morphology and connectomic data and are deeply embedded in the analysis patterns of the neuroscience domain. However, they are prone to suffer accuracy loss from choice of background threshold value, normalization of the image, and high-frequency, pixel-level noise which induce erroneous path breaks.

NN methods are a natural complement to mitigate the issues of graph methods, due to incorporating larger spatial context cues at various window sizes and resolutions and training. While this removes sensitivity to free parameters such as background threshold, it requires task-specific model selection, a robust training set and still must handle the conversion to the tree-like SWC format for downstream algorithms such as in Figure 1. In a recent 27 method comparison(6), APP2 tied for third place in reconstruction accuracy of mouse neurons. All top mouse algorithms were graph based(7)(8)(9)(10), with the exception of Advantra(11) which uses the Monte Carlo method at the initial stage. The top performing software Neutube(7) as well as many off the shelf reconstruction tools (12) use a fastmarching algorithm similar to APP2. The connected component, SDF and pruning stages of this paper therefore compare to the algorithms in APP2 as baselines for performance accuracy metrics.

Algorithmic Efficiency. Existing graph based reconstruction methods have comparatively low data access and computation. In conventional high-resolution light microscopy methods such as confocal (see details of our imaging pipeline in section 2.7), the *coverage topology* foreground is a tiny proportion of the total image. In such situations, graph approaches can avoid excessive computation since they avoid the dense and redundant access patterns found for example in convolution. Graph methods are therefore a strategic starting point to base further optimization effort upon. We refer to the metric of computation with respect to data size (e.g. image voxel count or n) as *algorithmic efficiency*. We compare algorithms based on their theoretic computation counts expressed in terms of n disregarding constant factors based on the big \mathcal{O} notation(13). This notation is fundamental in understanding and comparing efficiency. For example, in the same comparison study(6),




Fig. 1. Reconstruction progressively compacts the data footprint across the pipeline stages: starting with 3D images, we convert the representation into a sparse graph then further partition into individual graphs for each neuron, then finally aggregate for morphological analysis. With compaction, Recut's throughput is limited by the DRAM bandwidth of a system. Therefore the uncompressed footprint in memory of the active voxel working set (n) is shown on the right at each narrowing. Numbers shown are for a 30x objective lens non-downsampled section of the striatum brain region.

APP2 had the best runtime performance due to utilizing the fastmarching algorithm. Yet even APP2's radius calculation and pruning stages have sub-optimal algorithmic efficiency which severely limits the data sizes they can feasibly run on.

Performance. Better algorithm choice is an entry point for faster software, in practice however, hardware-specific op-

timizations yield far greater speedup factors(14). Neuroscience pipelines demand unprecedented scale exposing a long-standing weakness in computing infrastructure: data movement\{amount2021\}. Graph methods often traverse image volumes of sizes beyond memory capacity. Smaller subregions are then streamed as needed at runtime to mitigate large reads. These subregions are often visited and therefore read repeatedly and a large fraction of these dense regions are background or unaccessed pixels. This equates to programs that spend several orders of magnitude more time reading background values from disk than performing the desired computation. Additionally, generic compression methods on disk do not leverage the inherent spatial sparsity of neuroscience data. Algorithms are tightly coupled with data structures (15) and a unified vertex and data representation for all algorithms is essential to prevent excessive movement related to data conversions and inefficient access patterns. An ideal data model supports all desired algorithms efficiently and is leveraged as early as possible in a multi-stage pipeline. We choose the Volumetric Dynamic B+ tree (VDB) format (16) since it is optimized for fast access and low memory footprint for spatially sparse volumetric data. VDBs have a hierarchy of cubic regions for exploiting sparsity at various granularities. When a subregion contains no active voxels or entirely uniform values, its representation can be collapsed into a single value as visualized in Figure 5.

Scalability. Neuromorphology literature cites reconstruction as the bottleneck(4, 17) for high-scale neuroscience. Graph kernels such as pruning are also cited in neuroscience research as the bottleneck to analysis(6). Graph reconstruction methods lack *scalability*—the ability to handle large image regions on common systems such as a laptop or even on modern servers. In practice, this is due to the combination of the efficiency and performance concerns discussed above. Adapting to modern hardware requires enabling concurrency throughout a pipeline which can mitigate such data concerns. APP2, like other available iterative reconstruction methods, is only single-threaded, which is why large-scale applications must start multiple APP2 programs for each soma to enable some concurrency (18). Poor scalability is such a concern that common practice involves *windowing* regions of an image which can create cutoffs, artifacts, and intensive human intervention.

Automation. Large barriers to automation in reconstruction still exist. Even trivial manual steps vastly lower throughput by creating forced stops and long down times which greatly degrades the value of performance improvements. Lack of algorithmic efficiency or scalable design itself creates new barriers to automation, for example by enforcing windows of neurons to be cropped by hand which takes about 2-3 days per brain even when accelerated an automated cropping program. Additionally, swc outputs of semi-automated tools are usually filtered by hand. Other semi-automated pipelines such as (18) must manually identify soma locations since they are so critical to accuracy, but requiring intervention so early can limit many benefits of automated reconstruction. Finally, reconstruction is sensitive to signal-to-noise of the reporter

signal, which can be altered by tissue preparation, imaging, and stitching.

Correctness. Vaa3D and the BigNeuron project have significantly advanced collaboration in the field by providing a framework to compare reconstruction method accuracies. However, reporting end-to-end accuracy of reconstruction paper is only a first step towards reproducible research and reusable software tool-chains. For long-standing software that is intended to grow with collaborative effort, individual components of functionality also need to have corresponding automated unit tests. In order to improve algorithms and modules, baseline behavior must be first quantified and verified with such tests(19).

Portability. The neural anatomy and morphology community has had inspiring success in aligning data formats (e.g. SWC) and registration atlases (e.g. CCF). However coordinating software development collaboration is far more complex but has even greater potential benefits for sharing extensible reconstruction, analysis, and visualization pipelines. Software is far more dynamic, usually with a complex set of shifting dependencies that aren't declared explicitly by authors. The implications towards reproducible results are a growing concern, and other large scale scientific projects such as at CERN have adopted sophisticated methods to robustly lock and archive analysis software and all of its dependencies(20).

Experiment-Specific Data Type. Mouse brain connectomics utilizes several microscopy modalities, however this work focuses solely on light microscopy (LM), whereby genetic or viral expressers register fluorescent labels under varying light wavelengths. LM is a common modality for imaging labeled neurons in sectioned or intact brains. Building on previous studies, our work uses a novel membrane bound label MORF3(21) which has several advantages for studying Huntington's type disease cells. Many other labeling modalities target cell cytoplasm structures, which creates comparatively broad but diffuse signal in the inner volume of cells. Membrane-bound labels instead have sharper signals on the cell surface allowing less background image signal and higher fidelity membrane texture and synapses. Thin structures like neural branches–neurites–can also be more apparent at arbitrary distances from the cell body. Membrane labeling has the drawback of making somas appear hollow which must be corrected for in reconstruction software. When properly exploited, the inherent sparsity of our data can provide high throughput pipelines with significant yield per animal.

The sheer size of LM data induce the challenges we discuss above, yet we do not yet have strong heuristics for degree of acceptable *quantization*–bit width or range pixel values can take–or resolution / downsample rate to adequately describe desired morphology and connectomics.

Recut is designed to solve issues of general reconstruction pipelines, however specific algorithms were still implemented to solve the particular needs of our imaging capabilities to demonstrate real usage in an automated pipeline. Broader applicability is discussed in Section 4. We refer to the provided

algorithms as *stages* since they must proceed in a specific forward order for a given image region. Early stages of the pipeline operate on dense regions, whereas downstream stages tend to operate on an abstracted and sparse tree representation. The CC, SDF, and TC stages of Recut have a direct mapping with algorithms or components of APP2 and the intended transformation semantics are matched as best as possible. However, since the implementations are different we use the standard names of the transformations as they are used in the literature.

Contributions. To evaluate our implementation we provide several features and principles:

- **Correctness.** Recut's automation or computational techniques produce results consistent with the original intent without unexpected or unknown behavior changes or failures.
- **Accuracy.** Our framework demonstrates comparable or higher accuracy to those established in the reconstruction community.
- **Algorithmic Efficiency.** Recut demonstrates high algorithmic efficiency by reducing unnecessary computation and data access.
- **Performance.** Recut's implementation includes hardware specific optimizations to improve metrics like latency and throughput in ways that are as portable as possible.
- **Scalability.** Our software design is scalable, streamable and concurrent, meaning the implementation natively supports and performs consistently when adding significantly larger data sizes or more processors.
- **Automation:** Recut automates soma detection and first pass reconstruction and single-neuron partitioning from clusters. Manual proofreading steps are pushed as far downstream in the pipeline as possible to enable the benefits above.
- **Portability.** Open source software implementations are increasingly critical to scientific endeavors but have complex dependency chains that can take even software experts days to set up and integrate properly. Recut is installable and usable as a standalone program or a C++ library with 2 commands from a Unix command line.

Pipeline

Image Preprocessing (IP). The pipeline takes as input stitched image volumes. Inputs can be further enhanced to improve the fidelity however for discussions in this work we simply segment the raw inputs.

Segmentation (SG). To mitigate issues of raw images and membrane labeling, our pipeline applies the artificial neural network based algorithm U-net(22) to segment all regions of

the image belonging to neural tissue. This segmentation stage outputs a dense volume the same dimensions of the original image representing whether a given voxel is neural tissue. Additionally, this inference step aggregates a list of the filled soma locations. Soma location is critical for downstream stages since it acts as a root point for the branch-like projections—*neurites*—characteristic of neuron morphology. Somas are used to establish the starting seed locations for the subsequent stages of the pipeline. Somas are generally visible in image space and so critical to reconstruction accuracy that many pipelines rely on manual human identification of these locations. However, we rely solely on our automated methods which is essential for exploiting parallelism and boosting throughput.

VDB Conversion (VC). After segmentation, we convert the dense NN outputs into the low memory footprint VDB grid. This step conducts a dense read $O(n)$ such that all later stages of the pipeline can have fast accesses on the sparse volumetric grid.

Connected Component (CC). Segmentation establishes contiguous regions of coverage topology based on contextual information in local regions of image space. However, only a subset of these labeled regions are biologically plausible. In order to be relevant, any labeled voxel must also be reachable from the determined soma locations. Given the VDB graph with filled and identified somas as starting locations, the connected component stage labels all reachable vertices as selected in a region growing scheme. This stage yields all components of the image and the encompassed somas. *Component* refers to a cluster of neurons that are connected by at least 1 voxel. This CC stage fills a similar purpose to APP2's fastmarching stage or other reconstruction algorithms that use a Dijkstra or single source shortest path algorithm (SSSP) to traverse foreground voxels. Contrary to these other algorithms our CC stage simply stores a parent for each selected vertex with no distance or salience field.

Signed Distance Field (SDF). After the CC stage, the working set—all active elements—is complete. In other words, all voxels belonging to the coverage topology are known. However, the CC stage selects a vertex for each foreground pixel which is highly inefficient for subsequent stages in the pipeline. To compact the redundancy in selected vertices, each vertex is considered as a 3D sphere with a radius that extends to the nearest background pixel (inactive voxel). Calculating these distances provides a coverage area of each vertex via an operation referred to as SDF.

Tree Compaction (TC). Given a set of selected vertices with radii, we can perform graph compaction by starting from somas or more salient nodes first and traversing the graph, removing vertices that are already within the covered radius of a previously visited vertex. This leaves a graph that covers the same active image volume but is described by far fewer vertices. In neural datasets, this reduces vertex counts as detailed in Table 6.

Tree Prune (TP). While it is possible to compact losslessly for a exact volume coverage of foreground, in practice this is not useful since it produces enormous numbers of spurious branches that obfuscate morphology metrics. These adjustments traverse the active set of vertices and apply refinements on the sparse graph. In our pipeline, we use TP to remove all branches of length 1 with a parent bifurcation point.

Graph Cut (GC). G-cut segregates components into individual neurons in a data-driven, automated fashion. It assigns all pathways or branches to identified cell bodies in the image using the Dijkstra shortest path algorithm between all reachable paths between pairs of somas(23). This information passes to a linear solver which finds a global assignment optimum. The partitioned graphs represent each neuron which are aggregated for morphological analysis.

Proofread. Recut compresses the full image bounding volume to the VDB format then reconstructs and outputs discrete SWCs with their exact corresponding bounding window. This window is reflated from the foreground VDB voxels and written to a TIFF image on disk for subsequent proofreading in editing GUI software. Recut can optionally output swcs with coordinates with respect to the original image for correcting path breaks that extend beyond the components bounding volume. The final validated proofread category is derived by at least two annotators correcting the automated output of the APP2 category.

Morphological Analysis. The SG and CC stages entirely determine the coverage topology accuracy of reconstruction in our pipeline since subsequent stages merely compact the graph representation. Persistence homology for graph comparison and similarity builds a feature space that can express the morphological measurements of Sholl analysis while retaining spatial embedding information(24)(25). Persistence images have been used to classify neuronal types with subtle morphological differences that are difficult for even human experts to discriminate (26). The Topological Morphology Descriptor (TMD)(24) provides a metric equivalent to Sholl analysis and is less sensitive to imaging and the accuracy of the SG stage than pixel based accuracy metrics such as those used in (6). The TMD feature space has particular bearing on the edits to a tree that occur during proofread such as fixing breaks, merging branches, etc.

Results

Accuracy. At this labeling density reported, APP2 can only reliably terminate on data sizes up to bounding volumes of about $2048 \times 2048 \times 512$ or about 2 gigavoxels before runtimes become excessive even on a modern CPU as noted in (6) and verified in our test suite. At $30\times$ objective lens with our particular labeling probe, this bounding volume can encompass about 1 neuron as shown in Figure 4. Thus all accuracy comparisons to APP2, are for 1 neuron windows.

As shown in Table 2 our parallel connected component and SDF stage converge to exact equivalence with sequential runs

and ground truth over synthetic test and real neural data as expected and confirmed by unit tests. Note that Recut’s exact matching with APP2 sequential results diverges from the SDF stage onward since APP2 uses a 2D simplification of distance to compensate for the slow performance of its algorithm. APP2 employs hiercharchical pruning, a heuristic strategy to encourage selection a pruned set of maximally covering minimally redundant final vertices whereas Recut employs an optimized compaction strategy based off of the Advantra prune method(11).

Table 1. A comparison of the absolute difference between the average persistence image of APP2 and Recut against corresponding proofread neurons. Recut’s lower scores indicate that the morphological characteristics are closer to final proofread ground truth.

	APP2	Recut
Image distance	205	80.6

To compare these significantly different approaches, table 1 compares the absolute difference in average persistence images between APP2 and validated ground truth (205) and Recut and validated (80.6) for $N=42$ sample windows with a single neuron. Note that the set of proofread validated samples were derived directly from the APP2 sample thus the difference indicates manual morphological edits that humans made. The persistence images were generated from path distance thus neurite diameter information is not included in these statistics. The ground truth persistence image shows that the majority of branches start at a similar distance from the soma (x-axis) and terminate even closer to the soma (y-axis) which corroborates the bush-like morphology of D1/D2 neurons. Persistence images for APP2 have higher values along the diagonal $y=x$ which indicates that it introduces erroneous short branches beginning and ending at similar distances from the soma. The left-shifted peak of the APP2 compared to ground-truth possibly indicates that long paths are either added or extending during proofreading. Additionally, the image for Recut indicates that it often fails by introducing unnecessary branches, these branches are short and particularly common closer to the soma. Although the recut image does seem to capture some of the peak of the validated ground truth, the image is dominated by small branches.

Table 2. An accuracy comparison between APP2 and Recut across stages on synthetic inputs. For assigned parent error rates, automated recut SWCs are compared to APP2 parent assignments that yielded the best final output with hand tuned background threshold.

Stage	Error (%)
APP2 CC coverage	0.0
Recut CC coverage	0.0
APP2 SDF	23.1
Recut SDF	0.0

Note that Recut’s CC stage can be optionally run using a priority queue (heap) changing the stage to a parallel fast marching (FM) algorithm. Choice of FM vs. CC effects the CC traversal, the final parent paths back to root and the compaction order of TC. However, we did not find significant effects on final accuracy by using a FM approach thus we omit it from




Fig. 2. Persistence images for ground truth, Recut and APP2. Persistence images are aggregates of persistence histograms(24) a value for each branch is plotted by the radial distance from the soma at its birth (terminal) on the x-axis, to its death (bifurcation point or soma) on the y-axis.




Fig. 3. Recut or APP2 average persistence images subtracted from ground truth. Positive normalized values in red negative in blue.

the end to end accuracy comparison. However, researchers should compare and evaluate reconstruction algorithms via the respective effects on the final biological metrics they study.

Table 3. A comparison of the algorithmic efficiency of Recut not accounting for parallelism over multiple processor cores. n indicates the total working set count of vertices at that particular stage, r indicates the radius distance of a particular node.

Stage	APP2	Recut
CC	-	$\mathcal{O}(n)$
FM	$\mathcal{O}(n \log n)$	$\mathcal{O}(n \log n)$
SDF	$\mathcal{O}(nr^3)$	$\mathcal{O}(n)$
TC	$\mathcal{O}(nr^3)$	$\mathcal{O}(nr^3)$
TP	$\mathcal{O}(n)$	$\mathcal{O}(n)$

Algorithmic Efficiency. Besides TC, GC, or the optional fastmarching, the algorithmic efficiency of the graph stages of Recut are $\mathcal{O}(n)$ as shown in table 3. This is ideal efficiency since each stage needs to update n nodes and those updates are not dependent on the order of n .

The FM algorithm is less efficient than Recut’s CC stage because it requires on the order of $\log n$ more work per visited node to run. Note that we include the term r —a node’s radius—even though it is independent of working set count n since it helps in understanding the performance of stages below. APP2’s calculation of radius SDF algorithm requires r^3 operations per node, since it checks a circle, r times, whereas Recut’s implementation is again on the order of 1 operation per node, due to the same implementation pattern as the previous FM stage. Recut’s TC stage checks a sphere of vertices

according to each node's radius, whereas the APP2's hierarchical prune checks a circle, r times. TP for both Recut and APP2 require a single traversal of the compacted graph.

Performance.

Table 4. Recut runtimes per stage on confocal half coronal section with 30x objective lens on Camk-MORF3, 24-cores.

Stage	Recut (s)
CC	78.1
SDF	1.89
TC	76.8
TP	1.92
Hemi-section total	159.
Striatum total	2540

Table 5. APP2 vs. Recut throughput per striatum brain region 30x objective lens, Camk-MORF3, 24-cores.

Pipeline component	Striatum/day
Imaging	.25
Stitching	.33
Masking (manual)	.14
Inference	.30
APP2	.38
Recut	.34.
Proofread (manual)	.18

Runtime. The performance is even better than the algorithmic efficiency would suggest since n undergoes a substantial narrowing at the CC and TC stage as shown in Figure 1 and the downstream stages only visit the new sparse active sets. All Recut stages have a parallelized implementation, thus the n operations at each stage are conducted in parallel across available cores of a single CPU. To our knowledge the 509x compression factor of an image via VDB is the largest compression factor found in the neuroscience community due to being uniquely suited to sparse volumes. These algorithmic efficiency improvements combined with the concurrent execution model translate to lower runtimes and higher pipeline throughput as shown in Tables 4 and 5.




Fig. 4. Runtime to find the radii of synthetically generated spheres of varying radius. Average neurites in datasets images have diameters between 2-8 whereas somas have diameters between 16-32.

Table 6. Data footprints for the striatum region of 1 hemi-section, 30x objective lens, Camk2a; MORF3 ~1% fluorescent labeling density, bounding volume of 10k,18k,468.

Metric	Size
Dense voxel count (G)	85.2
Active voxel count (G)	.045
Active voxel %	.053
FIFO count (G)	.001
Max FIFO depth	512
Uncompressed HDF5 (GB)	170.
Disk coverage topology (GB)	.335
Disk compression factor	509.
Disk read + uncompress time (s)	2.89
In-memory coverage topology (GB)	2.85
Max FIFO footprint (GB)	4.01
Total runtime memory (GB)	6.86

Memory footprint. The graph topology footprints are shown at the vertex scale in Tables 7 and 8 and at the macro scale in Table 6

Scalability. Combining these principles we get consistent runtimes proportional to active voxel counts (instead of bounding volume) without hard to predict timeouts. We also achieve lower runtimes on 16x larger images than manually managed APP2. To remain responsive at interactive time scales, Recut must have adequate system memory (DRAM) such that the full runtime footprint can fit. For the data sizes used in Table 6 this equates to about 6 GB, which easily fits in memory of most modern laptops. The data footprint is so low that our reconstruction graph workloads are no longer disk-bound. While Recut is currently only supported for common PCs and CPUs, by comparison the average DRAM in newly purchased phones is almost 5 GB according to the Mobile Handset Sell-through Tracker. However to truly take advantage of these compression advantages at scale, a GUI tool for neuron proofreading should support visualization VDB grids. The core approach of Recut may be critical in future efforts for placing interactive reconstruction and manual proofreading in more hands. At a project scale, the accurately reconstructed neurons per animal—the *yield*—is a fundamental determinant of throughput. Yield is proportional to labeling density. Effectively leveraging sparsity in this way is the key to unprecedented image bounding volume scales.

Automation. Taken as a whole, these results have possible implications. Lossless high resolution imaging can be less costly to store or transfer than even heavily downsampled images. We can have productive and interactive workflows at brain region vs. single neuron scale on more common devices. We can also have predictable tradeoffs between label density and proofreadability, thus placing the reconstruction bottleneck back on yield per animal as opposed to image bounding volumes. Unified designs and simple implementations allows greater automation with less software infrastructure.

Future Steps. This work establishes a basic computational foundation for reconstruction pipelines that are automated

and amenable to hardware optimization, and guides decisions and trade offs between accuracy, automation and yield which is critical for the hardware software co-design(32) that has boosted fields like genomics.

Conclusion

This work aims to place reconstruction automation, hierarchical parallel infrastructure and data-oriented programming as central concerns in the design of future neuroscience pipelines. Given the right abstraction, we have demonstrated that a high performance and scalable framework is possible even in unique and challenging data settings. Though beyond the scope of this report, this framework is modular and extends to several graph-based neural kernels in our own pipeline, providing similar performance and scalability gains. Having correctness, accuracy and performance isolated and tested automatically allows Recut to enforce established standards and prevents regressions on these key principles by new external collaborators. Through simple design, we can provide a hardware-agnostic scientific community an open source framework for performance, thus advancing the possibility of making real-time analysis or interactive data science frameworks the common workflow.

Acknowledgments. We would like to thank the following authors for their contributions to this work. MZ, CC, HD, CP, WY, YC, ZC, and JC advised on design and provided helpful discussions and feedback. CP, CC and WY also provided the mouse imaging data. CC led the reconstruction proofread efforts. MZ designed and implemented the backend image IO library and the original sequential SG, CC, GC, and windowing output stages of the pipeline. KM designed, implemented and analyzed the Recut parallel framework and wrote the manuscript. We also thank Tony Nowatski for his suggestions on an early draft of the paper. This research is funded under NIH Grant No.: (U01MH117079-01).

References

- Hongkui Zeng and Joshua R Sanes. Neuronal cell-type classification: challenges, opportunities and the path forward. *Nature Publishing Group*, 18, 2017. doi: 10.1038/nrn.2017.85.
- R.C Cannon, D.A Turner, G.K Pyapali, and H.V Wheal. An on-line archive of reconstructed hippocampal neurons. *Journal of Neuroscience Methods*, 84(1-2):49–54, 1998. ISSN 0165-0270. doi: 10.1016/S0165-0270(98)00091-0.
- Katrin Amunts and Thomas Lippert. Brain research challenges supercomputing. *Science*, 374(6571):1054–1055, 2021. doi: 10.1126/science.abl8519.
- Chiara Magliaro, Alejandro L. Callara, Nicola Vanello, and Arti 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*, 7(AUG):1–8, 2019. ISSN 22964185. doi: 10.3389/fbioe.2019.00202.
- Ludovica Acciai, Paolo Soda, and Giulio Iannello. Automated Neuron Tracing Methods: An Updated Account. *Neuroinformatics*, 2016. ISSN 15392791. doi: 10.1007/s12021-016-9310-0.
- Jian Yang, Ming Hao, Xiaoyang Liu, Zhijiang Wan, Ning Zhong, and Hanchuan Peng. FMST: an Automatic Neuron Tracing Method Based on Fast Marching and Minimum Spanning Tree. *Neuroinformatics*, 17(2):185–196, 2019. ISSN 15392791. doi: 10.1007/s12021-018-9392-y.
- Linqing Feng, Ting Zhao, and Jinhyun Kim. neuTube 1.0: A new design for efficient neuron reconstruction software based on the SWC format. *eneuro*, 2(1):ENEURO.0049–14.2014, jan 2015. doi: 10.1523/eneuro.0049-14.2014.
- Hang Zhou, Shiwei Li, Anan Li, Qing Huang, Feng Xiong, Ning Li, Jiacheng Han, Hongtao Kang, Yijun Chen, Yun Li, Huimin Lin, Yu-Hui Zhang, Xiaohua Lv, Xiuli Liu, Hui Gong, Qingming Luo, Shaoqun Zeng, and Tingwei Quan. GTree: an open-source tool for dense reconstruction of brain-wide neuronal population. *Neuroinformatics*, 19(2):305–317, aug 2020. doi: 10.1007/s12021-020-09484-6.
- Hang Xiao and Hanchuan Peng. APP2: Automatic tracing of 3D neuron morphology based on hierarchical pruning of a gray-weighted image distance-tree. *Bioinformatics*, 29(11): 1448–1454, 2013. doi: 10.1093/bioinformatics/btt170.
- Dezhe Z. Jin, Ting Zhao, David L. Hunt, Rachel P. Tillage, Ching-Lung Hsu, and Nelson Spruston. Shuto: Open-source software for efficient and accurate reconstruction of dendritic morphology. *bioRxiv*, 2019. doi: 10.1101/226548.
- Miroslav Radojević and Erik Meijering. Automated neuron reconstruction from 3d fluorescence microscopy images using sequential monte carlo estimation. *Neuroinformatics*, 17(3): 423–442, December 2018. doi: 10.1007/s12021-018-9407-8.
- Cameron Arshadi, Ulrik Günther, Mark Eddison, Kyle I. S. Harrington, and Tiago A. Ferreira. SNT: a unifying toolbox for quantification of neuronal anatomy. *Nature Methods*, 18(4): 374–377, apr 2021. doi: 10.1038/s41592-021-01105-7.
- Donald E Knuth. Big omicron and big omega and big theta. *ACM Sigact News*, 8(2):18–24, 1976.
- Mike Acton. Data-oriented design and c++, 2014.
- Niklaus Wirth. *Algorithms + Data Structures = Programs*. Prentice Hall PTR, USA, 1978. ISBN 0130224189.
- Ken Museth. Vdb: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.*, 32(3), July 2013. ISSN 0730-0301. doi: 10.1145/2487228.2487235.
- Michael N Economo, Nathan G Clack, Luke D Lavis, Charles R Gerfen, Karel Svoboda, Eugene W Myers, and Jayaram Chandrasekar. A platform for brain-wide imaging and reconstruction of individual neurons. *eLife*, 5:1–22, 2016. doi: 10.7554/eLife.10566.
- Shengdian Jiang, Yimin Wang, Lijuan Liu, Sujun Zhao, Mengya Chen, Xuan Zhao, Peng Xie, Liya Ding, Zongcai Ruan, Hong-Wei Dong, Giorgio A. Ascoli, Michael Hawrylycz, Hongkui Zeng, and Hanchuan Peng. MorphoHub: A platform for petabyte-scale multi-morphometry generation. *bioRxiv*, 2021. doi: 10.1101/2021.01.09.426010.
- Michael Feathers. *Working Effectively with Legacy Code*. Prentice Hall PTR, USA, 2004. ISBN 0131177052.
- Chris Burr, Marco Clemencic, and Ben Couturier. Software packaging and distribution for LHCb using nix. *EPJ Web of Conferences*, 214:05005, 2019. doi: 10.1051/epjconf/201921405005.
- Matthew B. Veldman, Chang Sin Park, Charles M. Eyermann, Jason Y. Zhang, Elizabeth Zuniga-Sanchez, Arlene A. Hirano, Tanya L. Daigle, Nicholas N. Foster, Muye Zhu, Peter Langfelder, Ivan A. Lopez, Nicholas C. Brecha, S. Lawrence Zipursky, Hongkui Zeng, Hong-Wei Dong, and X. William Yang. Brainwide genetic sparse cell labeling to illuminate the morphology of neurons and glia with cre-dependent morf mice. *Neuron*, 108(1):111–127.e6, 2020. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2020.07.019>.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. 2015.
- Rui Li, Muye Zhu, Junning Li, Michael S. Bienkowski, Nicholas N. Foster, Hanpeng Xu, Tyler Ard, Ian Bowman, Changle Zhou, Matthew B. Veldman, X. William Yang, Hourii Hintiryan, Junsong Zhang, and Hong-Wei Dong. Precise segmentation of densely interweaving neuron clusters using g-cut. *Nature Communications*, 10(1), April 2019. doi: 10.1038/s41467-019-09515-0.
- Lida Kanari, Paweł Dłotko, Martina Scolamiero, Ran Levi, Julian Shillcock, Kathryn Hess, and Henry Markram. A topological representation of branching neuronal morphologies. *Neuroinformatics*, 16(1):3–13, 2017. doi: 10.1007/s12021-017-9341-1.
- Yanjie Li, Dinggang Wang, Giorgio A. Ascoli, Partha Mitra, and Yusu Wang. Metrics for comparing neuronal tree shapes based on persistent homology. *PLOS ONE*, 12(8):e0182184, August 2017. doi: 10.1371/journal.pone.0182184.
- Lida Kanari, Srikanth Ramaswamy, Ying Shi, Sebastien Morand, Julie Meystre, Rodrigo Perin, Marwan Abdellah, Yun Wang, Kathryn Hess, and Henry Markram. Objective morphological classification of neocortical pyramidal cells. *Cerebral Cortex*, 29(4):1719–1735, January 2019. doi: 10.1093/cercor/bhy339.
- Yuze Chi, Guohao Dai, Yu Wang, Guangyu Sun, Guoliang Li, and Huazhong Yang. Nxgraph: An efficient graph processing system on a single machine. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*, pages 409–420, 2016. doi: 10.1109/ICDE.2016.7498258.
- Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédéric Durand. Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics (TOG)*, 38(6):201, 2019.
- Jianming Yang. An easily implemented, block-based fast marching method with superior sequential and parallel performance. *SIAM Journal on Scientific Computing*, 41(5):C446–C478, January 2019. doi: 10.1137/18m1213464.
- Hanchuan Peng, Zhi Zhou, Erik Meijering, Ting Zhao, Giorgio A. Ascoli, and Michael Hawrylycz. Automatic tracing of ultra-volume of neuronal images. *bioRxiv*, 2016. doi: 10.1101/087726.
- Yimin Wang, Qi Li, Lijuan Liu, Zhi Zhou, Yun Wang, Lingsheng Kong, Ning Zhong, Renjie Chai, Xiangfeng Luo, Yike Guo, Michael Hawrylycz, Qingming Luo, Zhongze Gu, Wei Xie, Hongkui Zeng, and Hanchuan Peng. Terav empowers precise reconstruction of complete 3-d neuronal morphology in the whole brain. *bioRxiv*, 2019. doi: 10.1101/621011.
- Yatish Turakhia, Kevin Jie Zheng, Gill Bejerano, and William J. Dally. Darwin: A hardware-acceleration framework for genomic sequence alignment. *bioRxiv*, 2017. doi: 10.1101/092171.




Fig. 5. VDB grid of sparse neuronal data. The coverage topology of all active voxels is shown in red. Active leaf node regions—those that contain an active voxel—are outlined in blue. Active internal nodes regions are outlined in yellow and active secondary internal nodes in green. Empty leaf nodes and internal nodes have no outlines and have tiny memory footprint overheads.

Supplementary Note 1: Implementation

A. Algorithmic Efficiency. As demonstrated in Figure 1 we have two basic types of algorithms: NN and graph-based. NN is currently only used for the SG stage. NN methods have dense and regular data access patterns and simple parallelization schemes, whereas graph algorithms tend to have sparse irregular data access with complex concurrency schemes. The graph methods are also particularly well suited to take advantage of sparsity.

The desired transformations are quite common and there are many alternate implementations to these stages such as pruning with sphere packing, however CC and SDF stages are graph-based using breadth-first search on an iterative, advancing wavefront. This allows us to use the same underlying graph representation, overhead data structures and traversal patterns, thus greatly simplifying the implementation and reducing memory communication. The efficiency of the graph stages of Recut are listed in 3.

B. Performance.

Table 7. Vaa3D, APP2 marker footprint in memory

Attribute	Type	Bytes
Parent	pointer	8
Radius	double	8
Flags	int	4
Coordinate	3× double	24
Total	<i>Marker</i>	44

Table 8. Recut *vertex* and *message vertex* footprint in memory

Attribute	Type	Bytes	Red. v. APP2
Parent offset	3× int8	3	2.7×
Radius	uint8	1	8×
Flags	uint8	1	4×
Message offset	3× uint8	3	8×
Total	<i>Vertex</i>	5	8.8×
Total	<i>Message Vertex</i>	8	5.5×

B.1. Vertex Representation. We use a customized vertex model that can support downstream stages while remaining as compact as possible as shown in Tables 7 and 8. Neuroimaging data has inherent geometric properties and those properties

tend towards predictable power law distributions. Upon statistical analysis we can enforce data value range cutoffs to quantize the types fields used at each vertex to constrain the vertex memory size. Recut uses two basic types to represent the finest granularity of an image: the *vertex* and the *message vertex*. Any updates occur at the granularity of a vertex. All vertex attributes are held within VDB grids, whereas message vertices are temporary and are replaced dynamically in each FIFO. Concurrency is achieved via messages between leaf nodes which are buffered in FIFOs. Message vertices are simply an extension of the vertex type with a slightly larger memory footprint since they must be tagged with an offset to distinguish their coordinate identity within a leaf.

B.2. Image Representation. Typically, acceleration or parallelization schemes of graph processing rely on a preprocessing step to partition or arrange nodes and edges according to connections. This is an effective data-oriented approach for kernels operating on static power-law graphs(27). Unlike most high performance graph kernels, the graph stages of Recut are dynamic: they prune and reassign parents (edges) and repeatedly prune substantial portions of the active vertex set. Due to their embedding in volumetric space, the graphs are also planar, with an edge count of 1, with highly local connections. Additionally, only a small fraction (see Table 6) of all image voxels are ever visited and selected during the CC stage. For these reasons, the common high performance graph arrangement schemes, would mostly be single-use requiring a costly rebuilding during or after each stage and this rearrangement itself would have a similar runtime as a stage. To deal with these issues, Recut’s internal image and graph format both use OpenVDB grids, which are optimized for fast access and low disk and memory usage for sparse volumetric data(16). VDB grids allow Recut to transition from an image abstraction to a graph representation seamlessly which is why we use the terms voxels and vertices interchangeably. VDB grids hierarchically partition the data at different flexibly sized cube regions. The smallest partition region in a VDB grid is termed a leaf node, which is a cube of vertices processed by a thread independently of others leaf nodes. The leaf node size introduces spatial and temporal locality by enforcing a thread to stay within a small region of data until it completes processing all possible updates. Vertex attributes are accessed or stored as needed at the granularity of leaves. Each active leaf node requires a FIFO to hold future vertices to process and receive incoming messages.

B.3. Coverage Topology. Recut is designed to operate on arbitrarily large images that represent contiguous volumes of brain space produced either by accurate stitching or inherent in the imaging technique used. For example, light sheet data produces whole brains which Recut ingests as a single image, efficiently combines or compare multiple channels, and translates to a single topology representation. We conceptually split the topology into two parts, *coverage* and *attribute* which corresponds in concept to the underlying OpenVDB data structure’s usage of bit masks and attribute arrays respectively. The SG stages produces an over-representation

of the topology *coverage*—the total segmented voxels of the original image—which the CC stage constrains to only vertices reachable from somas. After the CC stage, the coverage is complete and the voxel based accuracy of the reconstruction is established (although paths can be reassigned to different neurons during GC). While *coverage* may be static after the CC stage, the *attribute* topology—the working set of all active vertices and their properties—gets reduced during TC. Despite the narrowing in memory footprint, the intended coverage is preserved abstractly via the radius property of the remaining vertices in the attribute topology.

B.4. Overhead. Before the CC stage, a FIFO is allocated for each of the active leafs output by SG. The subset of these FIFOs that are reachable from somas, have a lifetime for the rest of the pipeline since they need to dynamically buffer vertex messages to any leaf to ensure thread safety. We reduce memory requirements required for message buffers further by leveraging an update grid which sparsely stores whether border regions have been updated such that messages can be generated safely in the integration method at the end of each iteration shown in Algorithm 1.

B.5. Memory Leakage. Mismanagement of the dynamic memory footprint—*memory leakage*—leads to poor utilization of available resources and could cause large drops in program performance (i.e. scalability issues). More subtle issues in design can lead to memory leakage, for example while a raw image buffer is used for segmentation keeping it in program memory for the remainder of the stages would severely limit the regions that could be processed in parallel. This leads to a fundamental design point, only the minimum data needed for computation is brought up the memory hierarchy and cleared from higher memories where needed. Deep pipelines with diverse algorithmic needs require tradeoffs between memory footprint, algorithmic expressivity and fast access. For example, Recut builds a tree graph before the GC stage by enforcing each vertex has 1 edge: its parent. If Recut were expanded to include algorithms requiring undirected graphs such as GC, it would substantially increase the memory footprint of a vertex as it would additionally need to hold antiparallel edges (its children).

B.6. Sparse Adaptability. While the above details current design choices for our pipeline, almost all aspects can be tailored to other needs. Fluorescent labeling sparsity can vary greatly but is generally fixed across many imaging rounds. The sparsity displayed in neuromorphology is not a random set of pixels, it has localized correlation and spatially identifiable coherence, a property known as *spatial sparsity* (28). Furthermore, exploitable spatial sparsity tends to increase with higher dimensionality, resolution, expansion factors or added imaging channels.

The expected labeling density can vary between .001% to .01% in order to exploit ultra-sparse scenarios Recut can avoid treating internal node as dense regions entirely and instead monitor the sparse set of active vertices in a data structure optimized for low memory footprint and fast access for sparse

volumes known as OpenVDB(16). User-specified compile-time constants control this varying behavior and linkage with the optional OpenVDB 3rd party library so that runtime latencies are avoided.

B.7. Downsampling Flexibility. It's common in visualization and graphics to encode images at various resolutions—full resolution, half resolution, quarter resolution etc. This technique is used in high performance applications and the HDF5 format, and is most commonly referred to as multi-resolution or mip-mapping. This method can reduce or speed accesses in high performance applications at the cost doubling the memory footprint. The VDB library provides method to build multi-resolution sparse grids, but the sophistication of VDB data structure allows far faster access and smaller footprints without needing to adapt algorithms to this paradigm. While the resolution of voxels in grids is fixed throughout Recut, it would be possible to simulate upsampling by interpolation or downsampling by access or traversal pattern at runtime using VDB's provided functionalities.

B.8. Memory adaptability. Depending on the spatial sparsity characteristics, VDB grids can be arbitrarily deep with adaptive grid sizes chosen specifically for a data sample or type. This can be leveraged to substantially reduce the memory footprint at the cost of slower access times. To explore this tradeoff further refer to (16).

B.9. In-Memory Computing. Combining sparse data structures and a minimal vertex representation allows substantial reductions in memory footprint as detailed in Table 6. These reductions can shift the performance bottleneck from being disk-bound to main-memory bound (DRAM) for our entire end-to-end pipeline on a reasonably equipped imaging workstation. System builders of imaging workstations for sparse data should therefore keep in mind that in-memory computing will have disproportionately large benefits for performance with Recut.

C. Scalability. Due to the low reliance on global memory structures, we designed a lock-free shared memory approach based on (29) that operates on VDB grids that represent the required graph data. The basic structure of this scheme is illustrated in Algorithm 1. Concurrency is foremost at the granularity of leaf nodes via the Intel TBB library which parallelizes leaf processing across threads on a single CPU. We evaluate this framework via the speedup factor with respect to the number of cores used. Performance of these kernels is highly data dependent, therefore we define scalability metrics in terms of cores since they are aligned and localized with the CPU cache hierarchy more than threads. Much like other works such as Ultratracer(30) and TeraVR(31), during processing Recut can traverse through VDB grids following neurite paths in 3D space. However, since vertex attributes are stored in distinct regions of memory (in struct-of-arrays fashion), only the attributes of the active leafs specifically needed by each stage are accessed.

The VDB implementation supports image regions with x,y,z coordinates between $-2e9$ to $2e9$. Full adult mouse brain

confocal images imaged with a 30× objective lens have a bounding volume of about 65k×28k×2.7k with anisotropic z downsampling. A single VDB grid bounding volume can encompass 1.8 e16 of such brains arranged side by side due to the reliance on indexing via 3-dimensional coordinates of type signed integer-32. VDB grids are specifically designed to store the spatial structure in volumetric data and compress uniform or background regions. Our data has complex spatial sparsity at multiple resolutions which would not be effectively exploited by more naive adaptive grid methods. We evaluate the cost of the VDB format via the footprint in memory with respect to active voxel count of the labeled segmented regions. This metric should scale linearly for suitable sparse data structures.

Algorithm 1: This pseudocode represents the critical update loop. Each graph processing stage executes this loop with stage specific traversal and vertex field contained within the function on line 8. Integrating the ghost cell updates contains the added overhead of the parallel algorithm. Both inner parallelized for loops implicitly induce synchronization overheads.

```

1 Function update( $G, Q, stage$ )
   Input: Image  $G$  with located and filled somas, active
           set of queues or FIFOs  $Q$  of initial  $v$  to
           process, stage can be any arbitrary graph
           kernel with no side-effects outside of its  $L_j$ 
   Output: Grid  $G$  with updated attributes of  $v$ 
           according to stage
2 while any internal node  $I_i$  active do
3   for each active internal node  $I_i$  do
4     while any leaf node  $L_j$  active do
5       parallel for  $L_j \in I_i$ 
6         if  $L_j$  is active then
7            $\lfloor$  march( $L_j, Q_j, stage$ )
8         // parallel overheads
9       parallel for  $L_j \in I_i$ 
10      if  $L_j$  is active then
11        // check updated
           borders in active
           adjacent  $L$ 
12       $Q_j =$ 
           gather_updates( $L_j, Q$ )

```

D. Automation. The segmentation stage automatically identifies soma locations and removes need to specify background threshold value or do other image processing. Alternatively, Recut can take raw image buffers and a desired foreground percent. While a desired percent is less fraught than specifying a background intensity value as with APP2, it can still be an trial and error process which is why the automated SG stage is recommended.

E. Correctness. To ensure basic correctness, modern application development handles implementation complexity by *test-driven development* a paradigm whereby a program's isolated functionalities or integrated behavior is verified by a test suite accompanying the software. Rather than rely on developers to run these tests by hand to ensure no bugs have been introduced, full test suites can be run automatically on code changes or on continuous schedules to prevent any correctness regressions. This automated process, known as *continuous integration* (CI), enables large teams to deploy changes to complex software systems rapidly with confidence. Recut was developed in conjunction with a set of fast tests (less than 2 s total runtime) that are run automatically by the hosting platform.

F. Portability. At this initial release, the Recut framework is written in modern C++ (C++17 and above) leveraging language supported threading, atomic variables and some lightweight third party libraries. Recut optionally links with different image reading software as image data types are varied and non-standardized in the neuroscience community. In this paper we link with an internal library to read TIF or Imaris (HDF5) images. Additionally we leverage python functionality for the segmentation and graph-cut step detailed in (23). Recut and all its dependencies are distributed via the Nix package manager such that installation, compilation, and development requires running only two commands on a Mac, linux, or PC with WSL installed.

G. Experimental Setup and Technical Details.

G.1. Data. This pipeline takes TIF files as single-channel gray scale 3D images at 16-bits per pixel. The gray scale intensity represents the fluorescent labeling density determined by wet lab fluorescent techniques. Measured labeled signal density % is detailed in Table 6. Labeling is a product of months or years of biochemical tuning and research efforts which vary by scientific problem. If a fluorescent label is designed with too low a labeling frequency, the yield of neurons per mouse brain or breeding population may be too low to produce statistically viable results. If a fluorescent labeling is too high, neurons become a mesh, erroneously connected from the perspective of the reconstruction technique.

We determine a threshold intensity value below which all pixels are ignored for the entire reconstruction process. We term this the pixel selection density. This value can be input for an entire brain to save computation or can be calculated heuristically and applied to full images. The neural data used in this report is membrane-labeled mouse brain data with about .05% pixel label density. For real data testing we use a 7680×8448×383 image ($\frac{1}{8}$ coronal section; 2.36e10 total pixels) of mouse cortex labeled using CAM-K2 knockout at full resolution. The z dimension is downsampled 5× compared to x and y (covering steps of 1 μ m vs .2 μ m respectively). This region contains 232 distributed somas with dense mesh-like, continuous branching patterns. 1 section is about 16× the window size allowed by most reconstruction software. To

challenge and test performance, other synthetic images are generated as detailed in the tooling section.

G.2. System. We tested performance on a workstation with an AMD Ryzen Threadripper 3960X 24-Core CPU at 3.79 GHz with 2 hyper-threads per core, 256 GB DRAM, and 3× Samsung SSD 860 EVO 4 TB configured as RAID0.