Title Page

Title:

Biophysical Simulation Enables Multi-Scale Segmentation and Atlas Mapping for Top-Down Spatial Omics of the Nervous System

Authors:

Lina Mohammed Ali[1,2], Aldrin Kay Yuen Yim[1,3], Thien Nguyen[1,2], Nicholas Tu[1,2], Emanuel Gerbi[1,2], Diana Grigore[1,2], Jason Waligorski[1,2], Colin Kremitzki[1,2], Robi Mitra[1,3], Jeff Milbrandt[1,3], William Buchser[1,2,3]

Affiliations, Middle Names, E-mails, OrcIDs

[1] Department of Genetics, Washington University School of Medicine, St. Louis, MO USA

[2] Functional Imaging for Variant Elucidation at the McDonnell Genome Institute, St. Louis, MO USA

[3] McDonnell Genome Institute, St. Louis, MO USA

1. Lina K. Mohammed Ali - linakali@wustl.edu 0009-0002-9388-4150
2. Aldrin Kay Yuen Yim - yim@wustl.edu 0000-0001-9562-2328
3. Thien Nguyen - tnguye11@trinity.edu 0009-0003-1282-4832
4. Nicholas K. Tu - nicholas.k.tu@wustl.edu 0009-0008-9522-618X
5. Emanuel B. Gerbi - egerbi@wustl.edu 0009-0005-6314-8402
6. Diana Grigore - d.grigore@wustl.edu 0009-0008-6743-589X
7. Jason E. Waligorski - jwaligor@wustl.edu 0000-0001-8422-6004
8. Colin L. Kremitzki - ckremitz@wustl.edu 0000-0001-9175-3136
9. Robi D. Mitra - rmitra@wustl.edu 0000-0002-2680-4264
10. Jeffrey D. Milbrandt - jmilbrandt@wustl.edu 0000-0002-5477-7689
11. William J. Buchser - wbuchser@wustl.ed 0000-0002-6675-6359

Abstract (187 words)

Spatial omics (SO) has produced high-definition mappings of subcellular molecules (like transcripts or proteins) within various tissue samples. Most applications of SO are molecule-driven i.e. the spatial distributions of transcripts are used to make distinctions between samples. However, such inferences do not automatically utilize brain atlas regions. Here, we present SiDoLa-NS (**Si**mulate, **Do**n't **La**bel – **N**ervous **S**ystem), an image-driven (top-down) approach to SO analysis in the nervous system. We utilize the biophysical properties of tissue architectures to design synthetic images mimicking tissue samples. With these *in silico* samples, we train supervised instance segmentation models for nucleus segmentation, achieving near perfect precision and F1-scores > 0.95. We take this a step further with generalizable models that can identify macroscopic tissue structures in the mouse brain (mAP_{50} = 0.869) and spinal cord (mAP_{50} = 0.96) and pig sciatic nerve (mAP_{50} = 0.995). SiDoLa-NS provides a framework in applying and analyzing SO data that leverages high-definition images that are taken alongside spatial omics pipelines. Notably, SiDoLa-NS provides solutions for common challenges in supervised model training, including but not limited to annotator bias, limited generalizability, and production of massive, high-quality training sets.

Short Summary

The SiDoLa-NS micro-, meso-, and macro-scale models are generalizable, supervised models for neuronal segmentation in cell to tissue-level contexts. SiDoLa-NS is novel in its combination of three core ideas: it is top-down (image first), trained solely on synthetic images, and it is multi-scale. The tool is validated on brain, spinal cord, and sciatic nerve for advanced segmentation tasks.

Abbreviations

SiDoLa-NS: Simulate, Don't Label – Nervous System

SiDoLa-NS-micro-CNS: Micro – Central Nervous System

SiDoLa-NS-micro-PNS: Micro – Peripheral Nervous System

SiDoLa-NS-Meso-pSN: Meso – porcine Sciatic Nerve

SiDoLa-NS-macro-mCB: Macro – mouse Coronal Brain

SiDoLa-NS-macro-mSC: Macro – mouse Spinal Cord

IoU: Intersection over union

mAP: Mean average precision

YOLO: You Only Look Once; image detection and segmentation models

1 Introduction

1.1 Images as Datasets

Optics have been the fascination of scientists since Ibn al-Haytham, and the data contained in even a single image of cells can spark whole fields of biology. Images of tissues are even more rich, and some of the most complicated include the nervous system, where each subtly distinct region enables specific functions[1–3]. The complexity arises fundamentally from the tissues' intricate structure, highly organized in the brain, spinal cord, nerves and support structures[4,5]. Understanding the nervous system in health and disease requires an in-depth exploration of the cellular and molecular patterns within their spatial context[6–9]. An ongoing challenge is to extract all the image features, especially ones that are usually only identifiable by pathologists or anatomists. New technology has promised to shed light on the molecular underpinnings of the tissue, but what about the intricate regions and patterns in the image itself?

1.2 Spatial Omics

Over the past decade, spatial omics (SO) technologies have revolutionized our ability to map cellular and molecular targets within histological sections, offering spatial localization insights at the tissue scale [10–16]. The immense throughput of these technologies has provided transcript- and protein-wide data across diverse tissues including the nervous system. For example, 10X Visium has been applied in the human spinal cord, enabling the exploration of cellular and molecular drivers of spinal cord development and the creation of a transcriptomic atlas both in mouse and human[17–20]. Additionally, numerous applications of these technologies have been observed in the mouse brain, an extraordinarily complex and dynamic region with exact architectural and physiological properties[21–24]. Access to such SO data is essential in understanding the interactions among diverse cellular and molecular components in both health and disease. Notably, 10X Xenium has also been utilized in the mouse brain, allowing for subcellular mapping of transcripts[25]. Addressing questions pertaining to spatial identities of transcripts and proteins necessitates precise image segmentation and registration to biologically defined tissues[16,26,27].

1.3 Bottom-Up Analyses

Single cell sequencing technologies have changed the way cells are examined in the nervous system[28,29]. Inclusion of spatial coordinates in omics datasets has led to analysis methods which attempt to uncover regional associations between molecules like RNA and protein. Many computational methods have been developed to surmise biologically meaningful clustering of spatially variable features (typically gene expression) from these datasets[10,30–32]. These 'bottom-up' approaches often employ clustering algorithms which learn locally and globally variable gene expression patterns. Transcript density is used to either identify distinct spatial domains with discontinuous changes in gene expression or assume continuous expression and fit a function to model the expression of a transcript across space[33–36]. GASTON models the transcriptomic space as a topographic map, fitting the spatial landscape based on expression density[37]. Tools like Baysor model cells as elongated distributions of their transcriptional profiles within a Gaussian distribution[38]. Alternatively, bottom-up approaches may also resolve spatial domains via the integration of multi-omics modalities, such as SpatialGlue and MIP-Seq[34,39], which require access to various omics datasets for a given slice. There are even techniques to integrate two different species with gene expression information[40]. Unfortunately, the more genes that are used to define regions, the less there are to analyze for changes in expression patterns, often leaving insufficient genes to discover unknown cell types[41,42].

1.4 Top-Down Segmentation

Top-Down approaches determine regional information from image features or by alignment/registration. This methodology allows researchers to ask questions about molecular changes in a defined region and for cross validation of bottom-up transcriptomic analyses. Identifying pathologies and biological structures within tissue is a common use case for this approach. Original tools aimed to use interpretable algorithms based on local pixel values[43,44] or registration[45,46]. Recent methods leverage deep learning[47,48], transfer learning[49,50], transformers[51,52], and combinations of all these to segment relevant structures from background and label cells or tissue types in immunofluorescence and histology datasets. To train and validate these tools, data must be labeled either by hand or autoencoder, each presenting their own strengths and drawbacks. Autoencoders allow nuanced embeddings, but remove human input from the process, degrading interpretability and noise sensitivity[53,54]. Human-labeled data enhances model interpretability but increases the time to produce training sets, can introduce errors/biases, and may hinder novel discoveries[55]. Both methods suffer from a lack of widely available data[56]. Despite this problem, there is a growing desire for tools that perform non-manual segmentation of biological data: untapped image data in spatial transcriptomic analyses offers a space for expansion of these tools' use cases. DeepSlice and FTUNet perform atlas mapping and lesion segmentation in macroscopic neural architectures (macro-scale). iStar and UniFMiR can enhance existing low-resolution, high-magnification data sets (micro-scale). At intermediate (meso-) scales, regional inferences provide key insight to bridge tissue structures and the cells defining them[57,58].

1.5 Multiscale Segmentation with SiDoLa-NS

In this manuscript, we present a tool for tissue segmentation across micro-, meso-, and macro- scales of cell and tissue organization. We use a biophysics engine to create image simulations to mimic tissue structurally at each level of interest; at the micro-scale to identify nuclei and cell boundaries, at the meso-scale to identify intermediate features like clusters or bands of cells, and at the macro-scale to identify regions of the nervous system tissue with atlas mapping. Our *in-silico* image simulations offer 1) full control over feature distributions, 2) imperfection and noise representation (yielding tolerance), and 3) access to a virtually infinite and annotator-error-free training set. This is distinct from diffusion and GAN-based approaches, which take, as input, real microscopic images for synthetic data generation. Our datasets are used to train instance segmentation models, resulting in SiDoLa-NS, which can then define cellular and regional identity based on image features alone and avoid biases from global gene expression and double dipping. We validate SiDoLa-NS with publicly available brain, spinal cord, and sciatic nerve images and show the versatility of our tool across various platforms, including 10X Visium HD and immunofluorescence sections. At the micro-scale, SiDoLa-NS performs cell and nuclei segmentation, at the meso-scale, it allows (for example) within- and between- fascicle analyses in the sciatic nerve. At the macro-scale, SiDoLa-NS defines brain regions based on the Allen Brain and Spinal Atlases[59–66]. SiDoLa-NS is a specialized top-down tool that automatically annotates nervous system tissue and gives researchers new ways to group the data across the scales and ask unbounded molecular questions.

2 Results

2.0 Overview of SiDoLa-NS

We present SiDoLa-NS (Simulate, Don't Label – Nervous System), a suite of tools designed for the detection and classification of image features within various contexts of the nervous system. SiDoLa-NS (pronounced See Doh La, think Do-Re-Mi) employs an image-based top-down approach, where labels are assigned to biologically defined regions of an image (**Fig. 1A**). This contrasts with now frequent use of bottom-up methods that cluster cells based on their gene or protein expression, which while biologically meaningful, often blurs heterogeneous cell or tissue interactions. SiDoLa-NS models are trained on sets of fully synthetic/simulated images, which are created with biophysical properties, *not* by generative AI (**Fig. 1B**). We use uniform distributions of parameters to build a variety of cells which then fill regions, then convert these three-dimensional physical representations into images using an optics engine. We create images that reflect different modalities like immunohistochemistry (IHC) and immunofluorescence (IF). Using simulated images has the advantage of providing truly massive datasets, where we can easily generate hundreds of thousands of images in a few hours. There is also no annotator bias, perfectly resolved ground truth segmentation masks and labels, and no memorization issues since the classifier never sees real images until inference. With the simulated images, we represent examples of various tissue and optical aberrations that may be present in a physical tissue section including 1) out-of-field cells, 2) broken tissue, 3) poor focus / contrast / resolution, and 4) distorted or warped tissue sections. We also over-represent examples that are difficult to existing segmentation algorithms, such as extremely high density. Of note, at no point during model training are real biological samples used; both training and optimization utilize simulations, reserving the biological samples solely for downstream testing and analysis.

While each model is employed at a specific scale, SiDoLa-NS streamlines the integration of data across multiple scales. This is useful for tissue such as the brain, where it is essential to understand the spatial extents of regions from a brain atlas, to find clusters, lamina, and striations of cells, and to carefully detect nuclei (**Fig. 1C**). In this manuscript, we examine the results of SiDoLa-NS against three publicly available datasets from mouse and porcine neuronal tissues, including the brain, spinal cord, and sciatic nerve (**Fig. 1D**). SiDoLa-NS is designed to work with a wide range of platforms, from classical IHC and IF tissue sections, to the full range of spatial omics technologies (**Fig. 1E**). Lower resolution spatial omics, like the original 10X Visium platform, still operate normally since the accompanying tissue section imaging is done at relatively high resolution, even if the transcript data is obtained at a lower one.

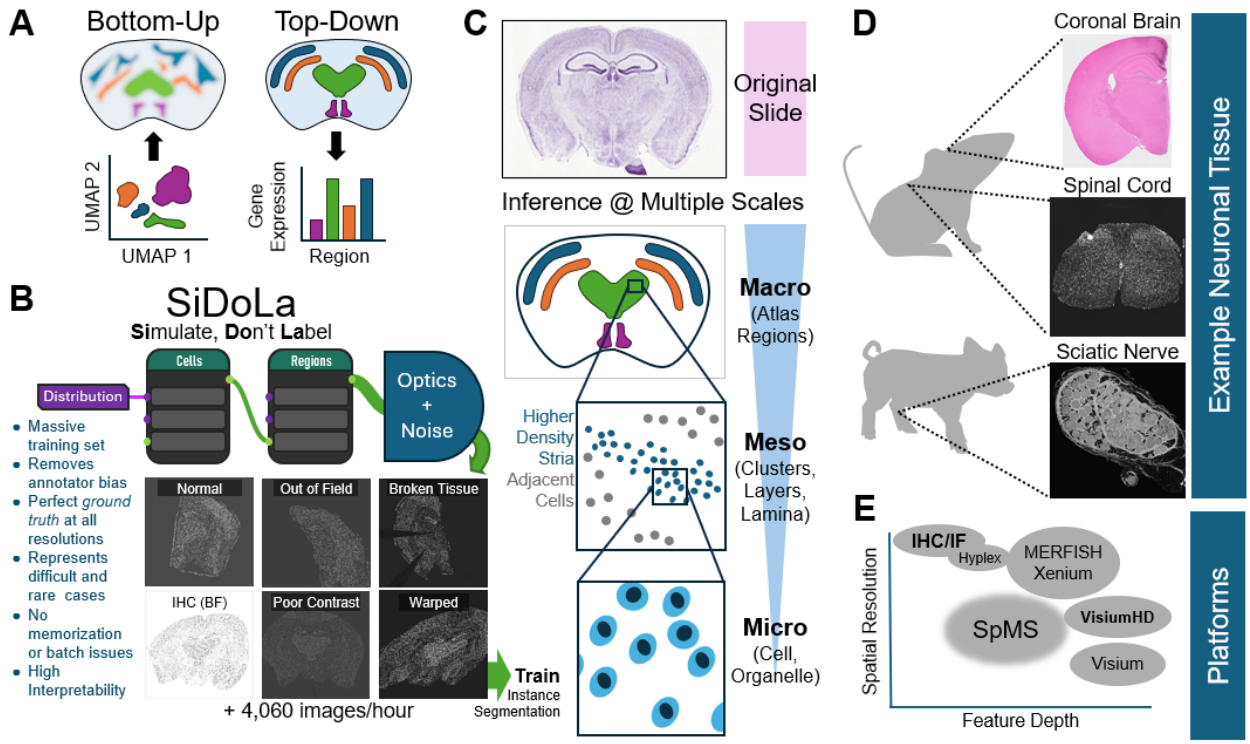

Figure 1. Simulate Don't Label: SiDoLa-NS for Top-Down Segmentation of Regions, Clusters, and Cells in Neuronal Tissue. A) Schematic illustrating the distinction between bottom-up and top-down methods. Top-Down approaches capture distinct image-based features and representations of a tissue. Bottom-Up approaches consider omics data (such as gene expression) to build molecular definitions of regions and cell types. **B)** Simulations create a massive synthetic dataset by stringing together biophysical nodes that build *Cells* and *Regions* of the image, and then transform them from 3D geometry into images with an *optics* engine. Synthetic micrographs are shown below the nodes, giving examples of variation and features for a mouse coronal brain slice. **C)** Mouse coronal brain Nissl section from the Allen Brain Atlas, shown as the "original slice", with cartoons below and the concept of multiple spatial scales where segmentation identifies individual instances of a variety of features. Specifically, this approach identifies atlas-based regions at a macroscopic scale, as well as cell density and other mesoscale patterning in the tissue and is also able to segment cells to incorporate subcellular microscale information, like cell type and structure. **D)** SiDoLa-NS was tested on three nervous system tissue types including mouse brain and spinal cord, as well as pig sciatic nerve (a human-scale nerve cross section). **E)** These tools are designed to work with the full spectrum of image-based datasets from classical immunohistochemistry to spatial transcriptomics platforms.

2.1 Case 1: SiDoLa-NS is Trained on Biophysical Simulations of Nuclei in Brain Slices to Segment a Histological Section of Mouse Brain

We first asked whether we could segment cells, layers, and atlas regions within a mouse brain slice. We constructed biophysical representations of mouse brain sections at multiple scales. The first fully synthetic image set (with 4,353 pairs and 1.96×10^6 objects) was at the cellular level and mimicked the DAPI staining phenomena observed in the nuclei of cells in the brain (**Fig. 2A 1**). These synthetic images represented neuronal, oligodendrocyte, and other nuclei of different size, shape, and intensity and were purposely rendered so that some were in better focus than others. We then trained an *instance*

segmentation model to find these nuclei, resulting in SiDoLa-NS-micro-CNS. We used a publicly available 6mm x 6mm H&E coronal cross section of mouse brain which was run on 10X Genomic's Visium HD platform. Zooming in to small tiles of the full brain image show varied morphology from a single channel of the H&E staining (**Fig. 2A 2**). The results of inferring on the biological examples with the model trained on simulation images is shown, where the brighter and dimmer nuclei are both detected, but given different class labels (**Fig. 2A 3**). Increasing the perspective to see larger numbers of cells, but still not the whole brain, we refer to as the "mesoscale". Here we observed two cellular populations with distinct densities, which we identified as 'striated' or 'laminar' regions. These striated regions are bands of high nuclear density, whereas the adjacent regions were observed with more diffuse nuclei (**Fig. 2B 1**). Striated regions include the granular cell zones (for example of the Dentate Gyrus or Hippocampal CA regions), which are composed primarily of neuronal cell bodies (**Fig. 2B 2**) and detected (**Fig. 2B 3**) in the Visium HD brain slice. Model performance was measured with the synthetic validation dataset. SiDoLa-NS-micro-CNS achieved high IoU (Intersection over Union) scores for detection and an mAP_{50} (mean average precision for IoU greater than 0.5) of 0.715 (**Fig. S1**).

We also asked whether the SiDoLa-NS-micro-CNS model could work beyond the nervous system, therefore testing its generalizability. We chose a publicly available DAPI stained image of lung adenocarcinoma (10X Genomics Visium HD), which is challenging for other segmentation algorithms due to regions of densely packed cells. We observed excellent segmentation and detection of nuclei, even in high-density regions of this tissue (no parameters changed from brain inference), supporting that the model is generalizable and unlikely to succumb to issues with batch effects (**Fig. S2**).

2.2 SiDoLa-NS is Trained on Biophysical Simulations of Brain Slices and used to Map Brain Regions in a Histological Section of Mouse Brain

At the macro scale (where the entire section of brain is visible), we asked whether we could assign named regional identities to the detected nuclei. To accomplish this, we created a dataset of fully synthetic coronal slices of the mouse brain (single hemispheres and full). These macro scale simulations were built from region labels based on the Allen Brain Atlas reference for the mouse brain (**Fig. 2C 1**). We used a set of 144,438 training image sets and 1,092 validation image sets (all synthetic) to train the models, which can be combined in an ensemble. Using the same coronal brain slice (**Fig. 2C 2**), the resulting regions inferred from the model clearly segment brain regions (**Fig. 2C 3**). The primary model (SiDoLa-NS-macro-mCB) achieved an excellent precision-recall curve and the confusion matrix on the validation set was a near-perfect diagonal, indicating close to zero class confusion (**Fig. 2D**). We measured a top mAP_{50} of 0.869 on the validation set (**Fig. S3**).

Another feature of the SiDoLa-NS models is their robust performance even with noisy or blurry samples. To demonstrate SiDoLa-NS-macro-mCB's ability to handle these imperfections, we inferred on images where noise was introduced in a calibrated fashion and found that SiDoLa-NS maintained an mAP_{50} above 0.783, even when most of the image was noise (up to 70% noise mix, **Fig. 2E**). Next, we simulated images where focus was reduced across the entire image (even) or in certain parts of the image (distorted). SiDoLa-NS-macro-mCB maintained an mAP_{50} above 0.803 and 0.722 across the even and distorted images respectively (where the aperture f-number was reduced to as low as 0.0013). While it was now clear that the ability for the model to obtain accurate results was robust to noise and blur, a more interesting test is on biological rather than synthetic images. We specifically sought to test versatility to *staining modality*, and we inferred on an unseen Nissl-stained brain samples with

corresponding true labels, provided by the Allen Brain Atlas (**Fig. 2F**). Here again, even though the Nissl image was a whole brain and not a single hemisphere, and has a bright background (the colors were not inverted prior to model inference), it was still able to find the brain regions and overlay them.

We were interested in seeing the differences from taking a top-down vs a bottom-up approach to cellular labeling. Specifically, we tested how SiDoLa-NS would look relative to Visium HD's Seurat pipeline for clustering and spatial overlay. We inferred SiDoLa-NS-macro-mCB on the whole mouse brain slice which automatically links these region labels to the cell-level segmentations and transcripts. Here, both individual cells and polygonal areas are assigned brain regions (**Fig. 2G**). We compared the atlas predictions (Top-Down) to the transcriptomic data (Bottom-Up) directly by overlaying the predicted region labels over the Seurat-generated UMAPs, or by overlaying the Seurat clusters over the brain image. As expected, the two methods are complementary, with the top-down excelling at finding boundaries and filling the whole area, while the bottom-up at finding different cell types within regions. For example, cells which SiDoLa-NS assigned to the Striatum-Like Amygdala, Lateral Hypothalamus, and Thalamus were difficult to differentiate by Seurat cluster. Using both together can help disambiguate numbered clusters, which can be matched up to named atlas regions (**Fig. 2H**).

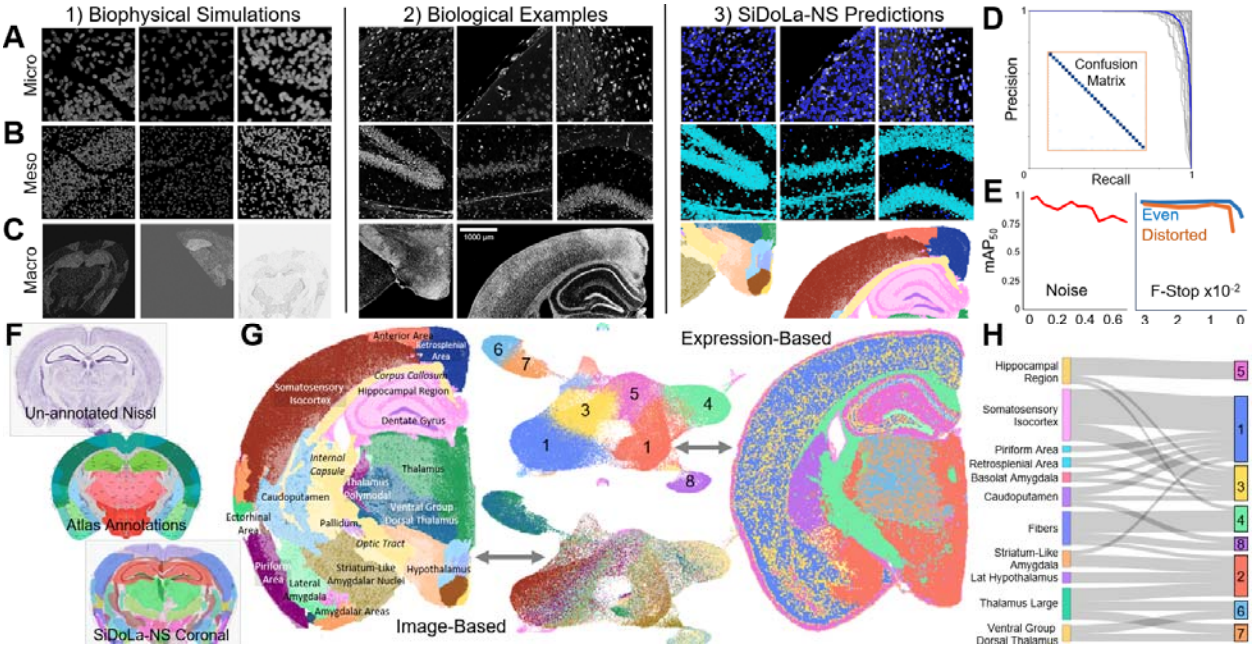

Figure 2. Multi Scale Segmentation and Classification of a Mouse Coronal Brain Section. A) At the micro scale, SiDoLa-NS trained exclusively on simulated data (1) featuring multiple cell types. Micrographs of real brain images (2) were used to run SiDoLa-NS-micro-CNS at this scale and identified nuclei (light and dark blue bounding boxes) from background (3). **B)** At the meso-scale, simulated images (1) featuring bands of higher density cells. Micrographs from brain slice, notice density patterns (2). SiDoLa-NS classes overlaid, which differentiated high density striated regions (cyan) and adjacent cells (blue) (3). **C)** At the macro scale, SiDoLa-NS was trained on simulated images (1) featuring full or hemi coronal sections of brain tissue. SiDoLa-NS-macro-mCB segmented real images of an H&E coronal brain slice (2), highlighting biologically defined regions (3). **D)** Precision-recall curve and confusion matrix for the macro atlas predictions. **E)** Line graph that depicts mean average precision (mAP_{50}) on the y-axis vs. different amounts of image manipulation on the x-axis. The left panel increases white noise proportion from 0 to

0.7. The right panel increases blur by lowering the aperture's f-Number in both even and distorted (blur was differential across the image) cases until ~0.01, where mAP_{50} dropped. **F)** Unannotated, Nissl-stained mouse P56 coronal slice, alongside Allen Brain Atlas annotations and SiDoLa-NS annotations (the Nissl brain is 9730 µm in width). **G)** Comparison of Top-Down and Bottom-Up approaches side-by-side. The left brain image and lower UMAP are SiDoLa-NS top-down, while the right brain and upper UMAP are bottom-up. The UMAP was constructed from transcriptomic data, and its colors are by Seurat clusters in the upper UMAP, and by SiDoLa-NS labeled brain regions in the lower. **h.** Sankey diagram mapping SiDoLa-NS-macro-mCB brain regions to the Seurat clusters in Figure 2G.

2.3 SiDoLa-NS for Spinal Cord Nuclei Detection and Atlas Mapping

We next asked whether a less-commonly studied neuronal tissue, the spinal cord, could be used with this system. SiDoLa-NS was tested on publicly available immunofluorescence (IF) sections of mouse spinal cord, stained with DAPI for nuclear localization, NeuroTrace (NT) for neuronal soma, and ChAT for motor neurons[65]. On a side note, this mouse spinal cord slice was not a spatial omics dataset, and therefore lacked transcriptional data. At the micro scale, we used the same model trained on simulated nuclei (**Fig. 3A 1**), even though the spinal cord IF section was lower resolution and slightly blurrier (**Fig. 3A 2**) than in the brain. Still, the SiDoLa-NS-micro-CNS with identical settings worked to segment the nuclei (**Fig. 3A 3**). Utilizing the SiDoLa-NS-micro-CNS model, we found sufficient generalizability to the mouse spinal cord nuclei images. This demonstrated the robustness and versatility of SiDoLa-NS across tissues and, more challenging, across platforms and image modalities.

At the macro scale, we simulated whole and hemi spinal cord histological slices and varied parameters including noise, warp, and tissue degradation (**Fig. 3B 1**). Spinal cord atlas mapping, like in the brain, previously required manual image registration; but the included variations in the training set allowed for streamlined application for this tissue section. Using this macro-scale simulated image set (12,609 training, 30 validation), we trained SiDoLa-NS-macro-mSC, for whole slice spinal cord atlas mapping in mouse. We predicted this model on the biological IF whole tissue sections (**Fig. 3B 2**), highlighting the various spinal lamina alongside other prominent regions in the white and grey matter (**Fig. 3C 3**). SiDoLa-NS-macro-mSC achieved a top precision of 0.96 at a confidence of 0.92 on the validation set (**Fig. S4**). We next examined this model's robustness to variations in focal distance and tissue rotation. SiDoLa-NS-macro-mSC was resistant to focal distance variations (zooming in to -300 µm and out to 4000 µm from a normal distance of 675 µm between the sample and the camera); the model maintained an mAP_{50} above 0.70 within a range of -200 µm to +1000 µm from the standard distance used in the training set images (**Fig. 3D**). Zooming in too tightly quickly degraded model performance presumable because it lost the spatial context provided by the boundaries of the spinal cord. Similarly, we tested the model's robustness to degree of rotation, and found that in our test case, SiDoLa-NS-macro-mSC maintained an mAP_{50} above 0.60 up to 40 degrees of rotation along the y-axis (tilting one side of the plane of the image 'into the camera').

We inferred SiDoLa-NS-macro-mSC on an unseen Nissl-stained image to examine the model's abilities to varying tissue and staining modalities. SiDoLa-NS-macro-mSC achieved a top precision mAP_{50} of 0.96 at a confidence of 0.929 (**Fig. 3E**). The mAP_{50} was class-variable, where larger and more well-bounded regions like the dorsolateral fasciculus achieved precisions of 0.995 and the thinner layers, like Lamina 1

and 2, had lower precisions (0.145 and 0.622 respectively). This is likely due to slight shifting of where the Allen brain polygons occurred compared to where the model predictions occurred.

To validate the regional assignments on the IF mouse spinal cord sections, we compared the staining intensities in the reported atlas regions. When inferring, SiDoLa-NS multiscale results provide quantitative measurements for each predicted object at the micro-, meso-, and macro scales, including area, channel intensity, and texture (measured as standard deviation of pixel intensities within the object). We combined the SiDoLa-NS-macro-mSC predicted regions to three generic groupings: white matter, dorsal, or ventral gray matter. As expected, regions labeled as white matter were located on the outer perimeter of the spinal cord, enclosing dorsal and ventral regions. Dorsal gray matter occupied the upper half of the spinal cord, and ventral regions fell in the lower half (**Fig. 3F**). We next quantified the variable feature distributions across these groupings. For example, the white matter regions had slightly larger nuclei area. After applying an intensity threshold, we observed a variable ChAT-positive cell distribution (**Fig. 3G**). We expected the ventral horn to contain the greatest fraction of ChAT-positive cells and indeed, the ventral horn was 38% ChAT+ compared to only 2% in the dorsal horn. Additionally, we observed less than 10% NeuroTrace-positive cells in the white matter (compared with 67 and 71% in the ventral and dorsal horns, respectively), as expected, and a uniform nuclear intensity.

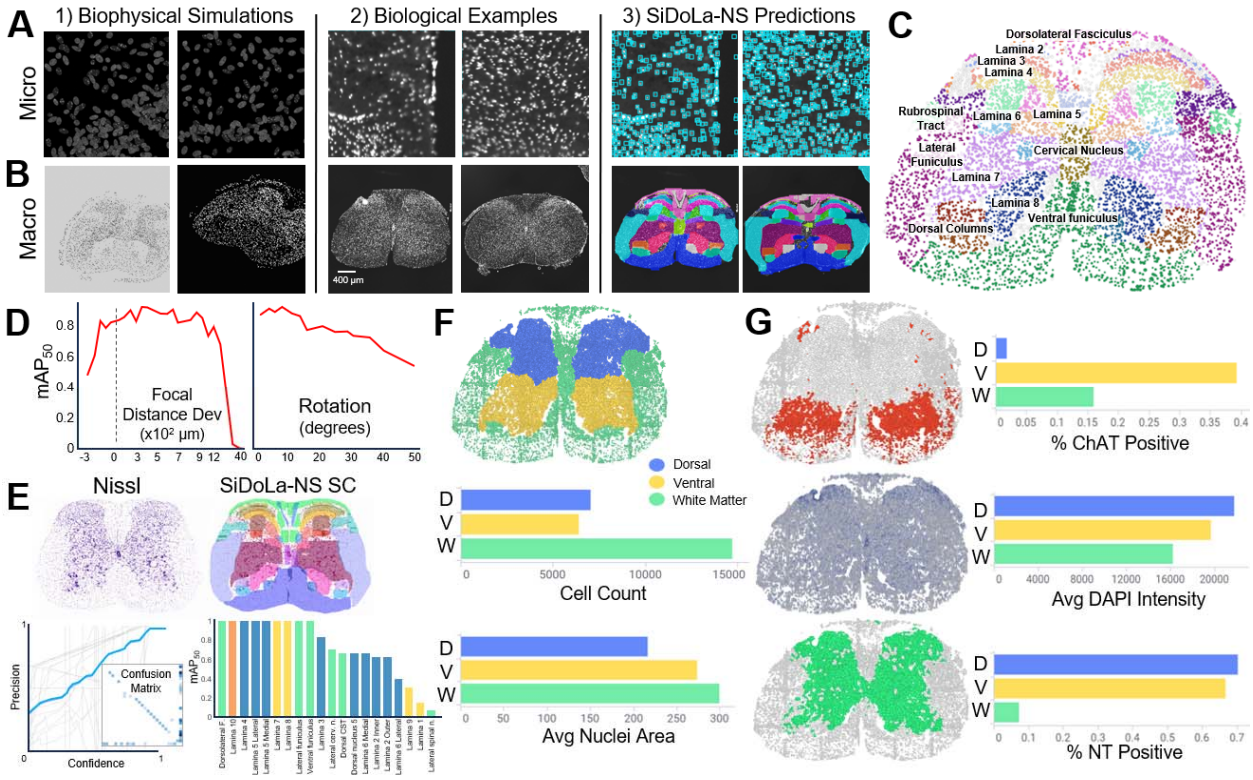

Figure 3. Segmenting Nuclei and Mapping Atlas Regions in a Mouse Spinal Cord Section. A) At the micro scale, SiDoLa-NS trains on simulated data (1). Micrographs of tiles from cross sectioned spinal cords from dataset[65] (2). SiDoLa-NS-micro-CNS predictions (blue squares) of nuclei in the spinal cord (3). **B)** At the macro scale, SiDoLa-NS trains on simulated images of spinal cord sections (1). Micrographs of mouse spinal cord cross sections. SiDoLa-NS-macro-mSC prediction on different spinal cord regions (different colors) (3). **C)** Cellular and regional results on the whole spinal cord section. Marker size is governed by the measured cell size, and coloration is by regions, some of which are labeled. **D**) Line

graph indicating performance of SiDoLa-NS-macro-mSC at different focal distances and tilts. Y-axis shows mean average precision at an IoU of 50%, with the left x-axis measuring the delta focal distance, where the dashed line is centered at 0, the normal camera distance, where the spinal cord takes up the majority of the frame. As the focal distance increases, the camera is backed away making the image smaller in the frame. The right x-axis displays degrees of tilt along the y-axis, rotating the camera around to the right side, skewing the sample. **E)** Micrograph of a Nissl section from the Allen Brain Mouse Spinal Atlas (section #875) on the left, and the same image with SiDoLa-NS-macro-mSC predictions overlayed (different colors are different regions). Below left, line graph of a Precision/Confidence curve, all classes at 0.96 precision with a confidence of 0.993. Inset, confusion matrix, and bottom left, a bar chart showing the mean average precision. All comparing SiDoLa-NS-macro-mSC to Allen Brain Atlas defined regions. **F)** Cellular representation of the same spinal cord section from C) but colored to define 3 main regions within the spinal cord (blue dorsal horn, yellow ventral horn, and green is white matter). Below are bar charts depicting the cell count and average nuclei area per region. **G)** The cellular spinal cord representation, but now the cells are colored to depict, on the top: ChAT+ cells in red, with the bar chart measuring the % of cells positive in each region, middle: the intensity of DAPI staining (higher intensity is a darker shading), and bottom: NeuroTrace (NT)+ cells in green.

2.4 Axonal and Nuclear Segmentation within Sciatic Nerve Fascicles

Few resources exist for the segmentation of peripheral nerves. Specifically, we sought to detect axons and Schwann cell nuclei as well as the fascicles that contain them within a cross section of a human-sized sciatic nerve. Here, we evaluated SiDoLa-NS's ability to detect axons, nuclei and meso-scale features in the nerve. In both H&E (hematoxylin and eosin) and DAPI staining, it is possible to discern a distinction between axons, which are larger, lower intensity structures, and the nuclei of Schwann cells, which are smaller and brighter. Thus, we created simulations (7,198 training and 612 validation images, 2.78×10^5 objects in the training set) at the cellular (micro) scale with simple geometry (**Fig. 4A 1**). We trained SiDoLa-NS-micro-PNS to explicitly segment these axons and nuclei and distinguish them as unique classes. The model achieved a perfect precision at a confidence of 0.894, and an F1 score of 0.9 at a confidence of 0.4 on simulated validation images (**Fig. S5**). We next used a publicly available cross section of a porcine sciatic nerve, which is composed of bundles of axons and glial nuclei (**Fig. 4A 2**). When inferring SiDoLa-NS-micro-PNS on the nerve, it detects the axons and nuclei within the tissue section (**Fig. 4A 3**).

While mouse, human, and pig have little difference at the cellular scale inside a sciatic nerve, the peripheral nerves in the larger mammals are composed of many fascicles (bundles) packed with axons and Schwann cells and are surrounded by perineural and epineural tissue. Most neurodegenerative diseases are specific to one class of neurons, for example sensory neurons, like in CMT2A[67] or motor neurons like in ALS[29] and the fascicles are impacted differentially due to their composition. We asked whether we could automate fascicle segmentation and then assign the cells inside to specific fascicles. We created meso-scale simulations for images of fascicles as ellipsoid objects with varying size, degree of overlap, and fascicle/background intensity contrast (**Fig. 4B 1**) with 6,588 training and 304 validation images (89,074 training objects). This was then trained with instance segmentation, and the result was termed SiDoLa-NS-meso-pSN. On simulated validation images, this model performed even better, with F1 reaching 1.0 for confidence values between 0.1 and 0.9 (**Fig. S6**). On the section of porcine sciatic nerve (**Fig. 4B 2**), single fascicle predictions were identified as ellipsoids in the tissue (**Fig. 4B 3**).

Next, we ran inferences over the porcine sciatic nerve section, on the micro and meso scales, where nuclei, axons, and fascicles are all fused in one resulting dataset. At the cellular level, we confirmed the expected variation in the cell area between axons and Schwann cells, where axons are larger on average (**Fig. 4C**). At the meso-scale, SiDoLa-NS-meso-pSN detected fascicles throughout the tissue slice, and unique fascicles were identified as numbered 'instances', where each cell within also had a feature indicating to which fascicle instance it belonged (**Fig. 4D**). We next asked what the variation in the number of axons or nuclei was in each fascicle and found that most porcine sciatic nerve fascicles had close to 100 axons, and about 70 nuclei, consistent with hand counting (**Fig. 4E**). Additionally, we noticed that fascicles on the lower left of the section had more nuclei, while the ones on the top right tended to have a lower percent of nuclei (data not shown). Just from an H&E nerve section, SiDoLa-NS-meso-pSN allowed for basic analyses of the variation between and within fascicles.

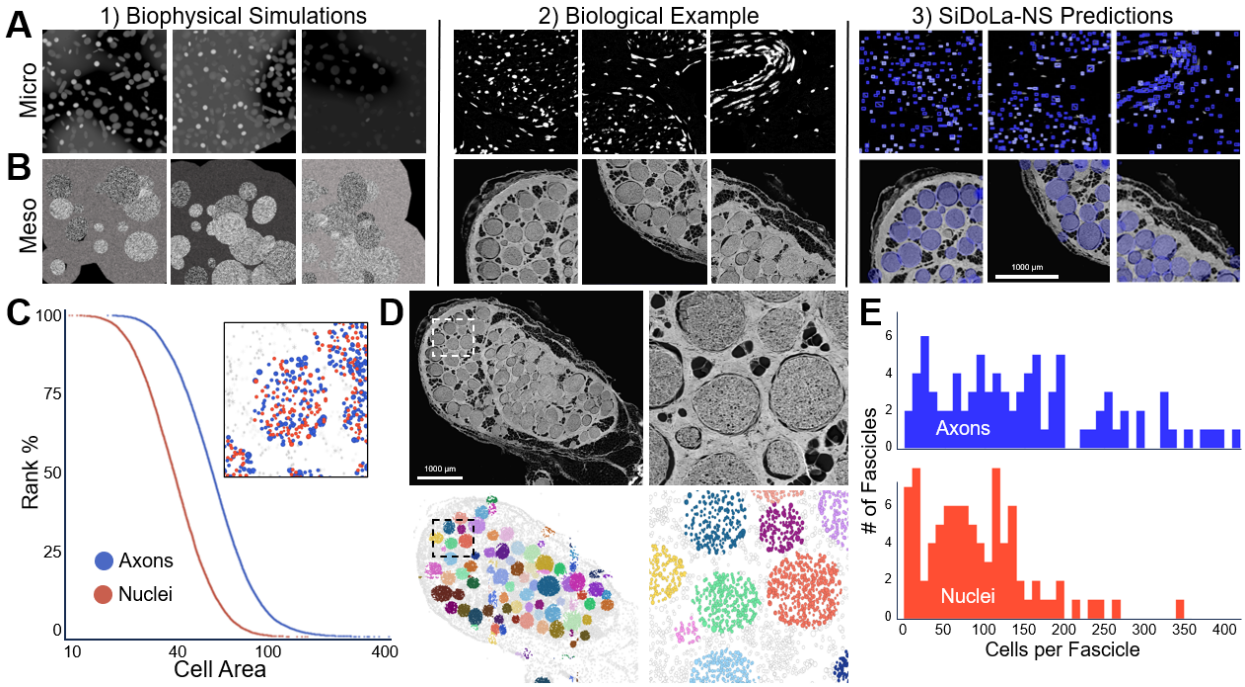

Figure 4. **Porcine Sciatic Nerve with Automated Fascicle Identification and Axon/Nuclei Discrimination.** **A)** Micrographs of a simulated cross section (1) of sciatic nerve with high intensity nuclei and dim axons. (2) Example tiles of a cross section of a porcine sciatic nerve and (3) SiDoLa-NS-micro-PNS prediction on axons (dark blue) and nuclei (light blue) on the same sections. **B)** Micrographs of simulated images (1) of textured bundles (fascicles) and perineurium using varying background contrast. (2) Micrographs of a cross section of a porcine sciatic nerve. (3) SiDoLa-NS-Meso-pSN prediction of fascicles on porcine sciatic nerve section (blue circles). **C)** Ranked histogram showing distribution of axon (blue line) and nuclei (red line) relative pixel areas within all regions of the porcine sciatic nerve section shown. Inset is a small region of the nerve showing non-fascicle-contained cells (gray), axons (blue), and nuclei (red). **D)** Micrograph showing overview of entire sciatic nerve, with SiDoLa-NS-Meso-pSN fascicle detection below. The dashed box indicates the area shown in higher magnification on the right. Circular markers are axons or nuclei, and their color indicates different fascicle instances. **E)** Histograms depicting the cell contents within fascicles. The blue histogram represents axons, and the red histogram represents nuclei. Fascicles with >= 10 are shown except for one outlier with 740 axons not shown.

2.5 Interrogating RNA Expression patterns and Cellular Morphology using Top-Down Analysis

Thus far, gene expression patterns have been used only to validate SiDoLa-NS's predictions at the macro (reference) scales. The top-down analysis after SiDoLa-NS allows for downstream investigations via measuring features across the SiDoLa-NS predicted regions. These features can be used, for example to define markers for distinct tissue regions. It can also be used to better compare gene expression changes within specific regions amongst cases and control sections.

When running the SiDoLa-NS inference, we can link gene expression data to add gene counts to the feature list per cell and region. In the mouse brain, the 10X Visium HD expression data was automatically linked to each cell, and we interrogated RNA expression patterns across cells and regions predicted by SiDoLa-NS. We identified a list of region-specific genes, ones that had the highest ratio ($region_x$ divided by the average counts of all the regions excluding *x*) and checked their expression across the atlas regions **(Fig. 5A**). Most of these genes were also quite specific (strong signal along the diagonal), but a few had some off-diagonal hotspots between related regions. For example, *Ltih3*, *Calb2,* and *Agt* were all differentially expressed in the hypothalamus and hypothalamic-adjacent regions. We also compiled a list of the top 4 highest ratio genes (**Fig. 5B**), highlighting known markers such as *Tshz2* in the Retrosplenial area and *Neurod6* in the hippocampus in addition to less commonly used marker genes like *Prkcd* in the ventral group of the dorsal thalamus. A larger, hierarchically clustered heatmap shows additional genes of interest per regions (**Fig. S7**). As expected, these genes maintain distinct regional localization in the coronal mouse brain (**Fig. 5C**). Additional gene localization patterns for *Tac1*, *Penk*, *Mobp*, *Acta2*, and *Prox1* are also visualized[68] (**Fig. S8**).

Next, we asked whether we could use morphological differences at the level of a cell (micro-scale), to find differences in the atlas regions within the mouse spinal cord section. The SiDoLa-NS inference automatically measures per-cell features for each available fluorescent channel. Indeed, we found that 'NT cell texture' (NeuroTrace intensity standard deviation amongst pixels per cell) arose, amongst others, as a strong segregator of spinal cord region (**Fig. 5D**, KW p~0, H-Stat = 6139). In particular, *Lamina 8* and the *lateral funiculus* have distinctly different NT textures (**Fig. 5E**). Therefore, with only image-based information processed at two different scales, SiDoLa-NS-micro-CNS and SiDoLa-NS-macro-mSC enabled spinal cord region distinction utilizing subtle morphological features.

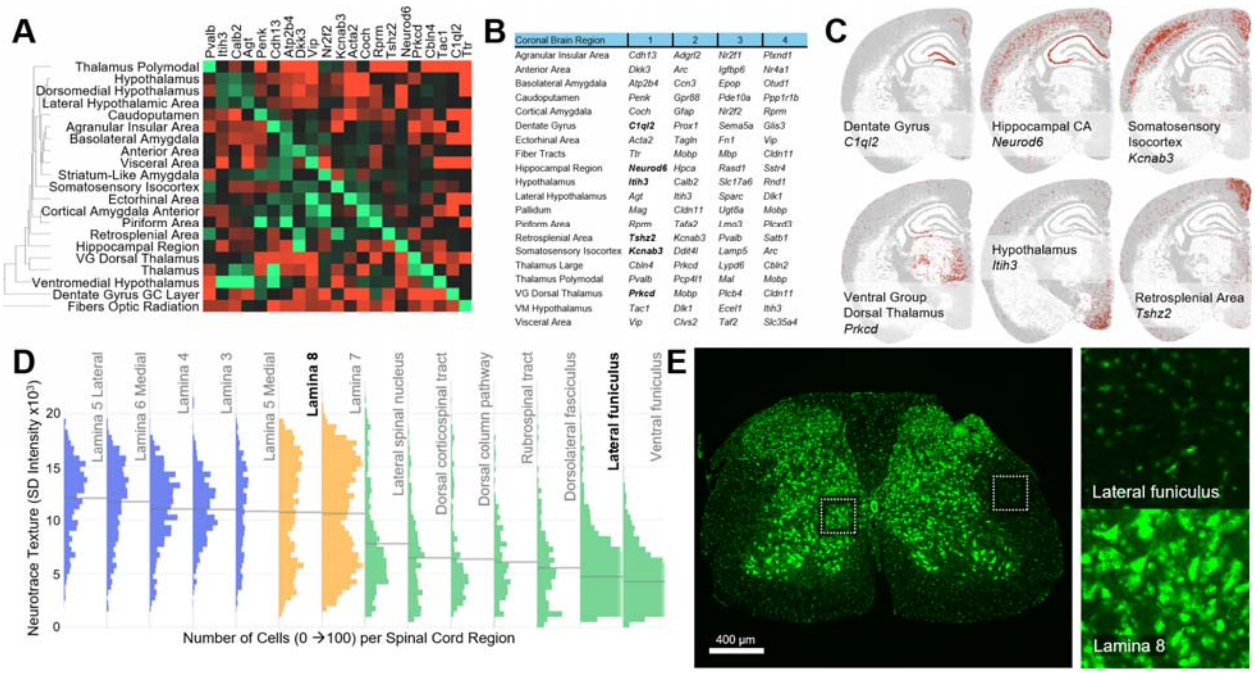

Figure 5. Top-Down Genes for Regional Marking and Utilization of Image Features to study Anatomical Regions in Mice. A) Heatmap showing the expression of the 'top ratio' genes for 21 regions from the mouse cortical Visium HD brain slice (green is high average counts, red is low). **B)** A list of the top four genes per brain region, where each gene was ranked by ratio of average count per region compared to the average counts in all other regions. Bolded genes are detailed in the next panel. **C)** Diagram of individual cells segmented with SiDoLa-NS-micro-CNS, colored by total counts per cell for the gene specified (gray = 0, black = mean, red = max. **D)** Histograms showing the distributions of NeuroTrace texture (measured as standard deviation intensity) per cell, then separated by Spinal Cord regions. Horizontal lines show the mean for that region. Regions with low cell numbers were excluded. **E)** Micrograph of Spinal cord stained with NeuroTrace, where insets show a region of white matter (lateral funiculus) compared with Lamina 8, each with different NeuroTrace textures.

2.6 Simulation Engine for Multi Scale Nervous System Image Generation

The exclusive use of simulated microscopy images allows for generalizable identification of the nuclei or cell bodies, organized cellular clusters, and reference regions from anatomical atlases. However, these systems are only as capable as the simulated images from which they are trained. We asked whether we could develop a framework to produce *biophysically informed* simulations that would be compatible with the spatial omics platforms discussed in Figure 1. The pipeline below was used to create the simulations used in this manuscript but could also serve as a baseline from which to build other synthetic datasets for anatomical regions of interest.

Continuing the multiscale theme, this pipeline operates at Micro-, Meso-, and Macro- scales (**Fig. 6A-C**). Each biological component is built as a 3D structure with a variety of properties defined and bounded by input parameters[53]. These components can be thought of as a function, and we represent them as a *node* with input parameters and outputs that return the geometry of that component (**Fig. 6D**). So, we first define the features of single cells (*Cell Node*), such as size, shape, and intensity, all of which can vary based on cell type and scan parameters and can be controlled through input parameters. The geometry for these (unique) individual cells is then used as an input to the next block (*Region Node*), which defines the makeup of a region or cluster (the details for the processes in the pipeline are defined in more detail below in Figure 6E and are indicated by a banner along the top). Cells are scattered throughout this region, with parameters for density and regional variation. The shape of these regions is either simulated generically, for example as laminar stripes, clusters, rings or is generated from a set of polygons which define a region from an anatomical atlas. Specifically, for sciatic nerve fascicles, we used rings of cells around a central ellipse of cells. For mesoscale inside the brain, we used higher density lamina that traverse across an image region as a stripe. Next, the geometry from each region is outputted and used to assemble the whole section (*Section Node*) that defines what will be a whole output image. Here, the parameters are randomized over a uniform random distribution as to change the atlas morphology, so that it includes normal physiological variation but also goes beyond to encompass a variety of pathological phenotypes.

We utilized the open-source CAD software Blender to produce the simulations, since it can be controlled through python scripting and with a visual generative workflow, called Blender Geometry Nodes (BGN), and has extensive optics and physics capabilities. The BGN pipeline for a micro-scale simulation focuses on cellular details (**Fig. S9**). A simplified version of the BGN pipeline that was used to create the macro mouse coronal brain dataset features transforms to distort the regions and overall section (**Fig. 6E**). The banner displayed over the pipeline defines the primary functions of each part of the pipeline: Regions, Cells within Region, Edge, Tears, Skew, Deform, and Empty. For the macro-scale reference regions, we defined their boundaries as exact copies of the reference polygons for one hemisection, and then had a switch to mirror those regions into a full set of hemispheres or just to use a single hemisection. The density of cells within the region was defined from a table (which we could manually set) or varied randomly between regions. Varying cell density randomly has the benefit of reducing the model's dependence on density in learning region labels. Each region was produced iteratively, indicated under 'Cells within Region'. The geometry from the regions and cells contained within were then joined together to define the whole section.

From here, the geometry of the simulated tissue section was manipulated. Specifically, we include within the BGN 1) Edge, which are extra cells and debris that are beyond the defined regions of the atlas (as there is often extra chunks of tissue that show up in a tissue section from errors in histological slicing), 2)

Tears, which are meant to simulate rips or tears in the tissue. These two features are disabled when generating ground truth, as we want to reconstruct the tissue as if these imperfections didn't exist. Additionally, we include two distortions, 3) Skew, which simulates different cut angles in sections and 4) Deform, which changes the relative size of the regions as would occur in samples from different individuals (the variation included purposely goes beyond common variation). Finally, there is a block that regulates 5) Empty, or background simulations, to reduce the influence of background variation and false positives in object recognition. More examples of these manipulations are shown in **Fig. S10, S11**. Although not shown in Figure 6, shading information is included to control how light is emitted from these objects and how it can interact with each object is also controlled using Blender Shader Nodes to enable variation and patterning.

At this point, the complex 3D geometry needs to be transformed into a 2D image, and the steps for that process are displayed below (**Fig. 6F**). To do this, we use optics and the rendering engine within Blender. Specifically, we employ a virtual camera with a 15.5 mm focal length and a low f-stop (0.01) to focus on the tissue section from about 1.4 mm away. Then the rendering engine 'Cycles' calculates the light being emitted or bouncing off the samples and uses it to produce an image. After the image layer is produced, we use a compositor layer to add varying amounts of white noise (0-95%), making the image slightly harder to interpret. This process is repeated (per 'frame') for the total number of images and produces the simulated 'microscope' images for our dataset. At the same time, the software also calculates the boundary coordinates of the various objects within the scene from the camera's point of view, and projects them onto a 2D plane so they are scaled to the rendered image. These coordinates are used to export the ground truth data, a set of polygon coordinates and class labels. For the ground truth construction, a switch is flipped that turns off some of the variation, such as the rips and the noise, so a clean reconstruction can be learned from the noisy microscope image. Also, during this process, all the parameters utilized are saved as metadata and are exported alongside the image and label files, so all the details can be tracked for different rendering runs and versions. In the macro mouse brain set, we rendered 4,065 images/hour on multiple desktops, with ~700 image sets an hour on a computer with only an 8th generation Intel CPU or NVIDIA 1000 series graphics card (i.e. standard, not specialty hardware).

The parameter input settings are constant for a single 'frame' of the BGN pipeline but are varied for different frames. Therefore, when producing the dataset, the software runs through one to hundreds of thousands of frames, creating the combinatorial variability of the training set. At the same time, we also produced a small, simulated validation dataset (hundred to thousand images). These images can then be fed into any number of image-based neural networks, such as a convolutional UNet architecture, instance segmentation architectures such as Mask-RCNN or YOLO[69], or used with Vision Transformers (ViT). In this manuscript, we presented data from YOLO instance segmentation, and neural architecture search (NAS) was used to find the optimal settings. The metadata exported during rendering is loaded during training so that it can be combined with model metadata and later utilized during inference.

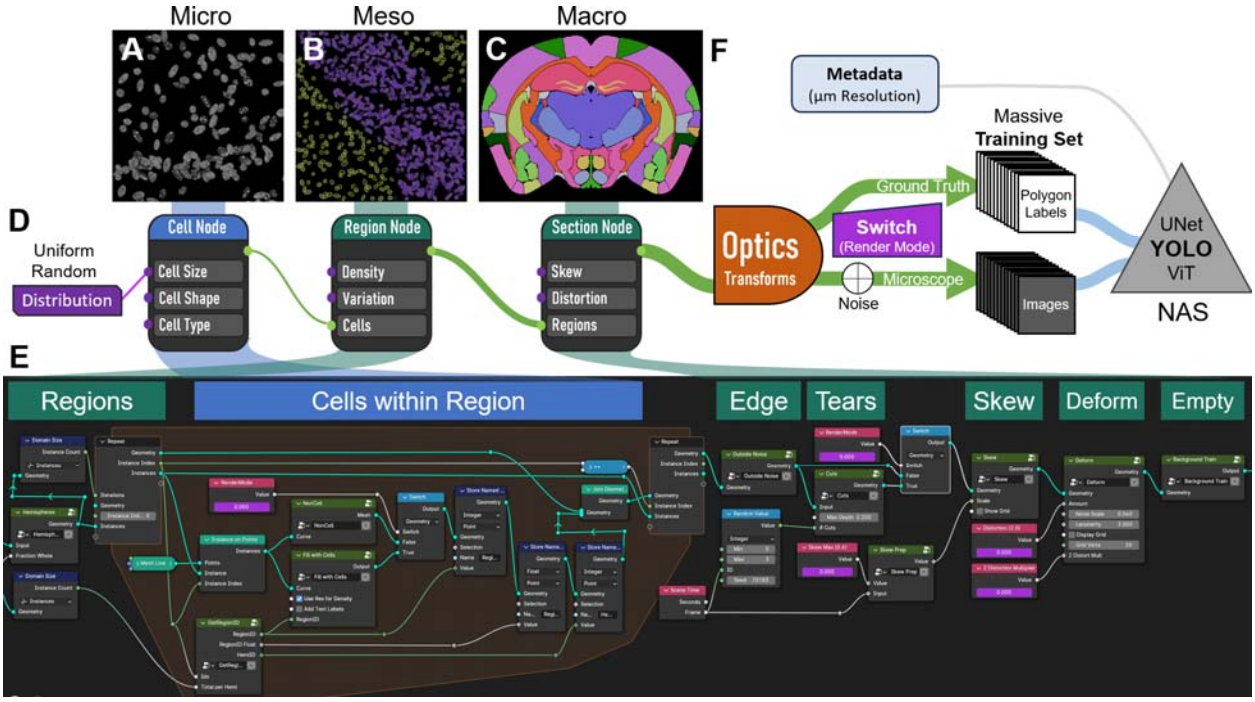

Figure 6. Simulation Engine for Creation of Synthetic Microscope Images and Corresponding Labels. The figure is organized to show the different modules used to produce the datasets for training, where **A-C** display micrographs and illustrations of the **A)** Micro-, **B)** Meso-, and **C)** Macro-scale representations of, as an example, a coronal section of a mouse brain. **D)** Three "nodes" that represent the main concepts within the engine. These connect to the images above and the pipeline within Blender Geometry Nodes that execute each of the scales below. The simulation pipeline in **D** starts with a series of uniform random distributions, each of which control input parameters that define *Cell*, *Region*, and *Section* properties. The random seed is different for each distribution input, so that combinations of parameters are achieved. Within each node, several example parameters are listed, such as Cell *size, shape,* and *type*, with a circle on the left-hand side, indicating it is an *input* node. On the right-hand side, there is a green circle indicating an output node that passes 3D *geometry* on to the next node. The connections between the nodes increase in diameter, indicating that the geometry is increasingly complex as the smaller components are assembled into larger ones. **E**) The Geometry Nodes pipeline within Blender that produces the coronal brain and spinal cord simulations. Polygon regions that define a section from an atlas are the input, and then the central repeat zone (with a translucent orange background) iterates through each region. The center of this iteration defines the geometry of the individual cells that are placed therein. Then, section parameters are set, including *Edge*: placing debris and cells around the outside of the section, *Tears*: which simulates rips or tears, Skew and Deform which distort the section, and finally *Empty*: which occasionally includes empty/background scenes. **F**) The created geometry is run through an optic and physics engine resulting in an image. For each image set, two versions are created, a *ground truth* version and a *microscope* version. The switch that controls these changes parameters throughout the workflow, including the addition of noise at the end. The images are rendered along with corresponding ground truth polygons and labels. Metadata is also saved alongside the images that includes information about the pixel resolution and all the settings used. This dataset is then used to train a neural network, usually YOLO instance segmentation using neural architecture search (NAS).

3 Discussion

Histological images, with or without omics data attached, are an incredible resource, and are themselves a dataset. Extracting useful features from these biological images has been the goal of countless projects since Leeuwenhoek and those like Jane Marcet that first started popularizing microscopic views of cells and tissues. Now, images carry vast additional dimensions, coupling the pixels with transcriptomics or proteomics, and they are also likely to expand to all types of omics and analyses[70]. The need to use features that have been defined for centuries by pathologists and histologists like Virchow, Malpighi, and Jacobi is of great importance and interest. This methodology aims to share a top-down approach to analysis that can be used in conjunction with all the amazing discoveries that spatial omics techniques afford.

Top-down simply means overlaying information that is gained from the image itself, rather than building up that knowledge from a lower molecular level (bottom-up). The heart of SiDoLa-NS is the concept that training a neural network on the available real images is insufficient. The variation in images that have been labeled is great, but the variation of a new section that needs to be analyzed is even greater. As with other advances in machine learning, this task can only be mastered by a neural network when it has the chance to observe vast amounts of examples and a high degree of variation. Therefore, this technique uses "**Si**mulate, **Do**n't **La**bel".

The vanguard models that we create with this approach are used in three distinct neurological cases, all with publicly available data. First, we infer cells, lamina, and atlas regions in a coronal slice of an adult mouse brain. While this works well, there are other approaches in the literature which can accomplish this task[48,71]. Next, we transitioned to two unique applications. A mouse spinal cord was segmented into its individual cells, and its reference atlas overlaid, identifying over 20 unique and identifiable regions accurately. In addition, we ran SiDoLa-NS on a cross section of sciatic nerve, which has distinctive morphology, including bundles of axons and Schwann cell nuclei organized into fascicles. Mice only have one or two of these fascicles, so we used a publicly available porcine section as a closer approximation of a human nerve. We demonstrate that these datasets can be used to identify marker genes when coupled with spatial transcriptomics. Additionally, we can use the multi-scale data to ask questions like how spatial distribution of neuronal cells in micro-scale correlate to the organization of spatial sub-regions of the corresponding tissue. Finally, we lay out the simulation process so that others can take advantage of the pipeline for their own applications.

The potential of SiDoLa-NS and other similar approaches is that it allows classifiers to become more and more general. Using machine learning out-of-context is fraught with problems, due to hallucinations and overfitting. One reason is because the model is able to 'memorize' the training set, and the other is related to variations in the real sample because of significant batch effects. Since we can simulate above and beyond the range of variation seen in existing samples, the neural network is more adept at handling variation in a completely new context. In fact, the only time the models in this manuscript are used is out-of-context since the training and validation consists *only* of simulated data. Another subtle but important note is that these generations are using an engine that is completely independent from training and validation. It is not using a GAN (generative adversarial network) or a diffusion approach to create the training set, which can lead to recursive generated collapse[72]. Other advantages of this approach include the generation of massive scales of training data, and the ability to overrepresent difficult and rare cases. For example, much of the mouse brain has nuclei spaced well apart, so existing

cell segmentation models can handle these regions easily. But some regions, like the granule cell layers have extremely high cell density, causing other models to fail. With simulations, we can overrepresent these phenotypes, making the SiDoLa-NS models robust to high cell density when appropriately large bounds in image parameters are set (**Figure S2**). Another important advantage of this approach is eliminating annotator bias, and that the quality of the ground truth annotations is much higher. For example, in the task of segmenting, the boundaries are now perfectly defined, since they are created from the simulation system at any resolution and the object bounds are defined before optical aberrations and sample imperfections, whereas manual annotations suffer from the vagaries of diffuse boundaries. Finally, since the simulation engine is based on biophysical properties (such as cells, clusters, and anatomically defined regions), its results are instantly and highly interpretable.

With these spatial-omics datasets, beware of double-dipping. Double dipping means using the data in a circular way, such as using molecular markers to find regions and then finding molecular signals in those regions. The issue this is that the assumptions from each approach are compounded and, although it is useful to enhance top-down image analysis with bottom-up clusters, the statistical significance of the results are falsely inflated. Using top-down regional definitions can help avoid these problems[73–75].

The power of SiDoLa-NS lies in the generation of the datasets. We chose to use those datasets with specific types of neural networks, but it is likely that additional exploration in that space would lead to even better inference. The models used in the manuscript are *instance* segmentation and/or *instance* classification models that are built on the YOLO[69] Ultralytics framework. Instance segmentation is powerful since it does not just find the region, but it understands the difference between different regions and their identities. We also worked with MaskR-CNN[76] and Feature Pyramid networks[77] but found YOLO to have the most active development and efficiency. Still, YOLO and instance segmentation may not be the right approach for all applications. For example, anatomical region definitions involve non-overlapping areas, which is not an architectural constraint for instance-based methods. Instead, convolutional UNets perform better at these tasks, since each pixel's location and identity are more consistent. In fact, we found that for some of these applications, a two-stage approach can enhance performance. First, UNets are trained to 'recode' the raw image with pixel-based classifications. Then, an instance segmentation model is trained from those, leading to better performance (data not shown).

We also have several improvements that are actively in development. During inference, the user must define the scale for which the various models should be run on the image. The next version of the inference system automatically checks the pixel resolution connected with the model and lines it up with the input image. Also, in terms of scale, additional simulations are being produced at varying focal distances to improve model performance with other microscope objectives or starting image resolutions. One major shortcoming of the current version is that the macro models are trained on only a single section from the reference atlas. While it is trivial to add additional sections into the training set, the next version will move beyond that. We are assembling all the slices in the atlas and using a system for virtual reconstruction into three-dimensional geometry[64]. Then, we can slice the atlas regions at any position, and at any angle. This will drastically improve the ability in multiple anatomical regions, in addition to helping deal with common histological issues like asymmetries in microtome slices. The current version has been extensively trained on both immunohistochemistry-like Nissl images, and immunofluorescence-like DAPI images. Although current models can handle H&E, we plan to explicitly include these stains. The mouse brain, spinal cord, and peripheral nerves were chosen as important nervous system regions that could benefit from SiDoLa-NS. The upcoming versions will include additional

nervous system tissue, including transverse sciatic nerve sections and dorsal root ganglia[78] (both at micro- and meso-scales).

4 Methods

Atlas Mapping

Resources from the Allen Software Development Kit (SDK) were utilized for the atlas mapping pipeline of the mouse brain and spinal cord. Each atlas image had a respective structure image with color-coded regions. The structure image for each atlas image was retrieved and a dictionary for regions and their respective color code was constructed. Next, binary masks in each color code were constructed over all structure maps to visualize the unique regions. Polygon coordinates defining the contours of each region were produced and assembled into a JSON file, organized by slice number and region name. These region outlines were used to generate macro-level brain and spinal cord simulations.

Simulations Overview

All simulations were done in the CAD software Blender 4.0+ using a feature called 'Geometry Nodes' which allowed for a visual workflow that encoded the geometry representing different scales of biology. Individual 'frames' were used to control variations and produce the different training and test image datasets needed for teaching the neural networks. While each tissue type and scale used distinct geometries, a few general rules were common: The basic concept was to build a plane or volume to fill with cells, then to 'instance' cell objects and randomly fill them into the region. Both the regions, the fill density/pattern, and the cells were all created dynamically, and random variations were added for each frame. The more variation that was shown to the neural network, the better it was able to generalize on real tissue.

SiDoLa-NS-macro-CNS Mouse Coronal Brain and Spinal Cord

These sections were both produced with the same pipeline explained in Figure 6. After atlas mapping, the regions were imported into Blender as curve objects using a custom script. These were fed into the Geometry Nodes pipeline described extensively in Figure 6E. Additional details and images are found in Figure S10 and Figure S11.

SiDoLa-NS-Meso-pSN Porcine Sciatic Nerve

Sciatic nerves were digitally generated using Blender with geometry nodes to simulate biophysical properties. The aim was to replicate natural structures within a sciatic nerve, including fascicles and Schwann cells. The Blender simulations were simplified to basic geometric forms, including ovular fascicles, circular Schwann cell nuclei, and circular epineurium. Additionally, varying sizes and counts of fascicles were distributed to ensure a manageable, yet flexible, simulation for training. This effectively prepared the model for real-world data. These simulated components were meticulously rendered into a set of .bmp files through Blender scripting, which produced the dataset.

SiDoLa-NS-micro-CNS Mouse Coronal Brain and Spinal Cord

The basis of this simulation was to instance ellipsoid nuclei on a large number of points across a plane. The nuclei varied in size, aspect ratio, rotation, and shading (which featured bright heterochromatin spots). We also introduced two other meso-scale features. The first of which was a bright laminar stripe or striation. These were bands of much higher density nuclei compared to the adjacent cells (**Fig. S9**). The other feature was a region devoid of cells, usually thinner than the high-density lamina, that also crossed the image. Finally, a slight amount of depth-of-field with a low f-Stop was used to simulate slight changes in focus.

SiDoLa-NS-micro-PNS Mouse Peripheral Nervous System

This version of the micro model is the original from which others were derived. It consisted of a plane that was filled with points that were split between two different cell type instances, each of which could be scattered in different ratios. One cell type was Schwann cell nuclei, which are more round, small and bright. The other type was the axons, which are slightly larger, often oblong, and dim to near invisible (by DAPI staining). The plane which contained these was circular shaped and had additional geometry outside of it representing the epineurium. The plane was also shifted around the frame of the camera so that occasional different edges of the perineurium were exposed.

Optics, Noise, and Ground Truth Polygons

A virtual camera within Blender was used to translate the three-dimensional geometry into an image. This was done with the rendering engine “cycles” or “EEVEE” and involved calculating where the emitted light from the fluorescent objects fell or where light from the background in an IHC setting was blocked by the objects. The camera had depth-of-field and a small aperture to give slight out-of-focus regions. The camera was also tilted and twisted to give different viewing angles and raised or lowered to adjust the overall field-of-view and zoom. White noise was added on top of the final image with the compositor and mixed at different ratios to make the model more robust to training. In addition to the standard rendering, we also extracted the coordinates of the vertices making up the 3D geometry, and projected them onto a plane from the camera’s point of view, converting them into camera x,y coordinates (world to camera). These vertices were then reassembled into polygons directly, or convex hull was applied. In more complex setups (atlas regions), we also downsampled the resolution of the polygons so that the exported segmentation files were more manageable. These files were jointly exported in Yolo.txt or coco.json format automatically as the images were rendered.

Model Training

The images created could be used with a variety of image-based models. All of the inferences in this document utilized Ultralytics YOLO v8 or v9 models trained in a python environment with CPU or GPU. For micro models, where the exact outline of the nucleus was not required, we trained ellipsoid nuclei segmentations into bounding boxes with YOLO v9c and v9e. For more complex shapes, like in the macro models, we trained with YOLO v8n-seg and b8l-seg models to extract the full polygon masks from the inferences. The number of images used in training is listed throughout the results. Generally, the training was for 10-200 epochs, the image size was between 448 and 640 pixels in width and height, and the batch size was between 8 and 32 images per batch. All these factors were varied, along with the random seed, for each NAS run. The results of the run were copied into a folder that was named by the resulting mAP_{75} on the validation dataset, and then inferred on biological examples as a quick reference.

Transcript Processing

All raw data was acquired from public sources. The 10X Visium HD Brain was provided at 2, 8, and 16 µm resolution in the binned outputs folder. We used pandas to export the tissue parquet files as CSV files, acquiring the SRT data in micron coordinates.

Data Collection and Image Preprocessing

Mouse Coronal Brain

The Visium HD Brain H&E image was acquired from 10X Genomics. The provided microscopic image was an H&E-stained section. Utilizing ImageJ, we split these images into its constituent channels and selected the channel that best highlighted nuclear features (red channel) for further processing. For both cases, we adjusted the image brightness and contrast, applying a linear normalization to the whole image. For Macro/Meso/Micro inference, these images were directly given to the program, without downsampling or tiling.

Mouse Spinal Cord

For the spinal cord analysis, we acquired the IF images from[65]. The spinal cord dataset included nine different cross sections of spinal cords, each containing three different channels of ChAT, Nissl, and DAPI stains. We selected *Slide1-6_Region0007_Channel555 nm,475 nm,395 nm_Seq0051* and labeled each spinal cord from one to nine starting from the top and moving horizontally. Taking the second spinal cord, we cropped it into a square and split it into three separate channels using ImageJ. The Nissl channel was then used for our SiDoLa-NS model prediction.

Porcine Sciatic Nerve

The pig sciatic nerve slide was obtained from Saarland University's histology site https://mikroskopie-uds.de/[79]. The sciatic nerve cross section had a scan area of 6.2 mm x 4.8 mm and H&E staining. The image was downloaded as a .zif file, which was then loaded into ImageJ and the highest resolution version of the sciatic nerve was exported as a 38885 x 30286 pixel 8-bit image. The channels were then split. Both channel 1 and channel 2 highlighted the fascicles within the sciatic nerve. Due to channel 2's more distinct contrast between the fascicles and the background, channel 2 was used for Figure 4. Channel 3 highlighted the axons within the fascicles. The pixels were then inverted to create a black background with white fascicles and adjusted to enhance the visibility of the fascicles and axons against the background and to mimic DAPI staining.

Inference

SiDoLa-NS must work on all parts of the image at different scales, and then that information needed to be fused together, along with any other omics data that corresponded with the image dataset. This was accomplished with the Macro/Micro inference package included in the code repository. The input allowed a single Micro-scale model to do the cell segmentation, and then allowed as many Meso- or Macro-scale models to be included as desired. Some parameters were set up for each model such as the location of the images, and then the software ran all the inferences and linked the data together. For the micro scale, the multiple channels available from the image were measured across three masks: Box, Poly, and Voronoi (which represent the bounding box, the full polygon outline if available, and the space all around the cell until the next cell's boundary is reached). Area, Size, Intensity, and Texture (SD) were measured for each mask and reported on a per-cell basis. If there was more than one Macro-scale model, then hard and soft-voting ensembles were generated and reported.

Analysis

Model Evaluations

Models were evaluated during training utilizing the Ultralytics *validation* tools. For all our models, the validation set used during training consisted of simulated images. For each validation image, many

metrics including precision, recall, IoU, and confidence were recorded. Precision and recall were used to generate an F1 score ($\frac{precision}{recall}$) as well as graphed across confidence values. Additionally, a confusion matrix was generated for visualization of classification accuracy, particularly in multi-class tasks like brain and SC atlas mapping.

Robustness Analysis

Various robustness analyses were conducted on the macro-level models for noise, degree of rotation, and focus. We generated a small set of 18 – 25 images with calibrated variations in one of the three ways. For noise, we altered the image resolution by reducing the pixel contrast. For degree of rotation and zoom, the camera was respectively rotated or moved along the z-axis relative to the reference object. We also adjusted the F-number for the camera's aperture for focus. For each of those images, we additionally generated a ground truth, which was used in model evaluation. The mAP_{50} was primarily utilized in assessing model performance.

Marker Gene Analysis

Because SiDoLa-NS-macro-mCB and SiDoLa-NS-micro-CNS were inferred directly on the VisiumHD mouse brain slice, we mapped the gene expression matrix at the 16 µm bin level to the data based on micron position. This resulted in each detected cell having a brain region assignment and a known gene count. For the list of the 'top ratio' genes, we ran a Kruskal Wallis test including all genes against brain regions to determine the highest rank of genes. From there, the ratio of each gene's count to the total region gene counts allowed us to rank these genes per region, yielding our list of the top 4 genes.

UMAPs and Clustering

The UMAPs were viewed using the Seurat embeddings from the VisiumHD dataset at the 16um bin level. The UMAP1 and UMAP2 data was plotted and colored by cluster. For the cluster definitions, we utilized the provided graph-based clusters, which we joined to the UMAP with the barcodes. The SiDoLa-NS UMAP in Figure 2G uses the same underlying projection data but was colored with the SiDoLa-NS-macro-mCB predictions for the respective barcode region.

Acknowledgements

Thanks to Jane Dodd at Columbia and Felix Fiederling for their paper providing the spinal cord images. Thank you to Rudolf Bock at Saarland University for their virtual microscope histology site that provided the porcine sciatic nerve cross section. We would also like to thank Serena Elia, Graham Bachman, Purva Patel, Jimin Lee, Josh Milbrandt, and Saul Weiss, for their help with aspects of the project. Additionally, we would like to thank the Milbrandt, DiAntonio and Mitra labs for their continued support. We would like to thank the **McDonnell Genome Institute**, and the rest of the FIVE@MGI lab members. Finally, we would like to thank the Genetics and MGI administration and our maintenance and cleaning staff.

Author Contributions

L.K.MA, A.K.Y.Y, T.N., N.K.T., E.G., J.E.W, C.L.K., and W.J.B wrote the manuscript. W.J.B., L.K.MA designed the experiments. Initial planning and ideas brought by W.J.B., J.D.M., R.D.M., A.K.Y.Y and L.K.MA. L.K.MA, A.K.Y.Y, T.N., N.K.T performed the experiments. L.K.MA, A.K.Y.Y, T.N., N.K.T, and W.J.B. analyzed the results. Key Software was developed by W.J.B, L.K.MA, and N.K.T. All reviewed the paper and gave suggestions.

Data and code availability

https://gitlab.com/buchserlab/sidola-ns

References

1. Okano, H., Hirano, T. & Balaban, E. Learning and memory. *Proceedings of the National Academy of Sciences* **97**, 12403–12404 (2000).

2. Roh, E., Song, D. K. & Kim, M.-S. Emerging role of the brain in the homeostatic regulation of energy and glucose metabolism. *Exp Mol Med* **48**, e216 (2016).

3. Eugene, A. R. & Masiak, J. The Neuroprotective Aspects of Sleep. *MEDtube Sci* **3**, 35–40 (2015).

4. Ganapathy MK, Reddy V & Tadi P. Spinal Cord Morphology. in *Neuroanatomy* (StatPearls Publishing, Treasure Island, 2022).

5. Kovacs, G. G. Cellular reactions of the central nervous system. in 13–23 (2018). doi:10.1016/B978-0-12-802395-2.00003-1.

6. Jaume, G. *et al. Deep Learning-Based Modeling for Preclinical Drug Safety Assessment*. (2024).

7. Schmauch, B. *et al. A Deep Learning-Based Multiscale Integration of Spatial Omics with Tumor Morphology*. (2024).

8. Piwecka, M., Rajewsky, N. & Rybak-Wolf, A. Single-cell and spatial transcriptomics: deciphering brain complexity in health and disease. *Nature Reviews Neurology* vol. 19 346–362 Preprint at https://doi.org/10.1038/s41582-023-00809-y (2023).

9. Sonawane, A. R. *et al.* Understanding Tissue-Specific Gene Regulation. *Cell Rep* **21**, 1077–1088 (2017).

10. Vickovic, S. *et al.* High-definition spatial transcriptomics for in situ tissue profiling. *Nat Methods* **16**, 987–990 (2019).

11. Ståhl, P. L. *et al.* Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science (1979)* **353**, 78–82 (2016).

12. Lein, E., Borm, L. E. & Linnarsson, S. The promise of spatial transcriptomics for neuroscience in the era of molecular cell typing. *Science (1979)* **358**, 64–69 (2017).

13. Maniatis, S. *et al.* Spatiotemporal dynamics of molecular pathology in amyotrophic lateral sclerosis. *Science (1979)* **364**, 89–93 (2019).

14. Asp, M., Bergenstråhle, J. & Lundeberg, J. Spatially Resolved Transcriptomes—Next Generation Tools for Tissue Exploration. *BioEssays* **42**, (2020).

15. Benjamin, K. *et al.* Multiscale topology classifies cells in subcellular spatial transcriptomics. *Nature* (2024) doi:10.1038/s41586-024-07563-1.

16. Chen, X. *et al.* Whole-cortex in situ sequencing reveals input-dependent area identity. *Nature* (2024) doi:10.1038/s41586-024-07221-6.

17. Zhang, D. *et al.* Spatial transcriptomics and single-nucleus RNA sequencing reveal a transcriptomic atlas of adult human spinal cord. *Elife* **12**, (2024).

18. Li, X. *et al.* Profiling spatiotemporal gene expression of the developing human spinal cord and implications for ependymoma origin. *Nat Neurosci* **26**, 891–901 (2023).

19. Russ, D. E. *et al.* A harmonized atlas of mouse spinal cord cell types and their spatial organization. *Nat Commun* **12**, (2021).

20. Yadav, A. *et al.* A cellular taxonomy of the adult human spinal cord. *Neuron* **111**, 328-344.e7 (2023).

21. Ortiz, C. *et al.* Molecular atlas of the adult mouse brain. *Sci Adv* **6**, (2020).

22. Yao, Z. *et al.* A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain. *Nature* **624**, 317–332 (2023).

23. Zhang, M. *et al.* Molecularly defined and spatially resolved cell atlas of the whole mouse brain. *Nature* **624**, 343–354 (2023).

24. Bhattacherjee, A. *et al.* Spatial transcriptomics reveals the distinct organization of mouse prefrontal cortex and neuronal subtypes regulating chronic pain. *Nat Neurosci* **26**, 1880–1893 (2023).

25. Salas, S. *et al*. *Optimizing Xenium In Situ Data Utility by Quality Assessment and Best Practice Analysis Workflows*. (2023).

26. Tian, L., Chen, F. & Macosko, E. Z. The expanding vistas of spatial transcriptomics. *Nat Biotechnol* **41**, 773–782 (2023).

27. Moses, L. & Pachter, L. Museum of spatial transcriptomics. *Nat Methods* **19**, 534–546 (2022).

28. Wang, P. L. *et al.* Peripheral nerve resident macrophages share tissue-specific programming and features of activated microglia. *Nat Commun* **11**, 2552 (2020).

29. Yim, A. K. Y. *et al.* Disentangling glial diversity in peripheral nerves at single-nuclei resolution. *Nat Neurosci* **25**, 238–251 (2022).

30. Russell, A. J. C. *et al.* Slide-tags enables single-nucleus barcoding for multimodal spatial genomics. *Nature* **625**, 101–109 (2024).

31. Marshall, J. L. *et al.* High-resolution Slide-seqV2 spatial transcriptomics enables discovery of disease-specific cell neighborhoods and pathways. *iScience* **25**, 104097 (2022).

32. Hu, Y. *et al.* Unsupervised and supervised discovery of tissue cellular neighborhoods from cell phenotypes. *Nat Methods* **21**, 267–278 (2024).

33. Hu, J. *et al.* SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network. *Nat Methods* **18**, 1342–1351 (2021).

34. Long, Y. *et al.* Deciphering spatial domains from spatial multi-omics with SpatialGlue. *Nat Methods* (2024) doi:10.1038/s41592-024-02316-4.

35. Pham, D. *et al.* Robust mapping of spatiotemporal trajectories and cell–cell interactions in healthy and diseased tissues. *Nat Commun* **14**, 7739 (2023).

36. Long, Y. *et al.* Spatially informed clustering, integration, and deconvolution of spatial transcriptomics with GraphST. *Nat Commun* **14**, 1155 (2023).

37. Chitra, U. *et al.* Mapping the Topography of Spatial Gene Expression with Interpretable Deep Learning. in 368–371 (2024). doi:10.1007/978-1-0716-3989-4_33.

38. Petukhov, V. *et al.* Cell segmentation in imaging-based spatial transcriptomics. *Nat Biotechnol* **40**, 345–354 (2022).

39. Wu, X. *et al.* Spatial multi-omics at subcellular resolution via high-throughput in situ pairwise sequencing. *Nat Biomed Eng* (2024) doi:10.1038/s41551-024-01205-7.

40. Zhang, B., Zhang, S. & Zhang, S. Whole brain alignment of spatial transcriptomics between humans and mice with BrainAlign. *Nat Commun* **15**, 6302 (2024).

41. Dries, R. *et al.* Advances in spatial transcriptomic data analysis. *Genome Res* **31**, 1706–1718 (2021).

42. Halpern, K. B. *et al.* Single-cell spatial reconstruction reveals global division of labour in the mammalian liver. *Nature* **542**, 352–356 (2017).

43. Bankhead, P. *et al.* QuPath: Open source software for digital pathology image analysis. *Sci Rep* **7**, (2017).

44. Takko, H. *et al.* ShapeMetrics: A userfriendly pipeline for 3D cell segmentation and spatial tissue analysis. *Dev Biol* **462**, 7–19 (2020).

45. Klein, S., Staring, M., Murphy, K., Viergever, M. A. & Pluim, J. elastix: A Toolbox for Intensity-Based Medical Image Registration. *IEEE Trans Med Imaging* **29**, 196–205 (2010).

46. Modat, M. *et al.* Fast free-form deformation using graphics processing units. *Comput Methods Programs Biomed* **98**, 278–284 (2010).

47. Dunn, K. W. *et al.* DeepSynth: Three-dimensional nuclear segmentation of biological images using neural networks trained with synthetic data. *Sci Rep* **9**, (2019).

48. Carey, H. *et al.* DeepSlice: rapid fully automatic registration of mouse brain imaging to a volumetric atlas. *Nat Commun* **14**, (2023).

49. Cisternino, F. *et al.* Self-supervised learning for characterising histomorphological diversity and spatial RNA expression prediction across 23 human tissue types. *Nat Commun* **15**, 5906 (2024).

50. Ma, C., Tan, W., He, R. & Yan, B. Pretraining a foundation model for generalizable fluorescence microscopy-based image restoration. *Nat Methods* (2024) doi:10.1038/s41592-024-02244-3.

51. Zhang, D. *et al.* Inferring super-resolution tissue architecture by integrating spatial transcriptomics with histology. *Nat Biotechnol* (2024) doi:10.1038/s41587-023-02019-9.

52. Wang, Y. *et al.* FTUNet: A Feature-Enhanced Network for Medical Image Segmentation Based on the Combination of U-Shaped Network and Vision Transformer. *Neural Process Lett* **56**, (2024).

53. Sekh, A. A. *et al.* Physics-based machine learning for subcellular segmentation in living cells. *Nat Mach Intell* **3**, 1071–1080 (2021).

54. Naseer, M., Bhatti, I. T., Hasan, O. & Shafique, M. Considering the Impact of Noise on Machine Learning Accuracy. in *Embedded Machine Learning for Cyber-Physical, IoT, and Edge Computing* 377–394 (Springer Nature Switzerland, Cham, 2024). doi:10.1007/978-3-031-40677-5_15.

55. Lang, O. *et al.* Using generative AI to investigate medical imagery models and datasets. *EBioMedicine* **102**, (2024).

56. Tudosiu, P.-D. *et al.* Realistic morphology-preserving generative modelling of the brain. *Nat Mach Intell* **6**, 811–819 (2024).

57. Tyson, A. L. & Margrie, T. W. Mesoscale microscopy and image analysis tools for understanding the brain. *Prog Biophys Mol Biol* **168**, 81–93 (2022).

58. Goubran, M. *et al.* Multimodal image registration and connectivity analysis for integration of connectomic data from microscopy to MRI. *Nat Commun* **10**, 5504 (2019).

59. Oh, S. W. *et al.* A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).

60. Harris, J. A. *et al.* Hierarchical organization of cortical and thalamic connectivity. *Nature* **575**, 195–202 (2019).

61. Daigle, T. L. *et al.* A Suite of Transgenic Driver and Reporter Mouse Lines with Enhanced Brain-Cell-Type Targeting and Functionality. *Cell* **174**, 465-480.e22 (2018).

62. Allen Reference Atlas – Mouse Brain [brain atlas]. Available from atlas.brain-map.org. Preprint at (2011).

63. Allen Mouse Brain Atlas [dataset]. Available from mouse.brain-map.org. Preprint at (2004).

64. Wang, Q. *et al.* The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **181**, 936-953.e20 (2020).

65. Fiederling, F., Hammond, L. A., Ng, D., Mason, C. & Dodd, J. Tools for efficient analysis of neurons in a 3D reference atlas of whole mouse spinal cord. *Cell Reports Methods* **1**, (2021).

66. Lein, E. S. *et al.* Genome-wide atlas of gene expression in the adult mouse brain. *Nature* **445**, 168–176 (2007).

67. Sato-Yamada, Y. *et al.* A SARM1-mitochondrial feedback loop drives neuropathogenesis in a Charcot-Marie-Tooth disease type 2A rat model. *Journal of Clinical Investigation* **132**, (2022).

68. Cembrowski, M. S., Wang, L., Sugino, K., Shields, B. C. & Spruston, N. Hipposeq: a comprehensive RNA-seq database of gene expression in hippocampal principal neurons. *Elife* **5**, (2016).

69. Redmon, J., Divvala, S., Girshick, R. & Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 779–788 (IEEE, 2016). doi:10.1109/CVPR.2016.91.

70. Zhang, Y. & Imoto, S. Genome analysis through image processing with deep learning models. *J Hum Genet* (2024) doi:10.1038/s10038-024-01275-0.

71. Stringer, C., Wang, T., Michaelos, M. & Pachitariu, M. Cellpose: a generalist algorithm for cellular segmentation. *Nat Methods* **18**, 100–106 (2021).

72. Shumailov, I. *et al.* AI models collapse when trained on recursively generated data. *Nature* **631**, 755–759 (2024).

73. Chen, H., Li, D. & Bar-Joseph, Z. SCS: cell segmentation for high-resolution spatial transcriptomics. *Nat Methods* **20**, 1237–1243 (2023).

74. Chen, K. H., Boettiger, A. N., Moffitt, J. R., Wang, S. & Zhuang, X. Spatially resolved, highly multiplexed RNA profiling in single cells. *Science (1979)* **348**, (2015).

75. Polański, K. *et al. Bin2cell Reconstructs Cells from High Resolution Visium HD Data*. (2024).

76. He, K., Gkioxari, G., Dollar, P. & Girshick, R. Mask R-CNN. in *2017 IEEE International Conference on Computer Vision (ICCV)* 2980–2988 (IEEE, 2017). doi:10.1109/ICCV.2017.322.

77. Lin, T.-Y. *et al.* Feature Pyramid Networks for Object Detection. in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 936–944 (IEEE, 2017). doi:10.1109/CVPR.2017.106.

78. Tavares-Ferreira, D. *et al.* Spatial transcriptomics of dorsal root ganglia identifies molecular signatures of human nociceptors. *Sci Transl Med* **14**, (2022).

79. Bock, R., Lang, I. & Schmitt, J. Virtual Microscopy. *University of Saarland*.