

Connectomic analysis of mouse barrel cortex and fly optic lobe

Dissertation der Fakultät für Biologie der Ludwig-Maximilians-Universität München

vorgelegt von Kevin Michael Boergens

München, 29. Januar 2016

Erstgutachter: Prof. Dr. Alexander Borst
Zweitgutachter: Prof. Dr. Gerhard Haszprunar
Tag der mündlichen Prüfung: 2. Februar 2018

Abstract

While neuronal circuits are thought to be at the heart of how brains function, their description still remains incomplete for most brain regions. Specificity patterns of innervation have not been properly mapped and even definitions of cell types remain partially ambiguous. Progress is hindered by the lack of readily available techniques that can efficiently record and annotate whole neuronal circuits.

The most advanced way to map neuronal circuits is using three-dimensional electron microscopy. A setup was created using an existing technique of this kind (serial block-face electron microscopy) and then improved to allow faster imaging. In a second step, the setup was considerably extended to remove overhead times between imaging steps. Additionally, an extraction technique was developed to create samples of mammalian neocortex that are centered on a region of interest. These technical improvements help creating larger and more targeted 3D electron microscopy datasets. Nevertheless, annotating such datasets is a challenge. A novel online tool called *webKnossos* was developed that allows large groups of annotators to collaborate on the analysis of 3D datasets, utilizing its efficient online transmission of raw data. Furthermore, *webKnossos*' new interaction mode, "Flight Mode" improves the speed with which neurites can be reconstructed up to 10 fold over existing tools.

Two datasets were acquired, one from layer 2/3 and one from layer 4 of mouse barrel cortex. Those datasets were used to measure how specific axons innervate target structures. Upon analyzing the layer 4 dataset, inhibitory axons were identified that in fact targeted structures with above-random specificity. These targets were apical dendrites, initial segments and the proximal part of spiny stellate dendrites. A group of axons was identified that innervated apical dendrites 27 times more frequently than what would be expected with random innervation. In layer 2/3, the high target specificity of chandelier axons was confirmed and new types of axons innervating pyramidal cell initial segments were described.

An existing dataset from the visual system of *Drosophila* was converted into a format that could be annotated with *webKnossos*. This dataset was used to identify lobula plate

tangential cells, a group of cells specialized on aggregating visual flow information.

Combining the reconstructions from many annotators, several lobula plate tangential cells were discovered that had not been described before in this species.

This study provides valuable groundwork for the field of connectomics to meet new challenges such as complete reconstructions of the visual system of flies and of whole circuits in mouse cortex. It delivers tools for creating larger datasets and circuit maps, and describes new and unexpected circuit rules in mouse barrel cortex and circuit elements in the *Drosophila* visual system.

Contents

1	Introduction	9
1.1	General motivation	9
1.2	Imaging requirements to detect connections between neurons	9
1.2.1	Spatial resolution requirements	9
1.2.2	Volume requirement.....	10
1.3	Electron microscopy.....	10
1.4	Imaging and staining	11
1.5	Sample preparation	12
1.6	Cell morphology reconstruction	13
1.6.1	The reconstruction challenge	13
1.6.2	Reconstructing axons is most critical.....	13
1.6.3	Manual versus automatic annotation.....	14
1.6.4	Semi-automated annotation.....	14
1.7	Redundant annotation.....	15
1.7.1	The RESCOP algorithm	16
1.7.2	Focused reannotation	16
1.8	Innervation specificity.....	17
1.8.1	Specificity calculations	18
1.9	Barrel cortex.....	19
1.10	Three-dimensional annotation tools	20
1.11	The lobula plate of <i>Drosophila melanogaster</i>	21
1.12	Speed considerations for SBEM imaging	24
1.12.1	Imaging.....	25
1.12.2	Cutting.....	25
1.12.3	Mechanical movement	26
1.12.4	Comparison of factors.....	26
1.13	Stitching and alignment	27

1.14	Structure of the thesis	28
2	Results.....	29
2.1	Innervation specificity in mouse barrel cortex	29
2.1.1	Datasets	29
2.1.2	Reconstruction reliability.....	30
2.1.3	Initial segment innervation	38
2.1.4	Specificity of apical dendrite innervation in L4.....	43
2.1.5	Spatial distribution of inhibitory innervation.....	46
2.2	webKnossos.....	49
2.2.1	Data transmission speed.....	49
2.2.2	Flight mode	51
2.2.3	Ease of teaching for different interaction modes.....	53
2.2.4	Speed advantages of flight mode	54
2.2.5	Redundant flight mode annotations.....	58
2.2.6	Fast synapse annotation in webKnossos	59
2.2.7	Fast spine reconstruction in webKnossos.....	60
2.2.8	Connectome strategies	63
2.3	Analysis of a <i>Drosophila</i> lobula plate SBEM dataset.....	65
2.3.1	Stitching and alignment	65
2.3.2	Annotation of lobula plate tangential cells.....	67
2.4	Continuous imaging	73
2.4.1	Tiling and continuous imaging	73
2.4.2	Piezo setup.....	74
2.4.3	Scanning pattern	75
2.4.4	Autofocus for large samples	76
2.4.5	Full barrel stack	76
3	Methods.....	82
3.1	Mouse barrel cortex analysis methods.....	82

3.1.1	Sample preparation	82
3.1.2	Electron microscopy.....	83
3.1.3	Cutting thickness.....	84
3.1.4	Annotation tools	84
3.1.5	Traceability.....	84
3.1.6	Synapse detection.....	85
3.1.7	Axon type detection.....	86
3.1.8	Initial segment detection	86
3.1.9	Apical dendrite detection	87
3.1.10	Inhibitory cell identification	87
3.1.11	Specificity calculations	88
3.1.12	Volume reconstruction and Merge Mode	90
3.2	Image stitching and alignment.....	92
3.2.1	L4 and L2/3 dataset.....	92
3.2.2	Improved L4 dataset for webKnossos experiments.....	93
3.2.3	The Lobula plate dataset.....	93
3.2.4	Speed improvements	95
3.3	webKnossos methods	96
3.3.1	Dataset	96
3.3.2	Data delivery	96
3.3.3	Data transmission test	96
3.3.4	Axon seeding.....	97
3.3.5	Speed measurement process reconstruction.....	98
3.3.6	Dendrite reconstructions	99
3.3.7	RESCOP algorithm	100
3.3.8	Error annotation.....	100
3.3.9	Synapse annotation mode	100
3.3.10	Time consumption for synapse annotation	101

3.4	Methods for continuous imaging.....	102
3.4.1	Setup of electronics	102
3.4.2	Control of the imaging setup	103
3.5	Methods for full barrel stack	106
4	Discussion.....	108
4.1	Axons in mouse barrel cortex target subcellular structures specifically	108
4.2	<i>webKnossos</i> is a novel tool for online 3D EM data annotation	111
4.3	Reconstruction of lobula plate tangential cells in <i>Drosophila</i> yields candidates for cells not previously identified	113
4.4	Continuous imaging reduces overhead time in serial block-face SEM dataset acquisition.....	114
5	Contributions	116
6	References	117
7	Acknowledgements.....	129

1 Introduction

1.1 General motivation

Brains of complex animals contain a large number of neurons (human: 86 billion neurons (Herculano-Houzel, 2009), mouse: 75 million neurons (Williams, 2000)). Even more staggering than the number of neurons within the brain is the number of synapses that these neurons can form with each other (for example every cortical pyramidal cell forms synapses with about 1000 - 10,000 partners (Braitenberg and Schüz, 1991; Helmstaedter, 2013)). Groups of neurons and their connections are the building blocks of neuronal networks, which are thought to be at the core of how the brain performs computations. While the field of neuroscience has gained enormous insights into the morphology and physiology of single cells (and the average pairwise connectivity and activity patterns of small sets of cells), only recently it has become possible to map connections within a neuronal circuit on a large scale. The key feature of these novel techniques is that they are not limited to statistical statements about connectivity but that for one tissue sample a large number of neurons and their synapses can be simultaneously observed. The field that is concerned with the measurement and analysis of such connectivity datasets has been named “connectomics”.

1.2 Imaging requirements to detect connections between neurons

Two neurons can only form a synapse if their neuronal processes touch. A way to measure if two cells are connected synaptically is to image the morphology of the two cells, to identify all spots where the two cells come into contact and to locally assess whether the contact spots are connections. If the morphologies of all cells and all synapses within a volume are being imaged simultaneously (“densely”), this method allows for the simultaneous identification of many synapses.

1.2.1 Spatial resolution requirements

The smallest processes encountered in mammalian cortex are about 40 nm (spine necks) and 50 nm (axons) (Helmstaedter, 2013). Other animals can have processes with even smaller diameters, e.g. 30 nm in *Drosophila* (Butcher et al., 2012). The processes are densely packed and can run next to each other (e.g. axon bundles). In many tissue types, there is also no preferred direction for thin processes.

Therefore the imaging requirements can be described as such: The resolution needs to be on the scale of the thinnest processes and such that two of these processes running next to each other can be discerned. A blurring of the image information, which would be unproblematic if only a single cell is labeled, e.g. (Markram et al., 1997), is not permissible (Helmstaedter et al., 2008). The neuropil is

usually isotropic and a lack of resolution in one dimension cannot fully be compensated for by better resolution in the other dimensions. Therefore in all three spatial dimensions the resolution has to be in the low tens-of-nanometers range.

1.2.2 Volume requirement

To build the full map of synaptic connections of a nervous system - its “connectome” - the nervous system has to be imaged entirely. At the necessary resolution, this has so far only been done once for adult animals: For the nematode worm *C. elegans* (White et al., 1986). Nonetheless, datasets of subregions of the brain can also be used to answer biologically relevant questions.

If all inputs onto a cell are to be mapped, its dendritic tree has to be imaged in its entirety. To evaluate if two cells form a synapse, one needs to image the whole volume where this synapse might occur, plus the parts of the neurites that lead to that volume. The volume that allows doing that for all pairs of cells in a circuit simultaneously is called its “minimal circuit volume” (Helmstaedter, 2013). (Kasthuri et al., 2015) imaged a single dendrite with its surround and were able to detect axons that form multiple synapses with the dendrite. The imaged volume was $40 \times 40 \times 50 \mu\text{m}^3$. Beyond the considerations of minimal circuit volumes, there are requirements for minimal volumes: To investigate how consistently an axon exhibits a certain innervation behavior, several of its synapses have to be within the dataset. Furthermore, neurons can have different types of dendrites, e.g. basal and apical dendrites. To be able to differentiate them by their morphology, a large enough stretch of neurite has to be imaged (see 3.1.9 for the criteria to identify apical dendrites).

1.3 Electron microscopy

Because the spatial resolution is so small and the volumes many orders of magnitude bigger, connectomic datasets are large. To acquire such datasets in reasonable timeframe, the microscopy has to be fast. Currently only electron microscopes can record data sufficiently fast. There are two types of electron microscopes, transmission electron microscopes (TEM) and scanning electron microscopes (SEM), both can deliver data at the necessary rates. These techniques are two-dimensional: TEM only operates on thin slices (typically 70-350 nm) because electrons traverse the sample. SEM only scans the surface of a sample and cannot penetrate substantially. The techniques have to be extended into the third dimension by other means. The techniques to do that are serial sectioning TEM, serial sectioning SEM and serial block-face SEM. For TEM, samples have to be sliced for imaging anyway and consecutive slices offer a 3-dimensional stack: This technique is called serial sectioning TEM (ssTEM) and has been used for several notable projects (White et al., 1986; Harris and Stevens, 1989; Harris et al., 1992; Fiala et al., 1998). Improvements to the technique have been

focusing on increasing the speed of imaging (TEMCA, (Bock et al., 2011)). These slices are manually transferred onto an electron-transparent tape. The slices are prone to rupturing and distortions (Briggman and Denk, 2006), so it is so far not possible to build stacks out of slices thinner than about 40-45 nm (Bock et al., 2011; Helmstaedter, 2013). As outlined in 1.2.1, this is not sufficient for correct reconstruction in some cases.

Because the manual handling is problematic, a novel approach called ATUM has been developed (Hayworth et al., 2006; Kasthuri et al., 2015) where the slices are transferred onto a tape automatically. Because this tape is not transparent to electrons, the imaging is done with an SEM, making this a serial sectioning SEM technique. Slice thicknesses down to 25 nm are possible.

Using an SEM, it is not necessary to image slices. Instead, the surface of the sample can be scanned without prior slicing. Then a thin layer of the top can be removed and the newly exposed surface can be imaged again. By repeating the process, over time a three-dimensional dataset is created. This technique is called Serial Block-face Scanning Electron Microscopy (SBFSEM or short SBEM). There are two options on how to remove the top layer of the block, either with a diamond knife (Denk and Horstmann, 2004) or with a focused ion beam (FIB-SBEM, (Knott et al., 2008)).

The main advantage of SBEM is that the top of the block is more mechanically stable than a single slice. Single slice techniques are prone to distortions, rupture and slice loss. Datasets from FIB-SBEM samples currently can only be acquired if the sample is smaller than 100 μm in the smallest dimension, but efforts are underway to remove this limitation (Hayworth et al., 2015).

1.4 Imaging and staining

Scanning electron microscopy requires a mechanism to give the images contrast. The two most important sources of signal in SEM are back-scattered electrons (BSE) and secondary electrons (SE). Back-scattered electrons are electrons of the imaging beam that have been Mott-scattered at atoms of the sample at such an angle that they return towards the column. They lose energy by creating secondary electrons. Back-scattered electrons are scattered back with their landing energy only slightly attenuated (Goldstein et al., 2012). Mott scattering is quadratically dependent on the atomic number (Z), therefore using high-mass atoms in staining gives better back-scattering contrast (Hennig and Denk, 2007).

Secondary electrons are electrons that are generated by the ionizing effect of the primary electrons. Secondary electrons have low energy, usually below 100 eV (Goldstein et al., 2012) and the image information is very superficial. Because secondary electrons are emitted more easily from exposed structures (corners, edges, ridges), the signal from secondary electrons is more sensitive to the

general shape of the sample (Goldstein et al., 2012). The signal from secondary electrons also contains information about the elementary composition of the sample, because strong scattering creates primary electrons that move laterally and superficially through the sample and elicit many secondary electrons that can escape from that superficial track (Titze, 2013). Secondary electrons are very susceptible to charging because of their low energy. Charging cannot be avoided in uncoated samples of neural tissue (Titze and Denk, 2013). Therefore in this piece of work imaging with back-scattered electrons is employed. To give differential contrast, the structures that shall be resolved are infused with atoms of high mass – for example osmium, lead and uranium.

1.5 Sample preparation

The high resolution of the datasets imposes strong requirements on the quality of the sample preparation. The first challenge is the protection of the sample from hypoxia. This can be achieved by intracardially perfusing the animal with a fixative. Furthermore, great care has to be taken when extracting the tissue. Even small forces can lead to visible ultra-structural damage (see 2.4.5.1). The actual sample preparation has to satisfy a number of requirements. The plastic (epoxy) in which samples are embedded must be hard enough so that the sample does not move away while the cutting force is applied and is sufficiently stable against electron beam damage, which changes the epoxy and makes it harder to cut (Starborg et al., 2013). A certain amount of beam damage is unavoidable and therefore the dose has to be limited to less than 20 electrons per nm², for more reliable cutting rather 10 electrons per nm² (Titze, 2013).

This can only be achieved with a sufficiently dense stain. The development of such a stain by combining several stains that existed before was therefore of critical importance (Briggman et al., 2011). The stain consists of several steps of heavy metal salt infiltration. The first step is a reduced osmium tetroxide infiltration, followed by a linking thiocarbohydrazide step, and a second non-reduced osmium tetroxide infiltration (Seligman et al., 1966; Karnovsky, 1971). The next step is an infiltration of uranyl acetate solution, followed by a lead nitrate step (Walton, 1979). Finally, the samples are dehydrated and prepared for epoxy infiltration (Briggman et al., 2011). The four steps of three different heavy metals together give sufficient contrast for imaging with 10 - 20 electrons per nm² at (10 - 15 nm)² pixel size, but this regiment remains technically challenging. For the data stacks presented in chapter 2.1 and chapter 2.4, several attempts were made before samples could be produced that combined sufficiently dense staining with good epoxy quality for reliable cutting. Also, this method yields a denser stain in mammalian cortex samples compared to *Drosophila* brain tissue samples, so a specialized *Drosophila* stain was developed by Ch. Kapfer for the dataset presented in chapter 2.3.

1.6 Cell morphology reconstruction

1.6.1 The reconstruction challenge

A cubic millimeter of brain matter contains about 4.5 km of neuronal processes (Braitenberg and Schüz, 1991)¹. Each process has to be followed (“traced”) to assign its volume to a specific neuron and to find its synapses. This constitutes a major challenge. If the process being traced is lost or misidentified (especially close to the soma), this can lead to the wrong attribution of thousands of synapses (Helmstaedter, 2013). If all neurites in the dataset have been traced, candidates for synapses between two processes can be identified by searching where the two processes come in close contact. Still, the image data at the contact spots has to be inspected to assess whether they are indeed synapses (Helmstaedter et al., 2013; Berning et al., 2015).

1.6.2 Reconstructing axons is most critical

The wiring consists (lengthwise) of 58% axonal wiring, 25% dendritic wiring and 17% glia processes (see 2.1.2.1, compare (Mishchenko et al., 2010)). To analyze connectivity, glia cells are not reconstructed. The dendritic reconstruction consists of two tasks: Annotating the shaft of the dendrite and annotating spines. The shaft reconstruction is relatively easy because the shafts have a large diameter. In mammalian cortex, the majority of dendritic wiring is formed by spines (Braitenberg and Schüz, 1991). Spine reconstruction also poses a challenge, because the spine neck diameters are so small (see 1.2.1) – but if the reconstruction of a spine fails, only the synapse connecting onto that spine is not detected.

For these reasons, reconstructing axons is the most critical part of the reconstruction challenge: They are very thin (at thin points down to 50 nm in diameter, see 1.2.1) and highly tortuous. Small and very local disruptions in imaging quality or small mistakes in the reconstruction process can immediately lead to an important branch of the axon being lost (Helmstaedter et al., 2011). Furthermore, axons generally do not become progressively thinner as they extend, so there is no built-in safeguard against losing important parts as there is with dendrites.

¹ If all neurites were circular, had the same diameter and would run densely packed in parallel, this would roughly correspond to a neurite diameter of 0.5 μm

1.6.3 Manual versus automatic annotation

Reconstructing a neurite not only includes following it through the dataset, but also measuring its variations of shape along the way. Ultimately this means assigning almost every pixel of the dataset to a neuronal process identity. This is called a “segmentation” mapping (Equation 1).

Equation 1
$$\Omega(x, y, z) = id$$

To assign voxels to neurons, tools have been created to manually annotate voxels (Kim et al., 2014; Harris et al., 2015; Kasthuri et al., 2015), but manually volume-annotating processes takes about 100 h of work per mm of neurite (Kasthuri et al., 2015), meaning that even a $(100 \mu\text{m})^3$ block would take 450,000 work hours.

An alternative approach is to develop an algorithm that assigns voxels to neuron identity without human help. To measure annotation quality by automated reconstruction, it is necessary to first define an error metric: Because the reconstruction is an attribution of voxels to processes, it would be suggestive to count the number of voxels that have been labeled wrongly. This metric however can be misleading: In the pursuit of reconstructing a neurite along its length, a certain amount of border voxels can fail to be assigned correctly without causing any problems (Turaga et al., 2009). Problems arise when a neurite is “split” into separate processes. Also problematic are “mergers”, which happen if two processes of different cells carry the same label. This rationale is efficiently captured in the split/merger metric. It measures the average distance along the centerline of a process before a split or merger error occurs (Turaga et al., 2010; Berning et al., 2015).

So far, attempts to build algorithms that can automatically segment whole cells have been unsuccessful. There have been several steps of improving the segmentation quality for various kinds of tissue and imaging method (Jain et al., 2007; Turaga et al., 2009; Jain et al., 2010; Turaga et al., 2010; Jain et al., 2011; Sommer et al., 2011; Kaynig et al., 2015). The one that is used in this thesis has an average inter-error distance for splits and mergers of 5.03 and 214.52 μm , respectively (Berning et al., 2015), which was a new record at the time of publication. Nonetheless, these split rates mean that an improvement of at least two orders of magnitude would still be necessary to reconstruct whole cells correctly.

1.6.4 Semi-automated annotation

Due to these constraints, neither a fully manual nor a fully automatic reconstruction is feasible at the moment in large datasets. However, the strengths of manual and automated annotation can be combined to achieve much longer error-free annotations (Helmstaedter et al., 2011). Human

annotation is slow because manually contouring the outlines of processes is time consuming (Berning et al., 2015). It is also unnecessary: As highlighted before, segmentation algorithms yield good results locally. Human annotators are most efficiently employed to traverse problematic splits where a continuation of the segmentation is not possible. To make use of this, a process for neuronal “skeletonization” was developed (Helmstaedter et al., 2011).

For skeletonization, the annotator follows the process through the dataset. The amount of annotation is greatly reduced; the annotator only marks the process in about every 10th layer by setting a single marker (“node”). Markers that are being set sequentially are connected with an “edge”. This annotation style is sufficient to reconstruct the path of processes and brings a speed improvement by a factor of 50 (Helmstaedter et al., 2011).

The skeletonization is a very sparse annotation and therefore not sufficient on its own to generate a list of places where the processes touch, but it can be combined with a computer-generated segmentation, especially a segmentation that is optimized to have very few mergers, like (Berning et al., 2015), by combining all segments that contain a marker of the same skeleton to create a full segmentation.

1.7 Redundant annotation

When several annotators reconstruct the same process (“redundancy”), the result will not always be identical. At times annotators connect parts that do not belong to the same process, end processes prematurely or miss branches that split off the main branch (Helmstaedter et al., 2011). When annotators are confronted with situations where their annotations differ, they are usually able to agree on one of their annotations. This makes it plausible that the majority of annotation errors are attention-related. Annotating dense neuronal tissue is strenuous and therefore the emergence of these errors is not surprising. It also means that multiple annotations can be pooled to create a better annotation quality.

With redundant annotations it is possible to consolidate multiple annotations so that errors are removed and only markers persist that correctly mark the process. However, the annotators are free to set markers anywhere within the process. For that reason two annotations that correctly mark the same process can look very different microscopically. This means that the consolidation process is more complex than a series of yes/no votes. Therefore, an algorithm called RESCOP was developed to automate this consolidation (Helmstaedter et al., 2011).

1.7.1 The RESCOP algorithm

The first part of the RESCOP algorithm looks at the set of edges in a given reconstruction and asks for each of those edges whether other annotators agree that the edge is correctly connecting two voxels of the same process. To do so, for each edge the number of annotators that followed the same path are considered agreeing votes and the number of annotators that annotated in the proximity but did not follow the path are counted as disagreeing votes (Figure 1). Annotators that did not set nodes in the vicinity at all are not counted one way or the other because their absence stems from an earlier disagreement.




Figure 1 Left: Overlay of neurite annotation done by 7 independent annotators. Right: Method for resolving disagreements. If neurites are annotated by placing points in them and connecting them with lines or “edges” (“skeletonization”), a consensus can be built from multiple annotations of the same neurite with the RESCOP algorithm. For each edge (here: dotted line in middle left panel) it is decided whether it belongs in the consensus by counting the skeletons that support its existence (red) or not support it (violet). Reprinted by permission from Macmillan Publishers Ltd: Nature Neuroscience, (Helmstaedter et al., 2011), copyright 2011

To determine which percentage of agreeing votes is necessary to consider an edge valid, the second part of the RESCOP algorithm compares the voting for a given edge to a baseline probability distribution that is specific to the dataset. For example, an edge can be considered valid if 2 out of 7 other annotators went along the same path if in this specific dataset correlated wrong edges (two annotators making the same mistake) are extremely rare. The baseline probability distribution can be also used to predict how fast the error rates would go down if the redundancy of annotation were to be increased. Because the second part of the RESCOP algorithm is independent of the specific annotation method it can be used for other redundant annotation setting where attention is a limiting resource for annotators, e.g. synapse annotation.

1.7.2 Focused reannotation

It was observed that for skeleton annotation the majority of decisions are very easy (Helmstaedter et al., 2011) and that the difficulty is concentrated in a few small regions. Therefore it is not very efficient to increase precision by reannotating the whole tree repeatedly. Instead it is preferable to

direct annotators specifically to places of high disagreement. This can be done efficiently in software that allows automatic task distribution like webKnossos (see 2.2).

1.8 Innervation specificity

A single axon has the ability to innervate multiple post-synaptic partners. In doing so, it can exhibit preferences for certain targets. One of the questions that this work addresses is how specific these connections are in the mouse barrel cortex. Innervation patterns are regularities encountered when analyzing neuronal connectivities. They can help to explain how neuronal circuits fulfill their task (Briggman et al., 2011; Borst and Helmstaedter, 2015). While some network motifs can only be detected when looking at a large number of cells, there are innervation patterns that can be detected by inspecting only a small number of cells in a network (Boergens et al., 2016).




Figure 2 Left: Possible innervation patterns, illustrated using example cell n4. Right: Matrix view of connections on cell n4 at different level of detail. A connectivity matrix is a compact way of displaying all innervation information of a cell. Depending on the details of the connectivity matrix, different patterns are visible: Neuron-to-neuron specificity, neuron-to-dendrite specificity and neuron-to-segment specificity seen in the respective connectivity matrices for example cell n4 (right panel, top to bottom). From (Boergens et al., 2016), reproduced by permission of Oxford University Press.

Innervation patterns are exemplified in Figure 2. The connectivity matrix $C_{N,N}$ shows one kind of innervation specificity. Here cell n10 (red axon) specifically innervates cell n4, but not cell n3 and cell

n5. This pattern can for example be observed for direction selective pyramidal cells in V1 that preferably connect to cells with the same direction preference (Ko et al., 2011).

A second kind of innervation pattern can be seen in the connectivity matrix $C_{N,D}$ (Figure 2, split up for dendrites): Axons can target specific parts of a neuron, e.g. soma or basal dendrites. For example, basket cells are known to preferably form synapses onto somata (Peters and Jones, 1984).

Splitting up the connectivity matrix further into the subregions of the dendrites ($C_{N,Dseg}$, Figure 2), a third kind of pattern becomes visible: Proximal and distal parts of the dendrite are targeted specifically. This pattern is implemented in the mouse retina, where bipolar cells of type 2 and 3 target the proximal and distal part of starburst amacrine cell dendrites (Kim et al., 2014).

1.8.1 Specificity calculations

The specificity of innervation patterns is not an all or none phenomenon - in reality, specificity measurements can be more nuanced. Some axons may preferably target a particular structure but not do all their innervations onto that structure: a baseline has to be found that defines “unspecific” behavior to be able to compare actually measured specificities.

This baseline can be expressed as Peters’ rule (Peters and Palay, 1991), a hypothesis that states that innervation is as unspecific as possible, that means that for all regions of interest the probability for a synapse between an axon type and target (= cell type) is proportional to how much wiring of the two structures is in the volume (Equation 2).

$$\int_{volume} p_{synapse}(x, y, z) \propto \int_{volume} WiringDensity_{axon}(x, y, z) \int_{volume} WiringDensity_{target}(x, y, z)$$

Equation 2

The validity of Peters’ rule may depend on the size of the volume. If it is assessed on a small volume, the probability of two processes forming a synapse will be proportional to their actual touching. This is roughly true for some tissue types (Helmstaedter et al., 2013) but it is not true in cortex (Takemura et al., 2015). If the volume is too large, the rule cannot be true because it gives non-zero connection probabilities for cells that do not even come in contact macroscopically. However, on intermediate scales – micrometers to hundreds of micrometers – Peters’ theorem generates testable hypotheses. It predicts that innervation should be totally nonspecific - whereas connectivity patterns would create violations of the rule.

1.9 Barrel cortex

In mammals different sensory modalities are represented in different areas of the cortex. A general organizing principle of each modality is the differentiation between early cortices where the information arrives and higher order cortices where the information is further processed. This is for example true for somatosensory cortex where these regions are called S1 and S2, respectively. Early sensory cortices retain the layouts imposed by the sensory modality, e.g. somatotopic maps, where areas that are close to each other on the body are represented close to each other in the cortex. This work focuses on the barrel cortex, a columnar structure within S1 found in rodents, which receives the primary sensory input from the whiskers on the animal's snout that help the animal navigate its environment. Within the barrel cortex, additionally a discrete organization exists. Units called barrels each primarily respond to information received from a specific whisker. It is remarkable that the layout of the whiskers on the snout of the animal corresponds to the same pattern in cortex, having the same number of rows and columns of barrels as rows and columns of whiskers. This region was first discovered using a Nissl stain (Figure 3, (Woolsey and Van der Loos, 1970)), which highlighted its segmented nature. Like the rest of the cortex, barrel cortex is organized into distinct layers that differ from each other in cell types and connectivity. These layers are numbered from 1 to 6, starting with layer 1 at the surface of the cortex. The barrel patterning is best visible in layer 4 of barrel cortex, where most of the fibers transporting the information from the whiskers (relayed over thalamus) terminate.




Figure 3: Nissl stain of layer 4 of mouse barrel cortex, showing how a soma density measurement can be used to distinguish single barrels, scale bar 100 μm . Reprinted from *Brain Research*, Volume 17, Issue 2, Woolsey and Van der Loos, The structural organization of layer IV in the somatosensory region (S I) of mouse cerebral cortex: The description of a cortical field composed of discrete cytoarchitectonic units, Pages 205-242, Copyright 1970, with permission from Elsevier.

The barrel cortex has become a model system for neuroscientific research. This is partly because of its existence in mice and rats, two of the most commonly used lab animals and partly because its

barrels are readily identifiable with several stains (Woolsey and Van der Loos, 1970; Land and Simons, 1985; Cooper and Steindler, 1986). The inputs are characterized (Agmon and Connors, 1991; Bruno and Sakmann, 2006; Meyer et al., 2010a), and there is a rich body of literature about connectivity statistics regarding cell types in barrel cortex (Feldmeyer et al., 1999; Beierlein et al., 2000; Koelbl et al., 2013) and single cell properties (e.g. (Korngreen and Sakmann, 2000; Waters et al., 2003)).

Many cortical neurons form synapses with cells that are several millimeters apart, thereby creating extensive networks. This makes connectomic analysis of cortical networks a task that requires very large datasets. Layer 4 of barrel cortex is a notable exception. Layer 4 of a single mouse barrel contains about 1650 excitatory cells ((Lefort et al., 2009), mostly “spiny stellate cells” that are morphologically distinct from the pyramidal cells of higher and lower levels, (Harris and Woolsey, 1979)) and 140 inhibitory cells (Lefort et al., 2009; Meyer et al., 2011). The majority of inputs those cells receive are either from thalamocortical axons or from each other (Lübke and Feldmeyer, 2007; Helmstaedter, 2013). The dendritic trees of the spiny stellate cells are much smaller than those of pyramidal cells, and they are oriented toward their respective barrel centers (Lübke et al., 2000). Furthermore, axons from barrel cortex spiny stellate cells leave their barrel and do not bend back to make further synapses. Because of this, barrel cortex is unique as a network having a small minimal circuit volume (see 1.2.2): only $300 \times 300 \times 300 \mu\text{m}^3$ have to be imaged to fully analyze this cortical circuit (Helmstaedter, 2013).

1.10 Three-dimensional annotation tools

Annotation tools for connectomic reconstruction face two challenges. Firstly, they have to deal with ever larger amounts of raw EM data. Secondly, the increasing demand for reconstructions and the associated costs call for improvements in the reconstruction software to allow annotators to achieve a higher throughput.

The first issue is becoming acute because datasets have continuously become larger and for many reconstruction scenarios it is difficult to devise strategies so that single annotators only have to work on subsets of the raw data: Cells usually extend over large portions of the dataset and it is hard to predict which parts of the dataset a cell projects to (Figure 4). Therefore, it is increasingly imperative not to store raw data locally on the annotator’s computer (as done in KNOSSOS, Helmstaedter et al 2011) but to download the data on the fly as needed (as done in CATMAID, Saalfeld et al 2009, Schneider-Mizell et al, 2016).

To decrease annotation costs, it is desirable to reduce training times, which for existing tools are usually above 10 hours (15 hours per tracer (Helmstaedter, 2013; Helmstaedter et al., 2011), 20 hours (Wanner et al., 2016)). Furthermore the interaction mode and image display should be streamlined so that the annotator can fully focus on the task at hand and not be distracted by technical limitation. This requirement can come into conflict with the on-the-fly delivery of data if the data delivery is not carefully optimized for slower connections (which are frequently encountered while traveling or in settings where undergraduate students aid with the reconstruction effort while working from home, see (Briggman et al., 2011; Helmstaedter et al., 2013)).

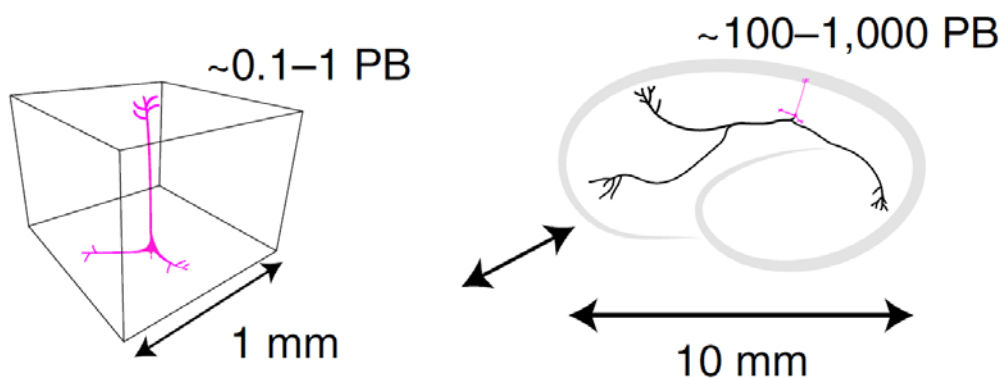


Figure 4: Dendritic arbor of cortical L5 pyramidal cell (left panel) and axon of the same cell projecting into the whole brain (right panel, black). Most cortical cells can innervate substantial parts of the brain - which means that even for single-cell reconstructions an annotator has to have large parts of the raw data readily available.

1.11 The lobula plate of *Drosophila melanogaster*

The optic system of the fly consists of the retina, lamina, medulla and lobula complex, the latter comprising the lobula and lobula plate (Borst, 2014). In each of these regions, the information is organized in a columnar retinotopic way, mirroring the layout of the compound eyes. Direction selective information is already present in the medulla and lobula. T4 cells in the medulla and T5 cells in the lobula respond in a direction selective manner to ON and OFF stimuli, respectively (Maisak et al., 2013). In that aspect they are similar to starburst amacrine cells in the mouse retina, which have direction selective properties and also exist in an ON and an OFF variety (Borst and Helmstaedter, 2015). The input from T4 and T5 cells is aggregated in the lobula plate, where information for each preferred direction is represented in a distinct layer. Going from anterior to posterior these layers are: front to back, back to front, upwards, downwards. Within the lobula plate exists a group of large cells called lobula plate tangential cells (LPTCs). In the blowfly, about 60 of them have been identified (Haag and Borst, 2008). Shared features of LPTCs are their large diameter

dendrites and that they collect information over many columns, playing a role as an aggregator of information. Due to their ability to quickly transport information about movement of the fly relative to its environment, they are essential for the flight control of the fly (Borst and Haag, 2002).

LPTCs all exist pairwise in the left and right hemisphere (some also project contralaterally) and two subgroups of them are cells reacting to horizontal and vertical optical flow. Horizontal cells are the horizontal system (HS) cells HSN, HSE and HSS (for north, equatorial, and south, respectively), which have been described in blowflies (*Calliphora*, (Hausen et al., 1980; Hausen, 1982)), house flies (*Musca*, (Pierantoni, 1976)) and fruit flies (*Drosophila*, (Heisenberg et al., 1978; Fischbach and Dittrich, 1989; Scott et al., 2002)); centrifugal horizontal (CH) cells which have been described in *Calliphora* and other blow flies (Hausen, 1976; Eckert and Dvorak, 1983; Egelhaaf et al., 1993; Gauck et al., 1997), *Musca* (Meyer et al., 1986) and potentially in *Drosophila* (Buchner et al., 1984; Rajashekhar and Shamprasad, 2004); the H1 and H2 cells, described in *Calliphora* (Hausen, 1984); and in *Drosophila* (Bausenwein et al., 1990; Schnell et al., 2010) and the Hu cell ((Haag and Borst, 2001), “called U in (Hausen, 1984)”) in *Calliphora*.




Figure 5: Light-microscopical images of the HS cells in *Drosophila*. The partial overlap of the dendritic trees (situated in the most anterior layer of the lobula plate) can be seen, as well as the somata which are connected to the main neurite by a very thin side branch. Scale bar 25 μm . From (Scott et al., 2002), with permission from John Wiley & Sons, Inc.

Cells that primarily react to vertical optical flow are the cells of the vertical system (VS), called VS1, VS2, VS3..., which have been described in *Calliphora* (Hengstenberg et al., 1982), *Musca* (Pierantoni, 1976) and *Drosophila* (Heisenberg et al., 1978; Scott et al., 2002), and the V1 and V2 cells as described in *Calliphora* ((Hausen, 1984), which were proposed to be “homologous to M-cells in *Drosophila*” (Heisenberg et al., 1978)). In *Calliphora* further tangential cells have been identified (Vi) and proposed (Vi2), for both see (Haag and Borst, 2007). Some thinner cells that bear a similarity to vertical cells have been described in *Calliphora* (Bishop and Bishop, 1981) and *Drosophila* (Rajashekhar and Shamprasad, 2004).

What all these cells have in common is that they react to a large field optical flow. There are other lobula plate tangential cells that are specialized in discriminating a moving object from the background. These are the FD (figure detection) cells, described in *Calliphora* (Egelhaaf, 1985) and *Musca* (Hausen and Wehrhahn, 1990). Their class potentially overlaps with the class of CI cells, described in *Calliphora* (Gauck and Borst, 1999).




Figure 6: Light-microscopical images of the VS cells in *Drosophila*. The arrows highlight features which are useful to distinguish between the cells, namely the orientation of the major dorsal dendrite of VS3, the curving of the same dendrite in VS4, two separate large dorsal branches in VS5 and the fact that all major dorsal dendrites originate from a shared stem in VS6. Scale bar 25 μm . From (Scott et al., 2002), with permission from John Wiley & Sons, Inc.

Disagreement exists over the number of cells in these classes. While the number of HS cells have been described as 3, the number of putative VS cells varies between 9 and 11 in *Calliphora* (Hengstenberg et al., 1982) and the number of VS cells in *Drosophila* is most commonly described as 6 (Scott et al., 2002; Raghu et al., 2007; Joesch et al., 2008). The initial study by (Heisenberg et al., 1978) identified 5-7 VS cells and recently up to three “Additional Vertical Cells” have been proposed (Rajashekhar and Shamprasad, 2004). Initially four H-cells (H1-H4) were described (Hausen, 1984), a number later increased to 6 (Douglass and Strausfeld, 1996), but only two of them have been characterized well enough to be included in recent models of the networks of lobula plate (Borst and Weber, 2011).

When optically counting the number and types of lobula plate tangential neurons, several problems present themselves. It is not possible to visualize all neurons at once. Only 6 VS cells were imaged simultaneously by (Scott et al., 2002) and even then it was not possible to assign all branches unambiguously to cells. Furthermore, no protocol is known that would selectively stain all lobula plate tangential cells, because they have been defined mostly anatomically and physiologically. This problem of separate molecular, anatomical and physiological classifications is not confined to fly research, compare e.g. (Ascoli et al., 2008). A GAL4-3A driver line exists that stains some LPTCs in *Drosophila* (Scott et al., 2002), namely 6 VS and 3 HS cells (see (Raghu et al., 2007) for the complete

expression pattern). Most likely this does not stain all LPTCs and perhaps not even all VS cells. For some multimodal cells it was explicitly shown that they do not colocalize with this Gal4-3A driver line (Levy and Larsen, 2013).

To reduce the number of visible cells, a random subsampling could be used. For the *Drosophila* GAL4-3A driver, such a system is MARCM, which labels a random subset of the cells marked in the line and has the advantage that cells that have a recent shared progenitor are often labeled together (Lee and Luo, 1999). If a purely anatomical reconstruction of cells is desired, a Golgi-silver stain can be employed (Fischbach and Götz, 1981; Rajashekhar and Shamprasad, 2004), which randomly stains about 1% of the neurons. For electrophysiological experiments, the visualization can be achieved by filling the cells with die after the conclusion of the experiment., This shifts the selection process to the onset of the recording. There cells can be selected by known anatomical or functional properties (compare for example (Warzecha et al., 1993)).

If not all cells have been imaged at the same time, datasets from different specimens have to be moved into a shared reference frame. However, aligning datasets from different animals, experiments, labs or species is challenging. Even for datasets from a highly controlled environment, aided by modern information processing technology, this remains an active area of research (Peng et al., 2011). Cells that exist in multiple datasets can be identified if the datasets have been successfully aligned and if the cells are sufficiently stereotypical. Scott and colleagues (2002) investigated the latter by comparing the dendrites of four VS1 cells from different animals. They found that while the dendrites were reasonable consistent with each other, one of the cells had a branched main dendritic shaft and two of the cells had more than one dorsal dendritic branch.

For all these reason, a dataset that contained all LPTCs in an unbiased fashion (marked in a way that they could be reconstructed) would have the potential to find new classes of neurons, to reliably count the number of neurons in each class and remove duplicate classes. 3D electron microscopy can do exactly that: record the morphology of all cells in a region of interest in a complete and unbiased fashion.

1.12 Speed considerations for SBEM imaging

Due to the large amount of data that has to be created for a successful 3D EM connectomics project, acquisition speed is one of the most important considerations. Three components play a role in the total acquisition time for a Serial Block-face EM stack: Pure imaging time (the time the beam spends scanning over the surface), cutting time and overhead time. The most significant contribution to overhead time is movement time: The practical field of view of a single beam scanning electron

microscope used for biological 3D EM is less than 100 μm . Therefore, in most projects, the field of view is extended mechanically. The most common mechanical way to extend the field of view is tiling (see Figure 50). Mechanical movement of the sample can introduce vibrations which persist after the movement has been completed and necessitates additional “cool-down time” after the movement before imaging.

1.12.1 Imaging

The imaging time depends on two factors. Firstly, the microscope has to be able to supply enough current to the rapidly moving small dot that does the scanning. Too few electrons hitting per pixel (charge density) means too little signal. To attain contrast that can be used for dense reconstruction (see 1.5), at least 1000 electrons have to be delivered to a region of about $10 \times 10 \text{ nm}^2$ (Binding et al., 2013). This means that the maximum pixel frequency is

Equation 3
$$f_{\text{pixel}} \approx \frac{I_{\text{beam}}}{1000 e}$$

where I_{beam} is the beam current possible at a spot size of 10 nm and e is the elementary charge. Microscopes that can deliver such currents at the low landing energy necessary for surface imaging in biological samples ($< 3\text{keV}$) are a relatively recent development (Crewe et al., 1968; Crewe, 1973; Hainfeld, 1977; Orloff, 1981, 1985; Tuggle and Watson, 1984; Tuggle et al., 1986). The second challenge is that the detection system needs to have a sufficient bandwidth to be able to separate the signal coming from two neighboring pixels. The back-scattered electrons (BSE, section 1.4) are usually detected with a large-area diode. While the detector itself is fast enough (GHz range), the signal from this setup is challenging to amplify: The diode has a high capacity and produces small currents (nA range). The first amplifiers used by, for example (Briggman et al., 2011), were based on amplifiers developed for electrophysiology (Sigworth, 1995), where a similar problem of high capacity and small currents had to be overcome.

1.12.2 Cutting

On a diamond-knife Serial Blockface SEM setup, the cutting process takes about 30 s. Therefore, for very small datasets where no mechanical extension of the field of view is necessary, cutting can constitute more than 50% of the overall acquisition time, but the problem becomes less acute for larger datasets. Because dataset size is a volume effect and cutting time only grows with the height of the sample block (and to a certain extent the cutting length), the bigger the dataset gets the less relevant cutting time becomes. For very anisotropic samples, it is possible to use one of the shorter edges as the cutting direction.

1.12.3 Mechanical movement

The fast and vibration-free movement of samples in an electron microscope is a topic of ongoing research (compare (Guo-Ying et al., 2013)). In the diamond-knife SBEM the situation is further complicated by the fact that the stage with the microtome on it is necessarily very heavy. Heavy objects have lower resonance frequencies and therefore longer cool-down times. An important consideration is the choice of motor for the movement. Geared motors have nearly unlimited range but are slow and induce vibrations. This can be improved with *piezo motors*, a type of motor that offers higher speed and creates the driving force by bending small piezo grips (Culp, 1991). Lastly, *piezo stacks* offer perfectly smooth movement, but have a limited range (the range of a piezo stack is about one thousandth of its length).

1.12.4 Comparison of factors

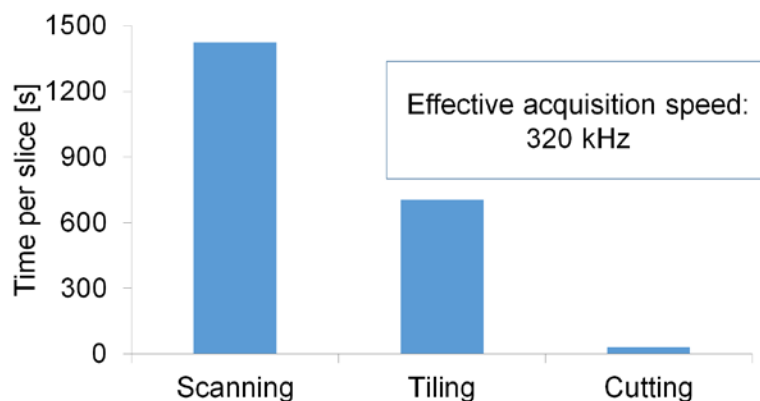


Figure 7: Comparison of time contributions to acquire a single $300 \times 300 \mu\text{m}^2$ slice in a SBEM setup, if the imaging speed is 500 kHz.

In a scanning electron microscope as used for (Briggman et al., 2011), a single image can take up to 12 seconds. Therefore the scanning is the dominant contribution to overall stack acquisition time, even if the tiling motion is slow or needs settling time after movement. Example acquisition times shown for single tile size of $30 \times 20 \mu\text{m}^2$.

The retina experiments that were the basis for (Briggman et al., 2011; Helmstaedter et al., 2013) were done with a FEI Quanta (FEI Company, Hillsboro, OR, United States) scanning electron microscope and a stage driven with conventional geared motors (P-227, Physik Instrumente, Karlsruhe, Germany). Because imaging speed was only 500 kHz, the distribution of acquisition time is dominated by actual imaging (Figure 7). This would be too slow for a whole-barrel experiment, which would take at least 250 days on such a setup. At the time when the work for this thesis was begun, a much faster FEI Magellan microscope had become available (FEI Company, Hillsboro, OR, United States), which created substantially higher beam currents at the same resolution. In

agreement with Equation 3, this allows imaging at a pixel frequency of 10MHz. Juergen Tritthardt and Winfried Denk developed an improved amplifier that overcame the limitations of earlier electrophysiology-based amplifiers and allowed imaging at up to this frequency.

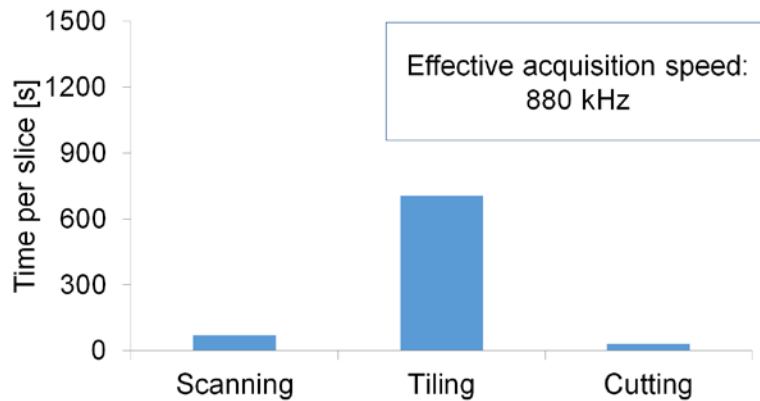


Figure 8: Comparison of time contributions to acquire a single $300 \times 300 \mu\text{m}^2$ slice in a SBEM setup, if the imaging speed is 10 MHz.

If the scanning speed is increased to 10 MHz (using a faster microscope and amplifier) and all the other parameters are kept the same as for the setup with 500 kHz imaging, the vast majority of time is spent for the tiling motion. Further increases to the scanning speed would not significantly improve effective speed. While the datasets from chapter 2.1 were recorded with this setup, the acquisition times for a whole barrel dataset would be excessive (more than 4 months).

In this setting, movement is the dominant contribution to stack acquisition time (Figure 8). Chapter 2.4 describes how to remove this movement overhead so that imaging is the dominant contribution to overall stack acquisition time again.

1.13 Stitching and alignment

Image stitching and alignment is the process by which a continuous three-dimensional dataset is created from a series of images. If the stack is tiled (i.e. each slice consists of several images) the images are stitched within a plane, forming a continuous image. To align the images in the z dimension, the images are positioned in a way that they are congruent with the image of the same position in the slice before it.

For this it is necessary to determine the shift between two images that would bring them into alignment. This can be done manually (Fiala, 2005), by determining the peak of the cross correlation between the two images (Preibisch et al., 2009) or by using the SIFT algorithm (scale-invariant feature transform, (Lowe, 1999); e.g. used in TrakEM2, (Cardona et al., 2012)).

SBEM data is created by imaging the top of a solid block and therefore the images do not have to be rotated to be aligned. For many existing datasets a translation-only alignment algorithm was used (namely (Preibisch et al., 2009), used in (Briggman et al., 2011; Helmstaedter et al., 2011; Helmstaedter et al., 2013))

The algorithm used there compared the overlap region of adjacent tiles. For these images of theoretically identical information, the cross correlation was calculated. Then a threshold was applied onto the correlation calculation result and the center of mass of the largest super-threshold area was marked. The shift of this marker from the null position (center-of-image) was interpreted as the local vote on which translation would have yielded perfect alignment. The same algorithm was used to calculate the relative shift of images being situated on top of each other in z direction. This created a set of local shift vectors. This set massively over-determined the positions that the tiles should eventually take (about $12n$ votes compared to $2n$ degrees of freedom for n pictures). This over-determined system was solved by a weighted least-square relaxation.

1.14 Structure of the thesis

Chapter 2.1 will report the acquisition of two 3D EM datasets in layer 2/3 and layer 4 of mouse barrel cortex and the identification of innervation patterns in these datasets²; chapter 2.2 will report a newly developed tool for collaborative 3D online data annotation called *webKnossos*³; chapter 2.3 will describe the processing and analysis of a large scale 3D EM dataset from *Drosophila* optic lobe; and in chapter 2.4 a 3D EM dataset will be reported that contains a whole barrel from mouse barrel cortex, and the methods will be highlighted that were used to acquire such a large dataset⁴.

² The “L4” dataset was used in the following publications: Berning et al., Neuron, 2015; Boergens, Berning et al., Nature Methods, 2017; Staffler et al., eLife 2017. The “L2/3” dataset and the results of chapter 2.1 are unpublished as of September 1st 2017.

³ This work was published as Boergens, Berning et al., Nature Methods, 2017

⁴ This work is yet unpublished. The method described in chapter 2.4 contributed to Schmidt et al., Nature, accepted.

2 Results

2.1 Innervation specificity in mouse barrel cortex

The analysis of innervation patterns can help elucidate the function of neuronal circuits (See chapter 1.8). For many parts of the nervous system (especially cortex), the extent of specific innervation of targets is open to debate (Peters et al., 1976; Braitenberg and Schüz, 1991; White, 2002; Sanes and Yamagata, 2009; Hill et al., 2012). Existing light-microscopical studies analyzing connectivity patterns in cortex rely on slice preparations and are susceptible to slice cutting artifacts (Song et al., 2005; Yoshimura et al., 2005). The existing 3D EM based studies investigate a rather small region of interest (Kasthuri et al., 2015) or are limited in their reconstructability due to insufficient z-resolution (Bock et al., 2011) and slice folding (Tomassy et al., 2014). Therefore an investigation was conducted looking for innervation specificity patterns in a sufficiently large, densely reconstructable 3D EM cortex dataset.

2.1.1 Datasets




Figure 9 Left: Surface view of the serial block-face electron microscopy stacks recorded from layer 2/3 and layer 4 of mouse barrel cortex (resolution $12 \times 12 \times 26 \text{ nm}^3$ (Layer 2/3) and $11.24 \times 11.24 \times 26 \text{ nm}^3$ (Layer 4)). Right: Light-microscopic image of coronal tissue slice from which the samples for the stacks were extracted (diameter of sample: 1 mm)

Two datasets of mouse barrel cortex were recorded, one in a sample extracted from layer 4, the other one in a layer 2/3 sample. Barrel placement was confirmed by cytochrome oxidase staining (Land and Simons, 1985). The datasets are $99 \times 66 \times 202 \mu\text{m}^3$ (Layer 2/3) and $96 \times 64 \times 89 \mu\text{m}^3$ (Layer 4) in size (Figure 9) and have a resolution of $12 \times 12 \times 26 \text{ nm}^3$ and $11.24 \times 11.24 \times 26 \text{ nm}^3$, respectively (see 3.1.3 for thickness measurement)⁵. The stitching and alignment of the datasets was done as described in 1.13, for further details see 3.2.1. The samples contain 187 and 117 somata, respectively. To aid exploratory discovery and to first test the tools for annotation, all 117 cells with somata in the L4 dataset were reconstructed. (Figure 10)




Figure 10: Layer 4 dataset of mouse barrel cortex from previous figure (outlines of the raw data stack in orange. All neurons with somata in the dataset (blue spheres) were reconstructed manually. The main effort of this annotation task is the faithful reconstruction of all dendritic spines. Strategies to reduce this workload are discussed in detail in chapter 2.2.7.

2.1.2 Reconstruction reliability

2.1.2.1 Neurite reconstruction

Single errors in manual process annotation can lead to tens or hundreds of synapses go undetected (compare 1.6). Because of this, it was important to measure how reliably processes can be reconstructed in the datasets. The datasets were recorded on the same setup just days apart and it was sufficient to assess reconstruction quality in one of them (the layer 4 dataset).

A team of annotators was asked to reconstruct all processes in a bounding box (size: $7.5 \times 10 \times 5 \mu\text{m}^3$, Figure 11) in the center of the dataset. To identify difficult spots where annotators struggled to

⁵ A subset of the L4 data is available here:
https://demo.webknossos.org/datasets/2012-09-28_ex145_07x2_demo/view

find the correct path, annotation speed was measured. The 2200 edges (7.3% of all edges) with the slowest annotation times were inspected by an expert annotator. It was found that 5 of the 2200 edges would need an increased amount of experience to be correctly judged. Thus, it was plausible that these spots would incur higher error rates from newly trained annotators. One spot was considered hard to decide even for expert annotators. This analysis was helpful for a first impression of annotation reliability in the dataset. However, annotation speed might be a weak proxy for annotation correctness and missed branches can correlate with faster annotation (further quantified in 2.2).




Figure 11: Evaluating traceability by densely reconstructing all neurites within a small volume of the layer 4 dataset ($7.5 \times 10 \times 5 \mu\text{m}^3$). Difficult spots were identified by finding loci with reduced annotation speed (scale bar $2 \mu\text{m}$). Snapshot from webKnossos annotation tool, see chapter 2.2. The red frame correspond to the coronal plane, the green frame is parallel to the pia.

A team of annotators then annotated a set of randomly chosen processes in a small bounding box, distributed all over the dataset (Figure 12, left panel, see 3.1.5 for methods). The first part of the RESCOP algorithm – which is used for identifying disagreements in redundant annotations (see chapter 1.7.1) – was run on these data (example: Figure 12, right panel) and yielded a vote matrix (Figure 13, left panel). All edges that correspond to entries on the decision boundary (black box Figure 13, left panel) were then manually inspected. Of those, only one position was not easily resolvable and left some ambiguity – which can be resolved with biological priors (e.g.: every axon

needs to be connected to a soma). Extrapolating from that one situation, the mean inter-error-distance can go up as high as 1.04 mm for high-redundancy annotation (see 3.1.5).




Figure 12: Measurement of reliability of reconstruction for layer 4 dataset. Left: To assess the ability to reconstruct all neurites throughout the dataset, random seed points were distributed and annotators were asked to trace them within a $(10 \mu\text{m})^3$ bounding box (outlines of dataset (orange) and dense cube of Figure 11 shown). Right: These reconstruction were done redundantly (7-fold) and it was measured how much they agreed by using the RESCOP algorithm (Helmstaedter et al., 2011), see Figure 1. (disagreements shown in red and yellow, scale bar $2\mu\text{m}$)

For process annotation, errors are attention related (Helmstaedter et al., 2011). The second part of the RESCOP algorithm can be employed to extrapolate how increasing numbers of annotators help improve the quality of the tracing. This was done for the redundant tracings here and yielded inter-error distances as good as or better than for the e1088 dataset (dataset used for (Helmstaedter et al., 2011; Andres et al., 2012)), (Figure 13, right panel), both for normal annotation and focused reannotation (see 1.7.2). This result was published as part of the comparison of staining techniques in (Hua et al., 2015).

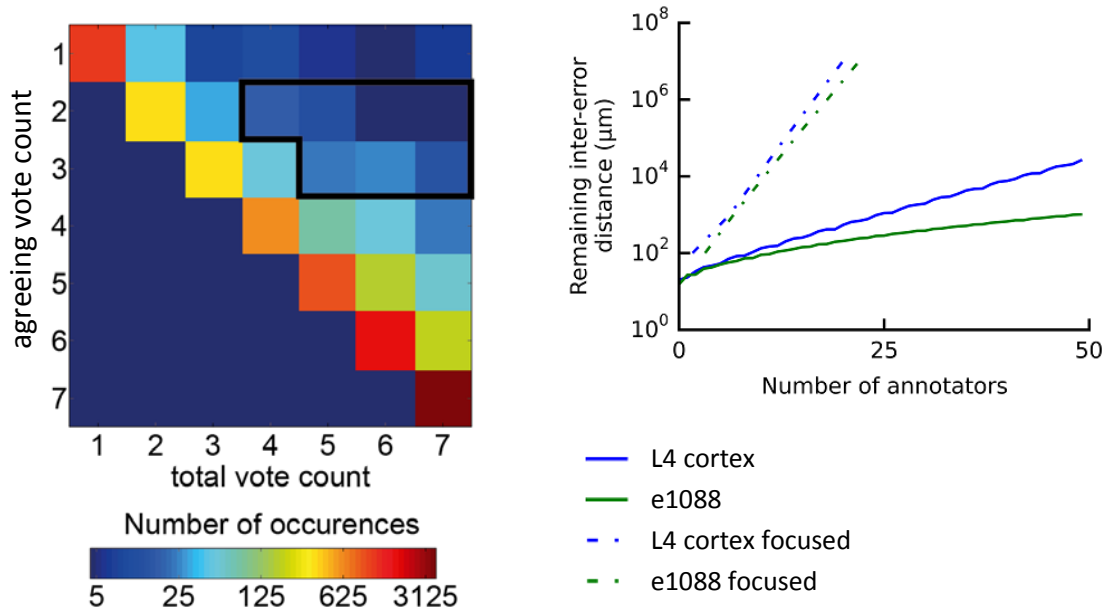


Figure 13: Analysis of neurite annotation consensus and prediction of annotation error rate. Left: vote histogram generated from annotations of the layer 4 dataset shown in Figure 12. Right: Prediction of annotation error rate for layer 4 and reference dataset. After the agreement for all neurite edges had been measured by RESCOP (Helmstaedter et al., 2011), they were accumulated into a vote matrix (left panel). Problematic edges are those that are close to the accept/reject decision boundary. Therefore those edges (within the black frame) were inspected manually. Only one of them was found to be critical for continuation of neurites, the others were related to spine-annotation. To estimate ease of reconstruction, it was measured which distribution of easy versus hard-to-reconstruct loci would generate a vote matrix as seen in the left panel. With this fitted distribution, it was possible to predict improvement of inter-error distance as the redundancy increases, both in a setting where new annotators reconstruct the whole neurite and where the equivalent work time is focused onto problematic loci (right panel)

2.1.2.2 Synapse detection

In cortex contact area is not a good proxy for synaptic connectivity (Takemura et al., 2015), therefore the next step was assessing synapse detectability. En-bloc synapse annotation is very time-consuming (6 min/ μm^3 , B. Staffler, personal communication). It is easier to search for synapses along axons. It is also more in line with the general task of measuring innervation specificity from a per-axon perspective.

On manual inspection, synapses looked clear and easily detectable in the data (Figure 14): Vesicle clouds are visible and docked vesicles can be seen close to the putative post-synaptic density. Post-synaptic densities are less clear for inhibitory synapses (Figure 14, right panel). Synapse detection profited from the three-dimensionality of the dataset: Synapses that were not clear in a single plane could be resolved when aggregating the information from several planes (compare (Gray, 1959),

traditional electron micrographs, where the image looks very sharp, but only a single layer is available to make the decision). Also, because the dataset is relatively isotropic, the concept of obliquely cut and thus hard to detect synapses was much less prevalent than, for example, in (Spacek and Harris, 1998).

To approximate error rates for manual synapse annotation, a group of annotators was asked to annotate vesicle clouds and synapses (Figure 15) along a shared set of axons (see 3.1.6 for methods). Then the agreement between the annotators was measured (Figure 15, top-right panel). When the original annotators were confronted with sites where they were disagreeing, they usually could agree on a consensus, reporting to have missed a given synapse or being temporarily unaware of the criteria on what should constitute a full vesicle cloud (see 3.1.6). Therefore this is a setting where most errors are lack-of-attention errors and thus the error rates could be estimated using RESCOP (1.7.1): The error rates are dropping quickly with an increased number of annotators (Figure 15, bottom-right panel).

This analysis also supplied an independent measurement for the synapse density along axons. Accepting synapses which 4 or more annotators agreed on yielded a synapse density of 0.23 ± 0.07 synapses/ μm in the three analyzed axons.

An alternative to manual synapse detection is automated synapse detection. It was shown that automated synapse detection on this dataset can operate with a precision and recall of 88% for single synapses and a precision and recall of 97% if the information from multiple synapses connecting the same pair of neurons is pooled (Staffler et al., 2017).




Figure 14: Detection of synapses in 3D EM data of mouse cortex. Left: Slices through excitatory synapse 5 slices (130 nm) apart. Middle: Volume model of excitatory and inhibitory synapse, with location of slices that are shown in the other panels. Right: Slices through inhibitory synapse 5 slices (130 nm) apart. Three-dimensional data helps with identification of synapses. In the left column from top to bottom, the presynapse, the synaptic interface and the postsynaptic spine are marked with arrows. At 26 nm cutting thickness, a synapse is present in 10-30 consecutive planes. Aggregating the information from many planes can help making identification less ambiguous, both for excitatory and inhibitory synapses (left and right panel, respectively; scale bar 1 μm ; in the images pia is to the top and the cuts were made in the coronal plane). This and all following cortex volume renderings have been created using a volume segmentation created by M. Berning, (Berning et al., 2015)

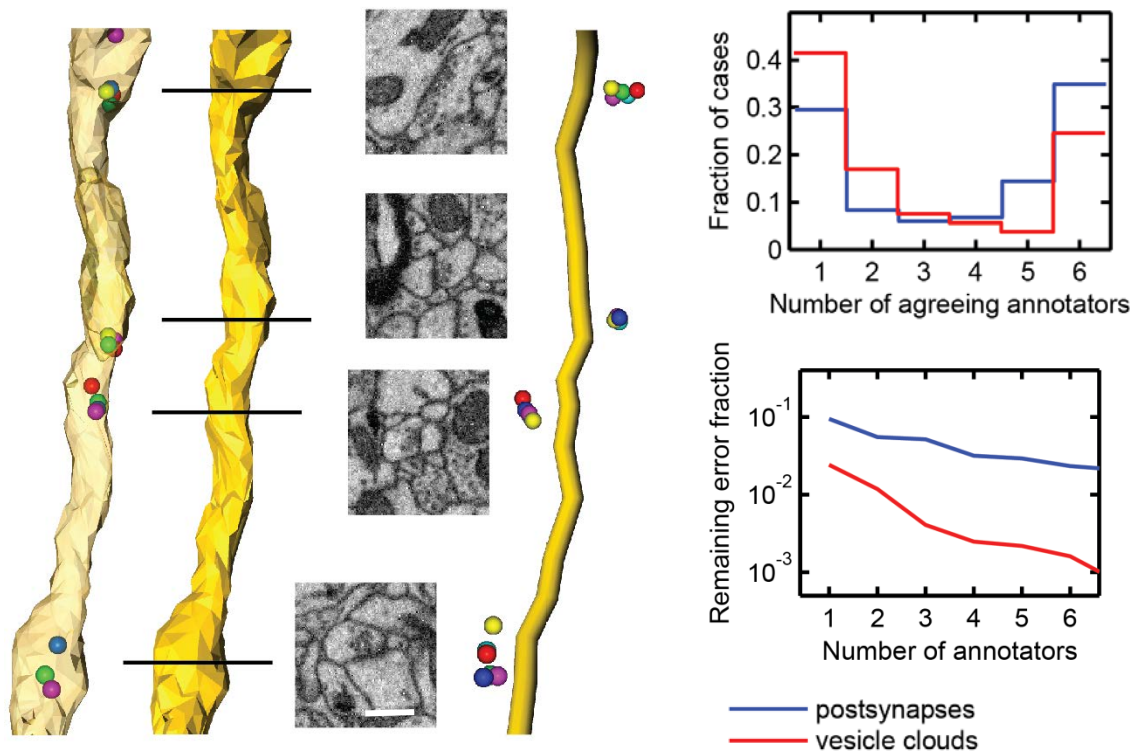


Figure 15: Method for identifying synapses along an axon. Left: Steps for synapse identification (from left to right). 1: Identification and marking of vesicle clouds (colored spheres show multiple annotators). 2: Reassessing raw data to identify postsynaptic partners. 3: Marking of postsynaptic partners. Right: Consensus quality and error rate prediction for scattered multi-annotator data shown on left.

In conclusion, synapse identification along an axon can be split into vesicle cloud and post-synaptic process detection. For both, annotator agreement is sufficient for efficient synapse detection. Few synapses yield annotator disagreement close to decision border (top right, between 3 and 4 agreeing votes). RESCOP (Helmstaedter et al., 2011) can be used to predict error rates as more annotators work on a given synapse (bottom right). Scale bar 500 nm

2.1.2.3 Axon type detection

Having established the ability to detect synapses, the next task was to categorize them as excitatory or inhibitory. Within the recorded data it was often difficult to categorize an isolated synapse. This issue has been previously reported for EM data (Gray, 1959). In the layer 4 and layer 2/3 datasets it is not necessary to categorize an isolated synapse: Most axons in the datasets carry multiple synapses and “Dale’s principle in Eccles’ interpretation” (Dale, 1935; Eccles, 1976) states that all synapses of an axon release the same neurotransmitters. Therefore analyzing several synapses along one axon can increase the confidence of the categorization. Whether a connection was excitatory or inhibitory could be determined based on the information gathered from adjacent synapses. In summary, innervation type was distinguished on a per-axon basis (as opposed to a per-synapse

basis). Expert annotators were able to distinguish excitatory and inhibitory axons based on PSD shape and coloration - but it was found that beginner-level undergraduate annotators were struggling to make the distinction. However, they were able to reliably distinguish spine and shaft synapses. There is a correspondence between synapse geometry and axon type, with spine synapses usually belonging to excitatory axons and shaft synapses usually being inhibitory. This is not a reliable indicator of axon type identity on a single-synapse basis (see (Chen et al., 2012)) and thus has to be pooled over several synapses of the same axon (Table 1).

<i>Excitatory vs inhibitory axons</i>	Distinction by PSD and color	Distinction by spine vs shaft
Decision per synapse	Not possible	Unreliable (Chen et al., 2012)
Decision per axon	In this dataset only possible for experts	Possible

Table 1: Comparison of axon type detection strategies

To measure the reliability of using only the shaft vs. spine synapse label to distinguish axon type, a pool of axons was traced completely through the dataset and all their synapses were annotated as either projecting onto a spine or a shaft (see 3.1.7 for methods). This showed that for those axons the synapses are either mostly onto spines or mostly onto shafts. It was then analyzed how many synapses per axon are necessary to reliably establish the axon type this way (Figure 16 left). For clusters of 10 synapses, the distribution of votes was found to be sufficiently bimodal and the majority of axons had a clear preference towards innervating either mostly spines or shafts (Figure 16 right). Therefore, for all type determinations in this chapter, 10 synapses were analyzed. If this yielded a very close result (4-6 synapses onto spines), another 5 synapses were annotated to break the tie.

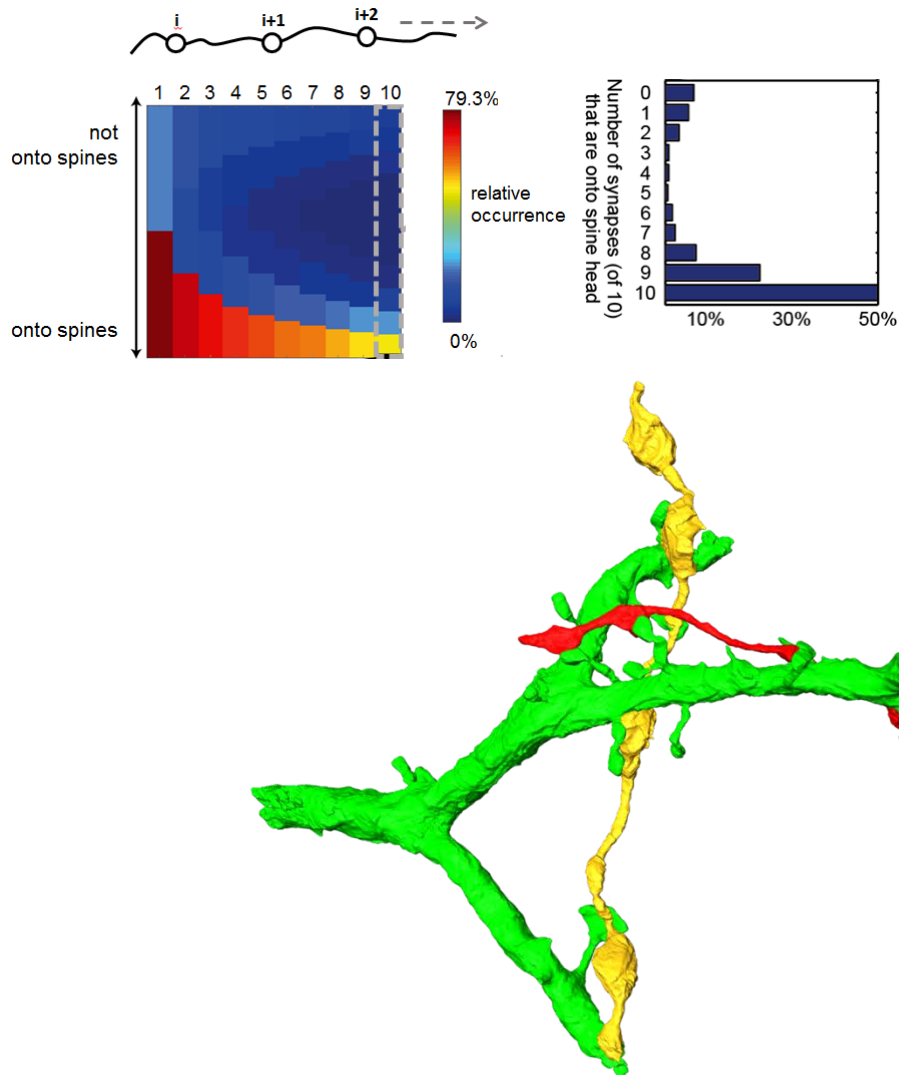


Figure 16: Axon type detection by counting synapses onto shafts versus synapses onto spines. Top left: Heat map of how consistent axons bear shaft or spine synapses, if n synapses of the axon are taken into account. Top right: Distribution of ratios of shaft vs. spine synapses for a surround size of 10.

The larger the analyzed surround (number of synapses n), the less frequently axons are encountered where roughly half the synapses are onto spines. For a surround of 10 synapses analyzed (top right panel), the distribution is sufficiently bimodal (split into spine-innervating axons and shaft-innervating axons) to allow a distinction between inhibitory and excitatory synapses. Bottom: example dendrite with inhibitory axon (yellow) forming three shaft synapses onto the dendrite and excitatory axon (red) forming three spine synapses onto the dendrite, an example case of high consistency of the spine-vs-shaft distinction.

2.1.3 Initial segment innervation

It is known that in all layers of cortex except layer 4, inhibitory chandelier axons innervate initial segments (IS) of pyramidal cell axons (Somogyi, 1977; Taniguchi et al., 2013). Because the action potential is usually initiated in the initial segment (Clark et al., 2009), this is a powerful spot to control firing rates. This innervation had been described as particularly specific, forming all synapses

onto initial segments, but the evidence remained inconclusive (Somogyi et al., 1982). There is only sparse information on innervation of initial segments in layer 4 (Marie and Peters, 1985). Therefore, it was measured how specific the innervation of initial segments in the layer 2/3 dataset actually is and whether there is a different kind of initial segment innervation in layer 4.

2.1.3.1 Layer 2/3 IS innervation

At first, one initial segment of a pyramidal cell was identified (see 3.1.8 for criteria to identify initial segments). Then all synapses onto that IS were annotated. It became clear that of those 10 synapses, 5 were from the same axon (Figure 17). This putative chandelier axon was then reconstructed throughout the dataset. Its overall morphology was consistent with the known shape of chandelier axons. All its synapses were marked and the identity of the postsynaptic target established. Of the targets, 39 were onto initial segments but 2 were not. This shows that the affinity of chandelier axons for initial segments is high, but not absolute.



Figure 17: Multiple branches of a chandelier axon (red) innervate an initial segment of a pyramidal cell (grey) in layer 2/3 mouse barrel cortex. Manually inspecting initial segments in layer 2/3 is an efficient strategy for identifying chandelier axons (scale bar 10 μ m)

The step of identifying all postsynaptic partners of the chandelier axon allowed for the convenient identification of further initial segments in the dataset. With the help of this set of initial segments, a second axon was identified that had multiple initial segment innervations and a chandelier-like

appearance (Figure 18, left panel). All of the synapses of that axon were also identified and all 80 of them were onto initial segments.

In conclusion, chandelier cells could be reliably identified and their specificity measured. It was then investigated whether this was the only type of innervation onto initial segments. To assess this, all synapses of 12 initial segments were annotated (marked with * in Figure 20). This yielded an average synapse density of 0.1232 synapses per μm . Of those 12 IS, 3 initial segments were selected and all 16 additional axons were reconstructed that innervated them (IS marked with ** in Figure 20). Of those axons, all postsynaptic partners were identified and it was assessed whether they were initial segments. This analysis yielded a much more varied picture than assumed before: While one more chandelier axon was identified, two more classes of axons were found: 2 high-specificity IS innervating axons that did not have the morphology associated with chandelier axons and 13 axons which targeted initial segments but also had other non-IS targets. What all axons had in common was that none of them was identified as excitatory (according to the criteria outlined in 2.1.2.3). The information of all axons was consolidated into a preliminary connectome of identified initial segments and their innervating axons (Figure 20).

The chandelier axons exhibit very high specificity. Therefore they are well suited to establish an upper limit on the specificity measurement. This specificity measurement was done in accordance with the ideas outlined in 1.8.1 and is described in detail in 3.1.11. Briefly, the wiring length of initial segments in the dataset was compared to the overall wiring length of postsynaptic structures. For random wiring, this ratio should be equivalent to the ratio of synapses targeting initial segments to total number of synapses. The factor in which the two ratios differ is the *specificity factor* (Equation 4).

Equation 4

$$specificity\ factor = \frac{\frac{num_{syn_target}}{num_{syn_all}}}{\frac{length_{wiring_target}}{length_{wiring_all}}}$$

For the chandelier innervation this factor is 234.6 (compare Figure 22) with error bounds of [+4.8 - 21.5] (see 3.1.11). The comparison of all wiring specificity factor of the layer 2/3 IS innervating axon shows the split into high- and low-specificity axons (Figure 18, right panel).




Figure 18: Multiple cells innervate a single initial segment. Left: A chandelier axon (orange) innervates an initial segment (dark grey). Additional synapses are interspersed (light grey). The orange axon exhibits a typical chandelier-like morphology. Scale bar 20 μm . Right: Distribution of axonal specificities (the preference of axons to innervate a target) for initial segment innervation. Initial segments are innervated both by axons with a very high and low specificity for initial segments.




Figure 19: Interneuron in layer 2/3 mouse barrel cortex dataset, dendrite red, axon blue, dataset bounding box orange (size: 99 x 66 x 202 μm^3). No innervation exists onto the initial segment of this interneuron.

All initial segments identified were initial segments of pyramidal cells. Therefore in a separate effort, a layer 2/3 inhibitory cell whose IS was within the dataset was identified (Figure 19, preliminary results indicate that one of ten cells in the dataset are inhibitory, see 3.1.10 for methods for inhibitory cell identification) and its IS was annotated. This IS did not contain any synaptic innervation at all. This means that even the low-specificity IS innervating axons had a specificity for pyramidal over interneuron initial segments.

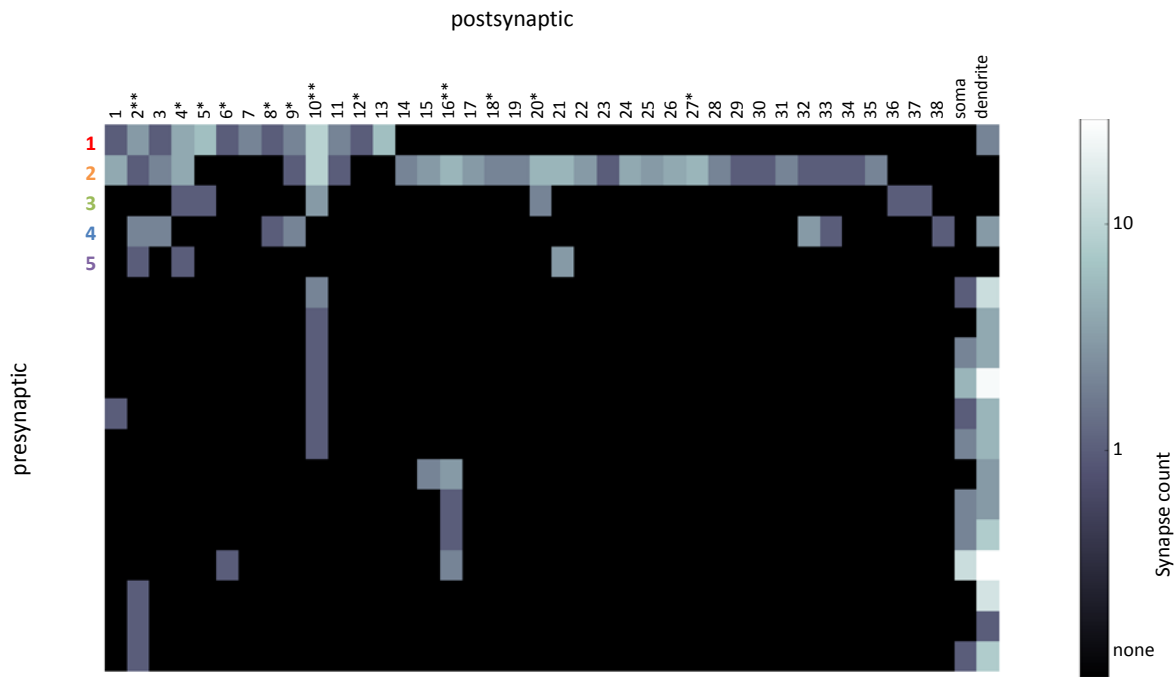


Figure 20: Connectivity matrix for innervation of initial segments in layer 2/3 of mouse barrel cortex. High-specificity axons colored, fully synapse-annotated initial segments are marked with a star, initial segments where all innervating axons have been fully reconstructed are marked with a double-star. All identified chandelier axons (line 1-3) have overlap with each other in their innervation targets. Some chandelier axons have non-initial-segment targets (line 1).

2.1.3.2 Layer 4 IS innervation

All initial segments of cells contained in the layer 4 dataset were annotated. The same criteria for identification were applied as in the layer 2/3 dataset (see 3.1.8). In the next step, all synapses projecting onto the initial segments were annotated. They were encountered at a rate of 0.22 synapses/ μm . This innervation density was higher than the value for layer 2/3 (0.1232 synapses/ μm , see 2.1.3.1), despite the lack of chandelier cells in layer 4 (Taniguchi et al., 2013).

In layer 2/3, the identified IS of an interneuron was not the target of any synapses. In the layer 4 dataset, only one inhibitory cell with an IS in the dataset existed at all. This made it plausible that the dataset was recorded in either the center of a barrel (the hollow) or the septal region between barrels (Lin et al., 1985; Kätzel et al., 2011). Unlike in layer 2/3, this one interneuron IS was innervated. The axon of one of those synapses was reconstructed and yielded a specificity of 1/3. While using the analysis from one cell cannot give conclusive results, this might hint at a separate network of inhibitory IS-innervating axons or innervation only by low specificity axons.




Figure 21: Rendering of an initial segments in layer 4 and an example innervating axon. Initial segments of spiny stellate cells in layer 4 of mouse barrel are innervated, but not by extremely-high-specificity axons (i.e. the kind of specificity chandelier axons exhibit in layer 2/3). As an example for increased specificity, an axon is shown (blue) that is innervating an initial segment with 4 synapses (dark grey volume model) and also two other initial segments (dark grey wire models). Scale bar 20 μm .

Of all synapses onto spiny stellate initial segments, 59 were chosen randomly and their axons fully reconstructed. All synapses formed by those axons were marked. Then automated contact detection was employed to generate candidates for synapses innervating pre-identified initial segments (see 3.1.12 for methods). Those candidates were manually inspected to discern synapses from mere contacts. The information was pooled into a connectivity matrix (Figure 23). It is important to note that this procedure did not mark all initial segment synapses. Only 61% were detected because the initial segment library only contained initial segments whose somata were in the dataset, but this was compensated for (see 3.1.11.1). Overall, the specificities encountered in the layer 4 IS innervating axons were lower and more varied than for layer 2/3 (Figure 22). Axons exist that have an elevated specificity for axon initial segments. A typical medium-increased-specificity axon is axon 51, which targets an initial segment with 4 synapses and then two further initial segments (Figure 21) out of a total of 144 synapses, which gives it a specificity factor of 10.6 ± 3.9 .

2.1.4 Specificity of apical dendrite innervation in L4

Another group of structures that can be investigated in the layer 4 dataset are the apical dendrites of layer 5 pyramidal cells which pass through layer 4 and form arborizations in higher layers. Therefore it was analyzed whether there is specific innervation onto layer 5 apical dendrites within

layer 4. First, a set of reconstructions was created that identified 117 apical dendrites within the dataset (as described in 3.1.9). Two apical dendrites were selected and nine axons innervating their shafts were randomly chosen. Those axons were completely reconstructed. According to the method detailed in 3.1.7, they were found to be inhibitory axons or, if they had too few synapses, they were at least consistent with the criterion. Then the innervation targets of those axons were identified and it was tested whether they were apical dendrites too. Of the 97 synapses tested in this way, 47 were targeting apical dendrites. Through this process 7 additional apical dendrites were identified that had not been part of the initial library. The innervation specificity of the axons was between 9.8 and 54.8 (Figure 22, excluding one outlier, the outlier is an axon with two synapses, hence the broad posterior – compare 3.1.11), with a mean specificity of 27.4 [+13.7 -3.3].

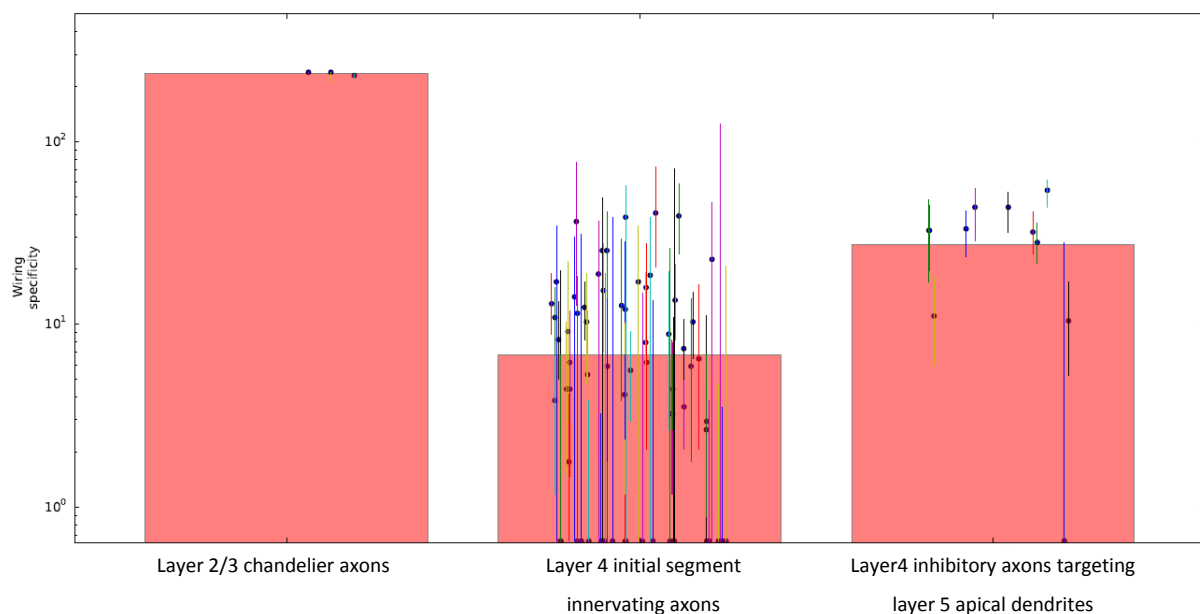


Figure 22: Overview of innervation specificities for structures in mouse barrel cortex, together with average specificity and error estimates for single axons. Mouse somatosensory cortex contains examples for extremely high specificity (chandelier axons onto initial segments in layer 2/3), low specificity combined with high inter-axon variance (initial segment innervation in layer 4) and medium high specificity combined with low variance (inhibitory innervation onto layer 5 apical dendrites within layer 4). Error bars were calculated using a Bayes simulation of an axon ensemble.

It was then investigated whether there are axons that innervate both initial segments and apical dendrites. For this an automated synapse detection algorithm was used (Staffler et al., 2017). No crosstalk between these two categories of axons was found (Figure 23, right panel). This can be described as an avoidance of a certain target class, which is also a type of specificity.

As discussed in the introduction (chapter 1.8.1), specificity statements can only be made when qualified by the volume over which the measurement takes place. The apical dendrites run through

the datasets in bundles. If the measurement volume is sufficiently small that it doesn't smooth out these bundles, an axon running parallel to a apical dendrite in a bundle would have a lower specificity than one running orthogonally despite having the same ratio of hit/non-hit synapses (see Figure 23, left panel for a stretch of axon that exhibits both behaviors). The degree to which availability of apical dendrites (as fraction of overall postsynaptic wiring) depends on the spatial scale of integration was measured for all axons (Figure 24). It can be seen that some axons have a lot of spatial overlap with apical dendrites because they run along their bundles. Other axons run across the dataset and only form synapses when they come close enough to apical dendrites.

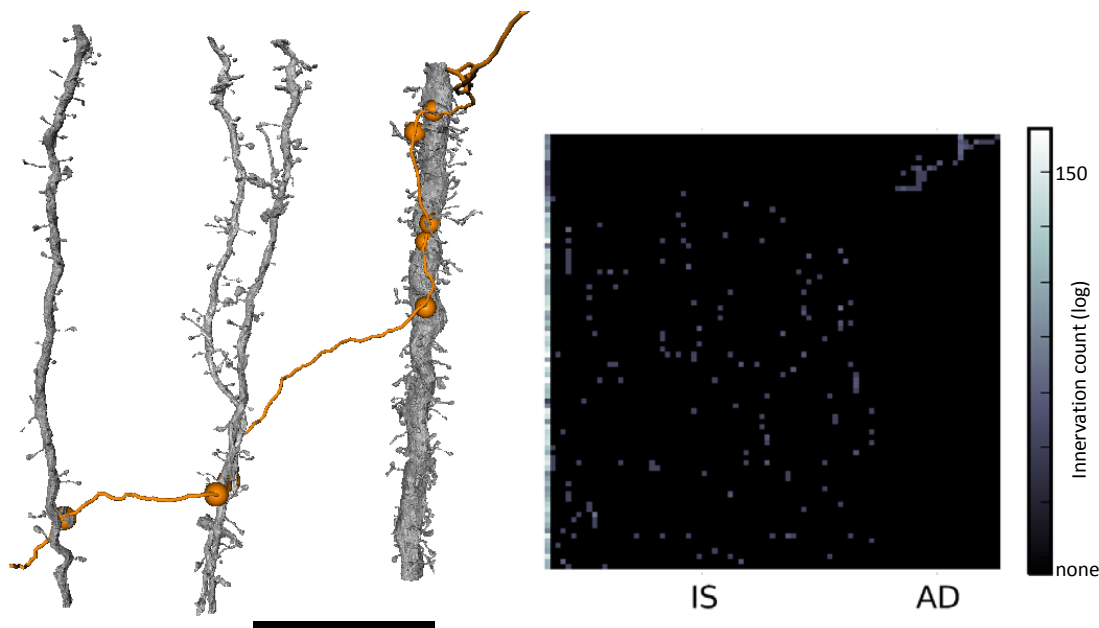


Figure 23 Left: Volume rendering of apical dendrites (grey) and an innervating inhibitory axon (orange) in layer 4 of mouse barrel cortex. Specificity in apical dendrite innervation is not solely caused by axons and apical dendrites running in parallel: The inhibitory axon crosses through the dataset to innervate multiple apical dendrites. Scale bar 20 μm . Right: Combined connectivity matrix for inhibitory innervation of apical dendrites (AD) and initial segments (IS) in layer 4 of mouse barrel cortex. Apical dendrite innervation is more specific. No axons were identified that innervate both apical dendrites and initial segments.

35 excitatory axons that innervate apical dendrites were reconstructed. This was done by choosing consecutive spine heads of random apical dendrites and asking annotators to trace the axons that were innervating those spines. Automated contact detection was used to detect potential synaptic sites. This requires all spines of the postsynaptic structure to be reconstructed. Therefore, a revised set of tracings of the apical dendrites was created where annotators were instructed to make sure to annotate all dendritic spines completely. This skeleton tracing was combined with a volume segmentation (compare 3.1.12 for methods) to create a contactome between the apical dendrites and the reconstructed axons (see 1.6.4). Then all contact sites were manually annotated by expert

annotators as synaptic or non-synaptic. This is an ongoing effort, but so far no specificity could be observed.

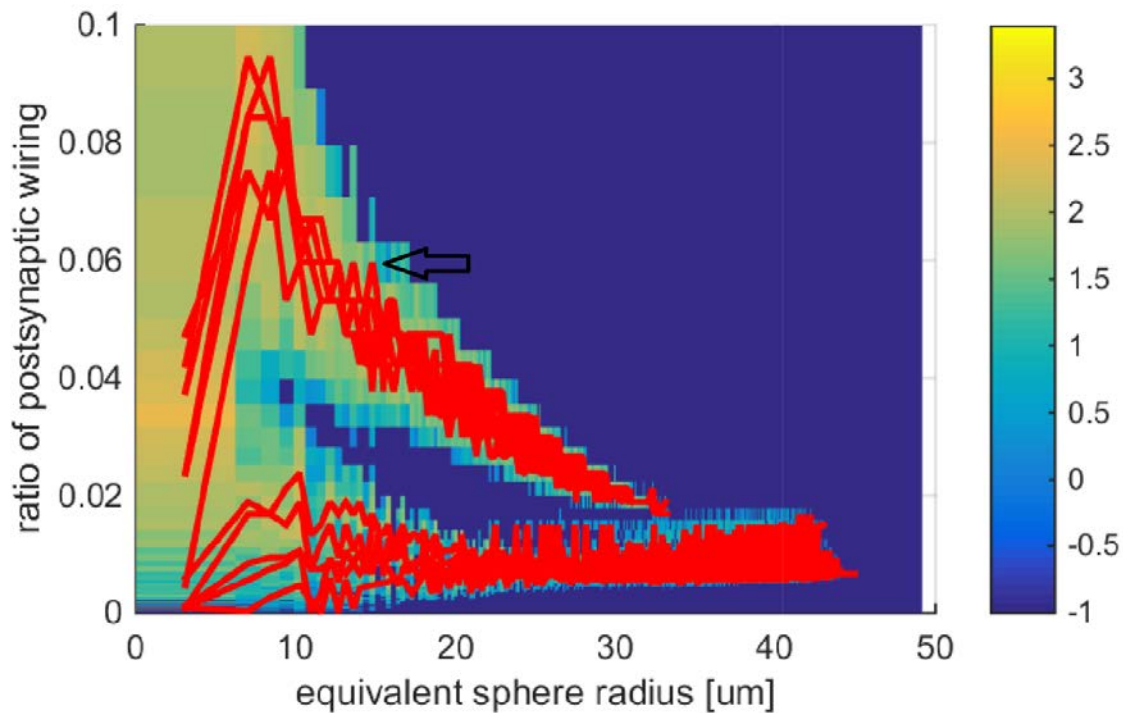


Figure 24: Availability of apical dendrites (as a fraction of total available postsynaptic wiring) for inhibitory apical-dendrite innervating axons, plotted against the volume in which the availability is measured. The red lines describe individual axons, the underlying heat map the whole population of axons. For example, the peak of the red axon plot marked with a black arrow corresponds to the measurement that the postsynaptic wiring in a tube with a radius of 15 μm around this axon contains on average 6% apical dendrites. This is an increase over the general prevalence of apical dendrites in the dataset, which is 1.5% of the postsynaptic wiring.

2.1.5 Spatial distribution of inhibitory innervation

As discussed in 1.8, one pattern of specificity is innervation that targets a specific region of a dendrite (see Figure 2). Competing predictions exist whether a proximal or distal cluster of inhibitory innervation would be most effective (Koch et al., 1983; Gidon and Segev, 2012). In the layer 4 dataset, the dendrites of the spiny stellate cells were used to answer for this cell type whether proximally clustered inhibition is present.




Figure 25: Measuring clustering of inhibitory innervation. Left: Set of 5 layer 4 spiny stellate cell main dendrites with spines reconstructed. Right: Example dendrite with all innervating axons.

This reconstruction of innervating axons was done for all dendrites. Each axon was reconstructed until enough of its synapses had been identified to allow a decision whether it is excitatory or inhibitory (usually 10 synapses, see Figure 16)

Five dendrites of spiny stellate cells with somata in the dataset were randomly chosen. Two of the dendrites were from the same cell. These dendrites were redundantly ($n=3$) reconstructed by expert annotators, including all spines. Branch points that were leading into side branches were ignored, resulting in just one branch with spines (Figure 25). If the main branch left the dataset, the biggest side branch was chosen instead.

All synapses onto those dendrites were marked and their axons were reconstructed into a 10-synapse radius (see 3.1.7) to determine the axon type. If more synapses were needed to make the decision, more were included. One dendrite had been used to establish the axon type detection (see 2.1.2.3) and its innervating axons and their synapses had already been reconstructed completely.

Initial visual inspection showed a strong clustering of inhibitory innervation in the most proximal region of the dendrite (example dendrite in Figure 26). It was then tested whether this pattern was significant. The excitatory and inhibitory synapses were split into two groups (less and more than 25 μm from the soma). On these groups a t-test was employed (Figure 26, inset), showing that there is a clustering of inhibition on the most proximal part of the dendrite (significant in 4 out of 5 dendrites).

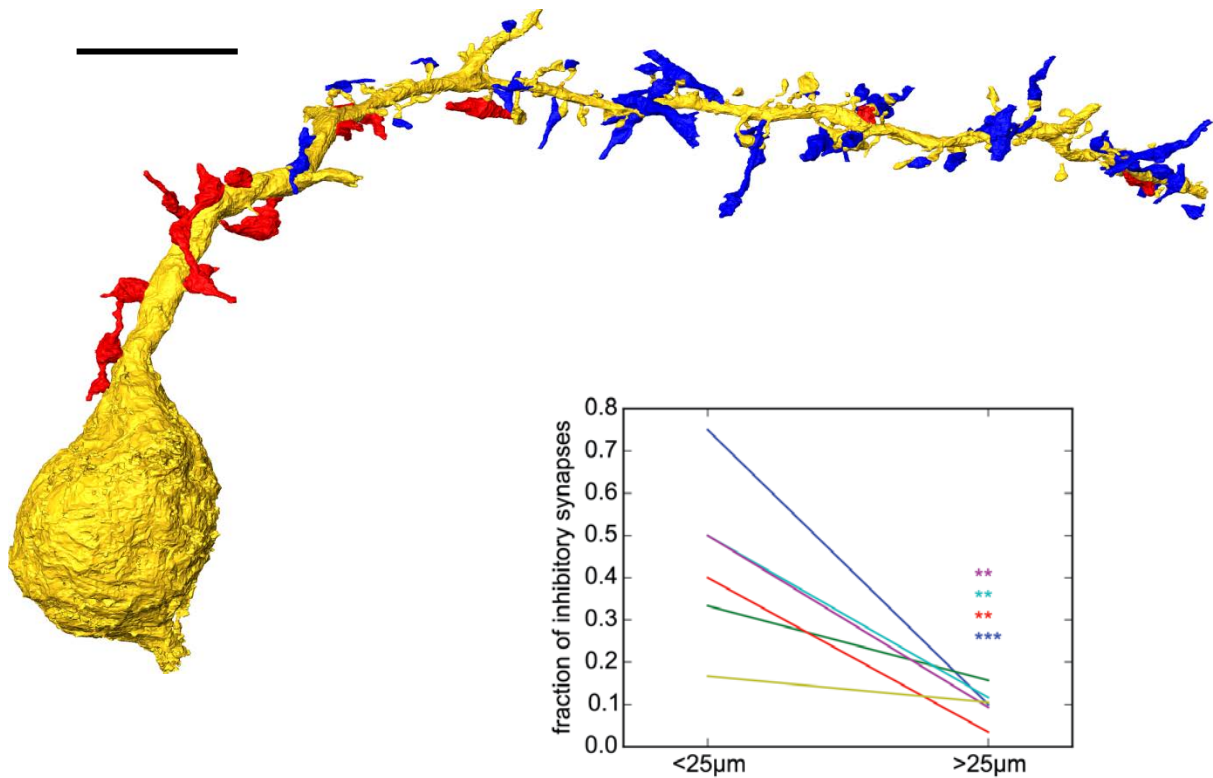


Figure 26: Example dendrite with its innervation (excitatory innervation: blue, inhibitory innervation: red). Inset: t-test for inhibitory clustering in the first 25 μm of each dendrite (**: $p < 0.01$; ***: $p < 0.001$). A bias for inhibitory innervation is visible in the proximal part of the dendrite. This clustering was significant in most cases (example dendrite: $p = 0.008$). Scale bar 10 μm

It was then investigated if this innervation specifically targeted a certain dendrite of the cell? For two of the dendrites, all inhibitory axons were reconstructed to the edges of the dataset and it was tested whether they innervated other parts of the cell. No hit was found targeting other main dendrites or the somata of the cells. If the somata or other dendrites had also been targeted, this might have meant that this inhibitory cluster was rather part of a larger cluster exercising tonic control over the cell (see 4.1 for an extended discussion).

2.2 webKnossos

webKnossos is a newly developed in browser cloud-based connectomic reconstruction tool, that also pioneers a novel data interaction mode which yields significant gains in reconstruction speed.

2.2.1 Data transmission speed

webKnossos has been optimized for non-optimal connection settings. To investigate how well it performs in telecommuting reconstruction setting where bandwidth and latency are worse than what staff scientists have available in a research institute, webKnossos was compared with the most widely cited cloud-based annotation tool, CATMAID (Saalfeld et al., 2009; Schneider-Mizell et al., 2016).

The webKnossos back-end transmits 3D EM data in packages called buckets sized 32 x 32 x 32 voxels. Additionally, webKnossos uses a heuristic to predict which buckets will be needed next and loads them preemptively (see 3.3.2 for details). CATMAID, on the other hand, stores and transmits the data as images (size configurable, e.g. 256 x 256 px, compare Figure 27), separately for each z layer. To compare the maximum achievable data rates of webKnossos and CATMAID, both tools were installed on separate dedicated servers (see methods). An expert annotator then reconstructed an easy axon, tracing as fast as the raw data rate allowed. The first measurement was performed while being connected to the server under optimal conditions (1 Gigabit connection to server, server in same country). There the expert annotator was able to achieve movement speeds of 2.032 mm/h for webKnossos and 0.470 mm/h for CATMAID (Figure 28; this and all following measurements were executed in a dataset with voxels sized 11.24x11.24x28 nm³, see 3.3.1).




Figure 27: Comparison of sequential 2D data loading strategies (as used in CATMAID) and true 3D loading strategies as used in webKnossos. In a 3D loading strategy, the lateral field of view can be smaller, laterally loading only the data that is necessary and then using the reduction of bandwidth demand to load the data that will be used as reconstruction movement proceeds along the neurite (pre-loading). Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

Then a speed restriction was activated in the browser, limiting the transmission to 100MBit, 20MBit, LTE-like speed, good 3G-like speed and normal 3G-like speed, respectively, and the test was repeated. The resulting speeds can be seen in Table 2 and Figure 27 and show a consistent advantage of webKnossos over CATMAID annotation-speed wise.

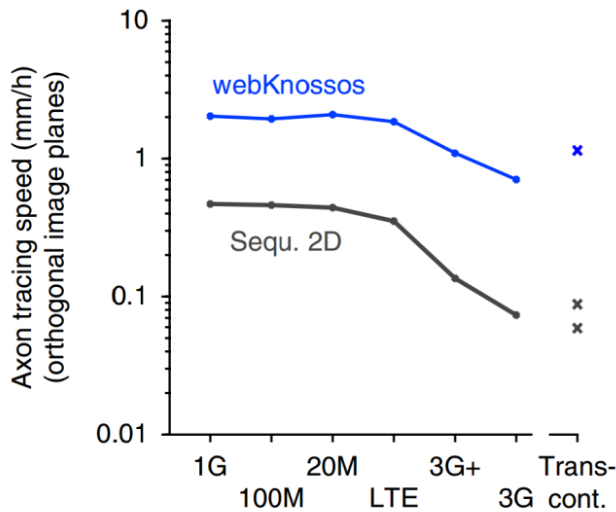


Figure 28: Comparison of maximally achievable axon reconstruction speeds for webKnossos and sequential 2D loading (as used in CATMAID). For high speed connections, webKnossos is about 4 fold faster than CATMAID. For connections with a higher latency and bandwidth, the difference becomes even larger. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

	webKnossos	sequ. 2D
1G	2,032 mm/h	0,470 mm/h
100M	1,939 mm/h	0,461 mm/h
20M	2,084 mm/h	0,442 mm/h
LTE	1,851 mm/h	0,353 mm/h
3G+	1,098 mm/h	0,135 mm/h
3G	0,705 mm/h	0,073 mm/h
Transcontinental	1,190 mm/h	0,088 mm/h
(openconnectome)		0,059 mm/h

Table 2: Comparison of maximally achievable axon tracing speeds in webKnossos and CATMAID

A stated development goal of online annotation software is to facilitate collaboration between laboratories worldwide (Saalfeld et al., 2009). Therefore the speed measurement was repeated in a transcontinental setting: Accessing the webKnossos and CATMAID server (located in Germany) from

a computer at the Instituto de Investigación en Biomedicina de Buenos Aires, Argentina. The resulting values were 1.190 mm/h for webKnossos and 0.08 mm/h for CATMAID (Figure 28). Additionally, the maximally attainable reconstruction speed using a public CATMAID server located in the United States (<https://www.openconnectomeproject.org>) was tested from Frankfurt. The resulting speed was 0.059 mm/h. In conclusion, these measurements highlight the data rate advantage that webKnossos enjoys over CATMAID.




Figure 29: Screenshot of the webKnossos user interface in ortho mode. Shown are the three orthogonal viewports, arranged along the three principal axes. The fourth viewport shows a 3D view of the reconstructed cells. On the right an abstract view of the active tree is rendered. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

2.2.2 Flight mode

In the “classic” orthogonal view employed by webKnossos (as introduced by KNOSSOS – (Helmstaedter et al., 2011)), the top left view (also called ‘viewport’) shows the original images as recorded by the electron microscope and two other viewports show the reslice of the data in the XZ and YZ plane (Figure 29, called “ortho mode”). This view might not be optimal in two aspects. First, anecdotally, it was a major threshold for new annotators to understand this arrangement and to be

comfortable interacting with it. Secondly, while reconstructing, annotators are instructed to reconstruct a neurite in the viewport where the neurite appears most round (i.e. the viewport to which the neurite runs most orthogonally). This is necessary to avoid missing branch points, but slows down annotators because both the decision process which viewport to use and the reorientation after switching viewports constitute a mental effort (compare Figure 30a). Furthermore, for neurites running obliquely to all viewports (e.g. along the (1,1,1) direction), all viewports are suboptimal.

These problems can be solved by using a single viewport that can be freely rotated. This idea was implemented in webKnossos in a mode called “flight mode”. The current view is defined by a position coordinate and a direction. The viewport is rendered by sampling the raw data in a hemisphere which is then projected into a rectangular frame (Figure 30b - 30d). The orientation can be adjusted with the mouse or keyboard, additionally keys are defined to move backwards and forwards. During the movement, the nodes that make up the annotation are placed at regular intervals without additional user interaction.




Figure 30: Interaction and data projection in flight mode. Moving along a neurite, ortho mode necessitates abrupt changes in the orientation with which the annotator views the dataset (a). Flight mode only has one viewport that can rotate freely around the annotator (b, c). This is done by projecting the voxels that intersect a sphere centered on the current position into a square view (c, d). Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

Both orthogonal and flight mode are equipped to handle branched neurites. The workflow by which this happens is that when an annotator encounters a branchpoint, they press B to set a marker and then select one of the two branches to continue. Once the annotator reaches the end of that branch, they press J and their view is moved back to the view they had when they pressed B. Then they can continue with reconstructing the other branch.

2.2.3 Ease of teaching for different interaction modes

To investigate whether flight mode indeed makes neurite annotation easier to learn, a team of ten undergraduates (which had previously shown interest in working as annotators but had never used webKnossos) went through a regime consisting of a short introduction and then a short test reconstruction. The ten newly recruited annotators were split into two groups of five. They each watched one of two versions of a 2½ minutes long introductory video that explained ortho or flight mode, respectively. Then they were asked to reconstruct an axon for 5 minutes (all annotators worked on the same axon). Afterwards their achieved speed and committed errors were measured and it was found that the flight mode annotators had significantly outperformed the orthogonal mode annotators (ortho: 74µm/h, flight 165µm/h; for error rates see Figure 31).

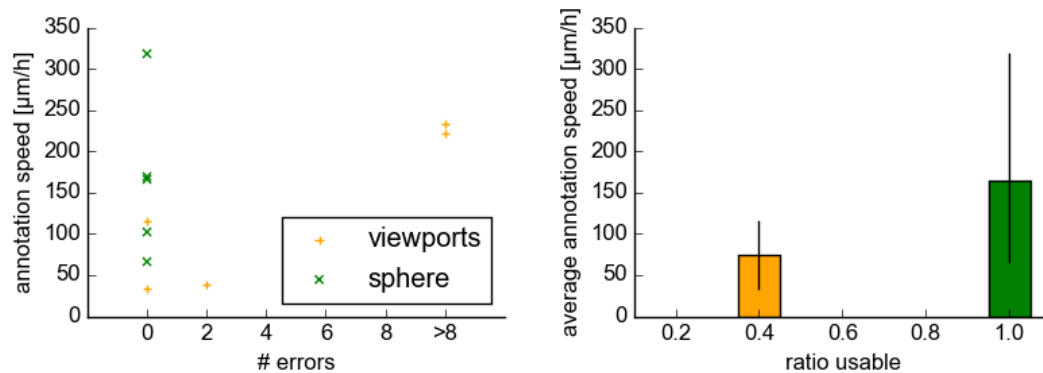


Figure 31: Comparison of almost naive annotators (2.5 minutes of training) performing in ortho and flight mode. It can be seen that all flight annotations were usable and the annotators working in flight mode outperformed their ortho counterparts in terms of speed. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

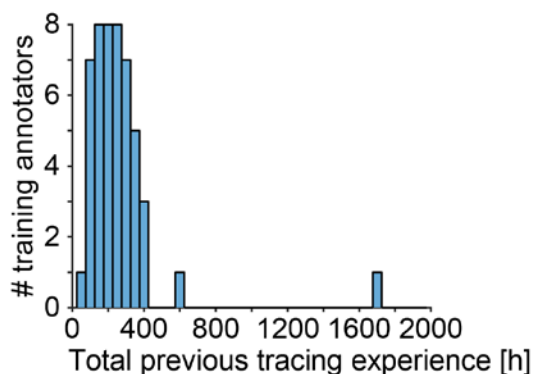


Figure 32: Distribution of previous annotation experience for annotators in the axon reconstruction tasks. Experience levels of annotators vary widely, from freshly trained to very experienced. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

2.2.4 Speed advantages of flight mode

It was then measured whether experienced annotators also profited from flight mode. From the center of the dataset, 40 neurites were chosen (32 axons) as training tasks and starting coordinates and orientations were selected for them. 51 annotators were randomly split into two groups (25 and 26 annotators, respectively). They came from a variety of experience (Figure 32). Both groups were asked to watch an 8½ minute long introductory video (different versions for flight and ortho mode). For the flight mode annotators, the flight mode video explained how this annotation mode works, the ortho mode video served as a refresher course on best annotation practices in ortho mode. Both videos emphasized the goal to speed up the annotation process as much as possible without sacrificing accuracy. Special focus was placed on the correct adjustment of the “move value”. The move value is the speed at which the annotator moves through the dataset if they keep the space bar constantly pressed and therefore the move value is the maximum speed an annotator can achieve. An excessive move value however can cause errors or inaccuracies while tracing.

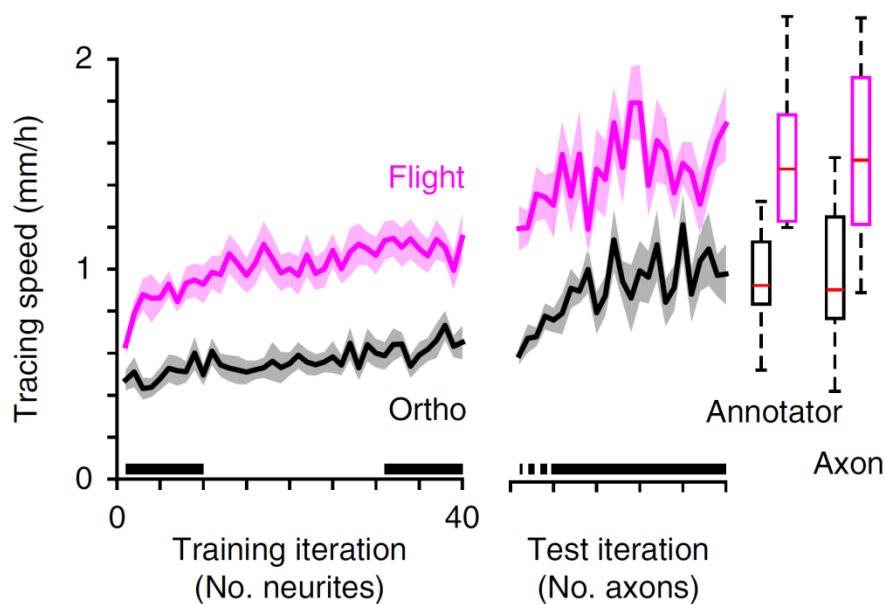


Figure 33: Speed development during training and test iterations. The black bars indicate the stretches for which the speed was evaluated, the broken bar the re-training iterations. Right: averages for flight and ortho annotators in the test iteration, averaged on a per-annotator and a per-axon basis respectively. Annotators in flight mode consistently outperform their counterparts in ortho mode. Both types of annotators increase their average speed during training and test iterations and the speed gains achieved during the training iteration are maintained over an 8 week gap before the onset of the test iteration. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

Then the annotators were presented with the 40 training tasks in a random order (random in a different fashion for each annotator). Their move value and average speed were measured on a per task basis. If the average speed exceeded 75% of their preset move value (i.e. if they moved more

than 75% of the time) their preset move value was insufficiently high and they received an email instructing them to increase their move value.




Figure 34: 20 axons selected for test iterations, all seeded from a $(2.5 \mu\text{m})^3$ bounding box roughly centered in the middle of the dataset. Random seeding from a densely reconstructed bounding box is superior to selecting random coordinates in the volume because the latter strategy has a bias to axons with a wider diameter. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

After all neurites had been reconstructed by at least 40 annotators each, 40 annotations of each neurite were randomly chosen (1600 annotations in total) and the speed of those 1600 annotations was measured. As can be seen in Figure 33 (left side), from the beginning the speed of flight mode annotations exceeded the ortho mode annotation speed ($0.84 \pm 0.05 \text{ mm/h}$ and $0.49 \pm 0.04 \text{ mm/h}$, $p < 10^{-4}$, Wilcoxon rank-sum test, averaged over the first 10 annotations). Additionally, the speed increased further during the duration of this training phase. The final speed was $1.11 \pm 0.07 \text{ mm/h}$

for flight mode and 0.63 ± 0.05 mm/h for ortho mode, again significantly higher (averaged over the last 10 annotations, $p < 10^{-5}$). The increase in speed was also significant ($p < 10^{-3}$, Wilcoxon signed-rank test).

After a pause interval of 8 weeks, it was investigated whether the flight annotators had retained their ability to outperform the ortho mode annotators in terms of speed. Out of the pool of the annotators that took part in the training iteration, 13 flight mode annotators and 13 ortho mode annotators were chosen (their speed performance was indistinguishable from the complete set of annotators in the last 10 iterations; $p = 0.246$, Wilcoxon rank-sum test). 20 axons were randomly selected from the center of the dataset (Figure 34, see methods for randomization procedure). These were shuffled (again on an individual per annotator basis) and paired with an initial 5 axons from the training set (randomly chosen but not shuffled) to give the annotators an opportunity to reacquaint themselves with the task. All annotators were instructed to complete the 25 tasks.




Figure 35: Overlay of ground truth and annotation tree. The two sides contrast the difference between an error that only caused the loss of a branchlet (which carried only one synapse, and therefore only induces a single synaptic connection error; left panel) and a continuity error that can have far-reaching consequences (the branch leaves the dataset - in a larger dataset this could potentially have meant losing hundreds of synapses; right). Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017




Figure 36: Comparison of error rates for axons in flight and ortho mode and dendrites in flight mode. Separately plotted are errors that affect the continuity of the neurite outside of its local surround (Cont. err.) and those that do not. The dashed line marked error rates normalized to the length of the overlap tree of the ground truth and the actual annotation tree, the full line shows the error rate normalized to actual annotation tree length. (Black: ortho axon, pink: flight axon, turquoise: flight dendrite. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

After they had finished, the average speed and the error rate for the test annotations were evaluated. Again, the flight mode annotators performed better (1.51 ± 0.04 mm/h for flight, 0.96 ± 0.03 mm/h, Figure 33, right side). The error rate was analyzed on a subset of 10 axons (randomly chosen). For the error rate annotation, a previously developed method was used (Helmstaedter et al., 2011). Briefly, an expert annotator created ground truth annotation for the 10 selected axons. Then overlays were plotted of the annotations and the ground truth, projected in the three cardinal axes. In these annotations all errors were marked; separately errors larger than $10\mu\text{m}$ and less than $10\mu\text{m}$ (Figure 35). Major errors can cause a whole branch of the axon to be lost, potentially affecting the correct identification of hundreds of synaptic contacts (therefore they are called continuity errors). Discrepancies smaller than $1\mu\text{m}$ were ignored. No significant difference between the error rate for flight mode and ortho mode annotations could be identified (all errors: $p > 0.34$, major errors: $p > 0.28$, Wilcoxon rank sum test). Both error rates were comparable to the error rates of published results (Helmstaedter et al., 2011; Hua et al., 2015).

2.2.5 Redundant flight mode annotations

Single annotations have errors at a rate that can hamper connectomic analysis. If several annotations of the same neurite are created, it is possible to use an algorithm called RESCOP to create a consensus set of annotation pieces ((Helmstaedter et al., 2011), compare Introduction chapter 1.7). To evaluate how much redundancy improves error rates over single annotator performance for the test annotations, the RESCOP algorithm was employed on the test axons for redundancies ranging from 2 to 13 and the same error rate annotation process was employed (Figure 36).




Figure 37: Relationship between error rate and speed for ortho and flight mode axons and flight mode dendrites: Plotting error rate vs speed shows both for continuity error and all errors that there is no decrease in accuracy as speed increases. This was evaluated on 30 flight axon, 30 ortho axon and 30 flight dendrite annotations. Color code as in Figure 36. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

Additionally, 497 dendrites were reconstructed in flight mode by the initial group of 51 annotators. The reconstruction speed exceeded the axon reconstruction speed (2.11 ± 0.16 mm/h, see lower panel in Figure 37, reconstruction without spines) A random subset of 10 dendrites was chosen and

the error rate was measured as described above (Figure 36, non-redundant annotations had an error rate of 2.7 +/- 0.69 per mm).

To investigate whether the error rates increased for annotations that had been executed with a higher speed, the speed and error rates for the axon and dendrite tracings were correlated. No positive correlation could be found (Figure 37; ortho: $r = -0.5$, $p = 0.007$; flight $r = -0.4$, $p > 0.04$ for test axons; Pearson's correlation), the data even suggest that the faster annotators committed fewer mistakes.

In conclusion, the instruction to annotators to speed up their reconstruction was successful in the sense that it did not increase error rates compared to previously published studies but did increase speed. On top of that, the introduction of flight mode created an additional speed gain and makes learning of webKnossos easier.

2.2.6 Fast synapse annotation in webKnossos

For the connectomic analysis of the set of reconstructed dendrites and axons, it was necessary not only to reconstruct the neurites, but also the synapses between them. So far, the focus had been on reconstruction speed, now it was investigated whether the synapses could be added without negating the speed improvement. To annotate synapses, two strategies were considered: Either to follow the axons and mark all the synapses on them and then identify the respective postsynaptic dendrites; or to follow the dendrites and mark all the synapses and then identify the respective presynaptic axons.

To reach a conclusion which strategy would be more efficient, it was measured which fraction of axons were close enough to reconstructed dendrites ($<5 \mu\text{m}$) to be able to form synapses onto them and which fraction of dendrites were close enough to axons to be innervated by them. If longer stretches of axons and dendrites run through areas where the opposite neurites are not present they wouldn't have to be synapse-annotated. However, this was not the case, the vast majority of axons or the majority of dendrites would actually have to be synapse-annotated. Because more reconstructed dendrites than reconstructed axons were available (93.6 mm dendrites, 4.55 mm axon), it was decided to annotate the synapses on the axons (see below for a more general discussion of this trade-off).

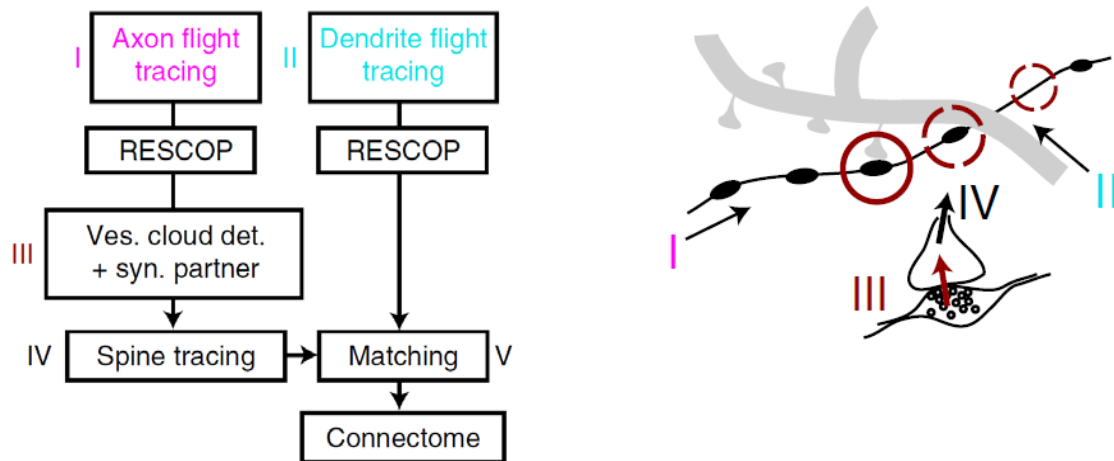


Figure 38: Work flow for whole-connectome axon-first synapse annotation. After axons and dendrites have been redundantly reconstructed (I and II) and a consensus version has been created using the RESCOP algorithm, the consensus axons are used as a template to search manually for synapses along them (III). Then reconstructions are started in the marked postsynaptic volume (IV). This leads to a reconstruction of the spines. In a last automatic step (V), the spine reconstructions are matched to the reconstructed dendrites and the connectome is assembled. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

A new interaction mode for webKnossos was developed where it was possible to follow an existing axon (“template” axon, see methods) and mark the postsynaptic neurite if one encounters a synapse. The output of a 6-fold-redundant RESCOP run of the 32 axons of the 40 training neurites was used as the template axons along which to annotate synapses. 10 annotators from our student team were trained for one hour on how to use the synapse annotation mode and how to spot inhibitory synapses. After that, they were asked to reconstruct the synapses onto the 32 axons. Inhibitory axons synapses present unique challenges and therefore the annotators were asked to mark suspected axon as inhibitory. For the axons marked as inhibitory, another annotator was asked to mark their synapses, with the instruction to prioritize precision over speed. The average time consumption for excitatory axon synapse annotation was 1.2 ± 0.5 h/mm, (including inhibitory axons 1.8 ± 1.0 h/mm).

2.2.7 Fast spine reconstruction in webKnossos

After that, markers were present in each postsynaptic volume that was relevant for the connectome, but the majority of them (90%) was in spine heads and not yet connected to their respective dendrite (i.e. the identity of the postsynaptic dendrite was unknown). Therefore a group of three annotators reconstructed the neurites starting at these markers (in ortho mode) until they encountered a dendritic backbone and then place three more nodes in that backbone. For shaft synapses this meant directly placing three nodes. After this task was completed, an expert annotator

measured the error rate in the spine tracings (3.3%; 1 of 30). The average spine annotation took 31.1 +- 28.0 seconds. Per length of axon from which these tasks were created, the process took 2.3 ± 1.3 h/mm.

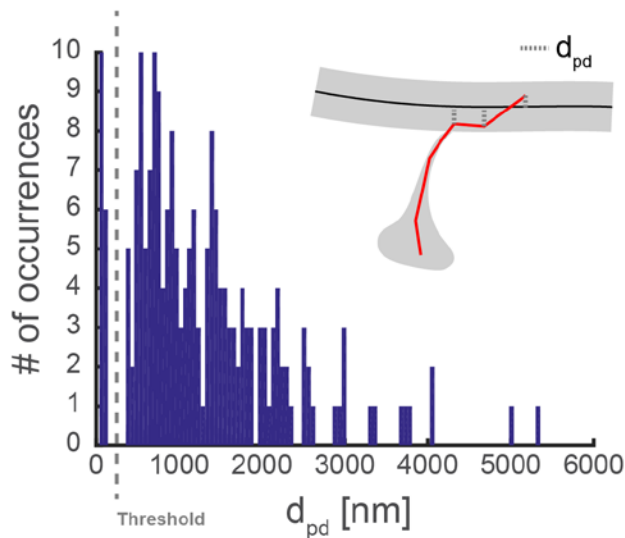


Figure 39: Distribution of d_{pd} values for spine attachment. The d_{pd} value measures the average distance between the last three nodes of the spine annotation and the dendritic backbone nodes (inset). For the training dataset (shown) a threshold of 250 nm splits the set correctly into attached and non-attached spines. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017




Figure 40: Resulting connectome matrix for 32 axons and 497 dendrites. Rows correspond to axons, columns to dendrites. Only dendrites innervated by the 32 axons are shown (Inh: inhibitory innervations, Exc: excitatory innervations). Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

The average distance d_{pd} of the last three nodes of each spine skeleton to each dendrites (dendrites at threefold redundancy) was measured and the dendrite that had the lowest score was selected as the candidate dendrite for that spine (Figure 39).

A training sample of 200 of these pairs of spine reconstructions and candidate dendrites was created. Within that training sample a threshold $d_{pd}^* = 250$ nm correctly split the training set into dendrites with their attached spines and spine annotations whose dendrite is not part of the set of 497 annotated dendrites, which therefore play no role in the connectome. An expert annotator then evaluated the performance of this classifier on a test set and found only one error (21 correctly attached, 1 wrong attachment, 178 correctly unattached).

Using this algorithm, the connectome of the 32 axons and 497 dendrites was created (Figure 40, Figure 41).




Figure 41: 3D display of resulting connectome. Dendrites in turquoise, axons in yellow. Synapses ($n = 104$) marked as spheres, excitatory synapses violet and inhibitory synapses red. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

2.2.8 Connectome strategies

With the connectome assembled, it was possible to investigate whether the initially used axon-first synapse annotation strategy was optimal (and more generally, for which densities of axons and dendrites this is the right strategy). The empirical measurements for axon synapse annotation and spine reconstruction speed was used and a dendrite-based synapse detection speed of 11.1 h/mm was assumed (including spine annotation and reconstruction). Additionally, a sparse and dense regime can be distinguished (as discussed above, a sparse regime is where parts of the synapse annotation can be omitted because there is no suitable synaptic partner close-by). In Figure 42 it is outlined for which density of axons and dendrites an axon-first and a dendrite-first strategy is faster and which densities lead to a sparse regime.




Figure 42: Ratio of connectome annotation time spent on synapse annotation (including spine reconstruction) compared to neurite reconstruction. The denser the reconstructed neurites fill the volume, the more time has to be spent on synapse annotation. The pink line shows the threshold below which it is more efficient to use the axons as a template to annotate synapses, above the line a dendrite-first strategy is more efficient. The dashed lines show the threshold below which not all stretches of template neurite have to be reconstructed because for some template neurite stretches the local surround contains no available partner neurites. The circle marks the situation for the local connectome from Figure 40, Figure 41. The cross marks the situation for a hypothetical L2/3-L4 cortical connectome as shown in the inset in Figure 43. The arrow in the color bar shows the ratio of time during the creation of the local connectome that was used for synapse annotation. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

As it can be seen, the decision to use an axon-first strategy for the 497 dendrites and 32 axons was correct (dendrites with redundancy 3, axons with redundancy 6). For a larger hypothetical annotation project consisting of 300 L2/3 cortical pyramidal cells (dendrites) and 100 L4 cortical spiny stellate cell axons (inset in Figure 43) the axon-first strategy is faster, too. That project is sparse

in both senses (not all dendrites have reconstructed axons close by as presynaptic partners and not all axons have dendrites close by as postsynaptic partners). What both the 497 dendrite-32-axon project and this have in common is that only a small fraction of the work load is spent on the synapse annotation (Figure 43). This means that the improvements in axon and dendrite reconstruction speed as reported here (1.5 and 2.1 mm/h respectively) are still valid even if synapse annotation time is taken into account.




Figure 43: Time comparison of relative reconstruction effort for the ‘local’ connectome (Figure 40, Figure 41) and for a hypothetical connectome of 300 L2/3 pyramidal cell dendrites and 100 L4 cortical axons. Shown in turquoise is the time consumption for dendrites, in pink for axons. The time demand for synapses is split up into synapse detection (crimson) and spine reconstruction (black). Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

2.3 Analysis of a *Drosophila* lobula plate SBEM dataset

2.3.1 Stitching and alignment

A protocol to stain *Drosophila* optic lobe for 3D EM was developed and a SBEM dataset was acquired by Ch. Kapfer containing the complete lobula plate, large parts of the lobula and a long stretch of tract medial of the lobula plate. At the time of the final stitching and alignment, this *Drosophila* lobula plate dataset was the biggest SBEM dataset that had so far been recorded (160 x 200 x 182 μm^3 at a resolution of 11 x 11 x 25 nm^3), a size which only recently has been surpassed (compare 2.4.5 for efforts in mouse barrel cortex). To achieve such dataset dimensions, massive tiling of the field of view was employed, mostly in a 7 x 11 images grid (see Table 3, page 93). This meant that afterwards over half a million images had to be stitched and aligned. While the techniques outlined in 1.13 could be used for this effort, additional changes were necessary due to the larger and more complicated layout of the stack.

2.3.1.1 Stack layout

The stack was recorded cutting from lateral to medial, starting about 24 μm laterally from the onset of the lobula plate. In this cutting plane, the lobula was already present. The onset of the lobula plate and the border line between lobula and lobula plate could be easily distinguished by a zone with wider processes and tracheae (Figure 44). Due to the specific anatomic nature of the sample, it was necessary to change the grid layout of the tiling motion several times during stack acquisition. This had to be compensated for by the alignment script (see 3.2.3)

2.3.1.2 Epoxy correction

The structure of lobula and lobula plate was surrounded by a layer of somata, which in turn was surrounded by epoxy (see Figure 44). Because the epoxy had a very low conductivity, strong charging effects could be seen. Furthermore, the epoxy had no intrinsic structure that could be imaged with a SBEM setup. In those cases the alignment algorithm could not make a meaningful approximation as to what should have been the correct stitching. Therefore these regions were detected and separated from the rest of the alignment (3.2.3).




Figure 44: Lateral view onto *Drosophila* lobula (anterior, right) and lobula plate (posterior, left), slice 1660. The plane has been stitched by combining 7 x 11 single images. The lobula and lobula plate are surrounded by a layer of somata, which in turn is surrounded by epoxy. The epoxy has very little structure that the stitching algorithm can analyze. Therefore epoxy stitching is not gapless, but due to decoupled optimization of neuropil and epoxy stitching the main structure is unperturbed.

2.3.1.3 Outlier detection

As described in chapter 3.2, the algorithm performs best if the shift detection receives cut-outs that are already well pre-aligned. Ideally the cut-out is selected in a way that the mean detected shift in the cut-outs is zero. Because the motor position jitter was small for this experiment, it was possible to pre-align images to a precision better than 200 nm. In the cortex stacks, the equivalent procedure made complete failures of the shift detection so rare that all remaining failures could be inspected manually (see 3.2.1). In this lobula plate dataset, there were two challenges to this approach: Firstly, the number of images was much higher. Even the larger one of the two cortex datasets only had a tenth of the images of this dataset - although the single images had a wider field of view. Secondly, contrast of the recorded tissue was not as good as for the cortex datasets, probably due to the challenges of staining fly nervous tissue in a way that is compatible with SBEM requirements (compare 1.4). Therefore, after a run where all cut-out positions had been optimized, still more than 50.000 shift value residual errors were more than 10 pixels. It was considered unfeasible to inspect

all problematic spots manually. Instead, a heuristic was defined that detected unreliable shift calculations and reduced the weight of that measurement in the least square relaxation (3.2.3).

These heuristics increased the alignment quality so that annotation was possible. A remaining issue was the set of very thin processes that connect the main branch of the LPTCs with their somata. Here the use of expert annotators was necessary to compensate for a remaining jitter in the alignment.

2.3.1.4 Mixing of alignment and annotation

An initial run of the alignment had been done before the acquisition of the stack had terminated. Due to the long duration of the stack acquisition it was deemed desirable to start with the annotation of cells with the data from the initial run. A challenge with this approach was that the least-square relaxation finds a global optimum and integrating the additional information of the later slice layers would move the position of all images. Therefore the results of the initial run were stored and were compared with the results from the final run that had included all slices.

This allowed creating an algorithm that assigned a voxel coordinate in the new output dataset to every voxel coordinate in the preliminary dataset. After the release of the final output dataset to all annotators, all skeletons that had been reconstructed in the first dataset were moved into this new coordinate system and no annotation effort was lost.

2.3.2 Annotation of lobula plate tangential cells

2.3.2.1 Explorative seeding by diameter

The first step to reconstruct cells in a 3D EM dataset is defining seeds from which the tracing can start. Processes of tangential cells in the lobula plate have diameters that exceed the diameters of other processes in the region. Therefore processes were selected from the lobula plate that had big diameters and annotators were asked to reconstruct the full trees of those processes. With this method, 3 cells were identified that corresponded to morphologies associated with HS neurons (Figure 45, left panel) and 3 cells were identified that corresponded to VS neurons (Figure 45, center panel). While this showed that the technique was viable, the large number of missed branches made it clear that one annotator per skeleton was not enough and redundant annotations would have to be employed (see 1.7). Furthermore it became apparent that to facilitate comparison to light-microscopic images (e.g. (Scott et al., 2002)) diameter information would be helpful in the skeleton reconstructions.

Therefore, a new annotation mode was developed, which used a feature that only KNOSSOS offered at that time: For each node, the annotator could set the diameter of the node freely. The skeleton of one VS cell was annotated redundantly ($n=3$) and a consensus skeleton was created from these multiple annotations. Then every 4th node along the branches was marked and a specifically trained annotator changed the diameters of the marked nodes so that they fit into the processes. It was ensured that all nodes around branch points were diameter-annotated, even if that locally broke the every-4th-node pattern. If the cross-section of the process was very anisotropic, the annotator was asked to aim for an equivalence of area. Then, a spline algorithm was used to interpolate the diameter between the marked nodes. The tree could then be converted into a HOC file and rendered in Amira (Figure 45, right panel).




Figure 45: Non-redundant reconstructions of 3 HS and VS cells (left and center panel, respectively). Volume reconstruction of VS cell (right panel) While many small branches have been missed in the non-redundant reconstructions, the overall morphology can be recognized. To increase ease of comparison with existing light-microscopical images, the diameters of nodes were reconstructed (exemplified for a VS cell, right panel). From this a rendering of a simplified volume model was created (scale bar 20 μm)

2.3.2.2 Transition of annotation effort to webKnossos

The first run of the alignment and stitching algorithm had created an output dataset that would fit onto a hard drive. Therefore it was possible to use KNOSSOS to distribute the dataset to all annotators. The second and final run of the dataset output was larger than 2.1 TB (compare 2.3.1.4) and would not fit on the hard drives that were deployed. Therefore this was considered an ideal use

case for the newly developed online annotation tool webKnossos (chapter 2.2). After implementing an extension that allowed to diameter-annotate nodes, it offered a much more convenient annotation environment, with automatic task distribution to annotators and remote dataset delivery.

2.3.2.3 Complete and bundle-based reconstruction

After the trees had been converted to the coordinate system of the final dataset (see 2.3.1.4), the skeletons could now be seen in a context where the bundle was present that carries processes from the LPTCs antero-medially. Therefore, bundle locations were included in the new seeding strategy.

To make sure that all large cells would be included, in 5 layers in the dataset (88 μm , 95 μm , 106 μm , 147 μm and 175 μm from the lateral end of the datasets) all processes present that had an increased diameter (first four planes) or were part of this bundle (two last planes) were marked for reconstruction. Processes were considered to have an increased diameter when they were still visible after a subsampling of the dataset by a factor of 16. All in all, 628 seeds were generated this way. This meant many annotations would be traced redundantly. As outlined in 2.3.2.1, this redundancy was desired - and due to the large number of available annotators not a problem. After the annotation was successfully completed, expert annotators created a library of all the cells thus reconstructed, joining redundantly-seeded trees and building consensus skeletons. This resulted in a set of 43 cells, which were then diameter-annotated. To increase volume quality, every node was then reannotated with diameter information.

2.3.2.4 The set of HS neurons

With this set of volume-reconstructed neurons, it was possible to draw analogies to the known horizontal cells in *Drosophila* (see Figure 5). A set of three cells was identified that corresponded to the known HSN, HSE and HSS cell (Figure 46). As can be seen, even at 4-5 fold redundancy, the annotators struggled with finding some smaller branches of the cell, probably due to low dataset contrast. In the top view of the reconstruction, the outlines of the lobula are marked in red, highlighting how the HS system cells branch in the most anterior layer of the lobula plate.

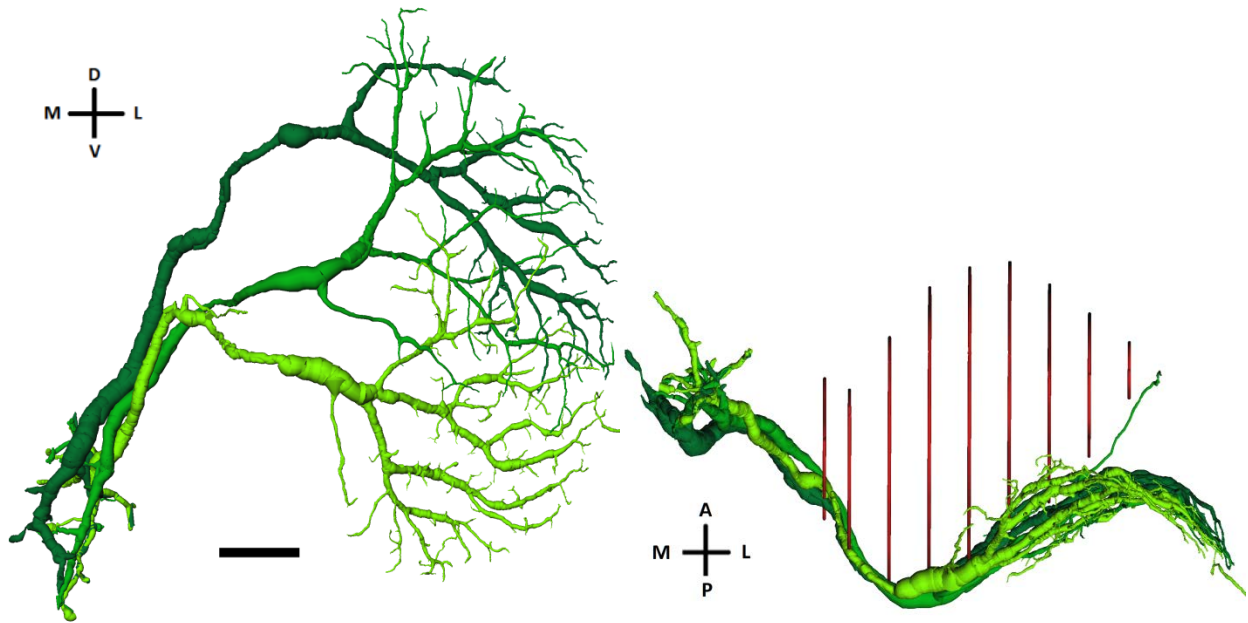


Figure 46: Volume rendering of redundant annotations of HS cells in posterior (left panel) and dorsal view (right panel). With redundant annotation and annotation of the projecting tract, morphologies of HSN, HSE and HSS can be seen more clearly. Outline of the lobula (red) shows stratification of HS cells close to lobula (most anterior layer of the lobula plate). The HS cells were contained in their entirety within the dataset, with no external projections that could not be reconstructed. Scale bar 20 μm

2.3.2.5 The set of VS neurons

The general structure of VS cells was already known ((Scott et al., 2002), Figure 6), it was also possible to identify the equivalent cells in the EM dataset (Figure 47). To make sure that identified structures were not projections from other regions of the brain, for the cells also soma reconstruction were attempted. The somata were densely embedded in the layer of somata around the lobula plate (right side in Figure 44) and were only connected by very thin processes with the rest of the cell. Because of alignment problems in this region (compare 2.3.1.3), this reconstruction could only be done by expert annotators. It was done separately and the annotations were joined afterwards.

Furthermore, there were three additional cells that were morphologically very similar to VS1-6 (Figure 48). If those were indeed VS cells, this would mean that the Gal4-3A driver line used in (Scott et al., 2002; Raghu et al., 2007) does not capture all VS cells. However, for one of them no soma could be identified so far.




Figure 47: Volume rendering of redundant annotations of VS cells in overlay (left) and single view (right). 6 candidates for VS cells were reconstructed, all redundantly by 4-5 annotators. Then, the diameter of all nodes was adapted to match the diameter of the neurite at the position of the node. Expert annotators then identified and reconstructed the somata and attached them to their respective cells. Left: scale bar 50 μm . Right: VS1-VS2 top row, V3-V4 center row, VS5-VS6 bottom row, size reduced 4x




Figure 48: Overlay of further VS-like cells with similar VS cells. Three VS-like cells were identified. The VS-like cells shown together with similar VS cells. They were identified as VS-like cells because their neurites ran in the same bundle as the medial part of the VS cells, because they were contained within the dataset and (optionally) because their soma could be identified in close proximity of the somata of the VS cells (scale bar 50 μm).

2.3.2.6 Candidates for CH cells

While several bodies of work were unable to identify CH cells in *Drosophila* (Scott et al., 2002; Raghu et al., 2007), this does not mean that they do not exist: The expression system used (Gal4-3A) might just not include them. Indeed, three candidates were found that might correspond to DCH or VCH. They roughly have the right morphology and form their dendrites in the right layer (Figure 49). In *Drosophila* more large somata that could belong to LPTCs have been identified than the number of cells that were visible in the Gal4-3A line (19 versus 9, A. Borst, personal communication). Therefore additional cells should exist and these CH-like cells are plausible candidates.

Together with the morphologies presented above, this highlights the possibilities of the dataset as well as the challenges: While correspondence of identified cells to classes identified in earlier work is not easy and cells need multiple annotations to yield high quality skeletons, the perspective to identify all lobula plate tangential cells in an unbiased fashion is exciting.




Figure 49: Known HS cells (left), candidates for CH cells (center) and overlay (right): while CH cells have not yet been conclusively described in *Drosophila*, many studies were able to identify them in *Calliphora* and other blow flies, e.g. (Hausen, 1976; Eckert and Dvorak, 1983; Egelhaaf et al., 1993; Gauck et al., 1997). Scale bar 50 μm

2.4 Continuous imaging

2.4.1 Tiling and continuous imaging

The easiest way to increase the field of view with a mechanical stage is called tiling. For this, images are taken that are smaller than the overall size of the region of interest. After one image has been acquired, the sample is moved mechanically so that the microscope is now centered on a region next to the one originally imaged. The movement is so configured that there is a small region of overlap between the two images that facilitates stitching (compare 3.2). This is continued in a grid-like fashion, until the whole region of interest is scanned (Figure 50, left panel).




Figure 50: Comparison of alternative methods for field of view extension: Tiling (left) and continuous imaging (right). Because the scanning electron microscope has a limited field of view, for large datasets the size of the overall field of view is extended by mechanical means. Using tiling, image acquisition is interspersed with abrupt movements to a new position (1, 2, 3, 4). This movement causes sample vibration and therefore necessitates long settling times (several seconds). For continuous imaging, the line feed of the scanning electron microscope is replaced with a smooth and continuous mechanical movement. This largely removes the need for tiling in one dimension.

This process limits the possible speed: The imaging beam is shut off during movement. Therefore the movement has to be fast. This means that substantial accelerating and decelerating forces are applied. This problem is especially acute in diamond-knife SBEM because the stage includes the cutting apparatus and therefore is heavy. For the stacks in chapter 2.1.1 the wait times had to be increased to several seconds between end of movement and start of image acquisition to prevent image artifacts.

Here a continuous imaging technique for scanning electron microscopy is presented. Continuous imaging moves the sample during imaging. This is routinely done for atomic force microscopy (Binnig et al., 1986), where the sample is moved beneath the cantilever. All scanning is done by a piezo stack mounted under the sample. A solution like this would not work for scanning electron microscopy because the line scanning of the scanning electron microscope is too fast: Even in a 1 mm sample, a row at 10 MHz would only take 10 ms (Equation 5).

Equation 5

$$t_{line} = \frac{width_{sample}}{f_{pixel}width_{pixel}}$$

The second dimension, the line feed, is much slower and therefore a continuous imaging solution was implemented that operated only in that dimension. This created a setting where long sweeps of the sample were imaged without interruption, and these long sweeps were put together by conventional tiling in the other dimension (Figure 50, right panel). The sweeping movement has to be very smooth to not to disturb the imaging and a tight integration between the movement and imaging software is necessary.

2.4.2 Piezo setup

Piezo stacks are comprised of piezo crystals that extend and compress driven by piezoelectricity. Therefore they do not have any discontinuity in their movement. Unfortunately, the biggest piezo stacks that were available to be fitted inside the chamber of the microscope had a range of 200µm. Therefore the piezo action was levered to extend their range.

A second concern was the stiffness of the movement system. The resonance frequency of a mechanical system like this can be approximated as the resonance frequency of a spring pendulum (Equation 6), where m is the mass of the stage and D is the stiffness of the motion system.

Equation 6

$$f_{spring} = \frac{1}{2\pi} \sqrt{\frac{D}{m}}$$

To prevent image artifacts, a stiff system is desirable. The stiffness depends on the levering, and the final stiffness is decreased by the levering factor according to Equation 7.

Equation 7

$$D_{final} = \frac{1}{levering^2} D_{raw}$$

Experimentally, a compromise between stiffness and range was found (see 3.4) that provided

sufficient stiffness and a range of 350 μm . This arrangement was built in a symmetrical fashion so that both axes were equipped with a piezo. The second axis was used to move the sample laterally between acquisition stripes (see Figure 50). To extend the field of view, both piezo actors were combined with geared motors (see Figure 62).

2.4.3 Scanning pattern

Combining the piezo stack and motor movement allowed fast acquisition over a large region of interest. The piezo system was used to create the stripes for the continuous imaging paradigm and to move laterally between the stripes. For stacks larger than 350 μm x 350 μm , the motors were used to split up the field of view into motor tiles (e.g. for a 600 μm x 600 μm whole barrel dataset 2 x 2 motor tiles). Motor movements did only add a couple of seconds to the overall per-slice time and were therefore considered acceptable.

A movement pattern was defined which the piezo actors would execute during imaging. The speed the stage has to move at during acquisition is given by Equation 8. Abruptly commanding the piezo actor to move at that speed would create forces that exceed the stiffness of the apparatus and therefore lead to shaking and imaging artifacts. Therefore no abrupt changes in movement speed were implemented. The final movement speed of the strip was reached by a linear ramp-up of speed and the scanning started as soon as the stage had reached the necessary speed (Equation 8)

Equation 8

$$v_{strip} = \frac{height_{pixel}}{t_{line}}$$

After the first strip had been acquired, the direction of strip movement was reversed so that the sample moved in the other direction, the absolute speed being the same. Furthermore, the stage was moved to the side. This was not done abruptly, but with a ramp-up and –down of the lateral speed. When these two processes were combined, the movement of the sample was semi-circular and without abrupt forces. Doing this over several strips resulted in a snake-like pattern (Figure 51).




Figure 51: Overall scanning pattern for large field of view stack acquisition, combining tiling and continuous imaging. This combination allowed for nearly unlimited fields of view acquired at high speed. Because the piezo system was built symmetrically in x and y , the piezo stack operating in the axis orthogonal to the continuous imaging was used to actuate the tiling for several stripes next to each other, thus saving about 50% of tiling time. Four of these patterns are combined to image the whole region of interest in the barrel dataset.

2.4.4 Autofocus for large samples

The electron microscope has a focal depth of several micrometers. While this was sufficient to allow for a global focus point for the whole surface of smaller sample (see 2.1.1), a sample that is several hundred micrometers wide becomes very sensitive to tilts of the cutting and movement apparatus. Therefore the autofocus that was described in 3.1.2 was split up to operate on four separate regions, searching for an independent focus point in all four motor tiles. This had the added benefit that the quality of the autofocus could be monitored by testing how stable these four focus points were in relation to each other (Because the difference between them is driven by stable geometric circumstances, the relative distance would not change even if the overall focus distance drifted).

2.4.5 Full barrel stack

2.4.5.1 Sample preparation

As discussed in 1.5 and 3.1, the process of preparing a sample for SBEM involves the infiltration of heavy metal compounds into the sample. For the samples for chapter 2.1 this was done as described in (Briggman et al., 2011), combining previously developed protocols (Seligman et al., 1966; Karnovsky, 1971; Walton, 1979). To assess placement of the small samples in barrel cortex, cytochrome oxidase staining was applied to the tissue (Figure 9). The cytochrome oxidase stain was applied to the same tissue position contralaterally. This was sufficiently accurate to confirm

placement in the barrel cortex but not to confirm placement of the stacks within a single barrel. Therefore a method was developed that allowed to extract a single barrel, stain it for SBEM and then report which barrel it is.

To create a stack that contained one whole intact barrel, the sample had to be at least 300 μm in diameter (if perfectly centered, see Figure 52, left panel). In a first attempt, a circular punch was used to extract a sample 1 mm in diameter (see 3.5). For geometric reasons it was certain that this sample contained at least one complete barrel (Figure 52, left panel). A post-hoc cytochrome oxidase stain proved successful placement (Figure 52, right panel) in the barrel cortex.




Figure 52 Left: Barrel field with 1 mm and 500 μm punch diameter (left panel), adapted from (Jan et al., 2008), permission granted through Creative Commons Attribution License 4.0. Right: Cytochrome oxidase stain of slice with 1 mm punch executed.

With a 1 mm punch it was certain that the sample would contain an intact barrel. With a 500 μm sample, the success rate was found to be much lower. Stereotactic targeting of a punch reliably created samples in barrel cortex, but cytochrome oxidase staining quality impairment due to traumatic punching was visible.

When attempting stains of 1 mm punch samples, the staining results were unsatisfying (Figure 53, left panel). The heavy metal staining agents did not fully penetrate the sample, resulting in an intermediate region with excess stain and ultrastructural damage and a center region with too little stain. It was found that without changing the staining protocol, 500 x 500 x 500 μm^3 was the maximum block size to allow homogeneous staining. Reasons for this behavior and remedies have in the meantime been described by (Hua et al., 2015). Therefore using a different kind of punch was attempted, which had a diameter of only 500 μm (see 3.5). The success rate with this kind of punch was very low. The punches had substantial wall thickness which caused tensile stress on the tissue, a problem that was worsened by the small diameter of the punch. The tissue was prone to tearing instead of getting cut (see Figure 53, right panel). The low yield rate was exacerbated by the fact that the probability of hitting an intact barrel with a circle of 500 μm was low; the few mechanically successful sample extractions didn't contain a whole barrel (see Figure 52, left panel). Furthermore,

the traumatic force exerted by the punch created a problem: the punch caused ultrastructural damage in the outer layers of the sample (visible in the EM) and a zone around the punch was unsuitable for cytochrome oxidase staining, making it hard to judge correct positioning (compare Figure 52, right panel, for blackening effect).




Figure 53 Left: Example of a center of an EM-stained sample generated with a 1 mm punch. Right: Suggested explanation for low success rate of sample extraction with 500 μm punch.

With the EM stain employed for this study, it was found that homogeneous staining of 1 mm punch samples was not possible. The outer region of the samples was stained correctly, but the center not sufficiently and an intermediate region contained excessive staining, while also showing tissue damage (left panel). On the other hand, extraction success rate was very low for 500 μm punches. The walls of the punch were very thick, quickly extending to 70 μm behind the cutting edge. This caused tensile stress which ripped the tissue apart before being cut (right panel).

It therefore became clear that a fundamentally new approach was needed. A stereotactical apparatus was constructed that employed a mounted scalpel blade to set a small cut at a defined position. All further cuts were done in a vibratome (see 3.5), because vibratome cuts did not reduce the quality of cytochrome oxidase staining even in close proximity to the cut (see Figure 54). After aligning the vibratome to the orientation cut that was done in the stereotactical apparatus, a 500 μm thick hemi-slice of barrel cortex was extracted by a series of rectangular cuts. This hemi-slice was expected to terminate at the medial end of barrel cortex (Figure 54). The remaining parts of barrel cortex were then sliced tangentially (parallel to pia) into 75 μm thick slices and a cytochrome oxidase stain was performed on them.

Due to the gentle nature of the vibratome cuts the cytochrome oxidase stain had sufficient quality in the vicinity of the region of interest and this could be used to target the sample directly onto a barrel: After the results from the cytochrome oxidase stain were available, it was decided how much tissue had to be removed on the medial side of the slice (in the example of Figure 54: 0 μm). The

medial piece was (if necessary) cut off with the vibratome, followed by another cut 500 μm more laterally. Lastly, the sample was rotated and two vibratome cuts were made at a depth of 150 μm and 650 μm relative to the pia. This left a cube with an edge length of 500 μm on a side, whose surfaces all had been cut with the vibratome.

This technique allowed a yield of about one targeted barrel sample per 5 animals. One of the successful sample extractions (Figure 54) was used for the stack described in 2.4.5.2.




Figure 54: Overlay of cytochrome oxidase stain of tangential slice of barrel cortex and posterior-lateral region, with manual barrel annotation and position of final barrel sample.

With the vibratome method, $(500\ \mu\text{m})^3$ samples could be generated that contained an intact non-border barrel sample. Here a 70 μm tangential slice stained with cytochrome oxidase is shown, which was created after the coronal 500 μm thick hemi-slice was removed. The posterior-lateral block was temporarily removed to gain access to the hemi-slice and the corresponding cytochrome oxidase images were reassembled. In this specific case, the most medial part of the hemi-slice was selected as the final sample, but this can be adjusted depending on the results of the cytochrome oxidase barrel pattern.

2.4.5.2 Stack acquisition

Using the methods for continuous imaging, a stack was recorded that contained the layer 4 of a barrel. The stack was acquired at an average data rate of 6.3 MVx/s. This was a 20-fold speed improvement compared to the retina stacks recorded (Briggman et al., 2011) and was only possible due to the reduction in tiling time with continuous imaging (Figure 55).

During stack acquisition, some issues were observed that limit reconstructability of the thinnest processes in the dataset (see Figure 56 for an overview and 3.5 for a detailed description).




Figure 55: Comparison of time contributions to acquire a single $300 \times 300 \mu\text{m}^2$ slice in a SBEM setup, if the imaging speed is 10 MHz and continuous imaging is employed.

For the layer 4 mouse barrel cortex stack, the continuous imaging can be used to dramatically reduce acquisition time. The tiling time had been limiting before, reducing the data rate from a raw 10 MHz imaging rate to an effective acquisition rate of 880 kHz. With continuous imaging employed, imaging time was restored at the dominant factor, but at a faster overall level, allowing for an effective acquisition speed of 6.3 MHz

Having finished the stack acquisition (after a total duration of 48 days), it was demonstrated that the newly developed method for barrel-centered sample extraction (see 2.4.5.1) indeed produced samples that contain a whole layer 4 of a barrel. To do this, a soma map of the dataset was created. It is known that somata density in layer 4 marks the barrels in barrel cortex (Meyer et al., 2010b). The quality in the dataset was sufficient everywhere to detect somata manually.

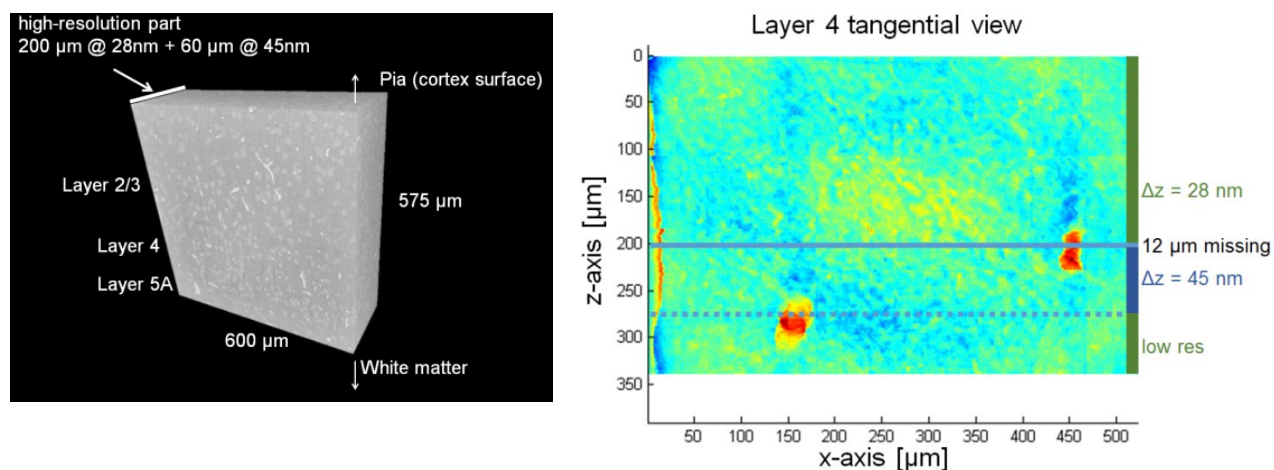


Figure 56 Left: 3D electron microscopy stack containing a whole barrel. Right: Tangential soma density map of dataset shown on the left (plus region imaged at low resolution).

The dataset includes the entire layer 4 and parts of layer 2/3 and layer 5. The cutting thickness for the first $200 \mu\text{m}$ was set to 28 nm and then was reduced to 45 nm for another $60 \mu\text{m}$. The rest of the sample was only cut to allow for the creation of a soma map. This soma map was created by measuring soma density in layer 4 and shows the successful targeting of the sample onto a barrel.

For automatic detection, at first the position of the layer 4 along the radial axis of the dataset was estimated. This was done by marking the sharp drop in soma density at the lower end of the layer 4 (Meyer et al., 2010b) and the transition between the larger somata of layer 2/3 pyramidal cells and the smaller spiny stellate cells of layer 4. Then a heuristic was developed to measure soma density automatically in the thus identified layer 4 (see 3.5). This showed the outline of a barrel very centrally in the dataset, validating the sampling method (Figure 56, right panel).

2.4.5.3 Image stitching

The recorded L4 dataset consisted of 6.1 million single images. Assembling them all together is a substantial challenge. The translation-only-script used for chapter 2.1 and 2.3 had failures in shift detection at about every 1000th image. It is believed that this rate could be reduced with a scale-invariant feature transform (Lowe, 1999) as the shift detection. Furthermore, a non-rigid approach would have more flexibility to counter imaging-induced distortion. Efforts to align this stack with such an algorithm are ongoing.

In the meantime however, a dataset from medial entorhinal cortex was recorded by H. Schmidt using the continuous-imaging technique which was then used to discover a novel mechanism by which synapses are sorted along axons of excitatory cells (Schmidt et al., accepted). The size of the dataset was (424x429x274) μm^3 , which highlights the increase in dataset volume that continuous imaging makes possible.

3 Methods

3.1 Mouse barrel cortex analysis methods

3.1.1 Sample preparation

Adult wild-type (C57BL/6) mice were used. They were transcardially perfused at postnatal 28 under isoflurane inhalation anesthesia with 2.5% PFA, 1.25% glutaraldehyde and 0.08M cacodylate buffer at pH 7.4. All procedures were approved by the local animal care and use committee.

The fixated brains were removed from the skull after 48 h of fixation and sliced coronally to a thickness of 1 mm. Two samples were manually taken from the 6th most frontal slice in the region, known to be barrel cortex from prior experiments. The samples were extracted with a 1 mm biopsy punch (Integra Miltex, Plainsboro, NJ, United States) – a hollow cylinder with a sharp front edge.

Contralaterally, the 6th most frontal slice was further sliced coronally into 70 μm sub-slices and a cytochrome oxidase staining was applied to them. With the help of an overlay, this confirmed that the samples were indeed from layer 4 and layer 2/3 of barrel cortex (Figure 57).

Afterwards the tissue was stained with a reduced osmium tetroxide solution (2% OsO_4 , 0.15 M cacodylate buffer, 2.5% KFeCN) followed by a 1% thiocarbohydrazide step and a 2% OsO_4 step for amplification (Briggman et al., 2011). After an overnight wash, the sample was further incubated with 1.5% uranyl acetate solution and 0.02 M lead (II) nitrate. The sample was dehydrated with propylenoxide and ethanol and embedded in Epon Hard (Serva Electrophoresis GmbH, Heidelberg, Germany) and hardened for 48 h at 60°C. This protocol did not allow for a homogeneous staining of the whole sample due to penetration issues. For both samples an imaging region was selected that was at the border of the sample.

Both samples were placed on aluminum stubs and were milled with a trimmer for EM preparation (EM TRIM2, Leica Microsystems, Wetzlar, Germany) so that all four sides of the sample had tissue exposed. The wall facing the direction from which the knife would approach and the contralateral side were smoothed with an ultramicrotome (UC7, Leica Microsystems, Wetzlar, Germany). The sides of the samples were covered with gold in a sputter coater (EM MED020, Leica Microsystems, Wetzlar, Germany). For imaging, the samples were placed separately into a custom microtome (Denk and Horstmann, 2004; Briggman et al., 2011) within a scanning electron microscope (Magellan, FEI Company, Hillsboro, OR, United States). The samples were oriented so that the radial axis was lying within the cutting plane.




Figure 57: Overlay of coronal slice of Figure 9 with samples extracted and cytochrome oxidase stain of contralateral slice. After the samples were extracted from the coronal slice with a punch (diameter 1 mm), a cytochrome oxidase stain was performed on 70 μm sub-slices of the equivalent slice from the other hemisphere. The overlay confirmed placement in barrel cortex (scale bar 1 mm)

In the layer 4 sample, the transition between layer 4 and layer 5A was identified by the sudden drop in soma density between the two layers (Meyer et al., 2010a). In the layer 2/3 sample, the transition between layer 1 and layer 2 was clear due to the very low soma density in layer 1. The region of interest was chosen so that it was situated in the transitory region between layer 2 and layer 3 using length approximations from existing literature (Lefort et al., 2009). Position and orientation of the transition zones agreed with the results from the cytochrome oxidase stain.

3.1.2 Electron microscopy

A region of $96 \times 64 \mu\text{m}^2$ within layer 4 was selected for imaging and a scanning pattern was established consisting of a grid of 3 by 3 separate image tiles with a pixel size of 11.24 nm. 3400 slices were cut within that region with a set thickness of 25 nm (see 3.1.3). The imaging was done at an acquisition rate of 10 MHz. Combined with overhead times for cutting and tiling, the effective data rate was 0.9 MB/s (compare 1.12.4).

A region of $99 \times 66 \mu\text{m}^2$ within the layer 2/3 sample was chosen for imaging. The pixel size was set to 12 nm and 7776 slices were cut. All the other parameters were equivalent to the settings for the layer 4 dataset.

The Magellan microscope (FEI Company, Hillsboro, OR, United States) was run with a fixed beam current of 3.2 nA, a landing energy of 2.5 kV and with the immersion field activated. An autofocus algorithm that had already been used for the (Briggman et al., 2011) study (a simpler version of the algorithm described in (Binding et al., 2013)) was customized to the Magellan microscope by adjusting gain parameters and adapting the calculated stigmatism axes to the one present in the

microscope. Overall, the acquisition of the layer 4 dataset took 108 h and the acquisition of the layer 2/3 dataset 247 h.

3.1.3 Cutting thickness

During stack acquisition, the cutting thickness was set to 25 nm. To check reliability of this setting, the cutting thickness was measured by manual inspection of an x-z-reslice, which yielded an approximate result of 28 nm. This value was used for dataset display during annotation. As an alternative metric, the cutting thickness was reevaluated by measuring the anisotropy of the nuclei of 18 randomly selected somata (in the plane orthogonal to the radial axis) and the cutting thickness was found to be $26.0 \text{ nm} \pm 0.45 \text{ nm}$.

3.1.4 Annotation tools

For all manual reconstructions, tools employing the skeleton annotation paradigm were used (1.6.4). Some tracings (most prominently the explorative layer 4 reconstructions from 2.1.1) were done using the software KNOSSOS (Helmstaedter et al., 2011), but the larger part of the reconstruction was done using a newly developed application called webKnossos (see 2.2). Briefly, KNOSSOS is a standalone application that necessitates that the annotators have a hard drive with the datasets available, while webKnossos is an online annotation tool that allows the data to be delivered over the internet. With this, larger annotator crowds can be put into action more quickly. Also of importance were the organizational features of webKnossos. Projects were created for the different questions that were investigated and the tasks belonging to these projects were automatically assigned to annotators. In both KNOSSOS and webKnossos, the adding of comments to nodes was used for synapse annotation. The newly developed merge-mode (3.1.12) was used to create the high-quality volumes in 2.1.

3.1.5 Traceability

For the bias-free traceability analysis (see 2.1.2.1), 90 pre-seeds were randomly chosen in the dataset. 30 seeds each were then surrounded by a flat bounding box ($10 \times 10 \mu\text{m}^2$) lying in the x-y, y-z and x-z plane, respectively. Annotators marked all the processes crossing these 2-dimensional bounding boxes. From this pool of processes, 90 were randomly chosen (out of 3198). Of those, 22 glial processes were eliminated because they cannot be efficiently skeleton-reconstructed. The remaining 68 processes were each assigned to 7 different annotators. These annotators reconstructed the processes in a $10 \mu\text{m}^3$ 3D bounding box. The resulting trees were fed into the first stage of the RESCOP algorithm, using the same set of parameters that was used for the retina study (Helmstaedter et al., 2011), except the test radius, which was set to 225 nm instead of 625 nm. The

test radius is the R.M.S. distance within which edges of other skeletons had to be in order to count as a vote in favor of an edge (see 1.7.1).

For the manual inspection of the output of the first RESCOP stage, all edges that had a vote distribution adjacent to the decision boundary (i.e. changing the vote by one would move the result into the other category, see (Helmstaedter et al., 2011)) were manually inspected. The majority of boundary cases were spines of dendrites that were not fully traced into the spine heads. These boundary cases would at worst lead to the loss of one synapse and probably none (see 3.1.12) and this was not considered a critical error. Only errors that would cause a break in the main dendritic or axonal stem were considered critical and this number was used for the approximation of maximally achievable average inter-error distance.

The second stage of the RESCOP algorithm used the vote distribution of the consensus trees to estimate remaining error rates. The assumption was that all edges have a ground-truth probability p_e of being detected. If this probability was lower than 0.5, the edge was assumed to have a ground truth of being not a biological connection in the neuron, if it was larger than 0.5 it connects two spots that are in fact part of the same neuron. To calculate how well n annotators would be able to reconstruct a given structure from the vote distribution it was estimated how many edges with intermediate values of p_e existed and how likely the consensus of n annotators would be wrong. The resulting value could be mapped for the number of annotators, estimating how many independent annotations are necessary to yield a certain mean error-free distance.

3.1.6 Synapse detection

For the test of manual synapse detection quality (see 2.1.2.2), two annotators traced the same set of three axons and compared the trees afterwards to make sure the tracings were without mistakes. At one point the two annotators differed and the annotation was corrected. Then a team of 6 annotators was assembled. They moved along the axon and annotated all vesicle clouds. In a vesicle cloud, at least 5 vesicles had to be visible and the putative vesicle cloud had to be in proximity to a membrane. In a second step, the annotators marked for each vesicle cloud all the putative postsynaptic partners. To be considered postsynaptic partners, the staining in the respective membrane had to be more intense and vesicles had to be visible presynaptically that were docked to the synaptic membrane. For each of the so identified partners, the annotators were asked to create a comment on how confident they were about their judgment on a scale from 0 to 10. This information was saved together with the partner annotation. One annotator misunderstood the instructions and labeled the confidence on a per vesicle-cloud basis instead. A colleague that helped with assembling data but who otherwise was not involved with the study noticed this mistake and

was able to instruct the annotator to correct the mistake. Then all the annotations were pooled anonymously and it was measured by manual inspection how well they agreed. Only 16 out of 488 synapses were between confidence level 3 and 7. Every synapse below confidence level 5 was rejected. Then the data was de-anonymized and annotators were allowed to discuss their results. Upon finding that majority of errors were attention-based, the error rate of manual synapse detection was estimated with the second part of the RESCOP algorithm (see 3.1.6). The prior was calculated and was used for the estimation of error rate over number of annotators.

3.1.7 Axon type detection

A dendrite within the layer 4 dataset was chosen and all synapses onto its main branch and its spines were annotated. These synapses were used as seeds and all seeded axons were reconstructed. On those axons, all synapses (3063 total) were categorized as either spine or shaft synapses. Additionally, it was indicated whether the synapse had multiple postsynaptic targets and/or mitochondria, but this information was not used further. Then, a custom MATLAB script found all possible clusters of n neighboring synapses along an axon. For each of those clusters it was measured which percentage of synapses was onto spine heads. In the $n=1$ case, this yielded the percentage of spine vs. shaft synapses, but for bigger n , this measure contained information on how consistent an axon targets spines or shafts. This information was normalized for axon length, so that long axons would not be overrepresented and plotted in a color-coded fashion over n (see Figure 16). The information for $n=10$ was extracted separately and plotted because it was decided that for $n=10$ the distribution was sufficiently bimodal that only a minority of axons would be ambiguous and would have to be further annotated after initial type detection.

3.1.8 Initial segment detection

To detect axon initial segments of layer 2/3 pyramidal cells, layer 4 spiny stellate cells and layer 4 interneurons, all the processes leaving the respective soma were inspected. The requirement for an initial segment was that its first $50\mu\text{m}$ did not have any further branching, that spine-like protrusions were rare (three or less per $50\mu\text{m}$) and that the diameter was relatively constant (less than a factor of two in diameter reduction over $50\mu\text{m}$). In all cells that were located sufficiently close to the pia-side border of the dataset (so that their initial segments could run for $50\mu\text{m}$ through the dataset), one and only one such process could be identified. If all but one process did contradict the requirements, this last process was assumed to be the initial segment even if less than $50\mu\text{m}$ of it were contained in the dataset. For the layer 4 spiny stellate cells it was permitted that a dendrite and an initial segment shared the very proximal part of the process (less than $5\mu\text{m}$). This exception was encountered in three cells. The validity of these criteria was further made plausible by the

occasional development of an initial segment into a branched and/or myelinated axon at the distal end of the initial segment (Layer 2/3: 4 times; layer 4: 18 times) and by the high consistency of their orientation - all pyramidal and spiny stellate initial segments were found to leave the cell towards the white matter, while the sole layer 4 interneuron initial segment exited towards the pia.

To discern initial segment synapses from soma synapses, the region of the axon hillock was identified where the curvature of the cell membrane was zero (neither the positive curvature of the soma nor the negative curvature of the narrowing initial segment). This was considered the border line between soma and initial segment.

3.1.9 Apical dendrite detection

Apical dendrites were detected by 3 independent annotators who marked all big processes traveling vertically through the dataset. This initial pool of candidates was then re-inspected and had to comply with the following requirements: The apical dendrites had to travel straight through the dataset, with a smaller entering diameter on the white matter side than the exiting diameter on the pia side. Only a 20% deviation of the direction of the process from the radial axis was deemed permissible. The dendrites had to have at least 1 spine every 10 μm . All candidates that were not myelinated axons fulfilled this requirement. Only one side branch was allowed, which had to be substantial smaller (less than half the diameter of the main branch) and branch off at an angle that was consistent with a growing direction towards the pia (not more than 10° angled downwards).

3.1.10 Inhibitory cell identification

For some cells that had their soma in the dataset the method outlined in 2.1.2.3 was not available because the axon of the cell was not sufficiently contained within the dataset. For those a different strategy was used to classify them as excitatory or inhibitory: The cell was completely reconstructed within the dataset, including all small protrusions. It was observed that cells show variation in how spiny their dendrites were. Namely, a subpopulation had a distinct pattern: Their dendrites had substantially fewer spines and the necks of those spines were stubby (Harris et al., 1992). Generally, dendrites had fewer spines closer to the soma, but the stubbiness was a consistent property per cell over regions of high and low spine density. The stubby non-dense spine bearing cells were assumed to be inhibitory cells. To classify cells of unknown identity, 20 spines in the most distal dendrites (edge of the dataset) were inspected. Spines that could completely and statically be seen in one of the three viewports (i.e. without moving the viewport, see 2.2) were classified as stubby. If 4 or more of the 20 spines were classified non-stubby, the cell was considered excitatory.

3.1.11 Specificity calculations

For the specificity calculations, it was necessary to determine which fraction of synapses of a given axon projected onto a target type of structure. The axons were not randomly chosen, but seeded by the fact that they innervated one of the sets of investigated target structures. Therefore Bayesian methods were developed to include this one seeding synapse in a neutral manner. The axon for which the specificity was calculated was assumed to be a member of an ensemble of m axons, all with the same properties: Each had n synapses within the dataset and a general predisposition p_e to form a synapse onto the target structure S . The size of the set of axons (set A) in the ensemble that confirmed with the observation of having k synapses onto S could be estimated with a binomial distribution (Equation 9). The ensemble would have a given number of synapses onto S (Equation 10), of which a subset would be from the axons in A (Equation 11). Therefore the probability to encounter an axon from A upon analyzing a random synapse onto S (which is what was done in these studies) could be expressed as a fraction of the two (Equation 12).

This can be turned into a Bayesian paradigm: Given a specific prior $p(p_e)$, what was the most likely p_e to explain the encounter? The $p_{\text{encounter}}(p)$ for all axons in all three settings (layer 2/3 initial segment, layer 4 initial segment, layer 4 apical dendrite) was calculated, using a flat p_e prior.

$$\text{Equation 9} \quad num_A(m, k, n, p_e) = m \binom{n}{k} p_e^k (1 - p_e)^{n-k}$$

$$\text{Equation 10} \quad num_{\text{synapses_all_to_S}} = m n p_e$$

$$\text{Equation 11} \quad num_{\text{synapses_A_to_S}} = k num_A$$

$$\begin{aligned} \text{Equation 12} \quad p_{\text{encounter}} &= \frac{num_{\text{synapses_A_onto_S}}}{num_{\text{synapses_all_onto_S}}} = \frac{k m \binom{n}{k} p_e^k (1 - p_e)^{n-k}}{m n p_e} \\ &= \frac{k \binom{n}{k} p_e^{k-1} (1 - p_e)^{n-k}}{n} \end{aligned}$$

Furthermore, a generalized best guess was made on what the predisposition of the innervating axons was to innervate S by multiplying the posteriors of all axons. It assumed that all the axons were following the same p_e value (no split in high-specificity and low-specificity axons), the probability of which was judged on a case-by-case basis (see 2.1.3.1 and 2.1.3.2). Importantly, because the posterior was much broader for axons with few synapses, this inherently gave more

weight to axons of which many synapses had been reconstructed. To calculate the margin of error of both the local and global posteriors, the smallest interval $[p_{e\ min}, p_{e\ max}]$ was chosen within which 67% of the probability spectrum was contained. To make statements about the absolute specificity of a given innervation type, the axons were compared to a hypothetical randomly innervating axon. Such an axon would target all structures in a blind fashion, therefore biasing itself towards structures with more wiring in the target region. To calculate the overall wiring length in the dataset, dendrites, somata and initial segments were considered viable targets. This was compared to the wiring length of the target structure to calculate a fraction of how much an axon should innervate the target structure if it were indeed random. The specificity factor was defined as the deviation from that random assumption.

3.1.11.1 Factor for initial segments

To calculate the total wiring length of initial segments of layer 2/3 pyramidal cells in barrel cortex, a literature value for the average soma density was used (Meyer et al., 2010b) and combined with the average length for initial segments observed in this body of work (50 μ m). While not all initial segments of all somata in the dataset were completely contained within the dataset, this was offset by other cells whose somata were not contained in the dataset but whose initial segments were. These effects cancelled each other out.

In the layer 4 dataset the contact detection was done on a reduced set of initial segments: Only the initial segments with somata in the dataset were included. The overall length of layer 4 initial segment wiring was calculated by summing the length of the initial segments in the library. Both the synapse detection and the overall length calculation operated on the same subset of initial segments and the effects cancelled each other out. To estimate how large this subset was, the following assumptions were made: Of the cells whose somata were in the top 14 μ m of the dataset, the whole 50 μ m of initial segment was contained. For the cells in the bottom 50 μ m of the dataset, on average 25 μ m of initial segment was contained. Furthermore, of cells whose soma was invisible and less than 50 μ m above the dataset, an average of 25 μ m of initial segment was contained within the dataset. This latter part had not been detected. With that, the percentage of detected initial segment length was calculated (Equation 13), which was assumed to be proportional to the number of synapses detected.

Equation 13

$$\frac{14\mu m \times 50\mu m + 50\mu m \times 25\mu m}{14\mu m \times 50\mu m + 50\mu m \times 25\mu m + 50\mu m \times 25\mu m} = 0.6094$$

3.1.11.2 Factor for apical dendrites

To calculate the number of apical dendrites in the layer 4 dataset, the number of pyramidal cells in layer 5 of mouse barrel cortex (Lefort et al., 2009) was normalized to the area of the dataset in a radial projection. Further it was assumed that all apical dendrites ran through the entire height of the dataset (64 μm).

3.1.12 Volume reconstruction and Merge Mode

For the automated contact detection used in 2.1.3.2, skeleton reconstruction was not sufficient (compare 1.6.4). The skeletons were combined with a segmentation that was optimized to be over-segmented (SegEM, (Berning et al., 2015), run 20130516T204040_{8,3}). Because the annotators had been instructed to set a node about every 100 nm and the average length of a segment was 4.56 μm , this created a model of the process that contained almost all the volume of the process, which could therefore be used for contact detection between processes. Two segments were considered to be in contact when they overlapped after being dilated by a 2 voxel spherical structuring element (compare (Berning et al., 2015)).




Figure 58: Two viewport panels showing user interaction in Merge Mode, with segments that the user selected because they belong to the neurite highlighted in color.

Precomputed segments can be collected into a larger volume: With this, the output from an automated segmentation can be used to generate a high quality volume models that contains all structural elements, e.g. the spine necks in Figure 16. This annotation mode was implemented outside of the normal webKnossos development as a plugin.

For visualization it was sometimes necessary to have high-quality 3D volume models of the involved processes (see Figure 26). For this again the segmentation described in (Berning et al., 2015) was used. The automated volume completion described above had small errors that reduced the

effectiveness of visualization (e.g. gaps in the volume reconstructions of spine necks). Therefore a new annotation mode was developed for webKnossos that allowed annotating a process in skeletonization mode while seeing the segmentation in a semi-transparent overlay (see Figure 58). With the creation of every node, the segmentation object within which the node was situated was added to the pool of segments that made up the process and the color of the segment was changed - so that it could be identified as part of the process. This allowed to manually collect all segments a process consisted of. As a last step, a custom MATLAB script used the final node list to collect all the segments thus marked, put together the volume, morphologically closed it with a structuring element of size 3, gauss-smoothed it and created a surface, simplifying that surface by reducing the number of vertices by a factor of 10 afterwards. This surface was then rendered using Amira (FEI Company, Hillsboro, Oregon, United States).

3.2 Image stitching and alignment

3.2.1 L4 and L2/3 dataset

After dataset acquisition, all files were inspected manually for imaging problems. The only problem was debris, which occurred when cut-off material remains at least partially on the sample and obscured the surface that is being imaged. Images thus affected were replaced by the images at the same position from the slice before or after. The most significant contribution that moved images from their fixed position in the grid was movement jitter of the x-y movement induced by the motors that moved the microtome. This effect could be compensated, i.e. there was a positioning of the tiles that removed the jitter effect of the motors. Additionally, images could be slightly distorted (e.g. due to charging from the electron beam), which could not be fully corrected by moving the images to a given position.

Because these latter effects were small, a rigid, translation-only type of alignment was chosen. This was based on the MATLAB script used for (Briggman et al., 2011), which followed closely (Preibisch et al., 2009). In the following the changes are described that were made to this approach.

The main challenge with this method was that occasionally the shift calculation failed entirely; returning values substantially more than 100 pixels off from the actual center. This did happen with image overlaps that contained very little structure. The least-square relaxation was not a good solution for this problem because outliers were given quadratically more influence. Therefore it was imperative to correct for these mistakes. For this, all local shift values were sorted by how big their residual errors after the least-square-fit were. This reliably identified overlaps where the shift calculation had failed. These usually corresponded to regions that had little structure (as mentioned before) or where the size of the overlap used for the cutout for the cross-correlation ($1\mu\text{m}$) and the actual overlap were too different. The second issue was fixed by setting the cutout size for the cross-correlation to a more appropriate value at the affected overlap. The remaining errors were iteratively fixed by manually reducing the weight of the corresponding entry in the least-square relaxation by a factor of 1000. The effort was considered finished when the highest remaining residual error dropped below 10 pixels. Because the shift calculation of subsequent images in cutting direction was found to be the most reliable measurement, it was weighted with a correction factor of 3 in the weighted least-square relaxation. Lastly, the data was written out with the shifts applied into a KNOSSOS hierarchy (see 3.3.2).

3.2.2 Improved L4 dataset for webKnossos experiments

The webKnossos flight mode (2.2.2) is sensitive to the quality of alignment. Therefore, an improved alignment was created of the layer 4 dataset. The prior alignment had reached a limit regarding how good a translation-only alignment of the raw data could get. To eliminate image distortions, it would be necessary to move different parts of an image in different directions. To do this, each image of the raw dataset was cut into smaller images sized 256x256 px. The delta calculation was run as described in chapter 3.2.1. The cutting procedure left no overlap between the cut images and their delta was set to zero. Additionally, a blood vessel and soma map was used and the delta results of images which mostly contained blood vessel or soma were given a decreased weight in the relaxation step. After the least-square relaxation, the cut images were discarded and their calculated shifts were used to create a smooth morphing of the original images. These morphed images were then used to write out the final dataset for the reconstructions.

3.2.3 The Lobula plate dataset

Number of slices	Height of sub-stack	Number of columns	Number of rows
284	7.1 μm	7	11
260	6.5 μm	9	11
356	8.9 μm	9	13
4276	106.9 μm	7	11
343	8.6 μm	7	11
1784	44.6 μm	7	11

Table 3: Overview of the Lobula plate substacks

Because the original script only allowed perfect cubic tiling patterns, it was changed to allow empty images and the grid was extended to the size of the sub-stack with the biggest lateral extent (third sub-stack in Table 3). For the “missing” virtual images the relative shift to the neighboring images was not calculated and they were not included in the least-square relaxation. To correctly follow the bundle of processes from the lobula plate tangential cells, the region of interest had been changed twice: Before starting the second to last sub-stack in Table 3, the region had been shifted 24 μm to the anterior. This shift was as wide as a single image in the grid and the software could represent that by shifting all following images in the grid one to the left. About 9 μm more medially, between

the second-to-last and last sub-stack, another readjustment had been necessary and the necessary shift (12 μm) could not be represented in the grid logic. It was found that the local shift detection performed very poorly when having to detect such big shifts in subsequent images without prior information. Therefore the algorithm was extended so that the shift detection received cut-outs in which the prior knowledge about the shift was already incorporated.

Neighboring images where the overlap region contained epoxy were isolated from the global alignment decision: As a first step, every 500th slice was inspected and tiles and overlap regions that contained epoxy were manually marked. It was sufficient to make that decision only every 500 slices (i.e. every 12.5 μm) because the border of the epoxy only moved slowly between slices. The relative shifts between images were calculated regardless of whether they had epoxy in them or not, but for the least-square relaxation, local shift vectors that were based on epoxy images were weighted significantly lower (10^{-4}). Consequently, the least-square relaxation primarily aligned the pictures that contained somata and neuropil using the relative information between them. Nonetheless, the information from the problematic images was not discarded and the least-square relaxation tried aligning them as well as possible too, but without disrupting the main region of interest.

To identify wrong shift measurements it was exploited that shifts were generally consistent between different slices: If a shift value differed more than 50 pixels from the average shift (regarding the same positions in all z slices), the weight of that measurement in the least-square relaxation was reduced by a factor of 10^4 .

Another heuristic was available to identify unreliable shifts: The shift detection operates by calculating the cross-correlation between the two cut-outs from the two adjacent images, applying a threshold to the result and detecting the center of mass of the largest super-threshold region. It was found that the size of that region correlated with the reliability of the shift detection: Failed shift detections often had significantly larger super-threshold areas. Therefore a graded weight reduction was implemented: If the area was larger than 600 px^2 , the weight of the local measurement was reduced by a factor of 10^7 , if the area was larger than 400 px^2 , the weight was reduced by a factor of 10^3 and if the size of the area was 200 px^2 or larger, the weight was reduced by a factor of 2. Combined, the outlier-sensitive weight reduction and the area-dependent weight reduction eliminated large-scale disruptions of the dataset

3.2.4 Speed improvements

Due to the large size of the datasets, efficient handling of data became a requirement. To fulfill this requirement, the datasets were not handled locally, but high performance compute clusters (located at the MPI for Medical Research and the Max Planck Computing & Data Facility, Garching) were used to do the necessary calculations, with up to 200 CPU workers processing the data simultaneously. The algorithm was modified to allow for that. Furthermore, it was not efficient to rerun the shift detection of the whole dataset for every proposed change of the algorithm. Therefore a modification was introduced that allowed to very selectively recalculate the shifts of overlaps that had proven to be problematic. Every image had to be used in up to 6 separate calculations, because it had up to 6 neighboring images. It would have unduly reduced the throughput to load each image 6 times. Therefore the stack was cut up in meta-slices of 9 slices each. The content of one meta-slice could be kept in memory completely and therefore every image was read only once (apart from the 12.5% of images in the border layer, which was read twice, once for each meta-slice they belonged to).

Initially, the output data was written slice-by-slice. Because each KNOSSOS cube incorporates data from 128 layers (see 3.3.2), every cube was opened 128 times and each time only 16 kBytes of data were added to the cube. This reduced the output data rate by a factor of 20. Therefore the algorithm was modified so that all necessary image data for blocks of 16 cubes (4 x 4 x 1 in x, y, z respectively) was read into memory and then entire cubes were created with one write command. With this final optimization, the total run time for the stitching and cubing could be brought down to 37 hours (for the Lobula plate dataset) and 20 hours (for the L4 dataset as it was used in the webKnossos speed tests).

3.3 webKnossos methods

3.3.1 Dataset

The test with naïve annotators was executed on the dataset 2012-09-28_ex145_07x2, described in chapter 2.1.1. For the other experiments of this chapter, the alignment of that dataset was improved (dataset name: 2012-09-28_ex145_07x2_new2, see 3.2 for details).

3.3.2 Data delivery

webKnossos data is stored on the server in cubes 128^3 vx in size, as previously developed for KNOSSOS (Helmstaedter et al., 2011). The data has a bit depth of 8 (values 0-255). These cubes are stored on original resolution and additionally downsampled by factors of 2 (up to 512-fold resolution in data volume and resolution). For transmission, these cubes are split up into smaller units called “buckets” 32^3 vx in size. Other values from 16^3 vx to 128^3 vx were tried as bucket size but 32^3 vx was found to be most performant for typical connections. The client at a given moment requests only buckets of one zoom level, depending on how far zoomed-in/out the view is.

In ortho mode, the buckets in the current plane of view (active viewport) are loaded first, prioritized by the Manhattan distance from the center of the viewport (higher numbers correlate to lower priority), the buckets in the next plane of view (in the direction of the last edge of the current reconstruction) with the distance multiplied by 2, the buckets of the next plane with the distance multiplied by 4 and so on.

For flight mode, all buckets are loaded that are at least partially contained in a rectangular frustum sized 5x5 buckets in the base and 4x4 buckets at the top, height 2.5 buckets, centered on the current position and rotated along the direction of view. The priority of the request is equal to the Manhattan distance from the current position. Also all the remaining buckets visible in the current viewport are loaded, as soon as all other cubes are successfully received.

All buckets are compressed in a lossless fashion for transport between server and client.

“4 bit mode” is an activatable setting that forces the server to only submit the 4 most significant bits of each voxel to be transmitted, thus saving about 50% of the bandwidth. It was used for all experiments.

3.3.3 Data transmission test

For the webKnossos speed test, a server at Hetzner Online GmbH (Gunzenhausen, Germany), was used (specifications: Intel(R) Xeon(R) CPU E3-1245 V2 (4 × 3,4 GHz); 32 GB RAM; 15 × 3 TB HWRaid

HDD). Annotation speed was manually controlled and the view was only advanced when the process was completely visible in the current viewport.

For the CATMAID speed test, a server was deployed at the Max Planck Computing and Data Facility in Garching (specifications: Xeon E5-2630 12 cores, 128 GB RAM, 10 Gb network, JBOD of 4× Intel DC S3500 240 GB SSDs). Version 2016.12.16 of CATMAID was used. The dataset was converted to a three series of JPEG images sized 256x256 images, corresponding to XY, YZ and XZ slices of the dataset. JPEG compression quality was set to 75%, and the headers of the files were stripped to further reduce image size. The viewport size was adjusted so that it would match the webKnossos viewport size. CATMAID has a function to auto-advance a layer once the current layer is loaded (this feature is also available as a script in webKnossos) and this was used to control the speed.

The test for webKnossos and CATMAID were conducted using version 56 of Google Chrome, working from the Max Planck Institute for Brain Research in Frankfurt and the Instituto de Investigación en Biomedicina de Buenos Aires, Argentina. The preset bandwidth settings of Google Chrome are detailed in Table 4

Name	Download	Upload	Latency
LTE	4 Mbit/s	3 Mbit/s	20 ms
Good 3G	1.5 Mbit/s	0.75 Mbit/s	40 ms
Normal 3G	0.75 Mbit/s	0.25 Mbit/s	100 ms

Table 4: Connection presets used by Google Chrome

3.3.4 Axon seeding

For the training tasks, a random box sized $4.5 \times 4.5 \times 4.2 \mu\text{m}^3$ was marked in the middle of the dataset. 68 processes were reconstructed within that box as putatively axonal. Of those 40 were selected randomly. Seeds were chosen for the 40 processes and for flight mode an initial orientation along the neurite was also provided.

For the test tasks, a bounding box sized $(2.5 \mu\text{m})^3$ was selected that was within the center $(15\mu\text{m})^3$. The bounding box was selected so that it was free of blood vessels, somata and large dendrites. All processes within the bounding box were reconstructed. Then for all neurites in it was decided whether they were axons. Both the initial marking of neurites and the subsequent axon-classification

was controlled by three independent annotators. Of the 41 axons thus identified, 20 were chosen randomly for the reconstruction.

3.3.5 Speed measurement process reconstruction.

To measure the process reconstruction speed (axons and dendrites), the length of the processes was measured and the time annotators worked on them. The time could be read directly from the webKnossos database (which also controls payout to the annotators). The way the time was measured is that whenever an annotator was actively working, every 30 seconds a token was sent to the server. If no activity took part in the last 30 seconds, the sending of the token was delayed until working was resumed. Each sending of a token added 30s of time to that annotators time account.




Figure 59: Comparing skeleton length between nominally identical flight and ortho tracings. As can be seen in the right subpanel, flight annotations are usually longer per actual neurite length. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

The naïve length measurement was done by measuring the length of all the edges in the annotation (compare chapter 1.6.4, skeleton based annotation). However, the naively measured length depends on the node density (more nodes reproduce the annotators' fluctuation more efficiently) - and flight mode annotation have a higher node density than ortho mode annotations (Figure 59). Therefore the skeletons were smoothed with the NURBS algorithm (ref: (Piegl and Tiller, 2012)). To calibrate the algorithm, a set of 10 axons was prepared that each had a matched pair of flight and ortho tracings that were otherwise identical. The node-order parameter (NO) of the NURBS algorithm controlled how much smoothing takes place. An attempted smoothing with an NO of 4 did not reduce the length difference between the paired skeletons below $14.93 \pm 1.08 \%$ (mean \pm SEM, compare Figure 60).

Therefore a variable NO approach was chosen, with

$$NO = \min(\text{ceil}\left(c_1 \left(\frac{D_s}{c_2}\right)^{c_3}\right), N_n)$$

where N_n is the number of nodes and D_s the density of nodes (in $1/\mu\text{m}$). The parameters c_1 , c_2 and c_3 were fitted so that the difference between the set of paired axons was below $1.79 \pm 1.16\%$ (mean \pm SEM; $c_1 = 50$, $c_2 = 5 \text{ }^1/\mu\text{m}$, $c_3=4$; Figure 60). Overall, this reduced the length of the skeletons by less than 20%.




Figure 60: Comparison of length of 10 reconstructed axons in flight mode and ortho mode. One can see that both for original tracings and tracings smoothed with NURBS with a fixed NO = 4 the length difference between flight and ortho persists. A variable NO approach abolished the problem. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

3.3.6 Dendrite reconstructions

The 497 dendrites were randomly chosen from a pool of over 2000 dendrites that had been previously reconstructed. The reconstructions were started at large diameter points of the dendrite. All annotations were conducted in flight mode. Annotators who had not seen the flight mode video before (because they had traced ortho mode before) were asked to watch it. All annotators were instructed to reduce their quality setting to medium (this only loads buckets that are downsampled once and thus reduces data rate demands). The annotators were instructed not to reconstruct spines.

3.3.7 RESCOP algorithm

The RESCOP algorithm (Helmstaedter et al., 2011) was used to make a consensus version of annotations that had been created redundantly. This necessitated the fitting of parameters of the RESCOP algorithm (namely, the edge probability prior) to the data on which it would be run. For this, the prior was calculated separately for ortho and flight modes annotations (Figure 61).

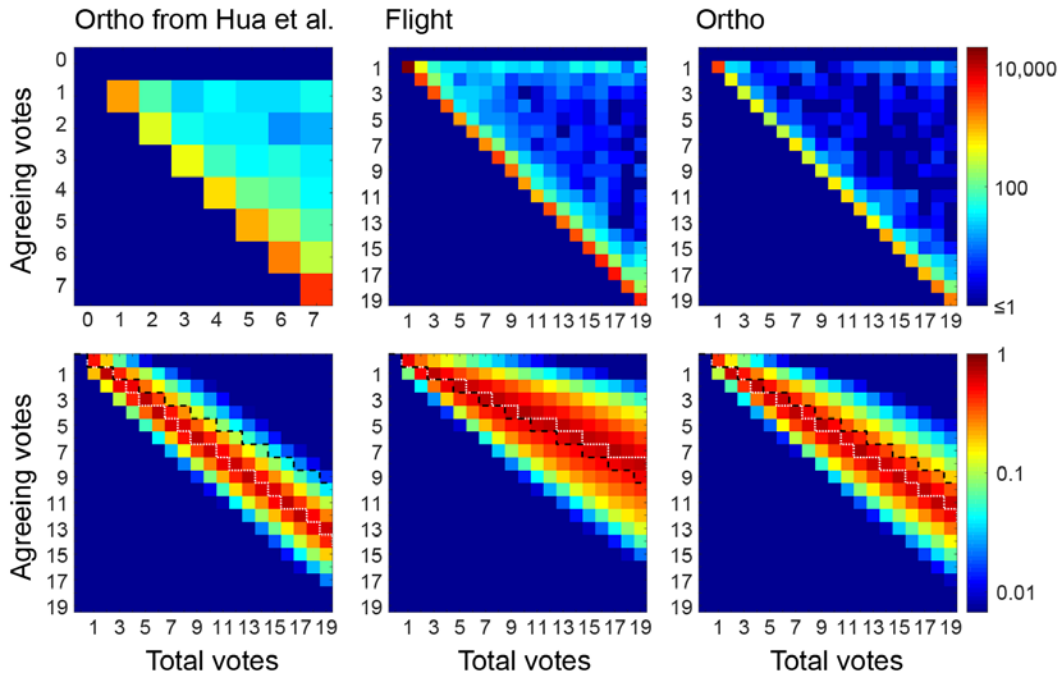


Figure 61: RESCOP algorithm output: Vote histogram (upper row) and optimized decision boundary (lower row; optimum vote: white line; majority vote: black line) for published data from (Hua et al., 2015) and for the axons from this study, calculated separately for flight mode and ortho mode reconstructions. Reprinted by permission from Macmillan Publishers Ltd: Nature Methods, (Boergens et al., 2017), copyright 2017

3.3.8 Error annotation

Redundancies 2, 3, 4, 5, 6, 7, 10 and 13 were used to measure the dependency of redundancy and error rate. For each setting, three random combinations were drawn for all redundancies and all measured neurites.

3.3.9 Synapse annotation mode

The synapse annotation mode combined flight-mode-style user interaction with manual node creation (in flight mode nodes are normally placed automatically). This mode was activated automatically when a task with the type “synapse annotation” was accessed. The task definition included a template neurite annotation so that the annotator could focus on the task of annotating synapses instead of neurite continuity. While moving along the template neurite annotation, right-

clicking the viewport at the position of a postsynaptic volume belonging to a synapse created a marker at the respective position in the dataset.

To determine error rates, three experts annotated the synapses on 4 test axons and compared this result to the ones supplied by the 10 annotators, measuring precision and recall. For the example connectome, the annotations of the most successful annotator (precision 0.96, recall 0.89) were used. The two axons that this annotator had marked as inhibitory were then reannotated by another student who had been instructed to focus more on correctness than on speed.

In these inhibitory annotations, 20 of the marked synapses were checked by an expert annotator and neither errors nor intermittent missed synapses could be identified.

3.3.10 Time consumption for synapse annotation

To predict the threshold between the sparse and dense regime of synapse annotation, the total volume of virtual tubes with a radius of $5\mu\text{m}$ around all reconstructed dendrites and axons was calculated (assuming non-overlapping tubes). The threshold was set to where this volume exceeds the dataset volume. For the axon reconstruction speed a value of 1.5 mm/h was assumed (with 6-fold redundancy), for the dendrite reconstruction speed 2.1 mm/h (with 3-fold redundancy). The time consumption for axon-based synapse annotation had been measured as 1.8 h/mm , the axon-based spine reconstruction as 2.3 h/mm . For the dendrite-based synapse-annotation a time consumption of 11.1 h/mm was approximated (including spine annotation and synapse annotation; using the measured times for spine reconstructions).

3.4 Methods for continuous imaging

For the movement in the x and y plane, internally levered piezo stacks were used, which had a stiffness of $0.4 \text{ N}/\mu\text{m}$ (P-602, Physik Instrumente, Karlsruhe, Germany), which were then down-levered again to a range of $350 \mu\text{m}$ and a stiffness of $3.2 \text{ N}/\mu\text{m}$ (down-lever factor $2\sqrt{2}$).

The piezo actors were combined with geared motors (M230-25, Physik Instrumente, Karlsruhe, Germany) to extend their range. These motors were able to carry the weight of piezo actors on their tip (see Figure 62). The mounting of the motors was set up so that they attached at the right position to allow for a down-lever factor as calculated above.




Figure 62: Mechanical construction that combined the range of a geared motor and a piezo stack to allow continuous imaging for large samples.

With the employed leveraging factor, the piezo stack that drove the continuous imaging was able to move the sample stage by $350 \mu\text{m}$. This was an insufficient range for the planned barrel cortex dataset. Therefore the piezo stack was coupled with a conventional geared motor. This motor employed a rack and pinion to retract and extend its tip. The piezo stack was mounted directly on the tip of the motor so that the movements of motor and piezo stack add together. Figure reprinted from (Schmidt et al., accepted)

3.4.1 Setup of electronics

Piezo stacks suffer from a hysteresis effect. There is no reliable relation between the applied voltage and the position of the piezo (Hall, 2001). Therefore the piezo actors used position sensors and control circuits to be able to exactly reach a desired position. This is was combined with amplifiers that generates the high voltages necessary for piezo extension (E-610, Physik Instrumente, Karlsruhe, Germany). The parameters of the closed-loop control circuit were fine-tuned to the applied load. However, it was not possible to find a set of parameters that combined fast

convergence time (<100 ms) with non-oscillatory behavior (without the mass of the microtome, this tuning had been possible).

The position was controlled by a signal voltage of 0-10 V to the amplifier/controller. This corresponded to 350 μm range. For the piezo position to have a stability of better than 5 nm, the control voltage had to have stability better than 140 μV . This was outside of the specifications of the amplifier used. Therefore a combined approach was taken: The E-610 amplifier executed the complete scanning pattern 100 times. The output voltages of these runs were recorded. Then these 100 voltage traces were averaged, removing the oscillatory behavior and the noise from the recording. Then the E-610 amplifier was removed and replaced by a different amplifier (ENV 40 SG, piezosystem jena, Jena, Germany). This amplifier had substantially better voltage stability. It executed the prerecorded voltage profiles in an open-loop setting.

If a piezo moves over a range of 350 μm in 10 nm steps, there are 35000 different intermediate positions. To correctly scan over all of them, a Digital-Analog-Converter with a resolution of 16 bit was used (NI6259, National Instruments, TX, United States). It was controlled by a custom MATLAB script that also did the curve generation, recording and averaging, outputting on two channels the trace for the two piezo stacks and on the third channel a trigger signal for the image acquisition (see 3.4.2).

To remove a mains artifact (image oscillations with 50 Hz and 20nm amplitude) from the system, the following steps were taken: The ground of the piezo system was separated from the vacuum chamber ground and a custom 50 Hz filters was patched into the lines that drove the piezo stacks. This had the added benefit of smoothing some of the oscillations induced by the PI amplifier, so that now rudimentary stack acquisition was possible without prerecording traces and changing the amplifier.

3.4.2 Control of the imaging setup

The microscope used (FEI Magellan) was not designed with a continuous imaging setup in mind. Therefore workarounds had to be found to make the microscope ready for the task. Ideally, the microscope would suppress the line feed generated by the microscope deflection system that is employed during normal image acquisition; and the whole strip would be read into one continuous image.

It was not possible to enable a mode where the whole strip was recorded as a single image. Initially, it was not even possible to record images directly after each other without pauses between the images that lasted more than 300 ms. Such a gap would correspond to a 10 μm gap between the

images. This was solved by a library that was supplied by FEI to the specifications of the project and allowed to record n consecutive pictures with gaps <30 ms between them. It was also not possible to completely disable the line feed, but the microscope was equipped with a mode called “Tilt mode” that allowed to correctly image tilted samples. When imaging a tilted sample the scanning rows have to be closer to each other. Tilting optically shrinks a sample in one axis, so bringing the lines closer to each other reverses that effect. For these reasons activating the “Tilt mode” reduced the strength of the line feed. The highest value the tilt correction could be set to was 85.9° , which meant that the line feed was reduced to 7% of its original value. The remaining line feed was used to compensate a remaining problem: The strip of pictures had small gaps (the aforementioned <30 ms gaps between the acquisition of consecutive images). The remaining line feed now caused the beam to act in synchronicity with the movement of the sample, thus stretching out the image frames. This closed the gaps and created a small overlap of about 80 pixels between the images that was used for stitching (see 3.2). For this to work, the sample had to move in the same direction as the line feed, but the direction of movement changed with every turn of the piezo snake (compare Figure 51). This reversal of scanning direction was compensated for by changing scanning *rotation* of the microscope by 180 degree for every strip.

During strip acquisition, the sample moved $1\ \mu\text{m}$ every 30 ms. Therefore the onset of the imaging had to be tightly controlled in synchronicity with the movement. It was found that triggering the acquisition start over a network connection was unfeasible because the jitter of the network connection exceeded the permissible jitter for imaging onset. Therefore a third channel was added to the control output of the movement generating script (see 3.4.1). This third channel was then used for triggering.

A program was created (written in C# in the CLR VM to be able to interface with the microscope library) that ran on the microscope computer and listened to the output of this analogue trigger channel with a Redlab 1208FS ADC (Meilhaus Electronic, Alling, Germany). On trigger it started the acquisition of the required number of images in rapid succession. The number of images per strip and the dwell time were sent to the C# program before the acquisition started (Figure 63). With this it was possible to attain a time-jitter between strips of less than 15 ms.

The number and height of the strips could be adjusted to the size of the region of interest. If the region of interest did not have a rectangular shape aligned with the axes of motion of the sample, the imaging during a strip could be started later or ended earlier for custom shape without having to change the scanning pattern. To further increase acquisition speed, the piezo stacks were not

returned to the original position after successfully moving through the pattern. Instead the next motor-tile was run rotated by 180° (compare Figure 51).

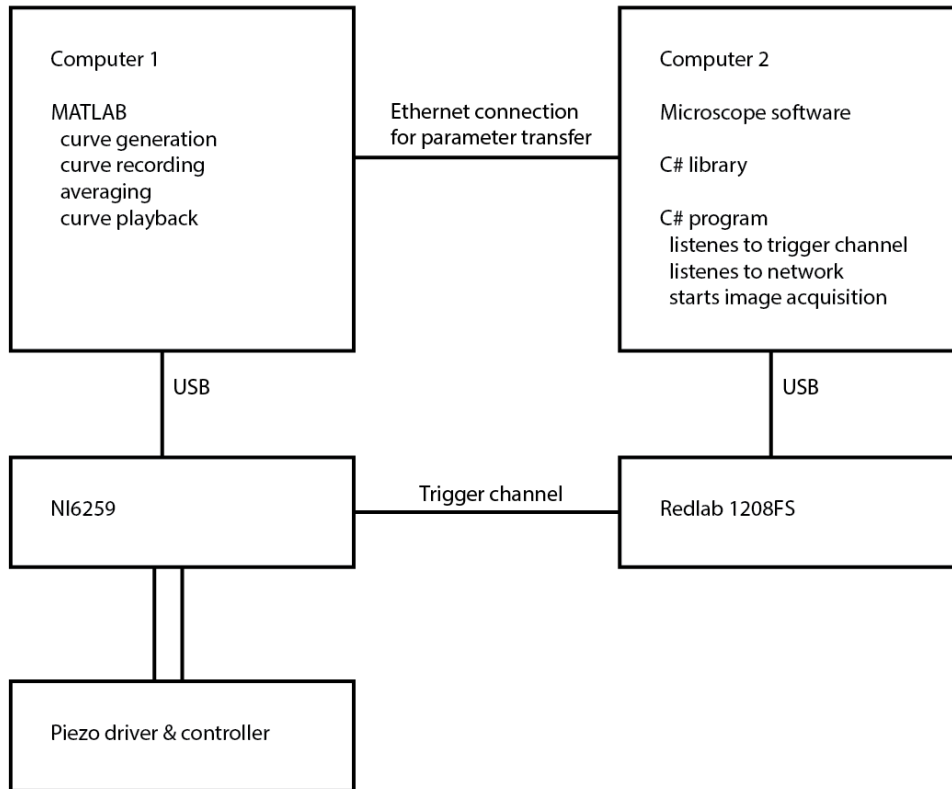


Figure 63: Layout of computers and electrical components for continuous imaging.

Two computers and two interfaces had to be combined to allow for a low-time-jitter synchronization of onset of imaging and piezo stack movement for continuous imaging. An Ethernet connection was found to be too unreliable (delay times could not be predicted between slices) and therefore the interface that generated the output ramps for the piezo stacks also generated a trigger signal that caused the other computer to start image acquisition.

3.5 Methods for full barrel stack

For the large sample extraction, 1 mm single-use circular punches were used (Integra Miltex, Plainsboro, NJ, United States). There were no single-use circular punches with smaller diameters available, therefore punches specifically designed for neuroscience applications were used for the 500 μ m diameter sample extraction (Leica Microsystems, Wetzlar, Germany).

For the vibratome-based extraction, a small cut was created in the stereotactic apparatus that marked the position 2 mm posterior of the C2 barrel. Then the brain was extracted from the skull and moved into a vibratome (HM 650 V, Thermo Scientific, MA, United States) that was modified to additionally allow placing cuts at a controlled 90° angle relative to the normal cutting direction.

All vibratome cuts were done while the sample was slightly submerged in cacodylate buffer to protect the sample from drying out.

After a successful extraction, the staining was applied to the sample as described in 3.1.1. The sample was not fully embedded in a block of epoxy after staining (as done in 3.1.1), because otherwise, it would have been necessary to liberate the sample from the block again and too much tissue of the 500 μ m sample would have been lost during this reexposure. Therefore, a “dry embedding” was developed. Before the sample was placed into the oven for epoxy hardening, it was removed from the bath of epoxy. Excess epoxy was removed with a tissue used for lens cleaning and the sample was placed onto a 1 cm pin with a tip that had been sand blasted to create a table for the sample 1mm in diameter. In this configuration, the sample was put into the oven for hardening. This in-situ curing of epoxy meant that the trimming step of sample preparation could be skipped and the sides of the sample could be directly smoothed with an ultramicrotome (UC7, Leica Microsystems, Wetzlar, Germany), greatly reducing the amount of lost material.

For stack acquisition, the field of view was 600 x 575 μ m² (Figure 56, left panel), consisting of 9 strips of 14 images (3072 x 2048 px²) in each motor tile (see 2.4.3 for description). The pixel size was 11.24 x 11.24 nm²; the cutting thickness was set to 28 nm.

Starting at slice 1400 (after 9 days), cutting became irregular, at times converging to a pattern where only every second cut took place. Furthermore, over time the autofocus (see 3.1.2) became increasingly unstable. This instability created a zone in one motor tile in the slices 900-920 that was out of focus. After slice 3000, the focus quality was suboptimal, but not prohibitive for reconstruction. Later a piece of cutting debris was found inside the electron beam column and subsequent experiments have struggled with similar problems, so it is likely that debris-induced focus instability was the cause.

At slice 7400, the electron column shut off spontaneously. It is known that the column does an emergency shut-off if it briefly encounters a non-standard condition – this is a reasonable behavior for most electron microscopy applications but problematic for long-term stack acquisition. This caused 430 slices to be imaged without any beam current and the information of them was lost (Figure 56, right panel). After this incident, the cutting thickness was increased to 45 nm, to allow for dendritic reconstruction in the distal part of the barrel. At slice 8660 (end of the barrel, see Figure 56, right panel), the cutting thickness was again increased (to 200 nm) and the acquisition was continued until the sample was completely cut away.

To automatically detect somata in the dataset, a custom MATLAB script was used. After the extent of layer 4 was marked, the data therein was normalized to compensate for brightness drift across the dataset. Then a series of filters was applied, namely Gaussian smoothing, upper thresholding to eliminate blood vessels, lower thresholding to eliminate neuropil, a binarization and a morphological opening to remove small items. The remaining signal clearly identified somata. This data was then projected radially and a tangential soma density map was created.

4 Discussion

The novel field of high resolution connectomics is still in the early stages of development but already yielding results. Neurons spread over great distances to form circuits and the neurites that do so are very thin, so two requirements have to be met. The first requirement is the acquisition of large data volumes that includes whole cells, and the second requirement is that imaging is done at a sufficiently high resolution to be able to identify even the smallest neurites. This creates large amounts of data and produces demands for high-throughput analysis. There are many open questions about connectivity of neural circuits, but even with regard to cell types and identities, meticulous analysis of cell morphology can yield surprises. This is already changing the way we think about the nervous system, e.g. (Helmstaedter et al., 2013; Morgan et al., 2016).

This work presents methodological development on two fronts. On one hand, to make acquisition of large datasets possible, an improvement of imaging speed was made, reducing necessary times by a factor of 30. If the sample is targeted on a region of interest no unnecessary data is recorded, which also makes acquisition of data less time-consuming. A protocol to achieve this for a barrel in mouse somatosensory cortex was developed. Jointly these advances make the reconstruction of the connectome of a cortical circuit a realistic endeavor. On the other hand, as data sets grow, more people have to be involved in the processing and analysis of them. Therefore it was necessary to devise a tool for distributed and scalable data annotation. In this work *webKnossos* is presented as such a tool, making creation and analysis of connectomes more feasible. Using these methodological advancements specific innervation patterns for inhibitory axons in layer 4 and layer 2/3 of mouse barrel cortex have been revealed. The search for innervation patterns made by excitatory axons is ongoing. Careful reconstruction of cells in a dataset of the optic system of *Drosophila* yielded new cells that previously had not been shown to exist.

4.1 Axons in mouse barrel cortex target subcellular structures specifically

As described in the introduction, the mouse barrel cortex is a highly organized structure that receives input from the whiskers of the rodent. An open question is whether additionally to the macroscopical organization a microscopical organization principle exists. This microscopical organization may manifest itself as specific innervation patterns. In the introduction, a number of possible patterns were described, namely a cellular specificity (targeting of innervation onto specific cells), subcellular specificity (targeting of specific dendrites or cell components) and spatially selective specificity (proximal versus distal innervation of dendrites). Furthermore, it was established how a non-specific scenario would look like. This non-specific scenario can be described as Peter's rule, stating that axons randomly innervate postsynaptic targets as they are available. To find out

whether wiring in cortex is following circuit patterns or is non-specific, two cortical datasets were recorded with Serial Block-face Electron Microscopy. One of them was located in layer 2/3 and the other in layer 4 of mouse barrel cortex, being radially aligned to each other. The size of the layer 2/3 dataset and layer 4 dataset was $99 \times 66 \times 202 \mu\text{m}^3$ and $96 \times 64 \times 89 \mu\text{m}^3$ respectively. It was established that processes could be reliably reconstructed, synapses could be detected and the axon type could be determined. Then, a method was developed to distinguish excitatory and inhibitory axons. It is difficult to infer the axon type from a single synapse. The number of synapses that innervate spines and/or shafts was counted; this helped determine the type of axon (either inhibitory or excitatory). Because all synapses on the axon had to be of the same type (Dale's law), information about single synapses (innervating a spine or shaft) could be statistically agglomerated to make a reliable decision regarding the type of axon.

To improve annotation quality, multiple annotators were asked to work in parallel. Afterwards, this redundant information could be pooled to remove errors. Loading preexisting annotations and reannotating spots where no consensus was apparent was a direct approach to improve annotation quality. For edge cases of synapse identification and axon type identification, further annotation effort was directed to them. This general principle of "Focused Reannotation" is powerful in settings like this where annotation may be insufficiently accurate but heuristics can be found to measure the accuracy. In the future a scheme could be implemented where disagreements or uncertainties in annotations automatically trigger the generation of new tasks to clear up the ambiguity.

The search for innervation patterns was started in layer 2/3. It is known that a class of interneurons called chandelier cells specifically target initial segments of pyramidal cells, but it was not clear whether chandelier cells also innervate other targets and/or whether other axons innervate pyramidal cell initial segments. After reconstructing 18 axons targeting 38 initial segments, it became clear that the situation is more nuanced than previously assumed. While chandelier axons target initial segments of pyramidal cells very specifically, there was a second class of axons that targeted pyramidal cell initial segments with a high specificity, but did not share the known morphology of chandelier cells. Furthermore, there were axons that innervated initial segments but did not exhibit an increased specificity for them.

The next step was to quantify the specificity of each axon. The percentage of synapses that the axon made with the target structure was compared to how much the target contributes to overall postsynaptic wiring in the dataset. The latter is the percentage of innervation the target would receive if the axon would innervate blindly in regard to target identity. The innervation specificity of the chandelier axons was found to be 235 times above the random threshold. Identifying initial

segments of pyramidal cell was not time-consuming, because once a chandelier axons had been identified, many synapses yielded new initial segments. It was challenging to find an initial segment of an interneuron. Nevertheless, one was successfully identified and surprisingly it was not innervated at all. It would be interesting to investigate this situation for more interneurons and determine whether this is a shared property of all of them.

It was investigated how spiny stellate initial segment innervation manifests itself in layer 4. For this, all synapses onto initial segments with somata in the dataset were reconstructed, which yielded the unexpected result that in layer 4, initial segments are more densely innervated than in layer 2/3. The specificity was measured for these initial segment innervations and was overall much lower. In contrast to layer 2/3, an interneuron initial segment was quickly discovered despite the fact that the interneuron density is lower in this dataset. No statistically significant difference could be found in the way its initial segment was innervated compared to the spiny stellate cells.

Apical dendrites of layer 5 pyramidal cells running through the layer 4 dataset were selected as another class of postsynaptic targets for analysis. Inhibitory axons were identified that innervated those apical dendrites and their specificity was measured. The specificity of them was more consistent within itself than for the other types of innervation. In absolute numbers it was roughly in the middle between the extremely-high specificity of the chandelier cells and the slight specificity of the spiny stellate initial segment innervating axons. This result suggests a scenario where pyramidal cells from layer 5, which receive excitatory inputs from higher layers, now also show the ability to receive coordinated inhibitory inputs in layer 4. This would allow for an interruption of information flowing down cortical layers. Inhibitory input onto apical dendrites has also been shown to be a potent mechanism for abolishing coincidence detection in apical dendrites (Larkum et al., 1999). The specificity was measured for excitatory innervation onto apical dendrites. This was done using automated contact annotation and so far has not shown an increased specificity: Excitatory axons seem to treat apical dendrites just like normal dendrites. This is ongoing research and rarer excitatory high-specificity axons might be found in the future.

A further question is whether the dendrites of inhibitory neurons receive specific innervation. These dendrites can be identified by the fact that they have fewer and shorter spines. This is an ongoing effort and results will add to the varied picture of specificities on cellular compartments that was already assembled. While the exact position of the synapses did not matter for the analysis of the initial segment and apical dendrite specificities, there are potentially patterns that use a spatial distribution of innervation along a dendrite. A likely candidate for that was inhibitory innervation onto spiny stellate dendrites. Competing theories postulate that innervation would be more efficient

on proximal or distal regions of the dendrite. Upon annotating a group of dendrites in the layer 4 dataset, it was indeed found that there is a strong and significant clustering of inhibitory innervation on the first 25 μm of the dendrites. It was then investigated whether this innervation was specific for that dendrite or whether a model of a more global tonic inhibition of the cell to which the dendrite belongs would be more fitting. For tonic inhibition, further innervation onto the soma and/or other dendrites of the same cell would have been likely, but no innervation of that kind was found, highlighting the role of the inhibitory cluster as a specific veto point for information flowing through that dendrite. Because of the limited size of the dataset, only the proximal and middle regions of the dendrites were visible. While being sufficient to prove a proximal cluster of inhibition, this did not rule out that distal clusters of inhibition exist. It would be interesting to assess whether such clusters can be found. Furthermore it should be investigated whether similar clusters exist on interneurons or layer 2/3 pyramidal cells.

4.2 *webKnossos* is a novel tool for online 3D EM data annotation

As datasets grow larger, it becomes increasingly less feasible to distribute the data physically to the annotator. Therefore an online annotation tool called *webKnossos* was developed that receives its 3D data over an existing internet connection. To ensure that the tracing speed was not limited by the speed of the internet connection, several optimizations were included that reduced the requirements in bandwidth and latency. Because annotation is an activity that requires to navigate the dataset in all three spatial axes, the data was transmitted in small cubes (called buckets) 32px on a side. This was shown to be superior to transmission as sequential 2D images and to transmission as cubes with a larger or smaller edge length. Furthermore, an algorithm was developed that predicted which buckets would be needed next and loaded them as soon as the data that made up the active view were completely loaded. With these optimizations, *webKnossos* was compared with the most successful existing online annotation tool, CATMAID and was found to be 4-13 times faster, depending on the quality of the connection (in-lab to mobile/transatlantic). It was especially relevant that the difference in performance became larger as connection quality became worse: In many settings, undergraduate annotators work from home over unreliable internet connections and collaborations often happen internationally, where an optimization for high-latency connection is especially relevant.

webKnossos offers an interaction mode that closely resembles the interaction mode of *KNOSSOS*, a standalone tool for skeleton reconstruction. There are three viewports, representing the three cardinal axes through the dataset and reslices of the dataset are shown in the respective views. When annotators are creating skeleton annotations, they are instructed to always follow the neurite

in the viewport to which it runs the most perpendicular. Failure to do so can lead to lost branches. As neurites run through the dataset, their orientation changes frequently, which makes it necessary for the annotators to switch between viewports. It had been long suspected that those context switches increase mental load while tracing and thus reduce the speed with which neurites can be reconstructed. Furthermore it had been observed that the conceptualization of this arrangement was a major obstacle in the process of learning how to annotate.

Therefore a novel mode called flight mode was included in webKnossos and evaluated. In this mode there is only a single viewport. The annotator sees the data as if they were in the center of a hollow sphere in the dataset and see the raw data that is just outside of that sphere. The annotator can now move forwards or change the direction where they are headed, like the pilot of a plane (which gave the flight mode its name).

It was investigated whether this mode was easier to learn and it was confirmed that after an introduction of only 2½ minutes, annotators could produce usable annotations with flight mode. Then it was investigated whether flight mode enabled experienced annotators to annotate faster. This yielded annotation speeds of 1.5mm/h for axons and 2.2 mm/h for dendrites, almost 10-fold faster than the fastest published annotation speeds. This was achieved without an increase in error rate and while annotating all branches of the neurite.

It was then investigated whether these axon and dendrite annotations could be combined into a connectome. A set of 32 axon and 497 dendrites were combined and the synapses between them were reconstructed. It was shown that depending of the density of reconstructed axons and dendrites in the dataset the ideal strategy to create a full connectome is to annotate the synapses along the axons or along the dendrites. In this specific case it was faster to reconstruct the synapses along the axons. This was done (separately for excitatory and inhibitory dendrites) and the result was combined with manual spine annotation and an automatic linking of spine and dendrite annotations. For all these steps the speed and error rates were quantified and it could be shown that it is possible to create a full connectome where the overall reconstruction time is still dominated by the achievable neurite reconstruction speed.

When reviewing videos of annotators reconstructing axons at speeds larger than 1.5 mm/h where no loading delays are encountered and the interaction mode introduces no effort into the process, it seems reasonable to assume that flight mode is so optimized that the measured reconstruction speed is limited by the speed with which an annotator can visually interpret 3D EM data and

therefore no further improvements in annotation speed through interaction tool development are possible.

Rather, the next frontier is to make the allocation of human annotation power more efficient. As documented, redundant annotations are necessary to achieve satisfactory precision in many tasks. An improvement over redoing the whole annotation with a large number of annotators is first doing the task with a subset of the annotators and then increasing the redundancy at the spots where the few annotators disagreed. As discussed before, this scheme is called “Focused Annotation” and it can be a massive improvement, especially when using a tool like webKnossos where there is little overhead to distributing small tasks that focus on a hot spot.

4.3 Reconstruction of lobula plate tangential cells in *Drosophila* yields candidates for cells not previously identified

The tangential cells in the lobula plate are a fascinating component of the optical system of flies. They provide important information for the control of flight and body position and therefore are of major interest to vision research. Their large diameters make them an ideal target for a diamond knife-based SBEM study. A dataset had been acquired that allowed to attempt an analysis of the cells in question. A first necessary step was the stitching and alignment of the data. For this, new methods for data handling and automatic error correction in the alignment algorithm had to be developed. The algorithm that detected the shifts between adjacent images had an error rate of about 0.1%. These errors created disruptions and severely limited the ability to reconstruct neurites in the output dataset. A manual repair of these mistakes – as it had been done in smaller datasets – was unfeasible due to the high number of images involved. Therefore two automated metrics were developed that detected failed shift measurements. One metric identified the deviation of the detected shift from the average shift, the second metric was an internal process metric from the shift measurement. Both could be used to remove the damaging effect of the wrong shift vector by reducing its weight in the least square relaxation. With this in place, a dataset was created where most details of the tangential cells could be reconstructed by non-expert annotators. This dataset was then used to reconstruct the tangential cells in the lobula plate. The selection of start points was done independently in several layers, which resulted in a high coverage and therefore high quality of the reconstructions. The inclusion of the tract which leaves the lobula plate in anteromedial direction was of special value, because most tangential cells have a process in this tract and could thus be identified.

The reconstructions were reannotated in a way that included the diameter of the processes, creating annotations that are easier to assess and compare to existing light-microscopical images. With this,

comparisons to neurons that had already been identified with a genetic expression system in *Drosophila* could be made. Interestingly, new cells were reconstructed that had not been seen in the genetic approach. Candidates for additional VS cells were found: three more than the six which were already described for *Drosophila*, making the situation more similar to blowflies, for which 9-11 VS cells exist. Also, two cells were identified that resembled CH cells, despite the fact that they had not been fully described in *Drosophila* so far. Lastly, many cells were found that did not match the known patterns but still had intriguing morphologies, these should be closer examined. Future work should focus on identifying synapses in the dataset in a systematic fashion.

This ongoing work shows parallels to recent progress in the field of mouse retina cell types (Helmstaedter, 2013). The unbiased inclusion of all bipolar cells within a volume revealed a novel cell type that had not been described so far (XBC). Using connectivity information, it was possible to split up the cells into subtypes that were not distinguishable morphologically (CBC5A/CBC5R). The unbiased inclusion of all these cells in the dataset shows that the 3D EM reconstruction of the *Drosophila* lobula plate was a necessary step to unravel the cell types of this intricate system. It has become a useful tool to understand the principles of fly vision.

4.4 Continuous imaging reduces overhead time in serial block-face SEM dataset acquisition

While small datasets can help gain valuable insights into the patterns of connectivity of neuronal tissue, acquiring stacks that contain whole circuits promises even deeper insights into how the brain performs its function. A layer 4 data set of a single barrel would take more than 4 months to image if recorded at the same speed as the small datasets described in this thesis. To be able to tackle such a big sample volume it is essential to improve the speed of acquisition. In this work, steps were implemented that allowed to increase the speed of acquisition by a factor of 30 compared to published SBEM studies. The first step was using a scanning electron microscope that had a much higher beam density. Beam density is essential for fast acquisition because the detector needs to be able to collect a sufficient amount of electrons per pixel, so fast pixel time can only be achieved with strong beams. The amplifiers for the first generation of SBEM setups were matched to the microscopes in which they were used and therefore could not correctly amplify the detector current on the new high-speed setup. A new amplifier had to be integrated with the new microscope to allow the fast imaging to take place.

With the new imaging setup in place it became clear that imaging was not the only bottleneck. When imaging a large sample, it is moved under the beam in a grid-like fashion (tiling). The movement was abrupt and therefore there was a settling time for the apparatus after each shift

causing major delays in image acquisition. This reduced the theoretically possible acquisition rate nearly by one order of magnitude. To solve this, a novel scheme was developed: while the microscope generated a line scan, the sample slowly moved underneath the beam, thus generating a two-dimensional image. This method was termed “Continuous Imaging”. Because there were no abrupt accelerations and decelerations of the sample stage, no vibrations were encountered and the imaging could take place almost without interruption. To be able to use this mode, the sample stage/microtome had to be modified mechanically. Geared motors were not able to move sufficiently smooth and a custom piezo-driven stage had to be designed. Secondly, a new software environment had to be created to accurately combine the smooth movement of the piezo stacks with coordinated image acquisition.

When trying to create a circuit-complete barrel cortex sample, it became clear that the EM staining had substantial limitations in penetration depth. Therefore a method was developed that allowed to extract a layer 4 single barrel sample in a targeted fashion, so that the sample did not exceed the penetration depth of the staining method. It was shown that this method generates high-quality samples that contain the whole layer 4 circuit of a barrel. While this was successful, novel staining methods were developed recently that have relaxed the volume restrictions for 3D EM staining. A dataset was recorded from a sample created by the method described in this work and it was shown that it indeed contains a whole barrel. With a size of $600 \times 575 \times 260 \mu\text{m}^3$ at a resolution of $12 \text{ nm} \times 12 \text{ nm} \times (28 \text{ nm} - 45 \text{ nm})$ it was the biggest SBEM dataset so far recorded. The stitching and alignment of it is an ongoing effort. While this barrel dataset is probably not fully reconstructable due to missing slices, it offers valuable ground work and an interesting starting point to analyze the microcircuit contained within a barrel. Furthermore, the fast imaging technique developed here can be used for large datasets of various tissues, providing a necessary improvement in 3D EM studies. Reducing acquisition times makes new classes of projects feasible and thus this method an important contribution to the field of connectomics.

All in all, it was demonstrated that connectomics with serial block-face electron microscopy can generate substantial insights into neuroscientific questions and that the method is advancing, with exciting new improvements being developed in this work.

5 Contributions

All experiments and analysis in chapter 2.1 were performed by the author except for contributions to staining improvement and samples creation done by B. Cowgill and contributions to MATLAB code for the proximal/distal analysis, volume reconstruction and contact detection supplied by M. Berning

The data from chapter 2.2 is the result of a collaboration with M. Berning (publication as shared first-authorship) and our statement regarding the contributions is as follows:

KMB was instrumental in the discovery of flight mode as an intuitive data mode; MB and KMB contributed to the testing and error measurements in an about 60% / 40% share, which constituted a major work load during the extensive revision phase; KMB contributed the user-scripting feature, and wrote many first user scripts, an important extension capability of the software package; software implementation was otherwise performed by the scalable minds team; the comparison to other tools was performed primarily by MB with contributions by KMB; the connectome reconstruction workflow and its results were contributed by KMB.

The dataset presented in chapter 2.3 was prepared and recorded by Ch. Kapfer, who also did the seed selection for the first and second phase and the curation for the consensus skeletons of the second phase. The author of this thesis stitched and aligned the dataset, led the team of annotators and developed the diameter mode and its Amira integration

The method for extraction of barrel-centered fully stained samples in chapter 2.4 was developed with contributions by P. Bastians (who also contributed to the creation of the final sample that was used to record the dataset). Furthermore, the fast amplifier used for the diode detector was developed by W. Denk and J. Tritthardt. All other experimental steps, most notably the development and implementation of the continuous imaging mode were performed by the author. This method has been successfully used by H. Schmidt to record the largest existing entorhinal cortex dataset (Schmidt et al., accepted).

6 References

- Agmon, A., and Connors, B.W. (1991). Thalamocortical responses of mouse somatosensory (barrel) cortex in vitro. *Neuroscience* *41*, 365-379.
- Andres, B., Koethe, U., Kroeger, T., Helmstaedter, M., Briggman, K.L., Denk, W., and Hamprecht, F.A. (2012). 3D segmentation of SBFSEM images of neuropil by a graphical model over supervoxel boundaries. *Med Image Anal* *16*, 796-805.
- Ascoli, G.A., Alonso-Nanclares, L., Anderson, S.A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsáki, G., Cauli, B., DeFelipe, J., Fairén, A., Feldmeyer, D., Fishell, G., Fregnac, Y., Freund, T.F., Gardner, D., Gardner, E.P., Goldberg, J.H., Helmstaedter, M., Hestrin, S., Karube, F., Kisvárdy, Z.F., Lambolez, B., Lewis, D.A., Marin, O., Markram, H., Muñoz, A., Packer, A., Petersen, C.C.H., Rockland, K.S., Rossier, J., Rudy, B., Somogyi, P., Staiger, J.F., Tamas, G., Thomson, A.M., Toledo-Rodriguez, M., Wang, Y., West, D.C., and Yuste, R. (2008). Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex. *Nat Rev Neurosci* *9*, 557-568.
- Bausenwein, B., Buchner, E., and Heisenberg, M. (1990). Identification of H1 visual interneuron in *Drosophila* by [³H] 2-deoxyglucose uptake during stationary flight. *Brain Res* *509*, 134-136.
- Beierlein, M., Gibson, J.R., and Connors, B.W. (2000). A network of electrically coupled interneurons drives synchronized inhibition in neocortex. *Nat Neurosci* *3*, 904-910.
- Berning, M., Boergens, K.M., and Helmstaedter, M. (2015). SegEM: Efficient image analysis for high-resolution connectomics. *Neuron* *87*, 1193-1206.
- Binding, J., Mikula, S., and Denk, W. (2013). Low-dosage maximum-a-posteriori focusing and stigmatism. *Microsc Microanal* *19*, 38-55.
- Binnig, G., Quate, C.F., and Gerber, C. (1986). Atomic force microscope. *Phys Rev Lett* *56*, 930.
- Bishop, C.A., and Bishop, L.G. (1981). Vertical motion detectors and their synaptic relations in the third optic lobe of the fly. *J Neurobiol* *12*, 281-296.
- Bock, D.D., Lee, W.-C.A., Kerlin, A.M., Andermann, M.L., Hood, G., Wetzel, A.W., Yurgenson, S., Soucy, E.R., Kim, H.S., and Reid, R.C. (2011). Network anatomy and in vivo physiology of visual cortical neurons. *Nature* *471*, 177-182.
- Boergens, K.M., Berning, M., Bocklisch, T., Bräunlein, D., Drawitsch, F., Frohnhofen, J., Herold, T., Otto, P., Rzepka, N., Werkmeister, T., Werner, D., Wiese, G., Wissler, H., and Helmstaedter, M. (2017). webKnossos: efficient online 3D data annotation for connectomics. *Nature Methods* *14*, 691-694.

Boergens, K.M., Berning, M., and Helmstaedter, M. (2016). Dendritic connectomics. In *Dendrites*, G. Stuart, N. Spruston, and M. Häusser, eds. (Oxford University Press).

Borst, A. (2014). Fly visual course control: Behaviour, algorithms and circuits. *Nat Rev Neurosci* *15*, 590-599.

Borst, A., and Haag, J. (2002). Neural networks in the cockpit of the fly. *J Comp Physiol A* *188*, 419-437.

Borst, A., and Helmstaedter, M. (2015). Common circuit design in fly and mammalian motion vision. *Nat Neurosci* *18*, 1067-1076.

Borst, A., and Weber, F. (2011). Neural action fields for optic flow based navigation: A simulation study of the fly lobula plate network. *PLOS ONE* *6*, e16303-e16303.

Braitenberg, V., and Schüz, A. (1991). *Anatomy of the cortex: Statistics and geometry* (Springer-Verlag Publishing).

Briggman, K.L., and Denk, W. (2006). Towards neural circuit reconstruction with volume electron microscopy techniques. *Curr Opin Neurobiol* *16*, 562-570.

Briggman, K.L., Helmstaedter, M., and Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. *Nature* *471*, 183-188.

Bruno, R.M., and Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* *312*, 1622-1627.

Buchner, E., Buchner, S., and Bühlhoff, I. (1984). Deoxyglucose mapping of nervous activity induced in *Drosophila* brain by visual movement. *J Comp Physiol A* *155*, 471-483.

Butcher, N.J., Friedrich, A.B., Lu, Z., Tanimoto, H., and Meinertzhagen, I.A. (2012). Different classes of input and output neurons reveal new features in microglomeruli of the adult *Drosophila* mushroom body calyx. *J Comp Neurol* *520*, 2185-2201.

Cardona, A., Saalfeld, S., Schindelin, J., Arganda-Carreras, I., Preibisch, S., Longair, M., Tomancak, P., Hartenstein, V., and Douglas, R.J. (2012). TrakEM2 Software for Neural Circuit Reconstruction. *PLOS ONE* *7*, e38011.

Chen, J.L., Villa, K.L., Cha, J.W., So, P.T.C., Kubota, Y., and Nedivi, E. (2012). Clustered dynamics of inhibitory synapses and dendritic spines in the adult neocortex. *Neuron* *74*, 361-373.

Clark, B.D., Goldberg, E.M., and Rudy, B. (2009). Electrogenic tuning of the axon initial segment. *Neuroscientist* 15, 651-668.

Cooper, N.G.F., and Steindler, D.A. (1986). Lectins demarcate the barrel subfield in the somatosensory cortex of the early postnatal mouse. *J Comp Neurol* 249, 157-169.

Crewe, A.V. (1973). Production of electron probes using a field emission source. *Prog Optics* 11, 225.

Crewe, A.V., Eggenberger, D.N., Wall, J., and Welter, L.M. (1968). Electron gun using a field emission source. *Rev Sci Instrum* 39, 576-583.

Culp, G.W. (1991). Piezoelectric actuator. US Patent 5,043,621.

Dale, H. (1935). Pharmacology and nerve-endings (Walter Ernest Dixon memorial lecture): (Section of therapeutics and pharmacology). *J R Soc Med* 28, 319-332.

Denk, W., and Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol* 2, e329.

Douglass, J.K., and Strausfeld, N.J. (1996). Visual motion-detection circuits in flies: parallel direction- and non-direction-sensitive pathways between the medulla and lobula plate. *J Neurosci* 16, 4551-4562.

Eccles, J. (1976). From electrical to chemical transmission in the central nervous system. *Notes Rec Roy Soc* 30, 219-230.

Eckert, H., and Dvorak, D.R. (1983). The centrifugal horizontal cells in the lobula plate of the blowfly, *Phaenicia sericata*. *J Insect Physiol* 29, 547-560.

Egelhaaf, M. (1985). On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly. 2: Figure-detection cells, a new class of visual interneurons. *Biol Cybern* 52, 123-140.

Egelhaaf, M., Borst, A., Warzecha, A.-K., Flecks, S., and Wildemann, A. (1993). Neural circuit tuning fly visual neurons to motion of small objects. II. Input organization of inhibitory circuit elements revealed by electrophysiological and optical recording techniques. *J Neurophysiol* 69, 340-351.

Feldmeyer, D., Egger, V., Lübke, J., and Sakmann, B. (1999). Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single 'barrel' of developing rat somatosensory cortex. *J Physiol* 521, 169-190.

Fiala, J.C. (2005). Reconstruct: a free editor for serial section microscopy. *Journal of Microscopy* 218, 52-61.

Fiala, J.C., Feinberg, M., Popov, V., and Harris, K.M. (1998). Synaptogenesis via dendritic filopodia in developing hippocampal area CA1. *J Neurosci* 18, 8900-8911.

Fischbach, K.F., and Dittrich, A.P.M. (1989). The optic lobe of *Drosophila melanogaster*. I. A Golgi analysis of wild-type structure. *Cell Tissue Res* 258, 441-475.

Fischbach, K.F., and Götz, C.R. (1981). Das Experiment: Ein Blick ins Fliegenhirn: Golgi-gefärbte Nervenzellen bei *Drosophila*. *Biologie in unserer Zeit* 11, 183-187.

Gauck, V., and Borst, A. (1999). Spatial response properties of contralateral inhibited lobula plate tangential cells in the fly visual system. *J Comp Neurol* 406, 51-71.

Gauck, V., Egelhaaf, M., and Borst, A. (1997). Synapse distribution on VCH, an inhibitory, motion-sensitive interneuron in the fly visual system. *J Comp Neurol* 381.

Gidon, A., and Segev, I. (2012). Principles governing the operation of synaptic inhibition in dendrites. *Neuron* 75, 330-341.

Goldstein, J., Newbury, D.E., Echlin, P., Joy, D.C., Romig Jr, A.D., Lyman, C.E., Fiori, C., and Lifshin, E. (2012). *Scanning electron microscopy and X-ray microanalysis: A text for biologists, materials scientists, and geologists* (Springer Science & Business Media).

Gray, E.G. (1959). Axo-somatic and axo-dendritic synapses of the cerebral cortex: An electron microscope study. *J Anat* 93, 420-433.

Guo-Ying, G., Li-Min, Z., Chun-Yi, S., and Han, D. (2013). Motion Control of Piezoelectric Positioning Stages: Modeling, Controller Design, and Experimental Evaluation. *IEEE/ASME Trans Mechatronics* 18, 1459-1471.

Haag, J., and Borst, A. (2001). Recurrent network interactions underlying flow-field selectivity of visual interneurons. *J Neurosci* 21, 5685-5692.

Haag, J., and Borst, A. (2007). Reciprocal inhibitory connections within a neural network for rotational optic-flow processing. *Front Neurosci* 1, 111-121.

Haag, J., and Borst, A. (2008). Electrical coupling of lobula plate tangential cells to a heterolateral motion-sensitive neuron in the fly. *J Neurosci* 28, 14435-14442.

Hainfeld, J.F. (1977). *Understanding and using field emission sources* (Brookhaven National Lab).

Hall, D.A. (2001). Review: Nonlinearity in piezoelectric ceramics. *J Mater Sci* 36, 4575-4601.

Harris, K.M., Jensen, F.E., and Tsao, B. (1992). Three-dimensional structure of dendritic spines and synapses in rat hippocampus (CA1) at postnatal day 15 and adult ages: Implications for the maturation of synaptic physiology and long-term potentiation. *J Neurosci* 12, 2685-2705.

Harris, K.M., Spacek, J., Bell, M.E., Parker, P.H., Lindsey, L.F., Baden, A.D., Vogelstein, J.T., and Burns, R. (2015). A resource from 3D electron microscopy of hippocampal neuropil for user training and tool development. *Scientific Data* 2, 150046.

Harris, K.M., and Stevens, J.K. (1989). Dendritic spines of CA 1 pyramidal cells in the rat hippocampus: Serial electron microscopy with reference to their biophysical characteristics. *J Neurosci* 9, 2982-2997.

Harris, R.M., and Woolsey, T.A. (1979). Morphology of Golgi-impregnated neurons in mouse cortical barrels following vibrissae damage at different post-natal ages. *Brain Res* 161, 143-149.

Hausen, K. (1976). Functional characterization and anatomical identification of motion sensitive neurons in the lobula plate of the blowfly *Calliphora erythrocephala*. *Z Naturforsch C* 31, 629-634.

Hausen, K. (1982). Motion sensitive interneurons in the optomotor system of the fly. *Biol Cybern* 45, 143-156.

Hausen, K. (1984). The lobula-complex of the fly: Structure, function and significance in visual behaviour. In *Photoreception and vision in invertebrates* (Springer), pp. 523-559.

Hausen, K., and Wehrhahn, C. (1990). Neural circuits mediating visual flight control in flies. II. Separation of two control systems by microsurgical brain lesions. *J Neurosci* 10, 351-360.

Hausen, K., Wolburg-Buchholz, K., and Ribí, W.A. (1980). The synaptic organization of visual interneurons in the lobula complex of flies. *Cell Tissue Res* 208, 371-387.

Hayworth, K.J., Kasthuri, N., Schalek, R., and Lichtman, J.W. (2006). Automating the collection of ultrathin serial sections for large volume TEM reconstructions. *Microsc Microanal* 12, 86-87.

Hayworth, K.J., Xu, C.S., Lu, Z., Knott, G.W., Fetter, R.D., Tapia, J.C., Lichtman, J.W., and Hess, H.F. (2015). Ultrastructurally smooth thick partitioning and volume stitching for large-scale connectomics. *Nat Meth* 12, 319-322.

Heisenberg, M., Wonneberger, R., and Wolf, R. (1978). Optomotor-blind H31—a *Drosophila* mutant of the lobula plate giant neurons. *J Comp Physiol* 124, 287-296.

Helmstaedter, M. (2013). Cellular-resolution connectomics: challenges of dense neural circuit reconstruction. *Nat Meth* *10*, 501-507.

Helmstaedter, M., Briggman, K.L., and Denk, W. (2008). 3D structural imaging of the brain with photons and electrons. *Curr Opin Neurobiol* *18*, 633-641.

Helmstaedter, M., Briggman, K.L., and Denk, W. (2011). High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nat Neurosci* *14*, 1081-1088.

Helmstaedter, M., Briggman, K.L., Turaga, S.C., Jain, V., Seung, H.S., and Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* *500*, 168-174.

Hengstenberg, R., Hausen, K., and Hengstenberg, B. (1982). The number and structure of giant vertical cells (VS) in the lobula plate of the blowfly *Calliphora erythrocephala*. *J Comp Physiol* *149*, 163-177.

Hennig, P., and Denk, W. (2007). Point-spread functions for backscattered imaging in the scanning electron microscope. *J Appl Phys* *102*, 123101.

Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. *Front Hum Neurosci* *3*.

Hill, S.L., Wang, Y., Riachi, I., Schürmann, F., and Markram, H. (2012). Statistical connectivity provides a sufficient foundation for specific functional connectivity in neocortical neural microcircuits. *Proc Natl Acad Sci USA* *109*, E2885-E2894.

Hua, Y., Laserstein, P., and Helmstaedter, M. (2015). Large-volume en-bloc staining for electron microscopy-based connectomics. *Nat Commun* *6*, 7923.

Jain, V., Bollmann, B., Richardson, M., Berger, D.R., Helmstaedter, M.N., Briggman, K.L., Denk, W., Bowden, J.B., Mendenhall, J.M., Abraham, W.C., Harris, K.M., Kasthuri, N., Hayworth, K.J., Schalek, R., Tapia, J.C., Lichtman, J.W., and Seung, H.S. (2010). Boundary learning by optimization with topological constraints (IEEE).

Jain, V., Murray, J.F., Roth, F., Turaga, S., Zhigulin, V., Briggman, K.L., Helmstaedter, M.N., Denk, W., and Seung, H.S. (2007). Supervised learning of image restoration with convolutional networks (IEEE).

Jain, V., Turaga, S.C., Briggman, K., Helmstaedter, M.N., Denk, W., and Seung, H.S. (2011). Learning to agglomerate superpixel hierarchies.

Jan, T.A., Lu, L., Li, C.-X., Williams, R.W., and Waters, R.S. (2008). Genetic analysis of posterior medial barrel subfield (PMBSF) size in somatosensory cortex (SI) in recombinant inbred strains of mice. *BMC Neuroscience* 9, 3-3.

Joesch, M., Plett, J., Borst, A., and Reiff, D.F. (2008). Response properties of motion-sensitive visual interneurons in the lobula plate of *Drosophila melanogaster*. *Curr Biol* 18, 368-374.

Karnovsky, M.J. (1971). Use of ferrocyanide-reduced osmium tetroxide in electron microscopy. *J Cell Biol* 51, 146.

Kasthuri, N., Hayworth, K.J., Berger, D.R., Schalek, R.L., Conchello, J.A., Knowles-Barley, S., Lee, D., Vázquez-Reina, A., Kaynig, V., Jones, T.R., Roberts, M., Morgan, J.L., Tapia, J.C., Seung, H.S., Roncal, W.G., Vogelstein, J.T., Burns, R., Sussman, D.L., Priebe, C.E., Pfister, H., and Lichtman, J.W. (2015). Saturated reconstruction of a volume of neocortex. *Cell* 162, 648-661.

Kätzel, D., Zemelman, B.V., Buetfering, C., Wölfel, M., and Miesenböck, G. (2011). The columnar and laminar organization of inhibitory connections to neocortical excitatory cells. *Nat Neurosci* 14, 100-107.

Kaynig, V., Vazquez-Reina, A., Knowles-Barley, S., Roberts, M., Jones, T.R., Kasthuri, N., Miller, E., Lichtman, J., and Pfister, H. (2015). Large-scale automatic reconstruction of neuronal processes from electron microscopy images. *Med Image Anal* 22, 77-88.

Kim, J.S., Greene, M.J., Zlateski, A., Lee, K., Richardson, M., Turaga, S.C., Purcaro, M., Balkam, M., Robinson, A., Behabadi, B.F., Campos, M., Denk, W., and Seung, H.S. (2014). Space-time wiring specificity supports direction selectivity in the retina. *Nature* 509, 331-336.

Knott, G., Marchman, H., Wall, D., and Lich, B. (2008). Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *J Neurosci* 28, 2959-2964.

Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöstrom, P.J., and Mrsic-Flogel, T.D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature* 473, 87-91.

Koch, C., Poggio, T., and Torre, V. (1983). Nonlinear interactions in a dendritic tree: Localization, timing, and role in information processing. *Proc Natl Acad Sci USA* 80, 2799-2802.

Koelbl, C., Helmstaedter, M., Lübke, J., and Feldmeyer, D. (2013). A barrel-related interneuron in layer 4 of rat somatosensory cortex with a high intrabarrel connectivity. *Cereb Cortex*, 263.

Korngreen, A., and Sakmann, B. (2000). Voltage-gated K⁺ channels in layer 5 neocortical pyramidal neurones from young rats: Subtypes and gradients. *J Physiol* 525, 621-639.

Land, P.W., and Simons, D.J. (1985). Cytochrome oxidase staining in the rat Sml barrel cortex. *J Comp Neurol* 238, 225-235.

Larkum, M.E., Zhu, J.J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398, 338-341.

Lee, T., and Luo, L. (1999). Mosaic analysis with a repressible cell marker for studies of gene function in neuronal morphogenesis. *Neuron* 22, 451-461.

Lefort, S., Tómm, C., Floyd Sarria, J.C., and Petersen, C.C.H. (2009). The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron* 61, 301-316.

Levy, P., and Larsen, C. (2013). Odd-skipped labels a group of distinct neurons associated with the mushroom body and optic lobe in the adult *Drosophila* brain. *J Comp Neurol* 521, 3716-3740.

Lin, C.-S., Lu, S.M., and Schmechel, D.E. (1985). Glutamic acid decarboxylase immunoreactivity in layer IV of barrel cortex of rat and mouse. *J Neurosci* 5, 1934-1939.

Lowe, D.G. (1999). Object recognition from local scale-invariant features. Paper presented at: Computer Vision, 1999 The Proceedings of the Seventh IEEE International Conference on.

Lübke, J., Egger, V., Sakmann, B., and Feldmeyer, D. (2000). Columnar organization of dendrites and axons of single and synaptically coupled excitatory spiny neurons in layer 4 of the rat barrel cortex. *J Neurosci* 20, 5300-5311.

Lübke, J., and Feldmeyer, D. (2007). Excitatory signal flow and connectivity in a cortical column: focus on barrel cortex. *Brain Struct Funct* 212, 3-17.

Maisak, M.S., Haag, J., Ammer, G., Serbe, E., Meier, M., Leonhardt, A., Schilling, T., Bahl, A., Rubin, G.M., Nern, A., Dickson, B.J., Reiff, D.F., Hopp, E., and Borst, A. (2013). A directional tuning map of *Drosophila* elementary motion detectors. *Nature* 500, 212-216.

Marie, R.L.S., and Peters, A. (1985). The morphology and synaptic connections of spiny stellate neurons in monkey visual cortex (area 17): A golgi-electron microscopic study. *J Comp Neurol* 233, 213-235.

Markram, H., Lübke, J., Frotscher, M., Roth, A., and Sakmann, B. (1997). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J Physiol* 500, 409-440.

Meyer, E.P., Matute, C., Streit, P., and Nässel, D.R. (1986). Insect optic lobe neurons identifiable with monoclonal antibodies to GABA. *Histochemistry* 84, 207-216.

Meyer, H.S., Schwarz, D., Wimmer, V.C., Schmitt, A.C., Kerr, J.N.D., Sakmann, B., and Helmstaedter, M. (2011). Inhibitory interneurons in a cortical column form hot zones of inhibition in layers 2 and 5A. *Proc Natl Acad Sci USA* *108*, 16807-16812.

Meyer, H.S., Wimmer, V.C., Hemberger, M., Bruno, R.M., de Kock, C.P.J., Frick, A., Sakmann, B., and Helmstaedter, M. (2010a). Cell type-specific thalamic innervation in a column of rat vibrissal cortex. *Cereb Cortex* *20*, 2287-2303.

Meyer, H.S., Wimmer, V.C., Oberlaender, M., de Kock, C.P.J., Sakmann, B., and Helmstaedter, M. (2010b). Number and laminar distribution of neurons in a thalamocortical projection column of rat vibrissal cortex. *Cereb Cortex* *20*, 2277-2286.

Mishchenko, Y., Hu, T., Spacek, J., Mendenhall, J., Harris, K.M., and Chklovskii, D.B. (2010). Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. *Neuron* *67*, 1009-1020.

Morgan, J.L., Berger, D.R., Wetzell, A.W., and Lichtman, J.W. (2016). The Fuzzy Logic of Network Connectivity in Mouse Visual Thalamus. *Cell* *165*, 192-206.

Orloff, J. (1981). A comparison of lanthanum hexaboride, cold field emission and thermal field emission electron guns for low voltage scanning electron microscopy. *Scan Electron Microsc*, 83-86.

Orloff, J. (1985). Thermal field emission for low voltage scanning electron microscopy. *J Microsc* *140*, 303-311.

Peng, H., Chung, P., Long, F., Qu, L., Jenett, A., Seeds, A.M., Myers, E.W., and Simpson, J.H. (2011). BrainAligner: 3D registration atlases of *Drosophila* brains. *Nat Meth* *8*, 493-498.

Peters, A., Feldman, M., and Saldanha, J. (1976). The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. II. Terminations upon neuronal perikarya and dendritic shafts. *J Neurocytol* *5*, 85-107.

Peters, A., and Jones, E.G. (1984). Classification of cortical neurons. *Cereb Cortex* *1*, 107-121.

Peters, A., and Palay, S.L. (1991). *The fine structure of the nervous system: Neurons and their supporting cells* (Oxford University Press, USA).

Piegl, L., and Tiller, W. (2012). *The NURBS book* (Springer Science & Business Media).

Pierantoni, R. (1976). A look into the cock-pit of the fly. *Cell Tissue Res* *171*, 101-122.

Preibisch, S., Saalfeld, S., and Tomancak, P. (2009). Globally optimal stitching of tiled 3D microscopic image acquisitions. *Bioinformatics* 25, 1463-1465.

Raghu, S.V., Joesch, M., Borst, A., and Reiff, D.F. (2007). Synaptic organization of lobula plate tangential cells in *Drosophila*: γ -Aminobutyric acid receptors and chemical release sites. *J Comp Neurol* 502, 598-610.

Rajashekhar, K.P., and Shamprasad, V.R. (2004). Golgi analysis of tangential neurons in the lobula plate of *Drosophila melanogaster*. *J Bioscience* 29, 93-104.

Saalfeld, S., Cardona, A., Hartenstein, V., and Tomančák, P. (2009). CATMAID: collaborative annotation toolkit for massive amounts of image data. *Bioinformatics* 25, 1984-1986.

Sanes, J.R., and Yamagata, M. (2009). Many paths to synaptic specificity. *Annu Rev Cell Dev Bi* 25, 161-195.

Schmidt, H., Gour, A., Straehle, J., Boergens, K.M., Brecht, M., and Helmstaedter, M. (accepted). Axonal synapse sorting in medial entorhinal cortex. *Nature*.

Schneider-Mizell, C.M., Gerhard, S., Longair, M., Kazimiers, T., Li, F., Zwart, M.F., Champion, A., Midgley, F.M., Fetter, R.D., Saalfeld, S., and Cardona, A. (2016). Quantitative neuroanatomy for connectomics in *Drosophila*. *eLife* 5, e12059.

Schnell, B., Joesch, M., Forstner, F., Raghu, S.V., Otsuna, H., Ito, K., Borst, A., and Reiff, D.F. (2010). Processing of horizontal optic flow in three visual interneurons of the *Drosophila* brain. *J Neurophysiol* 103, 1646-1657.

Scott, E.K., Raabe, T., and Luo, L. (2002). Structure of the vertical and horizontal system neurons of the lobula plate in *Drosophila*. *J Comp Neurol* 454, 470-481.

Seligman, A.M., Wasserkrug, H.L., and Hanker, J.S. (1966). A new staining method (OTO) for enhancing contrast of lipid-containing membranes and droplets in osmium tetroxide-fixed tissue with osmophilic thiocarbohydrazide (TCH). *J Cell Biol* 30, 424-432.

Sigworth, F.J. (1995). Electronic design of the patch clamp. In *Single-channel recording* (Springer), pp. 95-127.

Sommer, C., Straehle, C., Koethe, U., and Hamprecht, F. (2011). *ilastik*: Interactive learning and segmentation toolkit (IEEE).

Somogyi, P. (1977). A specific 'axo-axonal' interneuron in the visual cortex of the rat. *Brain Res* 136, 345-350.

Somogyi, P., Freund, T.F., and Cowey, A. (1982). The axo-axonic interneuron in the cerebral cortex of the rat, cat and monkey. *Neuroscience* 7, 2577-2607.

Song, S., Sjöström, P.J., Reigl, M., Nelson, S., and Chklovskii, D.B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol* 3, e68.

Spacek, J., and Harris, K.M. (1998). Three-dimensional organization of cell adhesion junctions at synapses and dendritic spines in area CA1 of the rat hippocampus. *J Comp Neurol* 393, 58-68.

Staffler, B., Berning, M., Boergens, K.M., Gour, A., van der Smagt, P., and Helmstaedter, M. (2017). SynEM, automated synapse detection for connectomics. *eLife* 6, e26414.

Starborg, T., Kalson, N.S., Lu, Y., Mironov, A., Cootes, T.F., Holmes, D.F., and Kadler, K.E. (2013). Using transmission electron microscopy and 3View to determine collagen fibril size and three-dimensional organization. *Nat Protoc* 8, 1433-1448.

Takemura, S.-Y., Xu, C.S., Lu, Z., Rivlin, P.K., Parag, T., Olbris, D.J., Plaza, S., Zhao, T., Katz, W.T., Umayam, L., Weaver, C., Hess, H.F., Horne, J.A., Nunez-Iglesias, J., Aniceto, R., Chang, L.-A., Lauchie, S., Nasca, A., Ogundeyi, O., Sigmund, C., Takemura, S., Tran, J., Langille, C., Le Lacheur, K., McLin, S., Shinomiya, A., Chklovskii, D.B., Meinertzhagen, I.A., and Scheffer, L.K. (2015). Synaptic circuits and their variations within different columns in the visual system of *Drosophila*. *Proc Natl Acad Sci USA* 112, 13711-13716.

Taniguchi, H., Lu, J., and Huang, Z.J. (2013). The spatial and temporal origin of chandelier cells in mouse neocortex. *Science* 339, 70-74.

Titze, B. (2013). Techniques to prevent sample surface charging and reduce beam damage effects for SBEM imaging. PhD thesis, University of Heidelberg.

Titze, B., and Denk, W. (2013). Automated in-chamber specimen coating for serial block-face electron microscopy. *J Microsc* 250, 101-110.

Tomassy, G.S., Berger, D.R., Chen, H.-H., Kasthuri, N., Hayworth, K.J., Vercelli, A., Seung, H.S., Lichtman, J.W., and Arlotta, P. (2014). Distinct Profiles of Myelin Distribution Along Single Axons of Pyramidal Neurons in the Neocortex. *Science* 344, 319.

Tuggle, D.W., Swanson, L.W., and Gesley, M.A. (1986). Current density distribution in a chromatically limited electron microprobe. *J Vac Sci Technol B* 4, 131-134.

Tuggle, D.W., and Watson, S.G. (1984). A low voltage field emission column with a Schottky emitter. Paper presented at: Proc EMSA.

Turaga, S.C., Briggman, K., Helmstaedter, M.N., Denk, W., and Seung, S. (2009). Maximin affinity learning of image segmentation.

Turaga, S.C., Murray, J.F., Jain, V., Roth, F., Helmstaedter, M., Briggman, K., Denk, W., and Seung, H.S. (2010). Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Comput* 22, 511-538.

Walton, J. (1979). Lead aspartate, an en bloc contrast stain particularly useful for ultrastructural enzymology. *J Histochem Cytochem* 27, 1337-1342.

Warzecha, A.K., Egelhaaf, M., and Borst, A. (1993). Neural circuit tuning fly visual interneurons to motion of small objects. I. Dissection of the circuit by pharmacological and photoinactivation techniques. *J Neurophysiol* 69, 329-339.

Waters, J., Larkum, M., Sakmann, B., and Helmchen, F. (2003). Supralinear Ca²⁺ influx into dendritic tufts of layer 2/3 neocortical pyramidal neurons in vitro and in vivo. *J Neurosci* 23, 8558-8567.

White, E.L. (2002). Specificity of cortical synaptic connectivity: Emphasis on perspectives gained from quantitative electron microscopy. *J Neurocytol* 31, 195-202.

White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos T Roy Soc B* 314, 1-340.

Williams, R.W. (2000). Mapping genes that modulate mouse brain development: A quantitative genetic approach. In *Mouse brain development* (Springer), pp. 21-49.

Woolsey, T.A., and Van der Loos, H. (1970). The structural organization of layer IV in the somatosensory region (S I) of mouse cerebral cortex. *Brain Res* 17, 205-242.

Yoshimura, Y., Dantzker, J.L.M., and Callaway, E.M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature* 433, 868-873.

7 Acknowledgements

I'm feeling very grateful to many people who played a role in the completion of this thesis, namely

To my supervisor Dr. Moritz Helmstaedter, who gave me the opportunity to do a PhD project in his lab, was a great teacher and inspiration and always pushed me to become a better scientist.

To Prof. Dr. Alexander Borst, for presenting my thesis to the Ludwig Maximilians University, for very helpful comments on the manuscript and for introducing me to the fascinating topic of the fly visual system.

To Prof. Dr. Winfried Denk, for invaluable contributions to the EM experiments, for being very generous with the equipment he made available and for highly inspiring discussions.

To PD Dr. Martin Hess, for giving me the opportunity to work on a very interesting project which combines FIB-SEM and SBEM acquisition in a fish retina dataset

To all of the above, for being a member of my thesis advisory committee and helping me with advice and feedback.

To Christoph Kapfer, Manuel Berning, Philipp Bastians, Beth Cowgill, Jakob Straehle, Heiko Wissler, Alessandro Motta and Benedikt Staffler for fruitful collaboration and my whole lab for a great work environment.

To Juergen Tritthardt for outstanding electronics support, to Niklas Neef and Martin Lukat for creating and discussing many mechanical parts used in this work and to Chris Roome and Mario Hilpert for IT support.

To my parents and sister, for encouraging me to follow my dreams, for always being there for me and for being an example in intellectual curiosity.

To Angeles Salles, for supporting me throughout the writing of this thesis and for making me a so much happier man.

This work was funded by the Max Planck Society.

Kevin Michael Boergens

Personal Information

Date of Birth	January 6 th , 1987
Marital Status	Married, no children
Citizenship	German

Education, Professional Experience

2014-present	PhD student at the Max Planck Institute of Brain Research, Frankfurt am Main, Germany Department of Connectomics, director Dr. Moritz Helmstaedter
2011-2014	PhD student at the Max Planck Institute of Neurobiology, Munich, Germany Laboratory of Dr. Moritz Helmstaedter
2010-2011	Diploma thesis Max Planck Institute for Medical Research, Heidelberg, Germany, Department of Biomedical Optics, director Dr. Winfried Denk. Development of FIB-SEM-based methods for neuroscience
2008-2011	Diploma in Physics, Ruprecht-Karls-University Heidelberg, Germany (specialization: Biophysics).
2008-2009	Pre-degree for Diploma in Mathematics (specialization: Group theory and Algebra) and Physics, Ruprecht-Karls-University Heidelberg, Germany
2006-2008	Undergraduate studies at Darmstadt University of Technology (<i>Technische Universität Darmstadt</i>) in Mathematics and Physics

Scholarships

2006-2011	German National Academic Foundation (<i>Studienstiftung des Deutschen Volkes</i>), full scholarship
-----------	---

Additional Information

Spoken languages	German (native), English (advanced), Spanish (basic) and French (basic)
------------------	---

Publications

Boergens KM, Denk W (2013) Controlling FIB-SBEM slice thickness by monitoring the transmitted ion beam. *J Microsc* 252(3):258-62

Berning M, **Boergens KM**, Helmstaedter M (2015) Efficient semi-automated image segmentation for high resolution connectomics. *Neuron* Volume 87, Issue 6, 1193 - 1206

Boergens KM*, Berning M*, Bocklisch T, Bräunlein D, Drawitsch F, Frohnhofen J, Herold T, Otto P, Rzepka N, Werkmeister T, Werner D, Wiese G, Wissler H, Helmstaedter M (2017) webKnossos: efficient online 3D data annotation for connectomics. *Nature Methods* 14, 691–694

Staffler B, Berning M, **Boergens KM**, Gour A, van der Smagt P, Helmstaedter M (2017) SynEM, automated synapse detection for connectomics. *eLife* 2017;6:e26414

Schmidt H, Gour A, Straehle J, **Boergens KM**, Brecht M, Helmstaedter M (accepted) Axonal synapse sorting in medial entorhinal cortex. *Nature*

Book chapters

Boergens KM, Berning M, Helmstaedter M (2016) Dendritic connectomics. In: Dendrites, 3rd Edition. Editors: Stuart G, Spruston N, Hausser, M, Oxford University Press

Invited talks

2017 FENS/IBRO Cajal Course “Interacting with neural circuits”, Champalimaud Centre for the Unknown: **Neuronal connectomics: serial EM approaches** (standing in for Dr. Helmstaedter)

2016 High-Resolution Circuit Reconstruction conference, Janelia Research Campus: **Local connectome statistics in mouse barrel cortex**

2015 Connectome Workbench, University of Cambridge: **Cellular and subcellular specificity in mouse barrel cortex** (standing in for Dr. Helmstaedter)

2015 Universidad de Buenos Aires, Laboratorio de Neurobiología de la Memoria: **EM-based connectomics in mouse barrel cortex**

2014 HHMI/Max Planck connectomics conference Berlin, **Novel methods for high-speed dataset acquisition in 3D electron microscopy**

Graduate Courses

Neural Systems and Behavior 2014, Marine Biological Laboratory, Woods Hole, MA, United States

Eidesstattliche Erklärung

Ich versichere hiermit an Eides statt, dass die vorgelegte Dissertation von mir selbständig und ohne unerlaubte Hilfe angefertigt ist.

München, den 27. Februar 2018..... Kevin Börgens.....

(Unterschrift)

Erklärung

Hiermit erkläre ich, *

dass die Dissertation nicht ganz oder in wesentlichen Teilen einer anderen Prüfungskommission vorgelegt worden ist.

dass ich mich anderweitig einer Doktorprüfung ohne Erfolg **nicht** unterzogen habe.

dass ich mich mit Erfolg der Doktorprüfung im Hauptfach
und in den Nebenfächern
bei der Fakultät für der
(Hochschule/Universität)
unterzogen habe.

dass ich ohne Erfolg versucht habe, eine Dissertation einzureichen oder mich der Doktorprüfung zu unterziehen.

München, den 27. Februar 2018..... Kevin Börgens.....

(Unterschrift)

*) Nichtzutreffendes streichen