

Increasing Depth Resolution of Electron Microscopy of Neural Circuits using Sparse Tomographic Reconstruction

Ashok Veeraraghavan¹, Alex. V. Genkin², Shiv Vitaladevuni², Lou Scheffer², Shan Xu²
Harald Hess², Richard Fetter², Marco Cantoni³, Graham Knott³, Dmitri Chklovskii²

¹MERL ²HHMI ³EPFL
Cambridge MA Ashburn, VA Lausanne
Contact: veerarag AT merl.com mitya AT janelia.hhmi.org

Abstract

Future progress in neuroscience hinges on reconstruction of neuronal circuits to the level of individual synapses. Because of the specifics of neuronal architecture, imaging must be done with very high resolution and throughput. While Electron Microscopy (EM) achieves the required resolution in the transverse directions, its depth resolution is a severe limitation. Computed tomography (CT) may be used in conjunction with electron microscopy to improve the depth resolution, but this severely limits the throughput since several tens or hundreds of EM images need to be acquired. Here, we exploit recent advances in signal processing to obtain high depth resolution EM images computationally. First, we show that the brain tissue can be represented as sparse linear combination of local basis functions that are thin membrane-like structures oriented in various directions. We then develop reconstruction techniques inspired by compressive sensing that can reconstruct the brain tissue from very few (typically 5) tomographic views of each section. This enables tracing of neuronal connections across layers and, hence, high throughput reconstruction of neural circuits to the level of individual synapses.

1. Introduction

The complexity of human behavior is matched by the complexity of underlying neuronal circuits, reconstructing which is an important step towards understanding brain function [11]. Such reconstruction poses a formidable challenge because of the large numbers and the wide range of spatial scales involved. The human brain contains 10^{11} neurons, each neuron receives and sends signals to nearly 10^4 other neurons. The signals are transmitted along neuronal processes called axons and across contacts called synapses, and received at dendrites. Axons may extend over many centimeters and yet are often only a few tens of nanometers in diameter.




Figure 1. ssTEM: Shown here is an illustration of the x-y and z-resolution of ssTEM for an example with 50 nm thick sections. While x-y resolution is sufficient to follow neural processes running in z-direction, z-resolution limited by physical section thickness precludes tracing processes running in x-y plane.

Because axons and dendrites of different neurons are tightly packed, tracing them requires resolution up to ten nanometers, or voxel size $5 \times 5 \times 5$ nm [11]. At such resolution imaging even a small brain produces astronomical amounts of data. For example, a tiny fruit fly brain contains 2×10^{14} $5 \times 5 \times 5$ nm voxels. Because such resolution is finer than the deBroglie wavelength of light, neuronal circuit reconstruction requires electron microscopy (EM). Several EM technologies differing in how tissue is sectioned and imaged have been proposed for circuit reconstruction [3, 11].

The oldest technology is serial section Transmission Electron Microscopy (ssTEM), which has been used to reconstruct the *C. elegans* nervous system [16, 6]. Processed brain tissue is sliced into a series of thin (50-75 nm) sections on a manually operated ultra-microtome. Then, each section is manually picked and placed on a thin support film and imaged under an electron microscope in transmission mode with a few-nanometer x-y resolution. Although this approach offers reasonable acquisition speed (several million voxels per second), its disadvantage is the low z-resolution determined by section thickness [3, 11]. Figure 1 shows an illustration of the problem of using ssTEM. While

the x-y resolution is sufficient to follow neural processes the z-resolution is insufficient.

There are two other approaches that use Scanning Electron Microscope (SEM) in reflection mode. First approach [10] uses sections that are cut by an automated ultramicrotome (Atlum), placed on a thick support tape, and imaged. Automation of slicing and section pick-up allows thinner (25-30nm) sections [11]. Second approach [7] images block-face and sections are cut by a custom ultramicrotome directly in the microscope at approximately the same thickness as Atlum but do not need to be handled. Although imaging can be done at speeds only somewhat less than ssTEM, the z-resolution is still insufficient to both trace the thinnest processes through the tissue and identify synapses [11].

There are two current approaches that have significantly better z-resolution (5nm) but these suffer from slow imaging speed. First approach called serial section transmission electron tomography images each section at several hundred tilts using ssTEM [14]. Then these hundreds of ssTEM images are used with backprojection, in order to reconstruct the original volume. In the second approach [13], focused ion beam (FIB) shaves brain tissue with precision exceeding that of diamond knife and the blockface is imaged by SEM in a reflection mode. Both these approaches are about hundred times slower than ssTEM.

In this paper, we propose a new EM technology that combines a high z-resolution of electron tomography and FIB with high speed of ssTEM by relying on computational processing. The main idea is to image thin sections, either in transmission or reflection mode, but only in a few, say five, tilt views. Such limited tomography can achieve required resolution while sacrificing imaging speed only by a small number of tilt views.

1.1. Contributions

The technical contributions of this paper are

- We study an important scientific challenge, neural circuit reconstruction and show that sparse representation and tomographic reconstruction can be used to enable neural reconstruction of large volumes.
- We show that the brain tissue can be represented as sparse linear combination of local basis functions that are like thin membranes.
- We develop reconstruction techniques inspired by compressive sensing that can reconstruct the brain tissue from very few tomographic views of each section.
- We show reconstructions on real EM images acquired a small brain tissue sample and show that the high depth resolution offered by the method significantly improves the ability to track neurons across sections.

2. Tracing Neural Processes: Z-resolution

The goal in neural imaging is to trace neural processes and identify synapses. Since, these processes may be less




Figure 2. TEM imaging of a 1D slice of a sample. The sample is 50nm thick. Vertical features are imaged sharply. Features that are slanted get blurred on the recorded image. Therefore traditional TEM will not be able to accurately image slanted membranes.




Figure 3. Illustration of process topology error due to thick sectioning. Shown are two different sections. The topology of the processes within the two sections are very different. Yet, the EM image obtained via a single 50 nm thick section is identical for the two sections. We propose using additional tilt-views to computationally resolve such ambiguities.

than 50nm thick and are delineated by thin membranes (as shown in Figure 1), the imaging voxel size should be around 5nm. While ssTEM achieves this resolution in the x-y directions, the depth resolution is limited to the thickness of physical sectioning (about 50-75 nm). This introduces two sources of errors that affect the ability to trace neurons. First, is the problem of blurring of the tilted membranes due to the fact that each ssTEM image is a projection of the section volume. Shown in Figure 2 illustrates the problem with identification of neural membranes that are not vertical. Shown in the image is a $X - Z$ section for illustration and the recorded 1D pixel intensities. While the membrane that runs vertically is recorded with high SNR on the captured image, membranes that are tilted get blurred. The amount of blur is dependent upon the tilt angle and this makes the identification, segmentation and/or localization of tilted membranes especially difficult. This is also illustrated in Figure 1. Further consider the two sections shown in 3. Though the topology of membranes for these two sections are radically different, the EM images obtained via thick sectioning are identical. This leads to membrane mergers which causes errors in the tracing of neurons. This necessitates the need to somehow enhance the depth resolution of electron microscopy.




Figure 4. System Diagram: The brain tissue is first physically sliced into thin 50-75 nm sections using a diamond knife. Each section is then imaged 5 times, with a different orientation of the section for each image. A computational reconstruction algorithm then uses the images for several sections together with priors about the brain tissue in order to reconstruct the volume. 3D segmentation is applied to the volume in order to trace neural processes and identify synapses.

3. System Description

A block diagram of the system to increase depth resolution for neural imaging is shown in Figure 4. The brain tissue is first physically sliced into thin 50-75 nm sections using a diamond knife. Each section is then imaged several times, with a different orientation of the section for each image. Typically, we imaged each section $N_{View} = 5$ times. One view was the central view, and the other four views were 45° tilt of the sample in each direction about the x and y axes respectively. A computational reconstruction algorithm then uses the images for several sections together with priors about the brain tissue in order to reconstruct the volume. 3D segmentation is applied to the volume in order to trace neural processes and identify synapses.

4. Modeling Neural Circuit Volumes




Figure 5. Modeling Neural Volume: A uniform grid with $15nm$ distance in x,y and z is placed in the volume. At each grid location basis functions are placed. The template basis function is a cuboid smoothed by a Gaussian. This template is then rotated around 12-20 different orientations that are sampled uniformly to create distinct basis functions at each grid location. The neural volume is modeled as being sparse in this basis.

In order to reconstruct the neural circuit volume with greater depth resolution from very few tomographic views of the section, one needs additional information about the volume being reconstructed. In this section we study the structure of neural circuits (as applicable to the imaging problem at hand) and model the volume using elongated

membrane-like structures that are both spatially and directionally isotropic.

Figure 11(a) shows an example of a reconstructed neural circuit volume with the different membranes encoded in color. This small volume was reconstructed from EM images acquired via FIB which has a greater depth resolution. We see from this volume that the neural circuits are made up of thin membrane like structures. Further these structures are isotropic both in space and in directionality over the course of the imaged volume (some exceptions exist eg. medulla). Therefore we model these neural circuits using basis functions that are made up of cuboids that are smoothed using a gaussian. The length, breadth and thickness of these cuboids are chosen so as to optimally represent the neural membranes. In practice, we set the length (X) of the cuboid to be about $55nm$, the breadth of the cuboid (Y) to be about $15nm$, while the thickness of the cuboid (Z) is set to be the thickness of the physical imaged section. In most of our real experiments the thickness of the physical section was about $50 - 60nm$.

Since the membranes could be present in any region in the brain tissue and its orientation could be arbitrary, we use the template cuboid basis function at every location on a subsampled grid. In practice the grid spacing is set to be around $15nm$ so that a membrane at any location can be approximated accurately. Shown in Figure 5 is a graphical illustration of the grid locations and some of the basis functions at each grid location. At each grid location we also create rotated versions of the cuboid basis in order to accurately represent membranes that have different orientations. In order to do this, we first uniformly sample Q points on the 2-sphere, generate the corresponding rotation matrices and rotate the basis function template accordingly. In practice we use about 12 – 20 uniform samples in most of our experiments. In order to account for other structures in the imaging volume that are not necessarily membrane-like, we also append to this basis, a constant intensity basis element and a non-oriented gaussian basis element at each grid location. This would accurately represent the constant gray scale regions and globule like structures present in the imaging volume.

Having created these basis functions, the original 3D volume of brain tissue can now be expressed as a linear combination of these basis functions. Let us denote by $B_{gx,gy,gz,\theta}$, the cuboid basis function at grid location gx, gy, gz with orientation θ . Now the original 3D volume can be represented as a linear combination of these basis functions as,

$$x(\cdot, \cdot, \cdot) = \sum s_{gx,gy,gz,\theta} B_{gx,gy,gz,\theta}(\cdot, \cdot, \cdot), \quad (1)$$

where x is the original volume and $s_{gx,gy,gz,\theta}$ are the weights of the individual basis functions. Further, since there are a small number of membranes within any small volume, and at any location within the volume, the mem-

brane has only one orientation (except at junctions), we also know that $s_{gx,gy,gz,\theta}$ is non-zero only for a very small set of basis functions. We exploit this property in the next section in order to reconstruct the volume from very few tomographic measurements.

5. Sparse Representation and Reconstruction

From the previous section, we know that the imaged volume x can be written as a linear combination of the Basis functions. Let us vectorize Equation 1 and write as

$$x = Bs, \quad (2)$$

where x is the unknown Volume, B is a matrix that contains all the vectorized basis elements and s is a $N \times 1$ vector that contains the weighting coefficients $s_{gx,gy,gz,\theta}$. N is the total number of basis elements which is the product of the total number of grid locations and the total number of basis vectors at each grid location. Since most of the elements $s_{gx,gy,gz,\theta}$ are zero, this implies that the vector s is sparse and contains very few non-zero elements. Since each tomographic measurement of the section can also be represented as a linear combination of the unknown Volume x , we can write

$$y = Ax, \quad (3)$$

where y is a long vector that contains all the varying view tomographic observations for several consecutive sections, A is a mixing matrix that represents the individual projection operators that simulate the tomographic views. From this, we see that the inversion problem that needs to be solved is

$$y = Ax = ABs, \quad (4)$$

where s is a sparse vector.

5.1. Compressive Sensing

Compressive sensing [4, 8, 1] is the concept of reconstructing sparse signals from under-sampled linear measurements. Consider a signal $x \in \mathbb{R}^V$, which is sparse in a basis B , that is, $s \in \mathbb{R}^N$ defined as $x = Bs$ is sparse. We call a vector K -sparse if it has at most K non-zero components, or equivalently, if $\|s\|_0 \leq K$, where $\|\cdot\|_0$ is the L_0 norm or the number of non-zero components.

We are interested in the problem of sensing the signal x from linear measurements. Ideally, with no additional knowledge about x , we would require N linear measurements of x , which would then form an invertible linear system. The theory of compressed sensing shows that it is possible to reconstruct x from M measurements even when $M \ll N$ by exploiting the sparsity of s . Consider a measurement vector $y \in \mathbb{R}^M$ obtained using a $M \times N$ measurement matrix A , such that

$$y = Ax + e = ABs + e \quad (5)$$

where e is the measurement noise. The components of the measurement vector y are called the *compressive measurements* or compressive samples. For $M < N$, estimating

x from the linear measurements is an ill-conditioned problem. However, when there exists a basis B such that s as defined above is K sparse, then compressive sensing allows for recovery of s (or alternatively, x) from $M = O(K \log(N/K))$ measurements.

5.1.1 Signal Recovery

Estimating K sparse vectors that satisfy the measurement equation of (5) can be formulated as the following L_0 optimization problem:

$$(P0) : \min \|s\|_0 \text{ s.t. } \|y - ABs\|_2 \leq \epsilon \quad (6)$$

where the L_0 norm, $\|\cdot\|_0$ counts the number of non-zero elements. This is typically a NP-hard problem. However, the equivalence between L_0 and L_1 norm for such systems [9] allows us to reformulate the problem as one of L_1 norm minimization.

$$(P1) : \min \|s\|_1 \text{ s.t. } \|y - ABs\| \leq \epsilon \quad (7)$$

with ϵ being a bound for the measurement noise e in (5). It can be shown that the solution to the (P1) is with a high probability the K sparse solution that we seek. There exist a wide range of algorithms that solve $P1$ to various approximations or reformulations [5][12][15][8]. Most of them note that the problem (P1) is a convex problem, and in particular, can be recast as a Second Order Cone Program (SOCP) for which there exist efficient numerical techniques.

In practice, we use a two step iterative shrinkage/thresholding algorithm (Twist)[2] which is a fast approximate algorithm that can handle very large problems to efficiently solve the sparse reconstruction. Further, since the data volumes that we consider are several hundred million voxels, it is not feasible to solve the entire system in one shot. Therefore, we subdivide the system into several overlapping blocks, solve the system for each of these overlapping blocks and retain the central portion of the solution for each overlapping block (in order to minimize edge effects). In particular, we use overlapping blocks of $150 \times 150 \times 27$ voxels, and retain the solution for the central $100 \times 100 \times 9$ voxels.

6. Experiments on FIB data

In order to evaluate the efficacy of the proposed method to reconstruct neural circuits we first collected high resolution data using FIB for a small volume of brain tissue. This provided us volume data for brain circuits at 5nm isotropic resolution. The total volume that we experimented was $1024 \times 1024 \times 1024$ voxels with each voxel being a cube of 5nm. We used this data in order to simulate thick sections as would be obtained using TEM and applied our reconstruction algorithm to reconstruct the volume from the observed tomographic views. Figure 6 shows the evaluation methodology.




Figure 6. Evaluation using high resolution FIB data: High resolution data for a small volume was acquired at $5 \times 5 \times 5$ nm voxel size using FIB which is a much slower acquisition method. We then simulated thick ssTEM slices using this data and reconstructed the volume using our reconstruction algorithm.

6.1. Comparison with Filtered Backprojection

In principle our reconstruction algorithm is different from traditional tomography in three significant ways.

- We use a computational reconstruction algorithm that exploits prior information about the brain tissue and this significantly reduces the impact of the 'missing cone' problem in traditional tomography.
- Traditional tomographic reconstruction requires several hundreds of projection per section in order to produce reasonably stable reconstructions. Imaging the same tissue hundreds of times would make the system throughput so slow to be practically infeasible. We propose a method that requires very few (typically 5) views for each section.
- The computational reconstruction algorithm that we develop exploits the continuity of membranes across sections which traditional tomography does not exploit. The presence or absence of a membrane at the bottom of one section provides information about the same at the top of the next section.

As a consequence of these differences, the quality of the reconstructions obtained using our algorithm is significantly better than those obtained by traditional tomographic reconstruction. Figure 7 compares the reconstruction signal to noise ratio (SNR) for both filtered backprojection and our reconstruction algorithm when the number of tilt views obtained per section is 3. The Figure clearly shows the superiority of our reconstruction algorithm over a wide range of thickness of the physical sections obtained.

6.2. Variation with section Thickness and NViews

Using the high resolution FIB data, we simulated obtaining the tomographic views of thick sections. We also corrupted these observations with 20dB noise drawn from an IID zero mean Gaussian distribution. We then performed our computational reconstruction which exploits sparsity of the reconstructed volume in the membrane basis. The SNR of the reconstructed volume for a wide range of section thicknesses and a wide range of number of tilt views per section is shown in Figure 7. The Figure shows that reconstruction quality degrades slowly as the thickness of the physical section increases. In practice it is fairly easy to




Figure 7. Performance of FIB Volume reconstruction with varying section thickness and number of tomographic views. The original volume was corrupted with 20 dB noise. In practice a physical section thickness of 50-60 nm can be obtained relatively easily. Notice that for this section thickness, the reconstruction SNR is about 19 dB.

obtain sections with thickness of about 50 – 60nm. Notice that even with very small number of views per section (NViews=3), the reconstruction quality is about 18dB when the section thickness is about 60nm. This indicates that our system is capable of performing high quality reconstructions with a very low cost in terms of throughput.

Since the eventual goal is to trace neural processes both within and across sections, we also quantitatively evaluate the reconstruction with a metric that better captures the ability to trace neural processes. The ability to perform 3D segmentation and tracing of neural processes is dependent on the quality and smoothness of the observed gradients G_x , G_y and G_z . The metric that we chose was the normalized dot product between the true gradients and the reconstructed gradients. The results are shown in Figure 8. Since the original acquisition via thick sections had different $x-y$ and z -resolutions, the reconstructions of the $x-y$ gradients are also significantly better than those for the z -gradients. We also compare the quality of the reconstruction with those obtained via traditional 3D interpolation on the acquired volume (which is non-uniformly sampled). The Figure clearly shows that with even as few as 5 views, the z -gradient reconstruction significantly improves the ability to trace neural processes across sections.

6.3. Effect of Noise and Registration Errors

In order to effectively fuse the observations of several tilt views across sections and create a single linear system, we need to perform registration both across views and across sections. In order to study the robustness of the reconstruction algorithm to errors in the registration, we simulated registration errors in the FIB data. Shown in Figure 9 is the SNR (in dB) of the reconstruction as the average registration errors varies. In practice, the registration algorithm that we use on real ssTEM data provides sub-pixel registration and the registration errors are typically utmost 1 – 2 pixels which is about 5 – 10nm.




Figure 8. Variable resolution of TEM: Normalized dot product of the true and reconstructed volume gradients. Notice that the thickness of the physical section affects the Z gradients much more significantly than the xy gradients. While simple interpolation schemes fail (dot product < 0.6) to capture the Z gradients at moderate section thickness, notice that our sparse tomographic reconstruction performs better. Gz is important in order to track membranes across sections in order to identify synapses and associate synapses with neurons and dendrites.




Figure 9. Effect of Registration Error on Volume Reconstruction: The volume reconstruction quality remains above 20 dB for average registration errors of less than 15 nm. In practice, we use a registration algorithm for registering across views that allows sub-pixel accuracy in most regions of the volume.

6.4. Increased Depth Resolution

Figure 10, shows some of the XZ and YZ sections from the reconstructions obtained by simulating both 50 and 75 nm thick sections and obtaining 5 tilt views for each section. The grayscale images are all shown in pseudocolor (colormap - jet) for visualization. The first column shows the original FIB data at 5nm isotropic voxel size. The second and fourth column shows the central view of the 50 and 75 nm thick sections that were simulated appropriately using the available high resolution FIB data. The third and fifth column shows the reconstruction obtained by our algorithm from five tilt views of the simulated 50 and 75 nm thick sections. The results clearly show greater depth resolution for the reconstructions. For each of the images notice the regions highlighted in a purple box showing the improvement

in z-resolution as a result of our reconstruction. We then applied a 3D segmentation algorithm based on watershed to segment the reconstructed volume into individual neurites. The results of the 3D segmentation are presented in Figure 11.




Figure 11. Segmentation of neurites from the high resolution FIB data and our reconstruction obtained via simulating 50 nm thick sections. Shown in (c) are some of the individual neurites in the volume.

7. Experiments on ssTEM

We collected real data from fly larva brain for evaluating the fidelity of our reconstruction algorithm. The tissue sample was cut into 50 nm thick sections using a diamond knife. Each of the 50nm samples was imaged at five different tilts using a Tecnai Spirit electron microscope at 4nm x-y pixel size. The first view was normal to the surface of the thick section. The next two views corresponded to tilting the section 45° and -45° about the y-axis. The last two views corresponded to tilting the section 45° and -45° about the x-axis. The five images within each section were registered using a simple affine registration. We acquired data for 30 consecutive sections and in each section we acquired five 4096×4096 pixel images. The total volume that was imaged was about $16\mu m \times 16\mu m \times 1.5\mu m$. Affine registration across sections was also performed using the central view of each section. We then divided the stack of thirty sections into overlapping blocks, with each block containing three sections. The reconstruction was performed independently for each overlapping block and the results for the central block was retained from each of the overlapping blocks. Shown in Figure 12 are some images from the original thick sections acquired by ssTEM and the corresponding images for the reconstructions. The reconstruction is shown in the second row of Figure 12. Notice the significantly higher z-resolution in the reconstructions compared to the original ssTEM data. Shown in Figure 13 is an illustration of the improvement in X-Z resolution and consequent improvement in the ability to segment neurites. Figure 13(a)(b) shows an X-Z section from the original ssTEM data and the corresponding X-Z section from the reconstruction respectively.

The segmentation of this section into neurites based on the reconstructed data is shown in Figure 13(c). We employ




Figure 10. Sample X-Z and YZ sections from the reconstructions. Zoom into the regions highlighted by purple box. The first column shows the 5nm resolution FIB data. The 2nd and 4th column shows a 50 and 75 nm thick sectioning while the 3rd and 5th column shows our reconstructions from 50 and 75 nm sections respectively.

a simple 3D hierarchical segmentation algorithm for reconstructing neuronal volumes in EM stacks. An initial dense over-segmentation is computed using 3D watershed. The watershed segments are clustered using an agglomerative segmentation algorithm based on boundary values between spatially adjacent segments. All stacks are pre-processed by histogram equalization in each section followed by median filtering with 3x3 kernel to reduce likelihood single-pixel breaches. The results of the 3D segmentation are shown in Figure 14. A 3D stack was created by concatenating the ssTEM images and the 3D segmentation algorithm was applied on this stack. In order to allow for misalignments in the membranes from one section to another, and because of coarse z-resolution, we tried several alternative preprocessing filters before preprocessing. These included performing median filtering by 3X3 and 8X8 kernel and also by dilating the boundaries with a disk of radius 4 pixels. Notice that the expected misalignment across sections is less than 3 pixels. The segmentation was tested manually for a sequence of several thresholds and the resulting segmentations were qualitatively evaluated. There did not exist a threshold using which we could reduce the number of false splits while still maintaining low mergers between processes.

8. Discussion and Conclusions

In this paper, we have proposed a method for increasing the depth resolution of ssTEM computationally while




Figure 12. Experiments on real data acquired via ssTEM. The top row shows an X-Z section of an original ssTEM stack obtained at 50 nm section thickness. The bottom row shows the reconstructed X-Z sections. Notice the significantly higher depth resolution in the reconstruction.




Figure 13. Illustration of segmentation in the X-Z section. (a) Original ssTEM data (b) Reconstructed data (c) segmentation into neurites based on the reconstructed data.




Figure 14. Two views of the 3D segmentation of neurites from our reconstruction in ssTEM data. Shown in (b) are some of the individual neurites in the volume.

imaging neural processes. The method requires only very few ssTEM images per section (5 or less) and therefore is only about a few times slower than traditional ssTEM while providing much higher depth resolution required for tracing neural circuits and identifying synapses. Other techniques that can provide similar depth resolution (vis-a-vis serial section transmission electron tomography and FIB) are both hundred times slower which is a huge handicap since the goal is to image and reconstruct about 10^{14} voxels. On the contrary, in contrast to FIB, we need to perform accurate alignment and registration of the captured images before they can be processed.

Acknowledgements

The first author would like to thank Prof. Rama Chelappa and Jay Thornton for their support and several members of MERL for their comments and suggestions.

References

[1] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008. 4

[2] J. Bioucas-Dias and M. Figueiredo. Two-step algorithms for linear inverse problems with non-quadratic regularization. In *IEEE International Conference on Image Processing ICIP*. Citeseer, 2007. 4

[3] K. Briggman and W. Denk. Towards neural circuit reconstruction with volume electron microscopy techniques. *Current opinion in neurobiology*, 16(5):562–570, 2006. 1

[4] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006. 4

[5] E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *COMMUNICATIONS ON PURE AND APPLIED MATHEMATICS*, 59(8):1207, 2006. 4

[6] B. Chen, D. Hall, and D. Chklovskii. Wiring optimization can relate neuronal structure and function. *Proceedings of the National Academy of Sciences of the United States of America*, 103(12):4723, 2006. 1

[7] W. Denk and H. Horstmann. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol*, 2(11):e329, 2004. 2

[8] D. Donoho. Compressed Sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006. 4

[9] D. Donoho. For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on pure and applied mathematics*, 59(6):797, 2006. 4

[10] K. Hayworth, N. Kasthuri, J. Lichtman, and E. Hartwig. Methods and apparatus for providing and processing thin tissue, June 5 2008. WO Patent WO/2008/066,846. 2

[11] M. Helmstaedter, K. Briggman, and W. Denk. 3D structural imaging of the brain with photons and electrons. *Current Opinion in Neurobiology*, 18(6):633–641, 2008. 1, 2

[12] S. Ji, Y. Xue, and L. Carin. Bayesian Compressive Sensing. *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, 56(6):2346–2356, 2008. 4

[13] G. Knott, H. Marchman, D. Wall, and B. Lich. Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *Journal of Neuroscience*, 28(12):2959, 2008. 2

[14] B. McEwen and M. Marko. The emergence of electron tomography as an important tool for investigating cellular ultrastructure. *Journal of Histochemistry and Cytochemistry*, 49(5):553, 2001. 2

[15] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 2008. 4

[16] J. White, E. Southgate, J. Thomson, and S. Brenner. The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philosophical Transactions B*, 314(1165):1, 1986. 1