

A Bayesian compressed-sensing approach for reconstructing neural connectivity from subsampled anatomical data

Yuriy Mishchenko and Liam Paninski

November 5, 2011

Abstract

In recent years, the problem of reconstructing the connectivity in large neural circuits (“connectomics”) has re-emerged as one of the main objectives of neuroscience. Classically, reconstructions of neural connectivity have been approached anatomically, using electron or light microscopy and histological tracing methods. This paper describes a statistical approach for connectivity reconstruction that relies on measurements of relatively easy-to-obtain fluorescent probes such as synaptic markers, cytoplasmic dyes, transsynaptic tracers, or activity-dependent dyes. We describe the possible design of these experiments and develop a Bayesian framework for extracting neural connectivity from compilations of such data. We show that the statistical reconstruction problem can be formulated naturally as a tractable L_1 -regularized quadratic optimization. As a concrete example, we consider a realistic hypothetical connectivity reconstruction experiment in *C. elegans*, a popular neuroscience model where a complete wiring diagram has been previously obtained based on long-term electron microscopy work. We show that the new statistical approach could lead to an orders of magnitude reduction in experimental effort in reconstructing the connectivity in this circuit. We further demonstrate that the spatial heterogeneity and biological variability in the connectivity matrix—not just the “average” connectivity—can also be estimated using the same method.

1 Introduction

The problem of the complete reconstruction of neural connectivity has, in recent years, re-emerged as a major research direction in neuroscience (Denk and Horstmann, 2004; Gray et al., 2005; Briggman and Denk, 2006; Broome et al., 2006; Svoboda, 2007; Haggmann et al., 2007; Smith, 2007; Wickersham et al., 2007; Petreanu et al., 2007; Helmstaedter et al., 2008; Bohland et al., 2009; Haggmann et al., 2008; Zimmer et al., 2009; Chiang et al., 2011). Since the seminal studies of Ramon y Cajal (Ramon y Cajal, 1904; Cajal, 1911), an enormous amount of information about the brain’s anatomical organization has been collected (Cajal, 1911; Westrum and Blackstad, 1962; Amaral and Witter, 1989; Braitenberg and Schuz, 1998). However, most details of fine-scale neural architecture remain obscure. Several large “connectomics” projects have recently been proposed, relying on such methods as electron microscopy (Briggman and Denk, 2006; Smith, 2007), large-scale injections of tracer viruses (Bohland et al., 2009), sparse expression of fluorescent cytoplasmic markers (Svoboda, 2007), or diffusion tensor magnetic resonance imaging (Haggmann et al., 2007; Haggmann et al., 2008). These developments have re-invigorated the field; however, the problem of practically affordable, routine circuitry reconstructions in neural systems remains far from being solved.

Classically, studies of neural connectivity have utilized conventional electron or light microscopy and histological techniques. For example, a dye can be injected into one or a group of neurons, and the projections of the stained neurons can be observed with conventional or fluorescent light microscopy. This approach is responsible for most of the information available today about the brain’s organization. However, only very sparse groups of neurons can typically be studied this way, and individual synapses usually cannot be observed (but see (Lichtman et al., 2008) for exciting recent developments). These limitations can be overcome using serial section electron microscopy (Westrum and Blackstad, 1962; White et al., 1986; Braitenberg and Schuz, 1998; Briggman et al., 2011; Bock et al., 2011), where dense groups of neurons can be examined and synapses between them can be observed directly by imaging large volumes of neural tissue at nanometer resolution (White et al., 1986; Mishchenko et al., 2010; Bock et al., 2011; Briggman et al., 2011). Such large scale serial electron microscopy semi-automated reconstructions in conjunction with fluorescent light microscopy were used recently to study the local organization of neuropil in rat hippocampus (Mishchenko et al., 2010), synaptic connectivity patterns of direction selective cells in mouse primary visual cortex (Bock et al., 2011), and amacrine cells in mouse retina (Briggman et al., 2011). However, data acquisition in this setting remains difficult, as is the analysis of the resulting very large imaging data sets (Burton, 1999; Briggman and Denk, 2006; Smith, 2007; Helmstaedter et al., 2008; Lu et al., 2009; Helmstaedter et al., 2011). The automation of data acquisition, image analysis, and neural process tracking in serial electron microscopy remains an area of active research (Denk and Horstmann, 2004; Smith, 2007; Helmstaedter et al., 2008; Mishchenko, 2009; Lu et al., 2009; Briggman et al., 2011; Bock et al., 2011; Helmstaedter et al., 2011).

In parallel with the anatomical histological approach, a novel paradigm for the reconstruction of neural connectivity has recently emerged (Svoboda, 2007; Micheva and Smith, 2007; Mishchenko, 2008; Bohland et al., 2009; Mishchenko, 2011). This paradigm relies on using large samples of relatively straightforward measurements obtained with fluorescent connectivity probes, such as tracer viruses (Wickersham et al., 2007; Bohland et al., 2009), cytoplasmic dyes (Svoboda, 2007), synaptic markers (Micheva and Smith, 2007; Mishchenko, 2008), or calcium dyes (Bureau et al., 2004; Sato et al., 2007; Mishchenko et al., 2011). In one example, a genetically modified tracer virus — that is, a virus that naturally propagates in the nervous system by jumping across synaptic connections and is genetically modified to cause infected neurons to produce a fluorescent protein such as Green Fluorescent Protein (Heim et al., 1995), hereafter GFP — can be injected into certain population of neurons, and secondary neurons infected via the virus’s synaptic propagation can be subsequently imaged (Wickersham et al., 2007). By repeatedly introducing such a virus in different primary groups of neurons, the detailed connectivity in a large neural system can be examined (Bohland et al., 2009). Similarly, the connectivity can be examined by aggregating data obtained with high throughput fluorescent imaging experiments performed with activity dependent calcium dyes (Mishchenko et al., 2011), sparsely expressed neural cytoplasmic markers (Svoboda, 2007), or fluorescent synaptic markers (Mishchenko, 2011). With recent developments in genetic systems for stochastic or specific genetic targeting (Luo et al., 2008), such large high throughput applications now become possible in practice.

While such new connectomics paradigms have been undergoing rapid development in the experimental community, the problem of analyzing the resulting data has remained largely unaddressed. This paper develops a general Bayesian framework for reconstructing neural connectivity from compilations of such fluorescent measurements obtained with various high throughput fluorescent connectivity probes. We show that the problem of reconstructing the

connectivity matrix in a large neural system from an ensemble of such measurements can be formulated naturally in terms of standard sparse regression techniques, known in the statistics literature as L_1 -penalized regression (Tibshirani, 1996) or in the applied math literature as “compressed sensing” (Candes and Romberg, 2005; Donoho, 2006; Candes et al., 2006; Candes and Wakin, 2008; Romberg, 2008). This provides a reformulation of previous work applying compressed sensing ideas to the connectivity analysis for connectivity data for high throughput GFP Reconstitution Across Synaptic Partners (GRASP) experiments (Mishchenko, 2008) or randomized photostimulation experiments (Hu et al., 2009). In addition, we demonstrate how similar ideas can be applied to a variety of other experimental settings currently considered as possible bases for large-scale future connectomics reconstructions.

As a concrete demonstration of the utility of our Bayesian formalism, we analyze simulated connectivity reconstruction experiments in *C. elegans*, a popular neuroscience model where a complete wiring diagram is known from long-term electron microscopy efforts (White et al., 1986). The connectivity in this circuit can, in principle, be inferred accurately with an orders of magnitude reduction in experimental effort, compared with classical approaches. We find that the methods proposed here significantly outperform the simpler inference techniques used in our previous work (Mishchenko, 2008; Mishchenko, 2011) in undersampled settings (i.e., where relatively few data observations are available), though the new methods are more computationally costly. Additionally, our approach allows the estimation of other neural system parameters, including the variability and spatial heterogeneity of the connectivity matrix, which would be difficult to assess with traditional approaches.

The paper is organized as follows. We begin in section 2.1 with a specific example of the ensemble of measurements obtained with randomly expressed fluorescent synaptic marker GRASP (Feinberg et al., 2008) or fluorescence co-localization markers (Mishchenko, 2010). We use this example to introduce the connectivity inference problem for an ensemble of high throughput fluorescent connectivity measurements, and then introduce and reduce the Bayesian formulation for connectivity inference from such data to a standard L_1 -penalized regression. Numerical methods for solving the resulting optimization problem are discussed in section 2.2. Results of numerical experiments and analysis of a realistic neural connectivity reconstruction problem in *C. elegans* are given in sections 3.1-3.2. In section 3.3 we discuss how spatial information obtainable from fluorescent light imaging can be incorporated into our framework to further dramatically improve connectivity estimates, and provide additional information about the spatial heterogeneity of the connectivity matrix. Conclusions and applications to other high throughput fluorescent microscopy connectomics paradigms are discussed in section 4.

2 Methods

2.1 Reconstruction of the neural connectivity matrix by sampling physical connections using fluorescent synaptic markers

We begin by describing our basic approach in a concrete example. In previous work (Mishchenko, 2008; Mishchenko, 2011) we proposed to exploit the experimental technique of stochastic GRASP to reconstruct a physical neural connectivity matrix: the idea is to sample synaptic connections by stochastically expressing the fluorescent synaptic marker GRASP (Feinberg et al., 2008), or pre- and post-synaptic co-localization fluorescent markers (Mishchenko, 2010), and then to use statistical inference techniques to estimate the full connectivity matrix. To

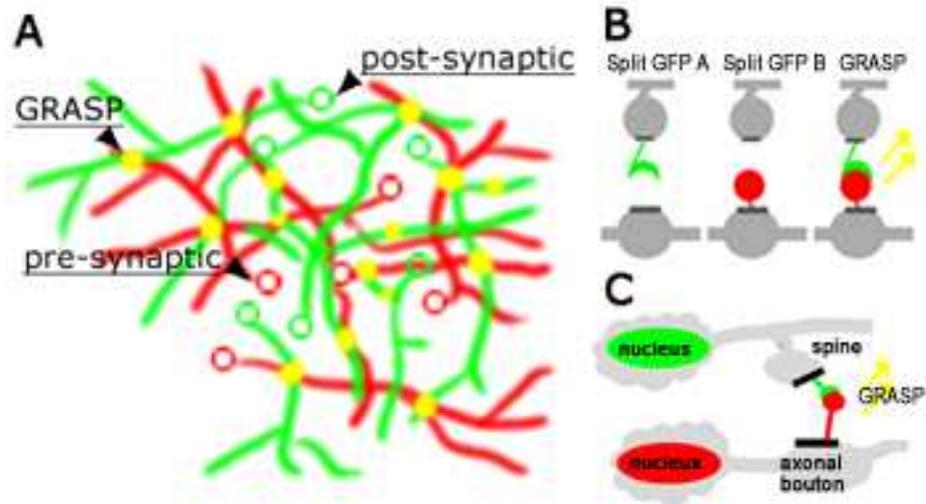


Figure 1: A cartoon schematic of high throughput GRASP sampling of neural connectivity (Mishchenko, 2008). A) A fluorescent synaptic marker GRASP (panel B) is randomly expressed in the neurons in a neural circuit using a genetic technique for stochastic gene expression such as Brainbow (Livet et al., 2007). Synapses between random pre-synaptic (red) and post-synaptic (green) neurons are thus labeled (yellow) and can be observed with fluorescent light microscopy. The pattern of labeled synapses in each animal is defined by the specific expression pattern of GRASP in different neurons and with high probability is unique in each animal. Since using current genetic techniques it is only possible to generate a single pattern per animal, only one experiment can be performed per animal. Nonetheless, experiments in a population of animals can be used to collect a statistical sample of such patterns sufficient for fully encoding the structural connectivity of the underlying synaptic circuit. B, C) GRASP (Feinberg et al., 2008) is a fluorescent synaptic marker that uses a complementary pair of proteins that can chemically recombine into a functional GFP whenever in each other’s proximity. One such protein (split GFP B) is introduced to the pre-synaptic region of selected neuron using genetic engineering techniques, and the other (split GFP A) is introduced to the post-synaptic region of another neuron. Wherever a synapse is formed by such neurons, the split GFPs can chemically interact to restore functional GFP, thereby tagging the synapse with a fluorescent puncta.

make our ideas explicit, we will explain the stochastic GRASP idea first, and then describe how to formulate the resulting inference problem in a Bayesian framework.

The GRASP technique (Feinberg et al., 2008) uses a synaptic marker composed of two proteins that chemically recombine into a functional fluorochrome (specifically, green fluorescent protein, hereafter GFP (Heim et al., 1995)) whenever in each other’s proximity (Zhang et al., 2004). Such proteins are expressed on the surface of the synapses of different neurons via a genetic engineering technique. At these synapses, such proteins recombine and form fluorescent puncta that can be observed with light microscopy and used to study the connectivity of these neurons (Feinberg et al., 2008) (see Fig. 1). The components in stochastic GRASP are expressed stochastically, so that the first component is expressed in one random set of neurons, and the second component in another. This allows the stochastic visualization

of synapses between different subsets of neurons with a standard genetic screen. Since this method relies on using a single fluorescent synaptic marker, all visualized synapses fluoresce with the same color, and therefore only coarse, gross measures can be assessed in a given experiment. (Mishchenko, 2008; Mishchenko, 2011) show that a simple measure such as the total count or total size (brightness) of labeled synaptic puncta in each pattern can suffice to exactly recover the connectivity matrix in the underlying neural circuit, given a sufficient number of such experiments.

While the problem of deducing the connectivity matrix from such data can be approached using compressed sensing techniques (Mishchenko, 2008; Mishchenko, 2011), this study considers the problem from a more standard Bayesian perspective. For concreteness, we develop this approach using the GRASP example described above, though as we discuss in Section 4, this approach can be applied in a number of other different high throughput fluorescent light microscopy connectomics paradigms.

We assume that we have a series of high throughput fluorescent light microscopy experiments producing some coarse measures of connectivity. We will denote different experiments in the series using the symbol k . For each experiment k we will denote the set of pre-synaptic and post-synaptic neurons affected in that experiment with the symbols A^k and B^k , respectively. We will denote the measurements, e.g. the total count of labeled synapses per experiment, with the symbol O^k . By introducing a synaptic connectivity matrix C_{ij} that describes the total count of synaptic contacts between pairs of neurons i and j in the circuit, we have that O^k , A^k , and B^k obey the equation

$$O^k = \sum_{i \in A^k} \sum_{j \in B^k} C_{ij}^k + \epsilon^k. \quad (1)$$

Equation (1) states that all synapses between neurons A^k and B^k that are “selected” in the experiment k are counted in O^k . Note that with i and j we can refer to both individual neurons as well as, more generally, entire classes of neurons, such as morphological or genetic classes, geometrically defined areas of the brain, etc., with respective changes in the interpretation of C_{ij} as describing the count of synapses between such groups of neurons. The error term ϵ^k describes the possible corruption of the measurements due to imperfections in the imaging process or noise; as a starting point, we can model these effects using the normal distribution, $\epsilon^k \sim \mathcal{N}(0, \sigma^2)$, though generalizations are possible.

If measurements are obtained from different animals, as discussed above, the measured connectivity matrix C_{ij}^k may vary in different experiments k . Thus it makes sense to consider C_{ij}^k as a random variable sampled from some distribution over possible neural circuit configurations, $\{C_{ij}^k\} \sim P(\{C_{ij}\})$, rather than an immutable constant.

To deduce the connectivity from a sample of such data, we introduce the Bayesian maximum a posteriori (MAP) estimate for the “average” connectivity matrix $M_{ij} = \langle C_{ij}^k \rangle$. The average connectivity matrix, specifically, represents the average count of synapses between different neurons or neural groups i and j as observed over a sufficiently large population of animals in a defined population (for example, for a particular genotype). As an example, consider the connectivity matrix in *C. elegans* (White et al., 1986). As was determined in extensive previous anatomical studies, the neurons in *C. elegans* are stereotypical and can be well identified by their position and morphological features. The connectivity matrix can be tabulated as the number of distinct anatomical synaptic contacts between different pairs of such neurons. For example, for the neurons identified in the *C. elegans* wiring diagram

as ADFL and RIAL, the count of synapses in the wiring diagram in (White et al., 1986) is $C_{RIAL,ADFL} = 12$. Note that this count was produced for the specimen imaged and analyzed in the above work; in other specimens such counts may further vary. If, after inspecting a sufficiently large number of similar specimens, we found that the count of synapses between these neurons ADFL and RIAL, on average, was 10.5, then we set $M_{RIAL,ADFL} = 10.5$.

By Bayes' theorem,

$$\hat{\mathbf{M}} = \arg \max_{\mathbf{M}} P(\mathbf{M}|\mathbf{O}, \mathbf{A}, \mathbf{B}) = \arg \max_{\mathbf{M}} P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})P(\mathbf{M}) \quad (2)$$

(note the bold notation for vector and matrix quantities, i.e., $\mathbf{O} = \{O^k \text{ for all } k\}$). As the first step in the MAP estimation, we want to evaluate $P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})$ given the ensemble of observations. At first glance, the determination of $P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})$ is a formidable problem, since $P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})$ is a result of a high-dimensional convolution with $P(\mathbf{C}^k)$, and involves integration over the ‘‘intermediate’’ connectivity variables \mathbf{C}^k , which are not observed directly:

$$P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B}) = \int d\{\mathbf{C}^k\} \prod_k P(O^k|\mathbf{C}^k, A^k, B^k)P(\mathbf{C}^k|\mathbf{M}). \quad (3)$$

$P(O^k|\mathbf{C}^k, A^k, B^k)$ represents the probability of observing a particular count of labeled puncta, O^k , for a given marker expression pattern, A^k and B^k , and connectivity between neurons, \mathbf{C}^k . Given Eq. (1) and the assumption of the Gaussian noise,

$$P(O^k|\mathbf{C}^k, A^k, B^k) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left\{ -\frac{\left(O^k - \sum_{ij} e_{ij}^k C_{ij}^k\right)^2}{2\sigma^2} \right\}. \quad (4)$$

Here σ^2 is the variance of ϵ^k in Eq. (1) and we have abbreviated $e_{ij}^k = a_i^k b_j^k$, where a_i^k and b_j^k are the indicator functions describing the count of neurons of kind i in A^k or B^k , respectively. The term $P(\mathbf{C}^k|\mathbf{M})$ represents the probability of observing a particular connectivity matrix \mathbf{C}^k in an animal given an average connectivity matrix \mathbf{M} . Then equation (3) defines $P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})$ by the law of total probability.

To simplify $P(\mathbf{O}|\mathbf{M}, \mathbf{A}, \mathbf{B})$ further, observe that each O^k is a weighted sum of a large number of random variables $C_{ij}^k \sim P(C_{ij})$. As long as the probability of labeling A^k and B^k are sufficiently larger than zero, so that the number of nonzero e_{ij}^k terms in the sum is sufficiently large, we can conclude by the central limit theorem that $P(O^k|\mathbf{M}, A^k, B^k)$ is approximately Gaussian, with mean $\sum_{ij} e_{ij}^k M_{ij}$, and variance

$$\begin{aligned} \text{Var} \left(O^k|\mathbf{M}, A^k, B^k \right) &= \text{Var} \left(\sum_{ij} e_{ij}^k C_{ij} \right) + \sigma^2 \\ &= \sum_{ij} (e_{ij}^k)^2 \text{Var}(C_{ij}) + 2 \sum_{ij \neq i'j'} e_{ij}^k e_{i'j'}^k \text{Cov}(C_{ij}, C_{i'j'}) + \sigma^2. \end{aligned}$$

Now we can argue as above to simplify these large sums, using the law of large numbers. For example, if we make the assumption that the C_{ij}^k are independent random variables, and that

the variables e_{ij}^k are independent of the connections C_{ij}^k , then the covariance terms are zero, and applying the law of large numbers we find that

$$\text{Var}(O^k | \mathbf{M}, A^k, B^k) \approx \left(\sum_{ij} E(e_{ij}^2) \text{Var}(C_{ij}) \right) + \sigma^2 \equiv \sigma_t^2,$$

For example, if the C_{ij} are Poisson random variables, then we can explicitly compute σ_t^2 as $\left(\sum_{ij} E(e_{ij}^2) M_{ij} \right) + \sigma^2$. More generally, we can expect these central limit theorem and law of large numbers arguments to apply under very weak conditions on the distribution $P(\{C_{ij}\})$, due to the very large sums involved; i.e., correlations in the synaptic weights C_{ij} will change the value of the overall variance σ_t^2 , but the Gaussian approximation for $P(O^k | \mathbf{M}, A^k, B^k)$ should hold broadly. Thus we arrive at the rather simple likelihood term

$$\log P(\mathbf{O} | \mathbf{M}, \mathbf{A}, \mathbf{B}) \approx - \sum_k \frac{1}{2} \frac{\left(O^k - \sum_{ij} e_{ij}^k M_{ij} \right)^2}{\sigma_t^2} + \text{const.} \quad (5)$$

Having defined the likelihood term, we now turn our attention to the prior term. In neurobiological settings, the connectivity matrix \mathbf{M} is known to be very sparse. For example, in the popular neuroscience model organism *C. elegans*, only about 3% of the 90,000 possible connections are realized (White et al., 1986). In *Drosophila melanogaster*, the total count of neurons is about 10^5 and the typical number of synapses per neuron is on the order of 100 (Chiang et al., 2011), putting the neural connectivity matrix sparsity at about 0.1%. Sparse priors are known to allow dramatic reductions in the number of measurements necessary for accurate signal reconstruction (Tibshirani, 1996; Tipping, 2001; Donoho and Elad, 2003; Ng, 2004; Candes and Romberg, 2005; Mishchenko, 2008; Hu et al., 2009). To account for sparseness in our setting, we introduce an exponential prior,

$$P(\mathbf{M}) \propto \exp \left(- \rho^{-1} \sum_{ij} |M_{ij}| \right), \quad (6)$$

where ρ is the expected density of \mathbf{M} . Specifically, ρ defines the expected average number of synapses between two neurons in the system. Such averages are well known for most neurobiological systems from past neuroanatomical studies, as noted above¹.

Combining the likelihood term (5) and the prior term (6), we write the full log posterior as

$$\log P(\mathbf{M} | \mathbf{O}, \mathbf{A}, \mathbf{B}) \approx - \sum_k \frac{1}{2} \frac{\left(O^k - \sum_{ij} e_{ij}^k M_{ij} \right)^2}{\sigma_t^2} - \rho^{-1} \sum_{ij} |M_{ij}| + \text{const.}, \quad (7)$$

¹We should emphasize that this simple sparseness prior should be considered just a first step; clearly additional prior information can be incorporated here. One obvious additional source of information could include the distance between neurons; in many cases the probability of finding a connection between two neurons shrinks as a function of the distance between the cells, and it is straightforward to incorporate this kind of information in our prior by including a distance-dependent term. Prior information about the connection probabilities between different cell types can be incorporated in a similar manner. More generally it should be feasible to incorporate other log-concave priors in the Bayesian maximum a posteriori approach followed here. For example, rank-penalizing priors such as the exponentiated nuclear-norm penalty have proven useful in a wide variety of matrix estimation problems (Candes and Tao, 2010; Goldfarb et al., 2010), and deserve further examination in the context of neural connectivity reconstruction problem.

and

$$\hat{\mathbf{M}}^{MAP} = \operatorname{argmin} \left[\sum_k \left(O^k - \sum_{ij} e_{ij}^k M_{ij} \right)^2 + 2\lambda \sum_{ij} |M_{ij}|, \right] \quad \lambda = \sigma_t^2 / \rho. \quad (8)$$

The nonnegativity constraint $M_{ij} \geq 0$ may be additionally enforced if the mean connectivity matrix M_{ij} is expected to be nonnegative, e.g. if it describes the expected count of (unsigned) synapses between neurons, which clearly can not be negative. Such a constraint can significantly simplify the task of numerically solving Eq. (8) but, in general, nonnegativity cannot be imposed if the connectivity matrix is not expected to be sign definite. For example, if M_{ij} describes the post-synaptic potential (PSP) strength of neural connections in an excitatory and inhibitory network, as discussed in section 4 below, the nonnegativity constraint cannot be enforced.

Equation (8) is an instance of a quadratic optimization problem known as the least absolute shrinkage and selection operator, hereafter LASSO (Tibshirani, 1996), a shrinkage and selection method for linear regression². The parameter $\lambda = \sigma_t^2 / \rho$ in Eq. (8) can be set using, for example, cross-validation techniques (Hastie et al., 2001), or with independent knowledge of the parameters σ_t^2 and ρ . As was emphasized above, excellent estimates for ρ are available a priori from past neuroanatomical studies. We discuss the problem of estimating σ_t^2 further below.

If σ_t^2 / ρ can be expected to be small, it may be also meaningful to consider Eq. (8) in the limit $\lambda \rightarrow 0$. In such a limit, certain bounds on LASSO reconstructions can be adapted from corresponding results from the compressed sensing literature, leading to theoretical predictions for the convergence speed and reconstruction accuracy with the LASSO. Specifically, we have the following propositions.

Proposition I: Let $\sigma_t^2 \rightarrow 0$ and let M_{ij} be a sparse matrix with \mathcal{N} elements, with support on an unknown set T of size N_s . Let $\{O^k\}$ be the set of the data samples in problem (8) and $K \geq K_{min} = C_\delta N_s \log \mathcal{N}$ be the number of samples, for a certain constant C_δ of order 1, depending on an accuracy parameter δ . Then under weak technical conditions on the matrix $\{e_{ij}^k\}$, with probability of at least $1 - O(\mathcal{N}^{-\delta})$, the minimizer of the problem (8) is unique and equal to the true M_{ij} .

In the case of the neural reconstruction problem, N_s is the number of nonzero elements in the connectivity matrix, that is, the number of neural connections in the circuit, and $\mathcal{N} \approx N^2$ is the total number of possible connections, where N is the number of neurons in the system. According to Proposition I, then, an accurate reconstructions of the connectivity matrix can be obtained with LASSO using only $K_{min} \propto N_s$ measurements O^k , much less than the expected $K_{min} \propto N^2$. In fact, the data sample with which LASSO can produce an accurate reconstruction is comparable to the minimal number of measurements needed to fully specify N_s nontrivial elements of \mathbf{M} had we known their positions a-priori.

Proposition II: Let $\sigma_t^2 \rightarrow 0$ and let the number of the data samples in problem (8) be K . Under similar weak technical conditions on the matrix $\{e_{ij}^k\}$, the minimizer of the problem (8), \mathbf{M}^* , obeys $\|\mathbf{M}^* - \mathbf{M}\|_1 \leq C_0 \|\mathbf{M} - \mathbf{M}_S\|_1$, for some constant C_0 , where \mathbf{M}_S is the matrix with all but the largest S components set to 0.

²To be precise, the LASSO problem corresponds to the case that the variance parameter σ_t^2 does not depend on the mean parameter \mathbf{M} . As discussed above, for some connectivity distributions $p(\{C_{ij}\})$ σ_t^2 might depend weakly on \mathbf{M} . For example, in the Poisson case, σ_t^2 corresponds to a one-dimensional projection of the very high-dimensional parameter \mathbf{M} . These cases can be handled with straightforward extensions of the LASSO approach; however, in this paper we will focus on the model with σ_t^2 independent of \mathbf{M} .

Proposition II is stronger than Proposition I. If \mathbf{M} is S -sparse, then $\mathbf{M} = \mathbf{M}_S$ and the recovery is exact. But if \mathbf{M} is not S -sparse, then the quality of the reconstruction is as good as if one knew ahead of time the location of the S largest values of \mathbf{M} and decided to measure those directly. In the recent applied math literature this statement had been proven for a wide range of distributions on $\{e_{ij}^k\}$ using the ideas of restricted isometry (Romberg, 2008; Candes and Wakin, 2008). If $K < K_{min}$ in Proposition I, then LASSO reconstructions still identify the strongest connections in the system, and the reconstruction can be continuously improved as more data is made available until $K = K_{min}$.

Proposition III: Let $\rho \rightarrow \infty$ while $\sigma_t^2 > 0$, so that $\lambda \rightarrow 0$, and let the number of the data samples in problem (8) be K . Under similar weak technical conditions on the matrix $\{e_{ij}^k\}$, the minimizer of the problem (8), \mathbf{M}^* , obeys $\|\mathbf{M}^* - \mathbf{M}\|_{l_2} \leq C_0 \|\mathbf{M} - \mathbf{M}_S\|_1 + C_1 \sigma_t$, for some constants C_0 and C_1 , where \mathbf{M}_S is the matrix with all but the largest S components set to 0.

Proposition III means that the reconstruction error of LASSO in the $\lambda \rightarrow 0$ limit is bounded by the sum of two terms. The first is the error which would occur if one had noiseless data. The second is proportional to the noise level. In particular, this means that the reconstruction even in the heavily undersampled regimes $K \propto N_s$ are stable: small amounts of noise will not immediately ruin the reconstruction but, instead, will introduce a proportionately small deviation from the truth.

These propositions are fully analogous to respective theorems from the compressed sensing literature — specifically, see (Candes et al., 2006) for Proposition I, (Candes and Wakin, 2008; Romberg, 2008) for Propositions II and III — and can be proved after establishing the equivalence between the $\lambda \rightarrow 0$ LASSO and the compressed sensing optimization programs. Note that the compressed sensing optimization program for the compressed sensing estimator $\hat{\mathbf{M}}_{CS}$ for a set of measurements \mathbf{O} and a sensing matrix \mathbf{E} is defined as follows: $\hat{\mathbf{M}}_{CS} = \text{argmin} \|\mathbf{M}\|_1$ s.t. $\mathbf{O} = \mathbf{E} \cdot \mathbf{M}$ (Candes and Wakin, 2008; Romberg, 2008). If the measurements \mathbf{O} and the matrix $\mathbf{E} = \{e_{ij}^k\}$ are such that there exist an \mathbf{M}^* such that $\mathbf{O} = \mathbf{E} \cdot \mathbf{M}^*$ identically for that \mathbf{M}^* , then we can equivalently re-write the compressed sensing optimization program as $\min \|\mathbf{M}\|_1$ s.t. $(\mathbf{O} - \mathbf{E} \cdot \mathbf{M})^2 = 0$. By using the standard Lagrange multipliers trick to convert the penalized LASSO optimization

$$\min [(\mathbf{O} - \mathbf{E} \cdot \mathbf{M})^2 + \lambda \|\mathbf{M}\|_1] \quad (9)$$

to a constrained optimization of the form

$$\min \|\mathbf{M}\|_1 \text{ s.t. } (\mathbf{O} - \mathbf{E} \cdot \mathbf{M})^2 \leq \epsilon^2(\lambda), \quad (10)$$

we can show that $\|\hat{\mathbf{M}}_{CS} - \hat{\mathbf{M}}_{LASSO}\|_2^2 = O(\epsilon^2(\lambda))$. $\epsilon^2(\lambda)$ can be subsequently bounded by using the equivalence of the problems (9) and (10) to show that $\|\hat{\mathbf{M}}_{CS} - \hat{\mathbf{M}}_{LASSO}\|_2^2 \leq B \|\mathbf{M}^*\|_1 \lambda$ for some constant B , so that $\hat{\mathbf{M}}_{LASSO}(\lambda) \rightarrow \hat{\mathbf{M}}_{CS}$ continuously as $\lambda \rightarrow 0$.

We further note that we can, in principle, estimate not only the average connectivity matrix $\hat{\mathbf{M}}$, but also the variability in the connectivity matrix. For example, Gibbs sampling methods have been applied successfully in smaller LASSO problems to sample from the posterior distribution of the parameter, given the observed data (Park and Casella, 2008). Unfortunately, the Gibbs sampling approach is too computationally expensive in the context of the very large neural connectivity reconstruction problem considered here. Thus, we developed a simple estimator for σ_t^2 ,

$$\hat{\sigma}_t^2 = \frac{1}{K - \hat{N}_s} \sum_k \left(O^k - \sum_{ij} e_{ij}^k \hat{M}_{ij}^{MAP} \right)^2, \quad (11)$$

where \hat{N}_s is the size of the MAP estimator support, that is, the number of elements $\hat{M}_{ij}^{MAP} \neq 0$. This estimator is equivalent to the standard unbiased estimator for the variance parameter in multiple linear regression settings, with \hat{N}_s playing the role of the number of “degrees of freedom” (Kutner et al., 2005). This formula makes sense if K is large enough so that the nonzero elements of M_{ij} have been identified correctly; as discussed above, in these sparse settings a relatively small value of K suffices. See (Zou et al., 2007) for a more detailed discussion of the degrees of freedom of LASSO estimators; as we discuss further below, we found that this simple estimator for σ_t^2 performs well in practice.

While Eq.(11) allows us to estimate the total variance σ_t^2 , recall that σ_t^2 is composed from σ^2 , which can be viewed as the noise contributed by the imaging process, and $\text{var}(C_{ij}^k)$, which can be viewed as the implicit “biological” variability in the structure of the inspected neural network. Given an estimate of σ_t^2 , then, the biological variability in the connectivity matrix can be estimated by subtracting an estimate of the noise in the imaging system, $\hat{\sigma}^2$, from $\hat{\sigma}_t^2$. The former can typically be measured, for example, by calibrating the experiment under controlled conditions.

2.2 Numerical approaches for solving MAP neural connectivity reconstruction

Problem (8) is a standard *quadratic program* and can be solved using standard interior point methods or active set methods (Vanderbei, 2001; Boyd and Vandenberghe, 2004). However, the size of this problem in neurobiological settings, with at least $N \sim 10^3 - 10^5$ neurons and $\mathcal{N} \approx 10^5 - 10^{10}$ possible connections, requires the development of specialized algorithms.

To make progress here, we have to take advantage of the specific structure of the Hessian matrix \mathbf{H} (which is the quadratic part of our objective function). Specifically, \mathbf{H} is a product of two low-rank $K \times \mathcal{N}$ matrices, \mathbf{e} and \mathbf{e}^T , which in turn are completely specified by the $2K$ sets A^k and B^k . It may be more beneficial then to define the linear operator corresponding to the Hessian matrix \mathbf{H} by its action on vectors, implicitly,

$$\mathbf{H} \cdot \mathbf{M} = \left\{ \sum_{i',j',k} e_{ij}^k e_{i'j'}^k M_{i'j'}, \text{ for all } i, j \right\}. \quad (12)$$

Each evaluation of Eq. (12) requires storing of only, at most, $2KN$ elements of $\{A^k\}$ and $\{B^k\}$ and performing $2KN$ arithmetic operations, significantly better than storing \mathcal{N}^2 elements of \mathbf{H} and performing \mathcal{N}^2 operations.

To numerically solve Problem (8), therefore, we use the interior point method based on the interior-reflective Newton method, as implemented in Matlab’s optimization toolbox, augmented with the implicit Hessian definition given by Eq. (12). The computational cost of such an algorithm scales with the number of measurements K and the size of the connectivity matrix \mathcal{N} as $K\mathcal{N}$. Even though the interior point method here does not maintain the sparse representation, the structure of \mathbf{H} as the outer product of two low-rank matrices can be used to the same effect to perform matrix–vector multiplication in $K\mathcal{N}$ rather than \mathcal{N}^2 operations.

Given that for a neural circuit with N neurons the size of the connectivity matrix $\mathcal{N} = N^2$ and $K_{min} \propto N_s \log N \propto N^2 \log N$ (assuming a constant connectivity sparseness), $K\mathcal{N} \propto N^4 \log N$ and the neural reconstruction computational cost should scale as N^4 . This scaling is observed in our calculations, described below (see Fig. 6). Since the computation is dominated by the multiplication of matrices, the algorithm, in principle, can be efficiently parallelized

and applied to very large problems. In our simulations, a laptop (Intel dual core, 2GHz, 3GB RAM) was sufficient to solve a neural connectivity reconstruction problem for the entire *C. elegans* circuit having $N \approx 300$ neurons and $\mathcal{N} \approx 10^5$ possible connections. Larger problems can be approached using parallelized algorithms on cluster- or super-computing infrastructure (Boyd et al., 2011).

3 Results

3.1 MAP connectivity inference from surrogate GRASP data generated using the *C. elegans* wiring diagram

In this section we apply the Bayesian reconstruction algorithm to a surrogate GRASP dataset in *C. elegans*. *C. elegans* is a popular model organism whose nervous system has been extensively studied using genetics, light microscopy, and electron microscopy. The neural wiring diagram for this organism was completely reconstructed in a heroic, decade-long electron microscopy effort (White et al., 1986) (Fig. 2). This wiring diagram contains approximately 300 neurons, 2500 distinct neural connections, and 6000 synapses. The strength of each connection is quantified as the count of distinct synaptic contacts formed between corresponding neurons. Such counts vary from one to approximately 40, with two to three contacts, on average, per connected neural pair. To date, *C. elegans* remains the only organism in which such a complete wiring diagram is empirically available.

We use this wiring diagram to generate a surrogate data set for an ensemble of GRASP measurements in *C. elegans*. For each measurement k , we first prepare the individual connectivity matrix \mathbf{C}^k . We chose the average connectivity matrix, \mathbf{M} , to coincide with the wiring diagram from *wormatlas.org*, and the variable part is modeled using a scaled Poisson random variable,

$$C_{ij}^k = \alpha M_{ij} + \text{Poisson}[(1 - \alpha)M_{ij}], \quad (13)$$

independently for different neural pairs (i, j) . The parameter α controls the degree of biological variability: $\alpha = 0$ corresponds to no variability and $\alpha = 1$ corresponds to the case where all synapses are formed completely at random, with M_{ij} synapses, on average. We varied α from $\alpha = 0$ to $\alpha = 1$ to inspect the impact of biological variability in the connectivity matrix on the reconstructions. At the same time, as was discussed in the previous sections, this impact, $\text{Var}(C_{ij}) > 0$, can be absorbed into the definition of the total noise, $\sigma_t^2 = \sigma^2 + \sum_{ij} E(e_{ij}^2) \text{Var}(C_{ij})$, so below we examine the impact of the total noise strength, σ_t^2 , on the reconstruction accuracy, where α is left implicit. The GRASP expression patterns A^k and B^k are generated from a uniform distribution over all neurons with a density of 50% (Livet et al., 2007; Mishchenko, 2008; Mishchenko, 2011). Each measurement O^k is then simulated using Eq. (1).

Figures 3 and 4 illustrate the result of our numerical experiments. As a simple measure of the reconstruction quality, we use the correlation coefficient r^2 between the original and reconstructed weights in the connectivity matrix. In low-noise regimes the Bayesian estimator produces nearly exact reconstructions, that is, $r^2 \approx 1$, with a data sample containing $K_{min} \approx 10^4$ measurements, while $K \approx 2000 - 3000$ measurements were found to be sufficient for reconstructions that recovered at least 50% of the variance in the connectivity matrix (that is, $r^2 \geq 0.5$). This data sample is significantly smaller than the $\propto N^2 \approx 10^5$ measurements that one can naively expect to be necessary to fully characterize the connectivity matrix of

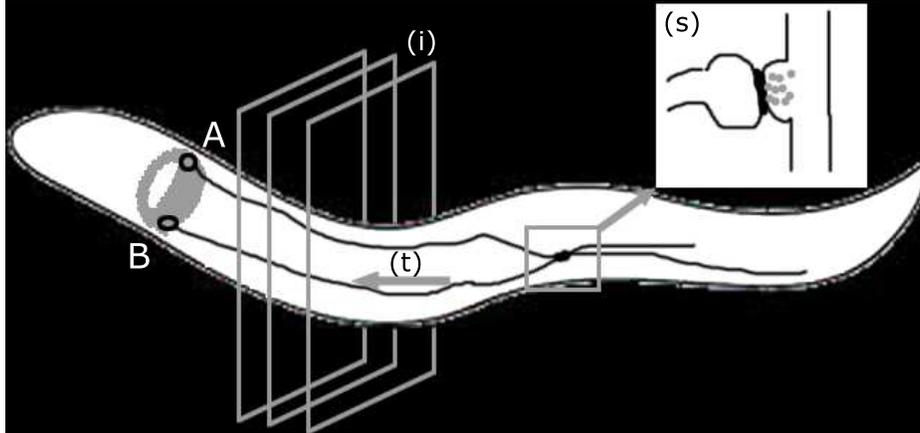


Figure 2: Schematic of the serial electron microscopy reconstruction of the wiring diagram in *C. elegans*, performed in (White et al., 1986). The entire volume of one *C. elegans* specimen was histologically prepared and imaged using serial electron microscopy as a stack of images (i) at very high resolution, several nanometers per pixel. In these images, synapses were manually located and associated with corresponding axons and dendrites (s). These axons and dendrites were then traced through the images until the corresponding cell bodies were found (t). A synaptic contact between the respective neurons A and B was then recorded. After examination of the entire volume of images, a table describing the number of synaptic contacts for all different pairs of neurons was compiled (see *wormatlas.org*).

$N = 300$ neurons. Instead, the number of samples required is much closer to the number of nonzero elements in the wiring diagram, $N_s \approx 2500$, as expected from the theory (Section 2.1).

We compare our results with the slightly cheaper compressed-sensing alternate projection algorithm used in (Mishchenko, 2008), as well as the much cheaper L_2 regularized estimator. The estimator in the alternating projection algorithm (Candes and Romberg, 2005) is operationally defined as the intersection of two convex sets, the hyperplane of constraints $O^k = \sum_{ij} e_{ij}^k M_{ij}$, $k = 1, \dots, K$, and the hypercube $\|\mathbf{M}\|_1 = s$, where the constant s is an estimate of the norm of the true solution $\|\mathbf{M}^*\|_1 = s$. By Candes' theorem (Candes et al., 2006), this intersection is unique and equal to \mathbf{M}^* with high probability in the regime $K \geq K_{min}$. For $K < K_{min}$ this intersection typically will not be unique and an arbitrary point on the intersection set is usually found depending on the starting point of the algorithm. The L_2 regularized estimator is defined as

$$\hat{\mathbf{M}}^{L_2} = \operatorname{argmin}_{ij} \sum |M_{ij}|^2, \text{ s.t. } O^k = \sum_{ij} e_{ij}^k M_{ij}, \text{ } k = 1, \dots, K. \quad (14)$$

This L_2 penalization is a more classical approach to the regularization of undersampled linear problems, largely because it is significantly simpler and faster to compute, since the solution \mathbf{M}^{L_2} can be obtained by solving a single system of ordinary linear equations.

Compared to the alternate projection algorithm, the MAP estimator approaches the true solution significantly faster as a function of K ; that is, more accurate estimates are produced with smaller data samples (Figs. 3-4). While in Section 2.1 we discussed that the compressed sensing and MAP estimators are in fact equivalent for small LASSO's λ , the al-

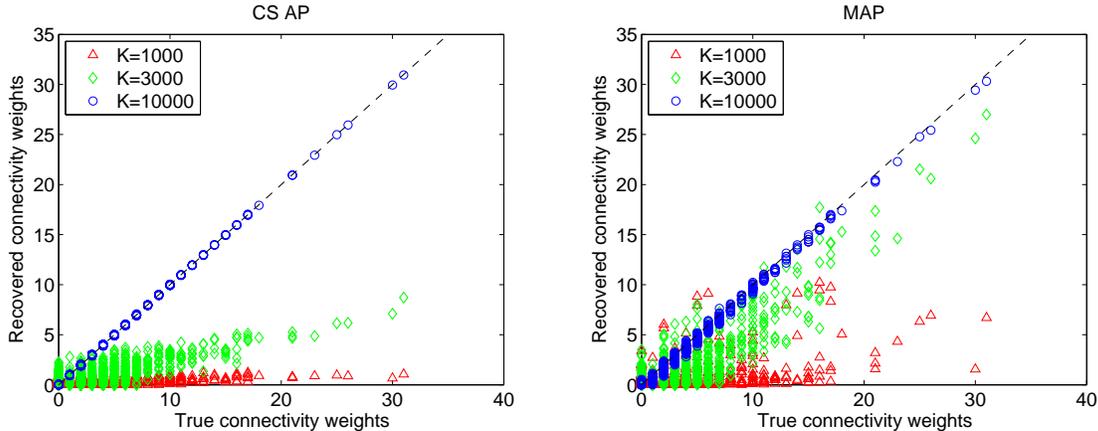


Figure 3: Results of the connectivity reconstruction experiment in *C. elegans* using surrogate GRASP data. Noiseless conditions are assumed. *C. elegans* contains approximately 300 neurons connected via 6000 synapses in 2500 distinct connections. The actual *C. elegans* wiring diagram from (White et al., 1986) was used to prepare the surrogate data. A total of K GRASP measurements in K different specimens were simulated, and MAP reconstruction of the structural connectivity matrix was subsequently performed. The scatter plots show reconstructed versus true connectivity weights for the MAP algorithm (right) and the alternating projection algorithm in (Mishchenko, 2008) (left). Dashes indicate identity line.

ternate projection algorithm is only an approximation to the compressed sensing estimator in the $K < K_{min}$ regime, as we discussed above. The L_2 -regularized estimator produces reconstructions that are incomparable to either the LASSO or the alternate projection algorithm: the L_2 -regularized estimator needed at least $6 \cdot 10^4$ data samples to produce a comparable reconstruction (data not shown).

We use Eq. (11) to perform estimation of the variance σ_t^2 from the simulated GRASP data. As can be seen from Fig. 5, the simple estimator given by Eq. (11) performs well in practice, and already with $K \approx K_{min}$ is able to estimate the variance accurately.

The computational cost of the MAP and alternating projections algorithms scales with the number of neurons N in our experiments as N^4 (Fig. 6). This is in agreement with the theoretical cost estimates in Section 2.2. Quantitatively, reconstructions of the wiring diagram in *C. elegans* with $N = 300$ could be obtained in approximately 50 hours of calculations with the MAP algorithm and five hours of calculations with the alternate projection algorithm on a 2GHz Intel dual core laptop. Therefore, to summarize, the alternate projection algorithm required substantially less computational effort than the MAP algorithm, but in the $K < K_{min}$ regime required larger data samples to achieve reconstructions of the same quality.

3.2 Robustness to noise and reconstruction stability

Next we examine the impact of elevated levels of measurement errors on the reconstructions. We find that the performance of the MAP estimator decays gracefully with noise, that is, the error in the MAP reconstructions grows proportionally to the noise variance and is not some discontinuous function that destroys the estimator instantly as $\sigma^2 > 0$. On the other hand, the proportionality constant is fairly large, and so the sensitivity to noise at this regime $K \approx K_{min}$ is quite substantial—only relative noise strengths of up to $\sigma/\langle O^k \rangle \approx 0.03 - 0.04$

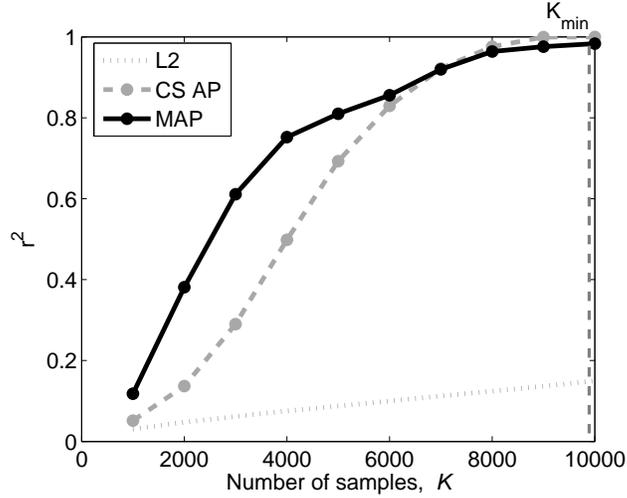


Figure 4: Results of the connectivity reconstruction experiment in *C. elegans* using surrogate GRASP data (details as in Fig. 3). The quality of the reconstruction for different sample sizes K is quantified using the correlation coefficient r^2 between the connectivity weights in the reconstructed and true connectivity matrix. Near-exact reconstruction, that is, $r^2 \approx 1$, is achieved with $K_{min} \approx 10^4$ samples, and meaningful approximate reconstructions ($r^2 \geq 0.5$) are achieved already with $K \approx 2000 - 3000$ samples, significantly less than the $N^2 \approx 10^5$ that could have been expected to be necessary to fully characterize the connectivity matrix of $N = 300$ neurons. The results of the reconstruction using a simple L_2 -regularized algorithm and alternate projection algorithm are also shown for comparison.

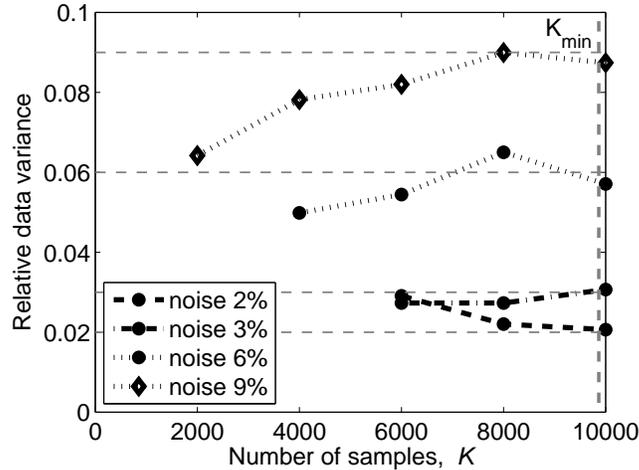


Figure 5: Results of the estimation of the biological variability in the simulated connectivity reconstruction experiment in *C. elegans*. Variability here is quantified in terms of $\sigma/\langle O^k \rangle$, used to characterize the variance in a dimension-independent way; $\sigma/\langle O^k \rangle = 0.02, 0.03, 0.06$, and 0.09 are shown. Already with $K \approx K_{min}$ samples the variability can be estimated very accurately. Simulation details are as in Fig. 3.

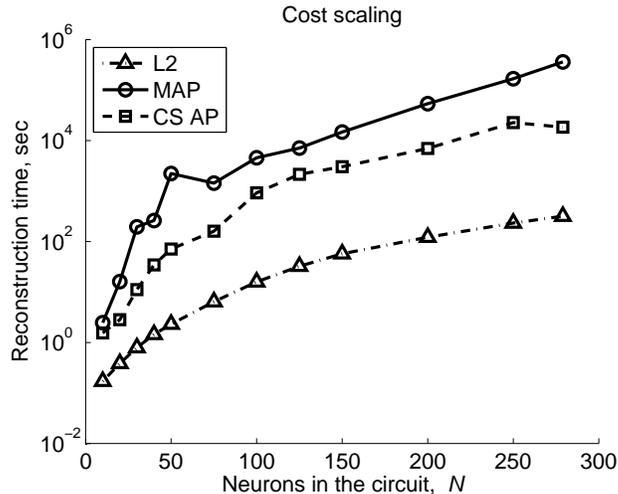


Figure 6: The computational complexity of the MAP, alternating optimization, and L_2 reconstruction algorithms for different neural circuit sizes N , assuming constant sparseness. The computational cost scales as $\approx N^{3.5-4.5}$, in good agreement with theoretical estimates in Section 2.2. Quantitatively, reconstruction of the complete wiring diagram in *C. elegans* ($N = 300$ neurons and $\mathcal{N} = 10^5$ unknown connections) could be obtained in approximately 50 hours of calculations using the MAP algorithm and five hours of calculations using the alternating projections algorithm on a standard 2Ghz Intel dual core laptop computer.

can be tolerated for this K (Fig. 7). Robustness increases with the dataset size, however, so that larger samples allow for more accurate reconstructions, even under heavier noise conditions. Compared to the alternating projection algorithm (Mishchenko, 2008), the MAP estimator performs better in the $K < K_{min}$ regime under noisy conditions as well. Similar results are obtained with respect to the sensitivity of the reconstruction algorithm to errors in the specification of the e_{ij}^k functions in different experiments, for example due to incorrect measurements of the expression patterns A^k and B^k (Fig. 8).

3.3 Increasing the efficiency of GRASP reconstructions by incorporating spatial information about the positions of the labeled puncta

Until now we have considered the situation where only one measurement could be obtained per one GRASP experiment, exactly as in (Mishchenko, 2008; Mishchenko, 2011). Indeed, as was discussed in (Mishchenko, 2008; Mishchenko, 2011), using GRASP only one expression pattern can be produced per one animal, for example, using gene fusion or the Cre-Lox genetic system. Since all synapses are labeled with the same fluorescent marker (GFP) in this setting, it was argued that only a single gross measurement — e.g., the total count of all labeled synapses — can be obtained. On the other hand, one can argue that the coarse spatial positions of synapses labeled and observed in different animals could serve as an additional source of information, allowing connectivity matrix reconstructions to be performed with improved efficiency and robustness. For example, if two neurons are known to make synaptic contacts in a certain region of the brain, then attributing the puncta observed in that region to these neurons should require less empirical evidence.

To formalize this notion, we consider spatially partitioning the neural circuit into a col-

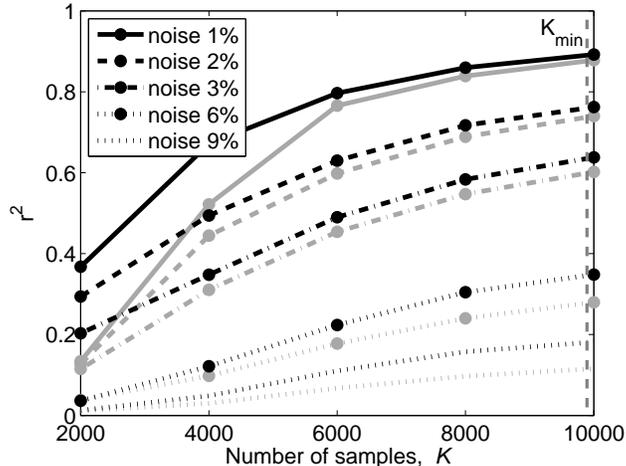


Figure 7: Sensitivity of MAP reconstructions to noise for different noise levels $\sigma/\langle O^k \rangle = 0.01, 0.02, 0.03, 0.06$, and 0.09 and sample size K . MAP reconstruction accuracy decays continuously with the additions of noise: the error grows proportionally to the amount of added noise, and is not some discontinuous function that destroys the estimator instantaneously as $\sigma > 0$, which one might have expected in the heavily undersampled regime $K \approx K_{min} \ll \mathcal{N}$. However, the proportionality constant is fairly large, so that the estimator is sensitive to noise—only relative noise strengths up to $\sigma/\langle O^k \rangle \approx 0.03 - 0.04$ can be tolerated at $K \approx K_{min}$. Robustness increases with the size of the dataset, so that larger sample sizes can be used to produce more accurate reconstructions. Similar graphs for the reconstructions using the alternating projection algorithm are shown in light gray from (Mishchenko, 2008), for comparison. Simulation details are as in Fig. 3.

lection of non-overlapping spatial regions. For example, partitions of the brain into known cytoarchitectonic parcellations (Kandel et al., 2000; Brodmann, 1994; Shepherd and Svoboda, 2005; Sato et al., 2007) can be utilized. Each region X is associated with a “local” connectivity matrix, $C_{ij}(X)$, defined to characterize the connections made by neurons in that region X . Note that here only the position of synapses in X is important for $\mathbf{C}(X)$; the corresponding cell bodies can be located anywhere in the circuit. If neurons i and j make a synaptic contact in region X , for example, then $C_{ij}(X) = 1$ even if the cell bodies are located elsewhere. On the other hand, if these neurons do not make a connection in X , then $C_{ij}(X) = 0$ even if they make a connection somewhere else.

Since the total count of synaptic contacts between any two neurons anywhere in the brain equals the sum of such counts for different X , the complete connectivity matrix is obviously related to the local connectivity matrices via summation over X :

$$C_{ij} = \sum_X C_{ij}(X). \quad (15)$$

Analogous to $C_{ij}(X)$, we can introduce the “local” puncta counts $O^k(X)$, defined to represent the counts of synapses labeled in experiment k in different regions. It is easy to see that for each X , for the local average connectivity matrix $\mathbf{M}(X) = \langle \mathbf{C}^k(X) \rangle$ we obtain an

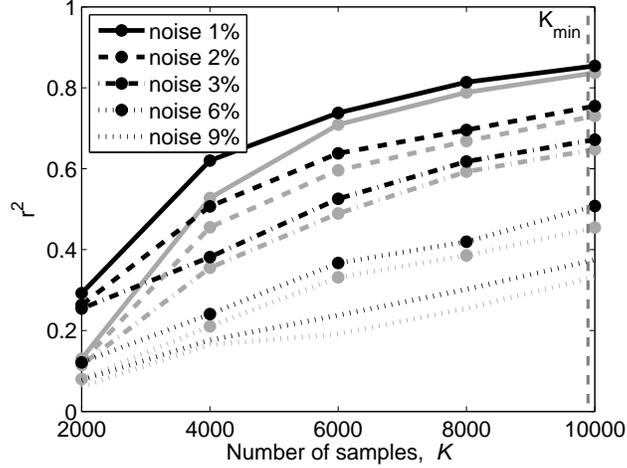


Figure 8: Sensitivity of MAP reconstructions to misspecifications in the measured expression patterns A^k and B^k , for different fractions of misspecified neurons in A^k and B^k , $e = 1\%$, 2% , 3% , 6% and 9% , and sample size K . MAP performance again decays continuously with noise, however the estimator is sensitive to noise—only up to $e \approx 5 - 6\%$ of errors can be tolerated at $K \approx K_{min}$. Robustness increases with larger dataset sizes, so that larger sample sizes can be used to obtain more accurate reconstructions. Similar plots for the reconstructions using the alternating projection algorithm are shown in light gray from (Mishchenko, 2008), for comparison. Simulation details are as in Fig. 3.

estimation problem similar to Eq. (8) in Section 2.1. That is, for all X ,

$$\hat{\mathbf{M}}^{MAP}(X) = \operatorname{argmin} \left[\sum_k \left(O^k(X) - \sum_{ij} e_{ij}^k M_{ij}(X) \right)^2 + 2\lambda \sum_{ij} |M_{ij}(X)| \right], \quad (16)$$

where

$$O^k(X) = \sum e_{ij}^k C_{ij}^k(X) + \epsilon^k(X). \quad (17)$$

Given Eq. (15), the complete $\hat{\mathbf{M}}$ can be written in terms of $\hat{\mathbf{M}}(X)$ as

$$\hat{M}_{ij} = \sum_{\mathcal{X}} \hat{M}_{ij}(X). \quad (18)$$

Observe that problems (16) for different X are completely independent and can be solved separately. In particular, to enter the compressed sensing regime, and thus to be able to produce an exact reconstruction, each sub-problem requires at least

$$K_{min}(X) \propto N_s(X) \log \mathcal{N} \quad (19)$$

measurements, where $N_s(X)$ is the total number of neural connections in the region X . If we use the spatial position information, then the number of separate experiments needed to completely reconstruct full connectivity matrix is reduced to

$$K_{min} \propto \max_X N_s(X) \log \mathcal{N}. \quad (20)$$

In other words, the number of measurements that must be taken from each spatial region X for reconstruction of the local connectivity matrix in each region $\mathbf{C}(X)$ is $N_s(X)$. However, in each experiment, the whole animal is imaged and synaptic puncta counts can thus be obtained from all different regions X . Thus, the measurements for all different X can be obtained simultaneously and in parallel. It is clear, therefore, that the total number of experiments required to reconstruct all the local matrices $\mathbf{C}(X)$ is determined by the region with the largest $N_s(X)$, so that the number of experiments required to reconstruct $\mathbf{C}(X)$ for all X , and thus the complete \mathbf{C} , is given by Eq. (20). Since $N_s(X) \approx N_s/N_X$, where $N_X = |\{X\}|$ is the total number of regions in the partition, this number of measurements will be significantly less than the $K_{min} \propto N_s \log \mathcal{N}$ required to reconstruct \mathbf{C} if no spatial data are used, as in (Mishchenko, 2008; Mishchenko, 2011).

3.4 MAP connectivity inference from surrogate data with spatial information generated using the *C. elegans* wiring diagram

We tested the above approach on a surrogate data generated using the *C. elegans* wiring diagram. We assume a partition containing 12×12 $30^\circ \times 0.1$ mm radially cylindrical segments that completely covered roughly cylindrical *C. elegans* body. This choice was somewhat arbitrary and should be considered for illustration purposes only. Of course, with fluorescent light microscopy imaging in *C. elegans* now routinely available at a single cell level (that is, $\approx 1\mu m$) (Chalasanani et al., 2007; Feinberg et al., 2008; Zimmer et al., 2009), enforcing such partitions should be entirely plausible. In fact, much finer partitions should be possible, conferring greater advantages for the reconstruction, according to our arguments above. In practice, the minimum partition width will be defined by the stereotypy of the positions of the synapses and the accuracy with which these positions can be practically resolved from the fluorescent light microscopy data.

To model the spatial distribution of synapses in *C. elegans*, Prof. Dr. David Hall of the Albert Einstein College of Medicine generously carried out a re-examination of the electron micrographs from (White et al., 1986) to estimate the density of the synapses over different longitudinal segments in *C. elegans* body. He found that the longitudinal distribution of synapses along the *C. elegans* body is roughly uniform, except for the segment of the nerve ring, where approximately 50% of synapses are concentrated. We used this model for the distribution of synapses in our numerical experiments.

The results of our experiments are shown in Fig. 9. As expected, complete reconstructions could be obtained with significantly fewer experiments if spatial information is exploited. Specifically, near-perfect reconstruction could be obtained with only $K_{min} \approx 600$ experiments, and practically reasonable reconstructions were obtained already with $K \approx 100 - 200$ experiments, a factor of 15–30 reduction as compared with (Mishchenko, 2008). The individual $\hat{M}_{ij}(X)$ can be also estimated, providing access to important information about detailed distribution of neural connections throughout the *C. elegans* body.

4 Conclusion and extensions

We have presented a statistical paradigm for structural connectivity estimation from compilations of fluorescent measurements using relatively easy to obtain high throughput fluorescent connectivity probes. Our Bayesian maximum a posteriori formulation can be reduced, using

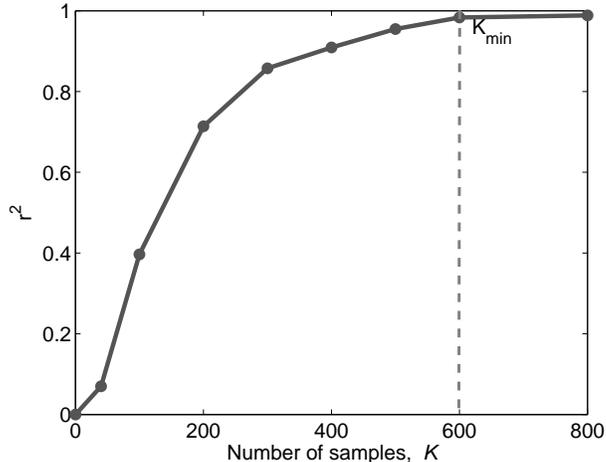


Figure 9: Results of the simulated connectivity reconstruction experiments in *C. elegans*, using GRASP with additional logging of the spatial positions of the labeled synaptic puncta. Partitioning into $144 = 12 \times 12$ radially cylindrical segments 30° wide and $0.1mm$ long is assumed, with uniform distribution of synapses throughout *C. elegans* body with the exception for the nerve ring, consistent with the available neuroanatomical data. Near-perfect reconstruction can be obtained already with $K_{min} \approx 600$ samples, and practically meaningful reconstructions can be obtained already with 100–200 experiments, a factor 15–30 reduction from (Mishchenko, 2008).

law-of-large-numbers arguments, to a surprisingly simple L_1 -regularized quadratic optimization problem. Standard numerical methods are capable of handling the large scale of this problem, after we exploit the specific factorization structure of the resulting quadratic form. We inspected the problem of connectivity matrix reconstruction in *C. elegans* using surrogate GRASP data generated from the empirically known wiring diagram. We showed that in this setting this wiring diagram can be inferred with an orders of magnitude reduction in experimental effort compared with classical neuroanatomical approaches, using fluorescent light imaging experiments in as few as ≈ 100 specimens (Fig. 9) and no highly-laborious axon and dendrite tracing. The resulting numerical reconstruction problem could be resolved on a standard laptop computer in under 50 hours of calculations. Along with the average structural connectivity matrix, we also showed how the variability in the connection strengths across multiple animals and the spatial distribution of the neural connections can be obtained using the same approach. Such information is difficult to obtain with other existing connectomics paradigms.

For clarity and concreteness, we have focused our attention on the specific case of the *C. elegans* wiring diagram and GRASP measurements (Feinberg et al., 2008). However, it is important to emphasize that the approach developed here can be extended and applied to data from a variety of different high throughput fluorescent light microscopy connectomics paradigms. The central assumption of our approach is that the measurements O^k are related to the connectivity matrix C_{ij}^k via a certain linear functional of C_{ij}^k :

$$O^k = f\left(\sum_{ij} e_{ij}^k C_{ij}^k + \epsilon^k\right), \quad (21)$$

for some coefficients e_{ij}^k and some function $f(\cdot)$. In that case, our approach allows one to accurately estimate the connectivity matrix from an ensemble of such measurements. Many current fluorescent light microscopy connectomics paradigms fall under this broad definition.

For example, in (Hu et al., 2009) a comparable approach was proposed for photoactivation imaging data. In photoactivation imaging, a preparation with caged glutamate (Callaway and Katz, 1993; Bureau et al., 2004) or photoactivated channelrhodopsin-2 (Petreanu et al., 2007) is used to enable optical excitation of different neurons or neural classes and subsequent patch clamp recording of the postsynaptic currents (Shepherd et al., 2003; Bureau et al., 2004; Petreanu et al., 2007; Nikolenko et al., 2007). When a group of neurons is excited in this setting, the resulting sub-threshold EPSP response can be described, to a first approximation, as a linear superposition of the EPSPs from individual pre-synaptic neurons. Observations in such an experiment then can be represented in the form of Eq. (21), and the connectivity extracted in a similar manner as described above.

A different high throughput fluorescent light microscopy neural connectivity mapping paradigm was recently proposed by Svoboda (2009). In this approach, a regular cytoplasmic fluorescent dye, such as genetically encoded GFP, is expressed in a sparse population of neurons and neural projections are observed in the entire brain using a high throughput fluorescent light microscopy imaging setup. Although a sparse set of neurons is visualized in any one such experiment, a complete coverage of the brain is expected from a large series of such experiments.

This connectivity inference problem also can be described in the language of Eq. (21). For concreteness, let us denote with D_v^k the density of labeled axonal projections in voxel v in experiment k . D_v^k is the sum of the axonal projection densities to voxel v from all individual neurons labeled in that experiment, whose collection we can identify with the symbol A^k ,

$$D_v^k = \sum_{j \in A^k} d_{v,j}^k. \quad (22)$$

Here, with $d_{v,j}$ we denote the projection maps of individual neurons or identified neural groups, which we can identify with the symbol j . Eq. (22) is in the form of Eq. (21), so the problem of efficient individual map extraction from this fluorescent imaging data can be approached using the same methods we discussed above.

Connectivity mapping using transsynaptic tracer viruses is another tool that has been recently proposed for reconstructions of synaptic connectivity in large neuronal networks (Kuypers and Ugolini, 1990; Wickersham et al., 2007; Bohland et al., 2009). In this setting, one uses a virus that can spread by crossing the synaptic junctions of synaptically connected neurons, for example, the Herpes simplex virus and rabies (Kuypers and Ugolini, 1990). The genetic code of such viruses is altered to make infected neurons also produce a fluorescent protein such as GFP (Heim et al., 1995). This allows infected neurons to be observed with fluorescent light microscopy. By modifying the virus further to ensure that the infection can cross only one synapse, only the neurons which are synaptically connected to the originally-infected cells can produce GFP (Wickersham et al., 2007). This makes it possible to observe synaptic connections in a neural circuit by using fluorescent light microscopy. It has been subsequently proposed that, by introducing such a virus into different populations of neurons in a series of experiments and imaging the populations of secondarily infected neurons, it may be possible to produce connectivity reconstruction for complete large neural circuits (Bohland et al., 2009).

We can cast this connectivity inference problem again in the language of Eq. (21). For the sake of clarity, the subsequent discussion assumes the retrograde propagation of the virus, that is, the virus infects neurons in the post- to pre-synaptic direction (Wickersham et al., 2007). In that case, denote the set of “primarily” infected neurons in experiment k with the symbol A^k . The set of secondarily infected neurons then will be denoted with B^k . A^k and B^k are obviously related to each other via the synaptic connectivity matrix C_{ij} . Specifically,

$$b_j^k = H\left(\sum_{i \in A^k} C_{ij}^k\right), \quad (23)$$

where b_j^k is the indicator vector for B^k and $H(\cdot)$ is the Heaviside step function, that is, $H(x) = 1$ if $x > 0$ and $H(x) = 0$ otherwise. The meaning of Eq. (23) is simply that a secondary neuron j becomes infected if and only if there is a synaptic connection to that neuron from any one of the primarily infected neurons, A^k .

Several interesting features of the resulting inference problem are worth noting. First, the connectivity matrix C_{ij} in the simplest case quantifies the unsigned counts of synapses, and therefore the nonnegativity constraint $C_{ij} \geq 0$ applies here. Second, since Eq. (23) is unchanged if we rescale $C_{ij} \rightarrow \eta_{ij} C_{ij}$, $\eta_{ij} > 0$, it is clear that we can only hope to reconstruct the presence or absence of a synaptic connection between a pair of neurons in this setting. Performing this inference reduces to a linear program (i.e., optimizing a linear objective function under a set of linear constraints), since Eq. (23) can be rewritten as a linear inequality: the right-hand side is positive if the left-hand side equals one, and nonpositive otherwise. Introducing a noise model in Eq. (23) — e.g., allowing the probability $P(b_j^k > 0)$ to depend on the magnitude (not just the sign) of the right-hand side — would lead to a more standard objective function which can be optimized using interior-point methods similar to those introduced here.

5 Acknowledgements

The authors are grateful to Prof. David Hall for his re-examination of the original print data from (White et al., 1986), and for providing the statistics about the spatial distribution of synapses in *C. elegans* used in section 3.4. This work was supported by an NSF CAREER grant, a McKnight Scholar award, and by NSF grant IIS-0904353.

References

- Amaral, D. G. and Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience*, 31(3):571–91.
- Bock, D. D., Lee, W. C., Kerlin, A. M., Andermann, M. L., Hood, G., Wetzell, 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*, 471:177–82.
- Bohland, J., Wu, C., Barbas, H., Hermant, B., Bota, M., Breiter, H., Cline, H., Doyle, J., Freed, P., Greenspan, R., Haber, S., Hawrylycz, M., Herrera, D., Hilgetag, C., Huang, Z., Jones, A., Jones, E., Karten, H., Kleinfeld, D., Kotter, R., Lester, H., Lin, J., Mensh, B., Mikula, S., Panksepp, J., Price, J., Safdieh, J., Saper, C., Schiff, N., Schmahmann,

- J., Stillman, B., Svoboda, K., Swanson, L., Toga, A., Essen, D., Watson, J., and Mitra, P. (2009). A proposal for a coordinated effort for the determination of brainwide neuroanatomical connectivity in model organisms at a mesoscopic scale. *PLoS Comp. Biol.*, 5:e1000334.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3:1–122.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Oxford University Press.
- Braitenberg, V. and Schuz, A. (1998). *Cortex: statistics and geometry of neuronal connectivity*. Springer, Berlin.
- Briggman, K. L. and Denk, W. (2006). Towards neural circuit reconstruction with volume electron microscopy techniques. *Current Opinions in Neurobiology*, 16:562.
- Briggman, K. L., Helmstaedter, M., and Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, 471:183–8.
- Brodmann, K. (1994). *Brodmann’s ‘Localisation in the Cerebral Cortex’*. Smith-Gordon, London, UK.
- Broome, B. M., Jayaraman, V., and Laurent, G. (2006). Encoding and decoding of overlapping odor sequences. *Neuron*, 51:467–482.
- Bureau, I., Shepherd, G. M., and Svoboda, K. (2004). Precise development of functional and anatomical columns in the neocortex. *Neuron*, 42(5):789–801.
- Burton, B. (1999). *Automated 3D reconstruction of neuronal structures from serial sections*. PhD thesis, Texas A&M University.
- Cajal, R. (1911). *Histologie du systeme nerveux de l’homme et des vertebres*, volume 2. Madrid, Instituto Ramon y Cajal.
- Callaway, E. and Katz, L. (1993). Photostimulation using caged glutamate reveals functional circuitry in living brain slices. *PNAS*, 90:7661–5.
- Candes, E. J. and Romberg, J. (2005). Practical signal recovery from random projections. *Wavelet Applications in Signal and Image Processing XI, Proc.*, page 5914.
- Candes, E. J., Romberg, J., and Tao, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509.
- Candes, E. J. and Tao, T. (2010). The power of convex relaxation: near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56:2053–2080.
- Candes, E. J. and Wakin, M. (2008). An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30.

- Chalasan, S. H., Chronis, N., Tsunozaki, M., Gray, J. M., Ramot, D., Goodman, M. B., and Bargmann, C. I. (2007). Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*. *Nature*, 450:35.
- Chiang, A.-S., Lin, C.-Y., Chuang, C.-C., Chang, H.-M., Hsieh, C.-H., Yeh, C.-W., Shih, C.-T., Wu, J.-J., Wang, G.-T., Chen, Y.-C., Wu, C.-C., Chen, G.-Y., Ching, Y.-T., Lee, P.-C., Lin, C.-Y., Lin, H.-H., Wu, C.-C., Hsu, H.-W., Huang, Y.-A., Chen, J.-Y., Chiang, H.-J., Lu, C.-F., Ni, R.-F., Yeh, C.-Y., and Hwang, J.-K. (2011). Three-dimensional reconstruction of brain-wide wiring networks in *Drosophila* at single-cell resolution. *Current biology*, 21:1–11.
- Denk, W. and Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLOS Biol.*, 2:e329.
- Donoho, D. (2006). Compressed sensing. *IEEE Trans. on Information Theory*, 52:1289–306.
- Donoho, D. and Elad, M. (2003). Optimally sparse representation in general (nonorthogonal) dictionaries via L^1 minimization. *PNAS*, 100:2197–2202.
- Feinberg, E. H., Vanhoven, M. K., Bendesky, A., Wang, G., Fetter, R. D., Shen, K., and Bargmann, C. I. (2008). GFP Reconstitution Across Synaptic Partners (GRASP) defines cell contacts and synapses in living nervous systems. *Neuron*, 57:353–63.
- Goldfarb, D., Ma, S., and Scheinberg, K. (2010). Fast alternating linearization methods for minimizing the sum of two convex functions. *Columbia University IEOR Technical Report*.
- Gray, J. M., Hill, J. J., and Bargmann, C. I. (2005). A circuit for navigation in *Caenorhabditis elegans*. *PNAS*, 102:3184.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, V., and Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):e159.
- Hagmann, P., Kurrant, M., Gigandet, X., Thiran, P., Wedeen, V., Meuli, R., and Thiran, J.-P. (2007). Mapping human whole-brain structural networks with Diffusion MRI. *PLoS ONE*, 2(7):e597.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer.
- Heim, R., Cubitt, A., and Tsien, R. (1995). Improved green fluorescence. *Nature*, 373:663–4.
- Helmstaedter, M., Briggman, K. L., and Denk, W. (2008). 3D structural imaging of the brain with photons and electrons. *Current Opinions in Neurobiology*, 18:633–41.
- Helmstaedter, M., Briggman, K. L., and Denk, W. (2011). High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nature Neuroscience*, 14:1081–1088.
- Hu, T., Leonardo, A., and Chklovskii, D. (2009). Reconstruction of Sparse Circuits Using Multi-neuronal Excitation (RESCUME). In *NIPS 2009, Poster*.
- Kandel, E. R., Schwartz, J. H., and Jessel, T. M. (2000). *Principles of Neural Science*. McGraw-Hill Professional.

- Kutner, M., Nachtsheim, C., Neter, J., and Li, W. (2005). *Applied Linear Statistical Models*. McGraw-Hill.
- Kuypers, H. and Ugolini, G. (1990). Viruses as transneuronal tracers. *Trends In Neurosciences*, 13:71–5.
- Lichtman, J. W., Livet, J., and Sanes, J. R. (2008). A technicolour approach to the connectome. *Nature Reviews Neuroscience*, 9:417–422.
- Livet, J., Weissman, T., Kang, H., Draft, R., Lu, J., Bennis, R., Sanes, J., and Lichtman, J. (2007). Transgenic strategies for combinatorial expression of fluorescent proteins in the nervous system. *Nature*, 450:56–62.
- Lu, J., Fiala, J. C., and Lichtman, J. W. (2009). Semi-automated reconstruction of neural processes from large numbers of fluorescence images. *PLoS ONE*, 4(5):e5655.
- Luo, L., Callaway, E. M., and Svoboda, K. (2008). Genetic dissection of neural circuits. *Neuron*, 57:634–60.
- Micheva, K. D. and Smith, S. J. (2007). Array tomography: a new tool for imaging the molecular architecture and ultrastructure of neural circuits. *Neuron*, 55(1):25–36.
- Mishchenko, Y. (2008). Strategies for identifying exact structure of neural circuits with broad light microscopy connectivity probes. *Nature Precedings Preprint*: <http://precedings.nature.com/documents/2669>.
- Mishchenko, Y. (2009). Automation of 3D reconstruction of neural tissue from large volume of conventional serial section transmission electron micrographs. *Journal of Neuroscience Methods*, 176:276–289.
- Mishchenko, Y. (2010). On optical detection of densely labeled synapses in neuropil and mapping connectivity with combinatorially multiplexed fluorescent synaptic markers. *PLOS ONE*, 5:e8853.
- Mishchenko, Y. (2011). Reconstruction of complete connectivity matrix for connectomics by sampling neural connectivity with fluorescent synaptic markers. *Journal of Neuroscience Methods*, 196:289–302.
- 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*, 67:1009–20.
- Mishchenko, Y., Vogelstein, J., and Paninski, L. (2011). A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics*, 5:1229–61.
- Ng, A. (2004). Feature selection, L_1 vs. L_2 regularization, and rotational invariance. *ICML*, 21.
- Nikolenko, V., Poskanzer, K., and Yuste, R. (2007). Two-photon photostimulation and imaging of neural circuits. *Nature Methods*, 4:943–950.

- Park, T. and Casella, G. (2008). The Bayesian Lasso. *Journal of the American Statistical Association*, 103:681–6.
- Petreaanu, L., Huber, D., Sobczyk, A., and Svoboda, K. (2007). Channelrhodopsin-2-assisted circuit mapping of long-range callosal projections. *Nature Neuroscience*, 10:663–8.
- Ramon y Cajal, S. (1904). *La Textura del Sistema Nerviosa del Hombre y los Vertebrados*. Moya.
- Romberg, J. (2008). Imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):14–20.
- Sato, T. R., Gray, N. W., Mainen, Z. F., and Svoboda, K. (2007). The functional microarchitecture of the mouse barrel cortex. *PLoS Biol*, 5(7):e189.
- Shepherd, G. M. and Svoboda, K. (2005). Laminar and columnar organization of ascending excitatory projections to layer 2/3 pyramidal neurons in rat barrel cortex. *Journal of Neuroscience*, 25:5670–9.
- Shepherd, G. M. G., Polgruto, T. A., and Svoboda, K. (2003). Circuit analysis of experience-dependent plasticity in the developing rat barrel cortex. *Neuron*, 38:277.
- Smith, S. J. (2007). Circuit reconstruction tools today. *Curr Opin Neurobiol*, 17(5):601–8. Journal Article Review England.
- Svoboda, K. (2007). Cajal 2.0 optical mapping project for mouse brain. In *Janelia Farm Meeting on Neural Circuit Reconstruction, HHMI, September 23-26, 2007*.
- Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B*, 58:267–288.
- Tipping, M. (2001). Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244.
- Vanderbei, R. J. (2001). *Linear Programming: Foundations and Extensions*. International Series in Operations Research and Management Science. Springer-Verlag, New York.
- Westrum, L. E. and Blackstad, T. W. (1962). An electron microscopic study of the stratum radiatum of the rat hippocampus (regio superior, CA 1) with particular emphasis on synaptology. *J Comp Neurol*, 119:281–309.
- White, J., Southgate, E., Thomson, J. N., and Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philosophical Transactions of Royal Society London. Series B, Biological Sciences*, 314(1165):1–340.
- Wickersham, I., Finke, S., Conzelmann, K.-K., and Callaway, E. (2007). Retrograde neuronal tracing with a deletion-mutant rabies virus. *Nature Methods*, 4:47–49.
- Zhang, S., Ma, C., and Chalfie, M. (2004). Combinatorial marking of cells and organelles with reconstituted fluorescent proteins. *Cell*, 119:137–44.

- Zimmer, M., Gray, J. M., Pokala, N., Chang, A. J., Karow, D. S., Marletta, M. A., Hudson, M. L., Morton, D. B., Chronis, N., and Bargmann, C. I. (2009). Neurons detect increases and decreases in oxygen levels using distinct guanylate cyclases. *Neuron*, 61:865–79.
- Zou, H., Hastie, T., and Tibshirani, R. (2007). On the “degrees of freedom” of the Lasso. *The Annals of Statistics*, 35(5):pp. 2173–2192.