ABSTRACT

Title of dissertation: SPARSE SIGNAL REPRESENTATION IN DIGITAL AND BIOLOGICAL SYSTEMS

Author: Matthew Guay, Doctor of Philosophy, 2016

Dissertation directed by: Professor Wojciech Czaja
Department of Mathematics

Theories of sparse signal representation, wherein a signal is decomposed as the sum of a small number of constituent elements, play increasing roles in both mathematical signal processing and neuroscience. This happens despite the differences between signal models in the two domains. After reviewing preliminary material on sparse signal models, I use work on compressed sensing for the electron tomography of biological structures as a target for exploring the efficacy of sparse signal reconstruction in a challenging application domain. My research in this area addresses a topic of keen interest to the biological microscopy community, and has resulted in the development of tomographic reconstruction software which is competitive with the state of the art in its field. Moving from the linear signal domain into the nonlinear dynamics of neural encoding, I explain the sparse coding hypothesis in neuroscience and its relationship with olfaction in locusts. I implement a numerical ODE model of the activity of neural populations responsible for sparse odor coding in locusts as part of a project involving offset spiking in the Kenyon cells. I also explain the validation procedures we have devised to help assess the model's simi-

larity to the biology. The thesis concludes with the development of a new, simplified model of locust olfactory network activity, which seeks with some success to explain statistical properties of the sparse coding processes carried out in the network.

SPARSE SIGNAL REPRESENTATION IN DIGITAL AND BIOLOGICAL SYSTEMS

by

Matthew Guay

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2016

Advisory Committee:
Professor Wojciech Czaja, Chair
Professor Radu Balan
Professor Alexander Barg
Professor John J. Benedetto
Doctor Richard Leapman
Professor Kasso Okoudjou
Doctor Mark Stopfer

© Copyright by
Matthew Guay
2016

Dedication

Dedicated to my grandmother Virginia Safford, *in memoriam*.

Acknowledgments

Thank you, mom, dad, and Jessy, for your love and support. The same for my extended family, too numerous to name. You are why I am where I am. A special thanks to my grandfather Albert Guay, for your support and inspiration over the past quarter century.

I am grateful to all of my mentors, especially Prof. Wojtek Czaja, Dr. Richard Leapman, and Dr. Mark Stopfer. Each of you helped me to unearth a new facet of myself, and to polish it.

Thank you to all of my friends, for helping me to grow as a person, not just an academic. Among them, a special thanks to all of my housemates past and present. You have kept me connected and happy throughout these past five years.

Last but not least, thank you to all of my colleagues at UMD, and collaborators at NIH. You are all responsible for the research environment which I have enjoyed so deeply during graduate school.

Table of Contents

List of Abbreviations

A	Ampere (unit of electric current)
AL	Antennal Lobe
BF	Bright Field
BPDN	Basis Pursuit Denoising
CS	Compressed Sensing
DF	Dark Field
ET	Electron Tomography
F	Farad (unit of electric capacitance)
FFT	Fast Fourier Transform
GGN	Giant GABAergic Neuron
IPI	Inter-Pulse Interval
KC	Kenyon Cell
LHS	Left-Hand Side
LLE	Locally-Linear Embedding
LN	Local Neuron
MB	Mushroom Body
NLDR	Nonlinear Dimension Reduction
NUFFT	Non-Uniform Fast Fourier Transform
ODE	Ordinary Differential Equation
OR	Olfactory Receptor
ORN	Olfactory Receptor Neuron
PN	Projection Neuron
PSTH	Peristimulus Time Histogram
RHS	Right-Hand Side
RIP	Restricted Isometry Property
RK4	Four-step Runge-Kutta
S	Siemen (unit of electric conductance)
V	Volt (unit of electric potential)

Chapter 1: Introduction

The theme of sparse signal representation has attracted broad interest over the past few decades, as researchers in information processing fields realize the importance and the utility of describing complex patterns in terms of a small number of meaningful components. Sparse signal methods demonstrate how proper choices of measurement, exploiting the statistical structure of signal ensembles, force signal representations to lie along low-dimensional manifolds within a larger measurement state space. Sparse representations are a useful intermediate step within larger signal processing pipelines. Collections of these representations will exhibit less overlap than dense representations, useful for categorization and classification. Sparse decompositions are a useful component of both digital and biological feature extraction systems, and can be used to improve a variety of other signal processing techniques such as image denoising or deconvolution.

The linear signal models predominantly used in digital signal processing are a primary focal point for the theory and practice of sparse representation, but the ideas are more broadly applicable as well. Abandoning linearity loses a great deal of the theoretical machinery commonly associated with sparse signal representation, but the phenomenon plays an important role in more complex systems such as neural

sensory processing. The importance of better understanding neural activity is itself reason enough to seek to understand a broader notion of signal sparsity, and this application is a key motivation for the work in this thesis.

In Chapter 2, I introduce the mathematical background required to understand sparse signal representation phenomena in the following chapters - compressed sensing and signal sparsity for Chapter 3, dynamical systems and the analysis of state space trajectories for Chapter 4. Chapter 2 contains numerical background as well, detailing the core algorithms that support the computational work done in this thesis. An important goal of my work over the past several years has been the implementation of efficient numerical techniques to quickly solve problems, and the fundamental algorithm design choices detailed in Sections 2.2.5 and 2.3.2 play an important role in that.

Chapter 3 details the results to date of an ongoing collaboration with Dr. Richard Leapman's lab within the National Institute for Biomedical Imaging and Bioengineering (NIBIB), an institute of the NIH. Working with this lab and its electron tomography data, we sought to investigate the effectiveness of compressed sensing for the 3D imaging of biological structures. We developed a software package to run compressed sensing reconstructions from this data, and applied it to several simulated and experimental datasets. We clarified the ways in which deterministic tomographic sampling is incompatible with the randomized measurement theories which are crucial to theoretical guarantees of signal recovery via compressed sensing, and we analyzed how the statistics of images of cellular structures deviate from the sparsity assumptions underlying compressed sensing reconstruction. Ultimately, we

addressed an important topic of interest within the electron microscopy community, and produced a numerical technique which matches or exceeds the state of the art in the field. The trifecta of scientific, mathematical, and computational knowledge required to develop this project is characteristic of the broad interdisciplinary skill set I have sought to cultivate in my time as a graduate student researcher.

My work with Dr. Mark Stopfer's lab within the National Institute for Child Health and Human Development (NICHD) sought to combine these skills for a different goal. This lab studies the fundamental mechanisms underlying sensory processing in insects, and one aspect of this sensory processing is a sparse neural code which is employed within the olfactory processing hierarchy. The Hodgkin-Huxley-type ODE dynamical systems, which describe the portions of the olfactory system that we studied, operate according to principles which differ radically from the digital signal processing models employed for CS-ET in Chapter 3. An important goal of my collaboration with this lab was to gain experience with the biological system being studied, in order to appreciate the subtleties of its operation which are responsible for important network-level activity features. To this end, the bulk of my research involved implementing a numerical ODE model for neurons and synapses in locust olfactory networks, and working with experimentalists to devise procedures to ensure the computational model's agreement with the biology. This was part of a project researching the offset spiking phenomenon in locusts. Participating in this project afforded me the opportunity to develop a new numerical protocol for investigating the network behavior associated with Kenyon cell offset spiking, and it was also important for understanding the temporal aspects of Kenyon cell activity

sparsity which are one focus for Chapter 4. My work in Chapter 4 concludes with the creation and analysis of a simplified network population model of the locust projection neurons and Kenyon cells, which demonstrates that important properties of the activity of these neurons may be accounted for by simple features of the neurons and their synapses.

Chapter 2: Preliminary Material

2.1 Overview

This thesis concerns itself with the forms which sparse signal representations take in vastly different systems - image processing on a digital computer using linear algebra and convex optimization in Chapter 3, and neural sensory processing in the olfactory system of locusts modeled as the evolution of a dynamical system in Chapter 4. To understand the commonality between these systems, it is necessary to define what is meant by "sparsity", "signal", and "representation". Therefore, the primary goal of this chapter is to define these terms in the contexts encountered in Chapters 3 and 4. For Chapter 3, this requires an explanation of linear signal models and the theory surrounding the development of compressed sensing. For the dynamical systems in Chapter 4, signals in network populations are nonlinear functions of their inputs. For systems of this form, sparsity amounts to a constraint on the number of neurons active at each time, and information transfer must be described without the machinery of linear transforms.

Additionally, this chapter explains the algorithms which underlie the numerical methods used in this thesis. The convex optimization routine used in Chapter 3 makes use of the split Bregman algorithm, a technique well-suited for the ℓ^1-

penalized least-squares problems associated with compressed sensing. In Chapter 4, the simulation of a system of ODEs is carried out using a 4-step Runge-Kutta method, a standard numerical ODE algorithm.

2.2 Compressed sensing and linear signal representation

The field of **compressed sensing** (CS) gained prominence starting around 2006, with seminal works by Candès and Tao [23], [22] indicating that sparse vectors could be recovered with high probability from a number of measurements far smaller than the bound prescribed by the classic Shannon sampling theorem [95]. Introductions to the general theory of CS can be found in e.g. [11], [35]. In this section, we highlight the aspects most relevant to the analysis in this thesis.

2.2.1 The linear signal model

A (finite) signal modelled as a vector $\boldsymbol{x} \in \mathbb{R}^M$ consisting of M data points may be decomposed as a linear combination [102] of other, simpler vectors using the tools of linear algebra. A basis decomposes $\boldsymbol{x}$ into a linear combination of M linearly-independent [102] vectors $\{\boldsymbol{\psi}_i\}_{i=1}^M \subseteq \mathbb{R}^M$. A frame generalizes the idea of a basis, and decomposes $\boldsymbol{x}$ into a linear combination of $N \geq M$ vectors $\{\boldsymbol{\psi}_i\}_{i=1}^N$ whose members span [102] $\mathbb{R}^M$ but cannot be linearly independent.

Definition 1. *A finite* ***frame*** *for a vector space* $\mathbb{R}^M$ *is a set of* N *vectors* $\{\boldsymbol{\psi}_i\}_{i=1}^M$ *which span* $\mathbb{R}^N$.

The elements of this frame can be stacked as rows of a change-of-basis matrix

Ψ.

Definition 2. *A **change-of-basis matrix** for a frame $\{\boldsymbol{\psi}_i\}_{i=1}^N$ for $\mathbb{R}^M$ is a matrix $\Psi \in \mathbb{R}^{N\times M}$ defined by*

$$\Psi = \begin{bmatrix} \boldsymbol{\psi}_1 & \cdots & \boldsymbol{\psi}_N \end{bmatrix}^T, \tag{2.1}$$

each frame element forming a row of Ψ,

and it is notationally convenient to identify the frame $\{\boldsymbol{\psi}_i\}_{i=1}^N$ with Ψ. This matrix can be used to compactly write the representation of a signal $\boldsymbol{x}$ with regards to the frame Ψ, which consists of the inner products [102] of $\boldsymbol{x}$ with each vector in the frame.

Definition 3. *A **representation** of a vector $\boldsymbol{x} \in \mathbb{R}^M$ with respect to a frame Ψ is the vector $\boldsymbol{c} = \Psi\boldsymbol{x} \in \mathbb{R}^N$. Note that $c_i = \langle \boldsymbol{x}, \boldsymbol{\psi}_i \rangle$.*

Inner products also form the basis for the measurement model used in this thesis.

Definition 4. *A **measurement** $y \in \mathbb{R}$ of a signal $\boldsymbol{x} \in \mathbb{R}^M$ is the inner product of a **measurement vector** $\boldsymbol{\varphi}$ with $\boldsymbol{x}$:*

$$y = \langle \boldsymbol{x}, \boldsymbol{\varphi} \rangle . \tag{2.2}$$

A vector of m measurements $\boldsymbol{y} \in \mathbb{R}^m$ can be produced by forming the inner product of $\boldsymbol{x}$ with a collection of m measurement vectors $\{\boldsymbol{\varphi}_i\}_{i=1}^M$. The measurement process may be written succinctly using the change-of-basis matrix Φ for these

measurement vectors:

$$\boldsymbol{y} = \Phi \boldsymbol{x} \tag{2.3}$$

2.2.2 Signal sparsity

The role of sparsity in CS theory is crucial to understanding CS-ET in Chapter 3 and gauging its limitations as a reconstruction method for biological data.

Definition 5. *A vector* $\boldsymbol{x} \in \mathbb{R}^N$ *is* ***k-sparse*** *if only* k *components of* $\boldsymbol{x}$ *are nonzero.*

The sparsity of a vector $\boldsymbol{x}$ is quantified by its number of nonzero entries, called the ℓ^0 "norm"

$$||\boldsymbol{x}||_0 = \#\{i : x_i \neq 0\}.$$

When contrasted with more general notions of signal sparsity, this value is also called *identity* sparsity.

Alternately, $\boldsymbol{x}$ may be sparse in another basis or frame.

Definition 6. *A vector* $\boldsymbol{x}$ *is* ***k-sparse in frame*** $\boldsymbol{\Psi}$ *if* $||\Psi \boldsymbol{x}||_0 \leq k$.

Under this condition, the mapping $\boldsymbol{x} \to \Psi \boldsymbol{x}$ is also referred to as a *sparsifying transform* for $\boldsymbol{x}$. For STEM tomography and other imaging fields, commonly-used constructions include wavelet bases and discrete cosine bases, though others exist. In this chapter, we use Daubechies' DB8 wavelet basis [34] as a putative sparsifying transform for our datasets.

One may attempt to push this approach further, and assume that $\Psi \boldsymbol{x}$ be sparse for certain nonlinear transformations Ψ. This is the case of interest for this chapter,

in which we assume a **TV sparsity** signal model. TV is the 2D total variation operator, a discrete analogue of the magnitude of the 2D gradient. The continuous gradient operator is approximated in the discrete setting by *forward differences*

$$\Delta_x^+ \boldsymbol{x}(i,j) = \boldsymbol{x}(i+1,j) - \boldsymbol{x}(i,j) \tag{2.4}$$

$$\Delta_y^+ \boldsymbol{x}(i,j) = \boldsymbol{x}(i,j+1) - \boldsymbol{x}(i,j), \tag{2.5}$$

where we have indexed $\boldsymbol{x}$ with two spatial coordinates for notational simplicity. A vector $\boldsymbol{x} \in \mathbb{R}^N$ may be visually interpreted as an $m \times n$ array for any $m, n \in \mathbb{N}$ such that $mn = N$, and vice-versa. Given this, we can define the isotropic total variation operator [49]:

Definition 7. *For a signal $\boldsymbol{x}$ interpreted as an $m \times n$ 2D array, the **isotropic total variation operator** $TV\boldsymbol{x}$ is defined for each $(i,j) \in [1,m] \times [1,n]$ as*

$$TV\boldsymbol{x}(i,j) = \sqrt{\left(\Delta_x^+ \boldsymbol{x}(i,j)^2 + \Delta_y^+ \boldsymbol{x}(i,j)^2\right)}, \tag{2.6}$$

where Δ_x^+ and Δ_y^+ are the forward spatial difference operators defined in Equations (2.4) and (2.5).

Letting $\Psi \boldsymbol{f} = TV\boldsymbol{f}$, corresponds to an assumption that a signal has a 2D spatial organization and its gradient is sparse, which is satisfied by signals whose dominant features are piecewise-constant regions separated by sharp boundaries. This model is chosen in Chapter 3 as representative of tomograms dominated by membranous features such as cell boundaries and organelle structures.

In application, the term *sparse* is often used to refer to vectors for which k/N is "small enough". The threshold for this qualification is vague and varies between application domains. Given a signal $\boldsymbol{x}^*$ to be recovered from measurements, a **sparse signal model** is an *a priori* assumption that $\Psi\boldsymbol{x}^*$ is sparse from some sufficiently small k and appropriate transform Ψ. In this chapter we investigate the use of identity, wavelet, and TV sparsity models.

For applications in which data is contaminated by noise or other sources of error, exact sparsity may not be feasible, but sparse approximations may suffice. To gauge this possibility, one may measure a signal's compressibility instead of sparsity.

Definition 8. *For a fixed $\epsilon > 0$, the **ϵ-compressibility** of $\boldsymbol{x}$, equal to the number of components of $\boldsymbol{x}$ with magnitude greater than ϵ.*

Equivalently, one may quantify compressibility in a relative fashion using compressibility ratios.

Definition 9. *The **p% compressibility ratio** of a signal $\boldsymbol{x}$ equals the proportion of components of $\boldsymbol{x}$ whose magnitude is at least p% of the largest magnitude among the components of $\boldsymbol{x}$.*

These more flexible notions of sparsity are useful in application, when noise and approximation error are inevitable. When signals are compressible, the sparse approximations reconstructed via CS have low error, and may be appropriate for imaging needs.

2.2.3 Sparse vector reconstruction

2.2.3.1 Mutual coherence

Given a collection of measurements $\{\boldsymbol{\varphi}_i\}_{i=1}^m$ and a basis for a sparsifying transform $\{\boldsymbol{\psi}_j\}_{j=1}^N$, the theoretical results of compressed sensing focus on the recovery of $\boldsymbol{x}$ from $\Phi\boldsymbol{x}$ with underdetermined measurements ($m << N$). Concrete results are possible when $\Psi\boldsymbol{x}$ is sparse and Φ and Ψ are *incoherent*.

Definition 10. *Given an orthogonal measurement matrix* Φ, $||\varphi_i||_2 = \sqrt{N}$ *for all* $i \in [1, m]$, *and an orthonormal sparsity basis* Ψ, *the* ***mutual coherence*** *of* Φ *and* Ψ *is defined as*

$$\mu(\Phi, \Psi) = \max_{i,j} |\langle \boldsymbol{\varphi}_i, \boldsymbol{\psi}_j \rangle|. \tag{2.7}$$

The mutual coherence of a measurement and representation system can be understood as a characterization of how "spread out" the measurement vectors are in the sparsity basis. When this value is small, CS theory guarantees recovery of sparse images: If $\boldsymbol{x}$ is k-sparse in Ψ, it may be recovered from m measurements where

$$m \geq C \cdot k \cdot \mu^2(\Phi, \Psi) \cdot \log N \tag{2.8}$$

for some (small) constant C [20], via a nonlinear convex optimization problem introduced in the next section. Proving that a CS recovery will work for suitably small m requires the use of representation and measurement systems with low mutual coherence - as we will see in Chapter 3, this assumption is often not met for real-world

measurement systems.

2.2.4 Compressed sensing and ℓ^1-regularized least-squares recovery

Under ideal conditions, a signal $\boldsymbol{x} \in \mathbb{R}^M$ can be recovered exactly from measurements by M linearly-independent measurement vectors - under these conditions, Φ is an invertible matrix, and $\boldsymbol{x} = \Phi^{-1}\boldsymbol{y}$. If there is any uncertainty or noise in the obtained measurements, a *least-squares* regression technique may be used instead [8], recovering an approximation $\boldsymbol{x}^*$ of $\boldsymbol{x}$ as

$$\boldsymbol{x}^* = \underset{\boldsymbol{z}\in\mathbb{R}^M}{\operatorname{argmin}} ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2. \tag{2.9}$$

When the number of measurements is smaller than M, Equation (2.9) will not have a unique solution, and in general it will not be possible to recover $\boldsymbol{x}$. However, when *a priori* information about the structure of $\boldsymbol{x}$ is available, it is possible to add cost functions to the RHS of Equation (2.9) which penalize improper reconstructions and can ensure better reconstruction accuracy from limited measurement data. These *regularized* least-squares equations take the form

$$\boldsymbol{x}^* = \underset{\boldsymbol{z}\in\mathbb{R}^M}{\operatorname{argmin}} ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2 + \lambda C(\boldsymbol{z}), \tag{2.10}$$

for an appropriate *cost function* $C(\boldsymbol{z})$ and *regularization hyperparameter* λ. The study of these equations is deep and wide, and an introduction to their general study may be found in [84]. Multiple penalization terms can be used concurrently,

so that for some $n \geq 1$ and $i \in [1, n]$, cost functions $C_i(\boldsymbol{z})$ and regularization hyperparameters λ_i can be used to formulate the minimization problem

$$\boldsymbol{x}^* = \underset{\boldsymbol{z} \in \mathbb{R}^M}{\operatorname{argmin}} \, ||\Phi \boldsymbol{z} - \boldsymbol{y}||_2^2 + \sum_{i=1}^{n} \lambda_i C_i(\boldsymbol{x}). \tag{2.11}$$

For CS, one wishes to use the prior knowledge that $\Psi \boldsymbol{x}$ is sparse to leverage reconstructions of $\boldsymbol{x}$ from measurements $\Phi \boldsymbol{x}$ where $m < N$. An ℓ^0 norm term can be added to Equation (2.9) to penalize non-sparse reconstructions, resulting in

$$\boldsymbol{x}^* = \underset{\boldsymbol{z} \in \mathbb{R}^M}{\operatorname{argmin}} \, ||\Phi \boldsymbol{z} - \boldsymbol{y}||_2^2 + \lambda ||\Psi \boldsymbol{z}||_0. \tag{2.12}$$

However, it has been proven that this is an NP-hard combinatorial optimization problem [48]. Its solutions will be computationally intractable in general, so that an alternate formulation is required. One approach is to consider a *convex relaxation* of Equation (2.12), replacing the non-convex ℓ^0 norm with an ℓ^1 norm term:

$$\boldsymbol{x}^* = \underset{\boldsymbol{z} \in \mathbb{R}^M}{\operatorname{argmin}} \, ||\Phi \boldsymbol{z} - \boldsymbol{y}||_2^2 + \lambda ||\Psi \boldsymbol{z}||_1. \tag{2.13}$$

This is the *basis pursuit denoising* problem (BPDN) [24]. The coherence bounds on measurement and representation systems in Equation (2.8) define conditions under which BPDN yields the same solution as Equation (2.12).

2.2.5 The split Bregman algorithm

Equation (2.13) is convex, but the $||\Psi \boldsymbol{z}||_1$ is not differentiable everywhere. This lack of differentiability poses a challenge for convex optimization techniques which rely on gradient descent to converge to local or global minima. The growth in interest in sparse signal reconstruction methods has fueled interest in convex optimization algorithms designed for cost functions which are sums of ℓ^2 and ℓ^1 norms. In Chapter 3, we make use of one such algorithm for our CS-ET work, the **split Bregman algorithm**.

The split Bregman algorithm, equivalent to the alternating direction method of multipliers discovered by Glowinski [46] and rediscovered and popularized by Osher and Goldstein in [49], seeks to compute the minima of objective functions of the form

$$J(\boldsymbol{z}) = ||F(\boldsymbol{z})||_1 + H(\boldsymbol{z}), \tag{2.14}$$

where both terms in the sum are convex. By introducing additional "dummy" variables into the optimization equation, one is able to split the objective function into separate ℓ^2 and ℓ^1 minimization problems. The ℓ^2 subproblem may be attacked by a number of standard methods, while the ℓ^1 problem can be solved by a so-called shrinkage operation introduced below.

This general functional $J(\boldsymbol{z})$ in Equation (2.14) includes Equation (2.13) for all of the Ψ investigated in Chapter 3 - the linear transforms, as well as the nonlinear TV transform. For these BPDN problems, the split Bregman works by decoupling

the ℓ^1 and ℓ^2 minimizations from each other. Equation (2.13) is trivially equivalent to

$$(\boldsymbol{x}^*, \boldsymbol{d}^*) = \underset{\boldsymbol{z}\in\mathbb{R}^M,\, \boldsymbol{d}\in\mathbb{R}^N}{\operatorname{argmin}} ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2 + \lambda||\boldsymbol{d}||_1 \quad \text{such that } \boldsymbol{d} = \Psi\boldsymbol{z}. \tag{2.15}$$

One can then convert this constrained optimization problem into an unconstrained one, with the equality of $\boldsymbol{d}$ and $\Psi\boldsymbol{z}$ enforced weakly:

$$(\boldsymbol{x}^*, \boldsymbol{d}^*) = \underset{\boldsymbol{z}\in\mathbb{R}^M,\, \boldsymbol{d}\in\mathbb{R}^N}{\operatorname{argmin}} ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2 + \lambda||\boldsymbol{d}||_1 + \frac{\gamma}{2}||\boldsymbol{d} - \Psi\boldsymbol{z}||_2^2. \tag{2.16}$$

The split Bregman iteration for solving this problem is

$$\left(\boldsymbol{z}^{k+1}, \boldsymbol{d}^{k+1}\right) = \underset{\boldsymbol{z},\boldsymbol{d}}{\operatorname{argmin}}\, \lambda||\boldsymbol{d}||_1 + ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2 + \frac{\gamma}{2}||\boldsymbol{d} - \Psi\boldsymbol{z} - \boldsymbol{b}^k||_2^2 \tag{2.17}$$

$$\boldsymbol{b}^{k+1} = \boldsymbol{b}^k + \left(\Psi(\boldsymbol{z}^{k+1}) - \boldsymbol{d}^{k+1}\right). \tag{2.18}$$

The minimization in Equation (2.17) can be rapidly solved by iteratively minimizing with respect to $\boldsymbol{z}$ and $\boldsymbol{d}$ individually:

$$\text{Step 1: } \boldsymbol{z}^{k+1} = \underset{\boldsymbol{z}}{\operatorname{argmin}}\, ||\Phi\boldsymbol{z} - \boldsymbol{y}||_2^2 + \frac{\gamma}{2}||\boldsymbol{d}^k - \Psi\boldsymbol{z} - \boldsymbol{b}^k||_2^2 \tag{2.19}$$

$$\text{Step 2: } \boldsymbol{d}^{k+1} = \underset{\boldsymbol{d}}{\operatorname{argmin}}\, \lambda||\boldsymbol{d}||_1 + \frac{\gamma}{2}||\boldsymbol{d} - \Psi\boldsymbol{z}^{k+1} - \boldsymbol{b}^k||_2^2. \tag{2.20}$$

Equation (2.19) is differentiable, and in this thesis we solve it using the conjugate gradient method [97], a standard technique for the minimization of sums of ℓ^2 norms of differentiable functions. Equation (2.20) is minimized using an componentwise

shrinkage operation. For each component d_j^{k+1} of $\boldsymbol{d}^{k+1}$,

$$d_j^{k+1} = \text{shrink}([\Psi \boldsymbol{z}]_j + b_j^k, 1/\gamma), \tag{2.21}$$

where the shrinkage operation $\text{shrink}(x, \tau)$ is defined by

$$\text{shrink}(x, \tau) = \frac{x}{|x|} \cdot \max(|x| - \tau, 0). \tag{2.22}$$

This operation is easily computed, as well. In all, the iterative variables $\boldsymbol{u}^{k+1}$, $\boldsymbol{d}^{k+1}$, and $\boldsymbol{b}^{k+1}$ are all computable using simple linear algebra, and Equation (2.13) may therefore be minimized efficiently using this method.

2.3 Signal processing in dynamical systems

In chapter 4, we model the activity of collections of neurons as a dynamical system. Dynamical systems may be formulated mathematically in multiple ways; for this thesis, the following definition suffices:

Definition 11. *For a given state space $\mathbb{R}^M$, a continuous-time **dynamical system** is a family of functions Θ^t on $\mathbb{R}^M$ which are the solution to the ordinary differential equation*

$$\frac{d\boldsymbol{x}}{dt} = F(\boldsymbol{x}), \tag{2.23}$$

for $\boldsymbol{x} \in \mathbb{R}^M$ and some suitably-smooth function $F(\boldsymbol{x})$.

A general introduction to dynamical systems may be found in [10]. Dynam-

ical systems can perform signal processing functions, in which the state of certain system elements reflects a constant or time-varying signal, and those states govern the evolution of other *downstream* system states over time. $F(\boldsymbol{x})$ may be nonlinear, and an input signal will generally not be a linear combination of the system's downstream states, but the neurons which are described by this model nevertheless carry out the information processing of neural systems. How can the activity of neuron populations be understood from an information processing perspective?

2.3.1 Low-dimensional activity

The definition of sparsity given in Section 2.2.2 may be applied to the state of a dynamical system at a given point in time. This activity may not carry the same meaning as the sparse coefficients in a linear signal representation, but sparse population activity in dynamical systems can convey similar functional benefits, as described further in Chapter 4. The tools of dimension reduction can be used to better make sense of the simultaneous activity of the elements of large dynamical systems governed by complicated or possibly unknown equations. Instead of predicting the evolution of the system based on its governing equations, we can measure the system state across time and analyze the volume of the state space which these measurements occupy. One hypothesis, popular in both machine learning and neuroscientific communities, is that "natural signals", i.e., signals generated by measurement devices exposed to ambient natural environments, occupy low-dimensional subspaces of the system's state space [87]. Activity states lying within low-dimensional sub-

spaces may be more easily distinguished from one another, aiding categorization and the extraction of features from an input signal.

2.3.2 Numerical simulation of dynamical systems

To simulate the evolution of a dynamical system (or other system of ODEs) on a computer, it becomes necessary to discretize the system in time. There exist many techniques to accomplish this goal, each suited to certain classes of problems. For the equations encountered in Chapter 4, the classic **four-step Runge Kutta** (RK4) method suffices [2]. The RK4 method can be used to compute an approximate solution to first-order initial value problems of the form

$$\frac{d\boldsymbol{y}}{dt} = f(t, \boldsymbol{y}), \quad \boldsymbol{y}(t_0) = \boldsymbol{y}_0. \tag{2.24}$$

Numerical solutions are computed as values $\boldsymbol{y}_n \approx \boldsymbol{y}(t_n)$ for a sequence of times $t_n = nh$ as

$$\boldsymbol{y}_{n+1} = \boldsymbol{y}_n + \frac{h}{6}(\boldsymbol{k}_1 + 2\boldsymbol{k}_2 + 2\boldsymbol{k}_3 + \boldsymbol{k}_4), \tag{2.25}$$

where

$$\boldsymbol{k}_1 = f(t_n, \boldsymbol{y}_n), \tag{2.26}$$

$$\boldsymbol{k}_2 = f\left(t_n + \frac{h}{2}, \boldsymbol{y}_n + \frac{h}{2}\boldsymbol{k}_1\right), \tag{2.27}$$

$$\boldsymbol{k}_3 = f\left(t_n + \frac{h}{2}, \boldsymbol{y}_n + \frac{h}{2}\boldsymbol{k}_2\right), \tag{2.28}$$

$$\boldsymbol{k}_4 = f(t_n + h, \boldsymbol{y}_n + h\boldsymbol{k}_3). \tag{2.29}$$

This method has $O(h^4)$ accuracy.

Chapter 3: Compressed Sensing Electron Tomography

3.1 Overview

Electron tomography (ET), as performed in the transmission electron microscope (TEM) or the scanning transmission electron microscope (STEM), has the unique capability of providing three-dimensional ultrastructure of cells and tissues in a native context, thus revealing important constituents such as membranes, cytoskeletal fibers, and protein complexes, on a macromolecular scale. These 3D visualizations are obtained from multiple 2D projections of a biological specimen when it is tilted through a wide range of angles relative to the incident beam direction. Currently, most ET reconstructions of cells are obtained using the weighted back-projection (WBP) algorithm, or the simultaneous iterative reconstruction technique (SIRT) algorithm. Electron tomography can be applied to specimens prepared by rapid freezing in a vitrified frozen hydrated state, or by fixation with cross-linking agents followed by plastic-embedding and staining with heavy metals to enhance ultrastructural contrast. The choice of specimen preparation depends on the structures that need to be visualized: whereas it is generally preferable to image bacteria and small eukaryotic cells in a frozen hydrated state using cryo-TEM, larger eukaryotic cells and tissues are often best visualized in specimens that have been stained

with heavy atom contrast agents, due to limitations in cryo-preparation of large samples. ET can be performed on both types of specimen but with different imaging modes: phase contrast TEM for frozen hydrated specimens; and amplitude contrast TEM, bright-field STEM, or dark-field STEM for stained specimens. In this work, we apply compressed sensing (CS) recovery algorithms to bright-field and dark-field STEM datasets, and detail their equivalence to certain regularized least-squares reconstruction problems.

With the continuing development of electron tomographic techniques and their widespread adoption by structural and cell biologists, there has been an impetus to enhance the quality of 3D visualizations to discern smaller structures within the complex milieu of the cell. Developing improved 3D reconstruction algorithms offers one potential route to better visualization of cellular ultrastructure. It has been well established that the quality of a tomographic reconstruction can be improved through the incorporation of prior knowledge about the specimen, i.e., through regularized image reconstruction. In this chapter, prior knowledge manifests through sparse signal models, which represent assumptions about image pixel statistics. Regularization takes the form of penalty functions on a least-squares optimization problem, designed to enforce sparsity in reconstructed images. More recently CS, which exploits sparse signal structure to reconstruct a signal from undersampled measurements via regularized recovery, has attracted increasing attention for a number of data processing applications. The recent success of CS and related mathematical techniques in medical imaging - particularly magnetic resonance imaging (MRI) [75, 27, 4, 58, 59, 28] and more recently x-ray computed tomography [112] - have

led to a growing interest in using CS methods in the field of electron tomography (ET). Most reports of CS applications in ET have involved the imaging of inorganic materials, e.g., STEM dark-field tomography of iron oxide nanoparticles [94, 73, 51]. There have also been studies aimed at applying regularization and CS methods to cryo-TEM [1, 99]. Here, we consider the application of CS to improve the quality of 3D cellular ultrastructure imaging in electron tomography of cells and tissues which have been prepared using heavy atom staining to enhance contrast. This type of specimen preparation is the technique that is most widely used among cell biologists because it enables the analysis of large eukaryotic cells.

To gauge what advantages CS might offer, we first examine how the sampling assumptions of CS relate to the mathematical models of tomographic sampling, and how CS image recovery relates to regularized least-squares recovery [108, 84]. From a computational perspective, the tools of CS and regularized recovery have much in common, and one of our aims is to clarify their relationship.

There remain important unsolved problems in rigorously grounding a theory of compressed sensing electron tomography (CS-ET). First, mathematical proofs of the efficacy of CS for undersampled recovery in biomedical imaging rely heavily on randomized Fourier sampling in order to obtain probabilistic reconstruction error bounds. The measurement process of STEM tomography can be interpreted theoretically as acquiring Fourier data, but the manner of acquisition violates CS sampling assumptions. In this chapter, we clarify the nature of this obstacle. A simple randomized variant of this procedure, in which ET projection angles are chosen randomly within a mechanically-feasible range, is a natural topic of investigation

for a new measurement procedure. However, numerical results from simulated data indicate that this strategy is inferior to traditional uniform sampling, as illustrated in Figure 3.9.

Second, the performance of sparse reconstruction methods depends on the structural complexity of the specimen being imaged. This may lead to different relative performance gains by CS-ET recovery depending on the application domain in question and the sparsity model chosen. In this chapter, we analyze the standard (*identity*) sparsity model as well as the *wavelet sparsity* and *TV sparsity* models [109, 108], which have previously been used in CS applications including work by [1, 75, 7, 51, 73, 63]. Simulated datasets and experimental datasets from nanoscale inorganic materials are often more sparse in common transform domains than experimental datasets from biological structures, with implications for the relative advantage of CS-ET recovery for each type of structure.

After reviewing pertinent aspects of numerical ET and CS theory, we examine CS-ET recovery for simulated membrane-bound compartments within cells; in these simulations statistical properties ensure a high degree of compressibility in the TV, wavelet, and identity domains. We also demonstrate that our CS-ET recovery on simulations of nanoscale inorganic materials gives results that are consistent with previous work. Under such conditions, the advantages of CS-ET relative to WBP or SIRT reconstructions are substantial. Reconstructions from a **membrane phantom**, a simulated dataset comprised of randomly-distributed ellipsoids and spheres at multiple scales, eccentricities, and orientations are used to characterize the performance of CS-ET under a range of noise and measurement conditions. In these

simulated membrane reconstructions, we find high correlation between sparsity and reconstruction error, significant performance advantages of CS-ET relative to WBP, and consistent superiority of uniformly-sampled CS-ET to random-angle CS-ET.

We then analyze experimental STEM tomographic tilt series acquired from plastic-embedded, heavy metal stained sections of fixed cells. Our simulations and experimental results indicate that although CS reconstructions substantially outperform the commonly-used WBP and SIRT methods for highly-sparse objects, CS generally provides a smaller or no advantage for reconstructions of biological structures. We attribute the difference this disparity in performance to the significant difference in sparsity between simple phantoms and complex biological structures. We discover that the relationship between data sparsity and the performance of CS reconstruction on the one hand, and between specimen structure and data sparsity on the other, can complicate a thorough understanding of the effectiveness of CS-ET under all potential imaging conditions. However, in application domains for which accurate sparse signal models can be established, there is increasing evidence for the efficacy of CS-ET in both fully-sampled and undersampled tomogram recovery.

3.2 Background

3.2.1 STEM tomography

In general, **tomography** refers to methods for constructing a 3D reconstruction of an object and its interior from 2D projections. These projections are created by measuring changes in penetrating waves or particles which are sent through the

imaging specimen. There are several imaging modalities which perform this function using many different projection geometries, producing 3D structural information at different spatial and temporal resolutions. With *electron tomography* (ET), 3D reconstruction of the electron density function $f(\boldsymbol{x}) : \Omega \to \mathbb{R}$, $\Omega \subseteq \mathbb{R}^3$, of microscopic structures are created from electron beams using a transmission electron microscope (TEM) [111] or scanning transmission electron microscope (STEM) [83]. A thorough introduction to the theory and practice of electron tomography may be found in [47], the most essential points of which are highlighted here. Both the TEM and STEM produce images by measuring changes in incident electron beams passed through the imaging specimen. STEM images are produced using a single beam emitter that is scanned over the surface of the specimen to acquire each spatial data point in a 2D projection.

In amplitude contrast imaging modes, structural information with roughly 1 nm spatial resolution is obtained by measuring the scattering of electrons along beam paths, comparing the number n_0 of incident electrons with the number n of electrons which pass through the specimen unscattered. Two imaging modes which fall into this category are explored in this chapter. The two modes, **bright field** (BF) STEM tomography and **dark field** (DF) STEM tomography, differ in how electron scattering measurements are acquired, but both can be analyzed with the same mathematical theory. In sufficiently-thin specimens, the quantities n and n_0 for a given beam L are related by the equation

$$n = n_0 \exp\left(-C\mu_L\right), \tag{3.1}$$

where C is a positive constant and μ_L is the integral of the specimen's electron density function $f(\boldsymbol{x})$ along the beam path L. As discussed in the next section, this quantity is the Radon transform $\mathcal{R}f(L)$ of $f(\boldsymbol{x})$ along L, so that

$$\mathcal{R}f(L) \propto -\log(n/n_0). \tag{3.2}$$

In this way, BF and DF STEM projections become related to the Radon transform of the specimen's density function $f(\boldsymbol{x})$, and this density function may be recovered from the computed Radon transforms.

3.2.2 The Radon transform

Mathematically, the problem of electron tomogram reconstruction can be cast numerically as an inverse problem for the **Radon transform**:

$$\mathcal{R}f(L) = \int_L f(\boldsymbol{x})|d\boldsymbol{x}|, \tag{3.3}$$

a line integral defined for any line $L \in \mathcal{L}$, the space of all lines in $\mathbb{R}^3$. The notation $|d\boldsymbol{x}|$ refers to the standard 1D Lebesgue measure on the line L. Though intrinsically a 3D problem, the parallel-beam measurement process used in STEM tomography samples the Radon transform along sets of parallel lines orthogonal to a fixed axis of rotation. This effectively decomposes pre-processed projection data into independent 2D sets, each containing samples of Radon transform data defined on the space of lines in $\mathbb{R}^2$. The sampling space, a proper subset of $\mathcal{L}$, can then

be explicitly parametrized by a *normal vector* α, a *distance coordinate* s equal to the minimum distance between L and the origin $(0,0)$, and a *depth coordinate* z indexing the planes normal to the axis of rotation:

$$\mathcal{R}f(\alpha, s, z) = \int_{-\infty}^{\infty} f(t\sin\alpha + s\cos\alpha, -t\cos\alpha + s\sin\alpha, z)\, dt. \tag{3.4}$$

After preprocessing, the data created by the microscope is interpreted as a $T \times S \times P$ array **sample set**

$$\boldsymbol{y}(i,j,k) \approx \mathcal{R}f(\alpha_i, s_j, z_k), \tag{3.5}$$

with a finite set D of Radon sampling parameters:

$$D = \{(\alpha_i, s_j) \mid i \in [1,T], j \in [1,S]\}, \tag{3.6}$$

as well as depth samples $\{z_k\}_{k=1}^{P}$. Approximation error stems from several sources: the most pertinent to our analysis is the shot noise inherent in the computation of y via discrete electron emissions [47]. In areas of ET such as cryo-ET where low electron doses are required, the deviation of this noise behavior from Gaussian noise models can be significant, and this presents a particular complication for CS - see e.g. [91] for a recent perspective on the problem of CS reconstruction in the presence of Poisson noise. For our analysis, we avoid this problem by using datasets obtained from larger electron doses, allowing the Gaussian noise model to suffice.

An additional complication arises from approximating the operator $\mathcal{R}$ digitally as a computational process. For iterative reconstruction methods such as CS-ET, it

is necessary to compute the forward and adjoint measurement operators. Multiple algorithms exist to tackle this problem, each performing a slightly different linear transformation on discrete images. For our analysis, the **digital Radon transform** R implemented using the ASTRA toolbox for MATLAB [88]. For the M and N values used for practical purposes, this approximation of $\mathcal{R}$ by R contributes negligibly to the error in (3.5). See [39] for further considerations on digital Radon transforms.

Despite these and other complications, the equation in (3.5) can be made sufficient for the purposes of numerical reconstruction. The measurement set $\boldsymbol{y}$ is received as a $T \times S \times P$ array of preprocessed electron counts. The 2D subsets for fixed z values:

$$\boldsymbol{y}_\kappa(i,j) = \boldsymbol{y}(i,j,\kappa),$$

where $i \in [1,T]$ and $j \in [1,S]$, are reconstructed independently. For the remainder of the chapter we will focus on these 2D reconstruction problems.

The solution to each such reconstruction is a vector $\boldsymbol{f}_\kappa^*$, a 2D slice of the full 3D **digital tomogram** $\boldsymbol{f}^*$ which itself is a finite $M \times N \times P$ approximation of f. Each reconstructed $\boldsymbol{f}_\kappa^*$ is an $x-z$ slice of the digital tomogram, with a height of P and a width of M pixels, and can thus be considered a vector $\boldsymbol{f}_\kappa^* \in \mathbb{R}^{MP}$. The choice of coordinate axis alignment, as well as a representative $\boldsymbol{f}_\kappa^*$, are illustrated in Figure 3.1.

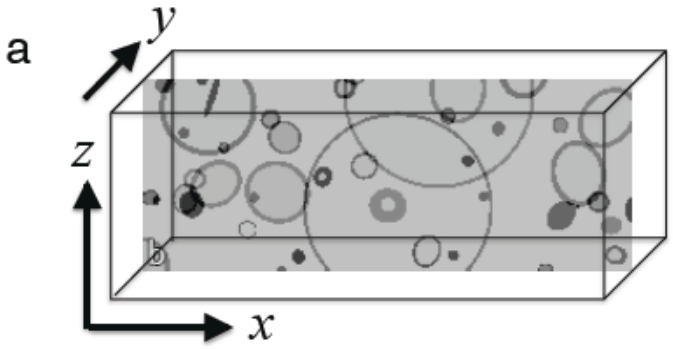

Figure 3.1: **Simulated tomographic reconstructions from a phantom consisting of stained membranes embedded in a uniform matrix**. **(a)** Slice through phantom in $x - z$ plane with coordinate axis orientation indicated.

3.2.3 CS-ET signal recovery

The problem of recovering $\boldsymbol{f}$ from tomographic measurements can be approached in many ways, but the approach most relevant here is to work explicitly with the digital Radon transform R and its linear algebra. For a fixed sample set D, the digital Radon transform R for $M \times P$ images can be realized as a $(TS \times MP)$ **measurement matrix** R_D. For any set D one obtains through tomographic measurement, R_D will not be invertible. This means that the problem:

$$\text{Find } \boldsymbol{f}_\kappa^* \in \mathbb{R}^{MP} \text{ such that } R_D \boldsymbol{f}_\kappa^* = \boldsymbol{y}_\kappa \tag{3.7}$$

is *ill-posed*, multiple $\boldsymbol{f} \in \mathbb{R}^{MP}$ satisfy this equation and the desired solution cannot be distinguished among them. To overcome this limitation, we recast the recovery of $\boldsymbol{f}$ using the regularized least-squares techniques introduced in Section 2.2.4. We seek a least-squares reconstruction of $\boldsymbol{f}_\kappa$ while enforcing sparsity in a combina-

tion of the TV, identity, and wavelet domains. For some choice of regularization hyperparameters λ_{TV}, λ_I, and λ_W:

$$\boldsymbol{f}_\kappa^* = \underset{\boldsymbol{f}\in\mathbb{R}^{MP}}{\operatorname{argmin}} ||R_D\boldsymbol{f} - \boldsymbol{y}_\kappa||_2^2 + \lambda_{TV}||TV\boldsymbol{f}||_0 + \lambda_I||\boldsymbol{f}||_0 + \lambda_W||W\boldsymbol{f}||_0. \tag{3.8}$$

As stated, however, this is an NP-hard problem, and we replace the ℓ^0 norm terms with ℓ^1 norms to obtain the objective function minimized by the CS-ET algorithm:

$$\boldsymbol{f}_\kappa^* = \underset{\boldsymbol{f}\in\mathbb{R}^{MP}}{\operatorname{argmin}} ||R_D\boldsymbol{f} - \boldsymbol{y}_\kappa||_2^2 + \lambda_{TV}||TV\boldsymbol{f}||_1 + \lambda_I||\boldsymbol{f}||_1 + \lambda_W||W\boldsymbol{f}||_1, \tag{3.9}$$

It is worth noting that in practice, useful mutual coherence bounds are not known for measurement and sparsity models of interest in ET. Radon analysis in particular remains limited, with respect to any sparsity model. Theoretically rigorous reconstruction guarantees depend delicately on the structure of the measurement matrix Φ and sparsifying matrices Ψ [35, 11, 21], and the conditions under which sparse reconstruction is guaranteed do not coincide with measurement schemes which STEM instrument limitations impose. Traditional tomographic sampling involves acquiring projections at a range of tilt angles, approximately uniformly-spaced within a mechanically feasible range.

This range is chosen deterministically, with small machine misalignments measured and reported after the fact. Therefore, for traditional tomography, R_D is a *deterministic* measurement matrix. The majority of CS recovery results rely on measurement matrices chosen from a random ensemble. See, e.g., [15, 41, 110] for

recent developments in deterministic measurement matrix design and analysis, but the theoretical results are as yet inadequate to address the R_D matrix structure. Moreover, work guaranteeing TV-sparse signal reconstruction is currently unable to say anything about solutions of equations such as (3.9) [19]. In a simple attempt to investigate the former issue, we perform random-angle tomography experiments on simulated datasets, choosing tilt angles uniformly at random within a mechanically-feasible range.

3.3 Methods

Pragmatically, the aim of this work is to implement a CS-ET algorithm for the regularized reconstruction of STEM tomograms and compare its performance to other reconstruction approaches. Evaluating this performance hinges on understanding how well the images encountered in biological STEM tomography satisfy the assumptions for sparse signal reconstruction problems; to this end we also attempt to quantify the sparsity of image specimens in multiple transform domains and relate these values to reconstruction performance.

3.3.1 Compressed sensing approach

Compressed sensing, strictly speaking, refers to the recovery of a signal from a small number of appropriately-chosen linear measurements. This achievement has become a popular focal point for a novel mathematical perspective on the recovery of signals from measurements. An understanding of the appropriateness of CS

methods for any physical measurement procedure must begin with a recasting of the measurement process into the language of CS theory. Recent work by [73] provides a comprehensive reinterpretation of the ET recovery problem, and here we therefore keep our development brief.

Any digital signal, including the tomogram images reconstructed by ET, can be understood as a vector of data points $\boldsymbol{f}$. A 2D digital tomographic image with N pixels is a (column) vector $\boldsymbol{f}$ with N entries, each entry containing a grayscale pixel value. This is a discrete approximation to the potential function $f(x, y)$ and is ultimately the mathematical object recovered via computational tomographic reconstruction.

Recall that for each 2D subproblem of tomographic reconstruction, tilt series data consists of T detector counts at each of S angles, as indexed by the set D in (3.6). There are therefore ST corresponding measurement vectors $\boldsymbol{\varphi}_i$, each evaluating to a particular tilt series pixel value when an inner product is taken against image $\boldsymbol{f}_\kappa^*$. The rows of the measurement matrix R_D are built from these vectors: for each $i \in [1, T]$ and $j \in [1, S]$, row $(i-1)S + j$ of R_D is a measurement vector $\varphi_{ij} \in \mathbb{R}^{MP}$, such that $\langle \varphi_{ij}, \boldsymbol{x} \rangle$ is pixel (i, j) of the tilt series data. Given these measurements and the sparsity models introduced in the previous section, we can seek to recover a 3D tomogram as a collection of 2D solutions to Equation (3.9).

3.3.2 Sparsity and incoherence for biological structures

In Section 3.4 of this chapter, we analyze the compressibility of an experimental biological dataset in the identity, wavelet, and TV domains. Though the appropriateness of each choice is dataset-dependent, all three are common in recent work in CS and regularized recovery [73, 94, 52, 75, 7], and our CS-ET implementation is capable of using any combination of the three. This analysis is performed by computing the 2.5% compressibility ratios of the experimental dataset in the three transform domains. See Section 2.2 for the definition of percent compressibility ratios. These values are expressed as a multiple ρ of the 2.5% compressibility ratios of the membrane phantom in each transform domain, for ease of comparison.

Additionally, we quantify the coherence of Radon transform measurements and the mutual coherence of Radon transform measurements with a DB8 wavelet representation system in order to gauge the discrepancy between this measurement system and the theoretical requirements imposed by CS theory for sparse signal reconstruction. The results of this can be seen in Figures 3.2 and 3.3.

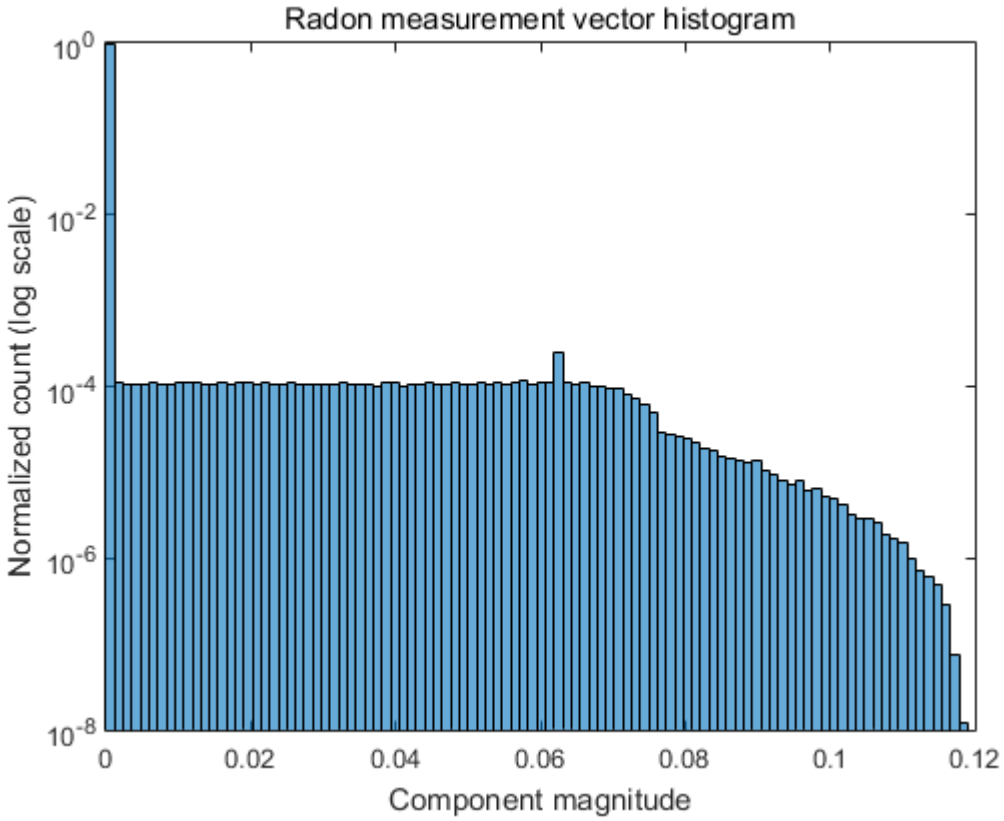

Figure 3.2: A histogram of the magnitudes of the components of Radon transform measurement vectors, taken from a Radon transform of a 256×256 image at angles from $-70°$ to $70°$ at $5°$ increments. The *coherence* of a measurement dictionary containing vectors of length M is defined as $\sqrt{M}$ times the largest magnitude among the vectors' components, and is used as a convenient but conservative bound for the restricted isometry property (RIP) of a measurement system. For the measurement dictionary displayed here, consisting of 7424 length-65536 vectors, the coherence value is approximately $0.1188 \cdot 256 = 30.4128$. This value is too large to be of use for a theoretical analysis, suggesting the need for new analytic tools for real-world measurement systems.

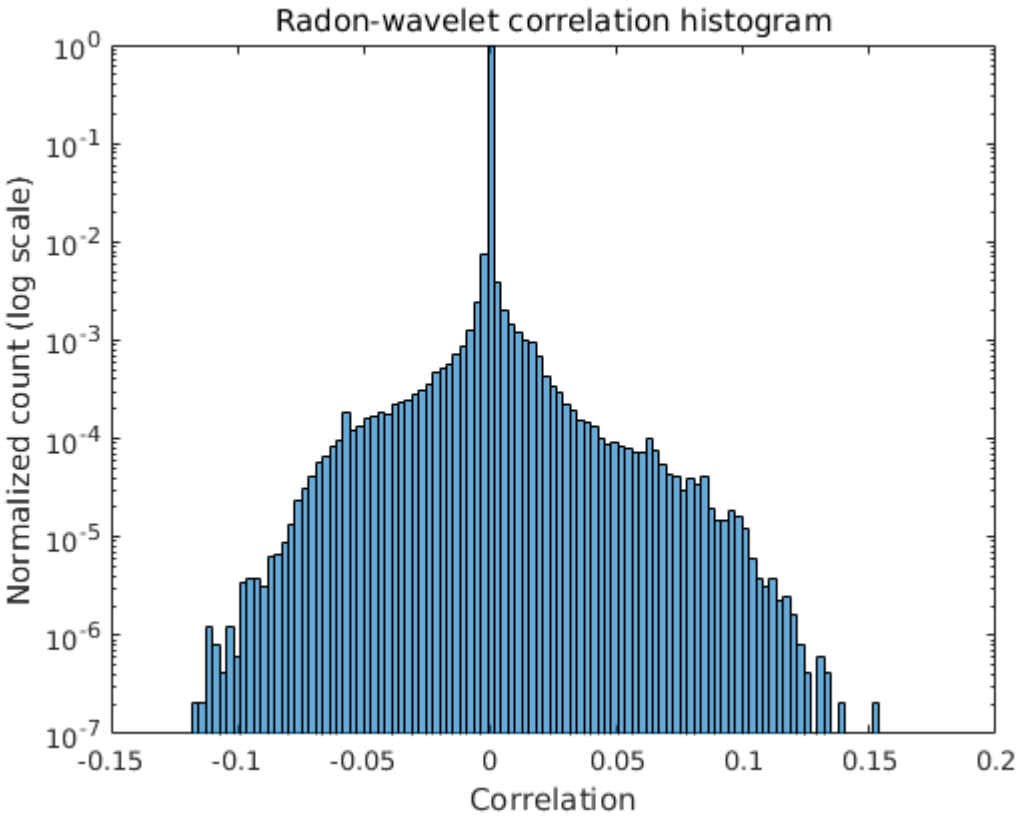

Figure 3.3: A histogram of the correlation values between normalized discrete Radon transform measurement vectors and elements of a DB8 wavelet basis, each for a 256×256 image and sampling 10% of the vectors in each dictionary. The Radon transform used consists of projections at angles from $-70°$ to $70°$ at $5°$ increments. The *mutual coherence* of a measurement dictionary and representation basis for vectors of length M is defined as the $\sqrt{M}$ times the magnitude of the largest correlation between their elements, and is used as an upper bound for the restricted isometry proprty (RIP) of a measurement system for signals with sparse representations in the given representation basis. For the dictionaries displayed here, a mutual coherence of approximately $0.15 \cdot 256 = 38.4$ is too large to be of use to a theoretical analysis of the optimization routine, suggesting the need for new analytic tools for real-world measurement systems.

Since the number of vectors to check grows linearly with image size, all calculations were performed using a test system which assumed a 256x256 image. The coherence of the Radon measurement vectors is equal to the largest component of any of the vectors, while the mutual coherence between Radon measurements and

wavelet representations is equal to the largest magnitude of any inner product between the two, with Radon measurement vectors all normalized to have Euclidean norm equal to the square root of the number of pixels in the test system, which in this case equals 256. The values obtained in this manner, approximately 30.4 for the Radon measurement coherence and 38.4 for the Radon-wavelet mutual coherence, are too large to be of use for theoretical justifications for the efficacy of CS-ET. To quantify the discrepancy, note that according to Equation (2.8) in chapter 2, we would require a mutual coherence of approximately 2.8 between our measurement system and wavelet representation system in order to provably recover a 256x256 image composed of 100 nonzero wavelet coefficients from 70 tomographic projections. The empirical efficacy of the CS-ET algorithm therefore suggests the development of improved theoretical tools for understanding when CS recovery does and does not succeed for deterministic measurement systems.

3.3.3 CS-ET reconstruction

Following the developments discussed in Section 3.3.1, tomograms were reconstructed in independent 2D slices as solutions to the equation

$$\hat{\boldsymbol{f}} = \underset{\boldsymbol{g}\in\mathbb{R}^N}{\operatorname{argmin}} \; ||R\boldsymbol{g} - \boldsymbol{y}||_2^2 + \lambda_{TV}||TV\boldsymbol{g}||_1 + \lambda_I||\boldsymbol{g}||_1 + \lambda_W||W\boldsymbol{g}||_1, \tag{3.10}$$

where matrix R has Radon measurement vectors as its rows, W is a DB8 wavelet transform [34], and $(\lambda_{TV}, \lambda_I, \lambda_W)$ is a collection of regularization hyperparameters. By exploiting the Fourier Slice Theorem [62], one may treat STEM tilt series data as

samples of either the Radon or Fourier transform, and choose to work in the Radon domain as we have, or in Fourier space [53]. The latter approach has been developed and used successfully for CS-ET, [73, 45, 78], but working in the Fourier domain requires numerical overhead in the form of non-uniform fast Fourier transform (NUFFT) operators [45] which may be avoided in the Radon domain. We pursue a Radon domain reconstruction for this reason, although the overall methodology closely parallels the development described by [73].

For various integer values of k, k-fold undersampled or simply k**x undersampled** tilt series were created by keeping every kth projection and additionally the last projection from a full tilt series.

3.3.4 Comparisons between reconstruction methods

Experimental dataset CS-ET reconstructions were compared against WBP and SIRT reconstructions of the same projections. WBP reconstructions were implemented in MATLAB using the MATLAB Imaging Toolbox's iradon command. SIRT reconstructions were written in MATLAB and Fortran and run on the high-performance Biowulf Linux cluster at the National Institutes of Health (`http://biowulf.nih.gov`). Simulated dataset CS-ET reconstructions were compared against WBP reconstructions performed in MATLAB with the imaging toolbox's iradon command.

3.3.5 CS-ET numerical implementation

Equation (3.10) is a convex optimization problem, but the ℓ^1 norm is not a globally differentiable function which presents an obstacle to efficient optimization. In recent years, several efforts have been made to create algorithms for rapidly solving ℓ^1-regularized least-squares problems. We choose to use the **split Bregman** algorithm [49], in contrast to the conjugate-gradient and interior-point methods used in prior studies [73, 52, 18]. See Section 2.3.1 for more details on the split Bregman algorithm.

In addition to the λ_{TV}, λ_I, and λ_W regularization parameters in (3.10), the split Bregman method requires the specification of a data fidelity parameter μ. A convenient method of parameter specification is to fix a value of μ and write $(\lambda_{TV}, \lambda_I, \lambda_W)$ as multiples of the μ value. In this way, μ can be viewed as controlling the size of each update of the iterative split Bregman procedure. This method allowed us to make simple changes to the hyperparameter sets used for each level of undersampling in the experimental reconstructions.

The bright-field dataset used the following parameter values. For the 1x reconstruction, $\mu_1 = 5 \times 10^{-6}$. The 3x, and 6x reconstructions had $\mu_3 = \mu_6 = 1 \times 10^{-5}$. All bright-field reconstructions had $(\lambda_{TV}, \lambda_I, \lambda_W) = \mu_i \cdot (1.2, 6, 4)$ for $i = 1, 3, 6$.

The dark-field dataset used different parameter sets. For all three reconstruc-

tions, $\mu_1 = \mu_3 = \mu_6 = 10^{-6}$. Then,

$$(\lambda_{TV,1}, \lambda_{I,1}, \lambda_{W,1}) = \mu_1 \cdot (4, 8, 10), \tag{3.11}$$

$$(\lambda_{TV,3}, \lambda_{I,3}, \lambda_{W,3}) = \mu_3 \cdot (6, 2, 6), \tag{3.12}$$

$$(\lambda_{TV,6}, \lambda_{I,6}, \lambda_{W,6}) = \mu_6 \cdot (4, 2, 4). \tag{3.13}$$

As an iterative method the split Bregman algorithm also requires the specification of a number of iterations, in this case divided into "inner" loop iterations and "outer" loop iterations. All experimental reconstructions used 12 inner loop iterations and 10 outer loop iterations.

The core CS-ET algorithm derives from the mrics.m MATLAB function developed by Tom Goldstein in conjunction with [49]. A repository containing this algorithm as well as supplementary utilities is available at https://github.com/norbert-wiener-center/cset. Our algorithm differs from mrics.m by using a Radon measurement matrix instead of a Fourier measurement matrix in the least-squares term of Equation (3.10). This modification bars the solving of the intermediate ℓ^2 minimization problem through Fourier methods, and so we use a conjugate gradients implementation tailored to ASTRA's Radon transform code.

Reconstruction of 3D volumes through independent 2D problems lends itself naturally to parallelization to improve reconstruction times. This was implemented in a simple way in MATLAB using the parfor loop structure, reducing the fully sampled reconstruction time to under 20 minutes. Experimental dataset reconstruction was carried out on a Windows 7 PC with dual 8-core Intel®Xeon®3.40GHz CPU's,

using a MATLAB parallel pool using 26 workers.

3.3.6 Datasets used

Experimental data

The three experimental datasets used in this chapter are tomographic tilt series acquired from a stained beta cell in a plastic-embedded mouse pancreatic islet of Langerhans. The three tilt series were acquired on an FEI Tecnai TF30 TEM operating at 300kV (FEI Inc., Netherlands). Images were acquired in the STEM mode, using a high-tilt tomography holder (Fischione Instruments, Inc., Export, PA), in conjunction with an on-axis bright field detector (Gatan Inc., Pleasanton CA) and in-column high-angle annular dark-field STEM detector (Fischione Instruments, Inc., Export, PA). For each dataset, a single tilt series was acquired automatically with a tilting range of $\pm 78^{\circ}$ and increments of approximately 2°, resulting in 79 projections. Each projection was 2048×2048 pixels, with a pixel size of 1.67nm (BF) and 1.4nm (DF). Two tilt series were acquired in the DF-STEM mode - one with a standard dose of 24 nA probe current and one with 10% of the standard dose, to test the reconstruction algorithms under noisy conditions. The tilt series were aligned using IMOD software [68] and binned by 2. An example of a binned projection is shown in Figure 3.4.

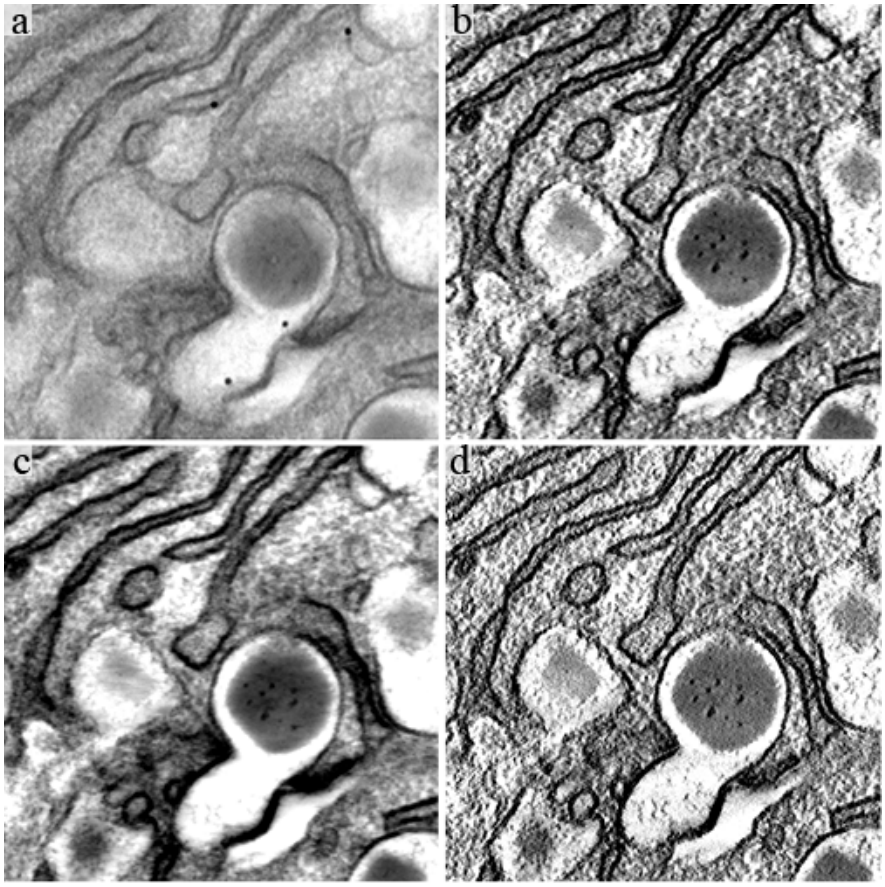

Figure 3.4: **Comparison of bright-field STEM projections and $x-y$ reconstruction slices**. **(a)** The STEM projection of a biological sample at 0° tilt after initial preprocessing in IMOD. This provides a "ground truth" that serves as a point of reference for the broad morphological structure of reconstruction volume data. Reconstruction $x-y$ slices from fully sampled data are displayed in **(b)** for the CS-ET volume, **(c)** from the SIRT volume, and **(d)** from the WBP volume. The reconstruction slices resolve small-scale features not apparent in the overhead projection alone, and are free of interference from the high-contrast gold nanoparticles present near the upper and lower boundaries of the sample.

Each reconstructed volume of 1024 independent reconstructions of 1024 $\times$ 100 2D images ($N = 102400$). Figure 3.1 demonstrates this chapter's choice of orientation for the coordinate axes, consistent with common nomenclature in the ET community. Each reconstruction generates $x-z$ data of the tomogram volume for a fixed y value, but the perspective most commonly used for examining imaged

specimens is the orthogonal $x - y$ view.

The use of the identity sparsity model with the experimental images requires background subtraction from the resulting reconstructions before evaluating the reconstruction sparsity. Background values are estimated by averaging over patches of manually-observed background areas in an initial WBP reconstruction.

Simulated data

A key difference between the imaging of biological and nanometallurgical specimens is the complexity of the resulting volumes. Tomograms of assemblies of nanoparticles may be highly sparse or TV-sparse, in comparison with the decreased sparsity of biological tomograms. In addition to the evidence from experimental reconstructions for the decreased utility of undersampling for biological CS-ET, we consider the problem numerically using a mock-nanometallurgical phantom shown in Figure 3.5.

The construction of this phantom is intended to draw on the statistical properties of datasets observed [73, 94], and others - high contrast, piecewise constant components with sharp edges. Projections of the nanoparticle phantom are simulated in MATLAB. Poisson noise with a rate parameter of 5500 is added to the projections, which are then corrupted with Gaussian noise with standard deviation equal to 10% of the projections' mean.

The membrane phantom is a $256 \times 100 \times 256$ array taking values in $[0, 1]$. On top of a non-zero background, multiple ellipses of random sizes, locations, eccentrici-

ties, orientations, and contrasts are superimposed to imitate observed arrangements of membranous bodies in heavy metal-stained cellular preparations. This phantom provides contrast and structural conditions more akin to those encountered in biological imaging, and sparsity levels between those found in the nanoparticle phantom and the experimental datasets. Both noiseless and noisy projection data are tested; the noisy projections are corrupted with Poisson noise and have a Gaussian noise level equal to 10% of the projection's mean after background subtraction.

3.3.7 Data analysis

Reconstruction of experimental tilt series

The experimental datasets were reconstructed using CS-ET, SIRT, and WBP algorithms from 1x, 3x, and 6x projection data. For the bright-field dataset, these contained 79 tilts, 27 tilts, and 14 tilts respectively. For the dark-field dataset, these contained 77, 26, and 13 tilts. Figure 3.11 illustrates the results of this procedure for $x-z$ slices of the reconstructions, comparing the same slice across multiple levels of undersampling.

Figure 3.12 gives a comparison of several fully-sampled $x-z$ slices from the high-dose DF reconstruction. Figure 3.13 shows representative overhead $x-y$ slices from reconstructions of the low-dose DF dataset at multiple levels of undersampling.

After reconstructions were performed, all tomogram volumes were whitened (scaled to zero mean and unit variance) and then uniformly scaled for better image contrast for visualization purposes.

All regularized reconstruction methods require a choice of hyperparameter "weights", and many open problems remain for obtaining good *a priori* estimates of weight values in experimental application conditions. Parameter combinations used in this chapter were found heuristically based on experience with the datasets. See [73] for a more in-depth discussion of hyperparameter tuning concerns.

The sparsity of the experimental datasets was quantified via the n%-compressibility ratio defined in Section 1.2 for $n = 2.5$ represented as a distribution of multiples ρ of the 2.5%-compressibility ratios of the 256 membrane phantom $x - z$ slices. This distribution gives an indication of the relative compressibility of the experimental and phantom datasets.

Reconstruction of simulated tilt series

The goal of reconstructing the nanoparticle phantom was to establish a baseline comparison between our CS-ET implementation and the Fourier-based method used in [73]. The full set of data analyzed in this work supports a conclusion that CS-ET does not substantially outperform alternative methods such as SIRT and WBP for less-compressible signals, so we seek to eliminate algorithm implementation errors as a source of performance degradation. See Figure 3.5 for a description of the nanoparticle phantom reconstruction procedure.

For the membrane phantom, reconstructions $\hat{\boldsymbol{f}}$ were analyzed in a standard manner via their *root mean squared error*: $||\hat{\boldsymbol{f}} - \boldsymbol{f}||_2^{1/2}/\sqrt{M \cdot N}$, calculated for each 2D $x - z$ phantom slice $\boldsymbol{f}$. Root mean square error (RMSE) was calculated for all reconstructions, and the CS-ET results were compared with the identity, wavelet, and TV sparsity of each slice of the 3D phantom to examine the correlation between

them. See Figure 3.8 for the results of this analysis.

Reconstructions are computed for the membrane phantom using both random and traditional tomographic sampling schemes, to test the performance of this mechanically-feasible randomized sampling strategy. The uniform tilt angles were chosen to coincide with the tilt series angles reported for the bright-field experimental dataset. Random tilt angle vectors were chosen to have the same number of projections as the uniform series, chosen uniformly at random from the same tilt range as the uniform angles. The results of this analysis are displayed in Figure 3.9.

To test the stability of the reconstruction algorithm, all reconstructions were repeated under two different noise conditions. First, reconstructions were made from projection data free of noise. Second, reconstructions were made from projections corrupted by Poisson noise with a rate parameter of 5500 and Gaussian noise with standard deviation equal to 10% of the projection mean. Since the noisy conditions are more relevant to experimental measurement conditions, images of those results were included in Figure 3.6. Noiseless reconstruction results can be found in Figure 3.7.

3.4 Results

In this section, we present the reconstruction results for the CS-ET and comparison methods on the simulated and experimental datasets. We also present our analysis of the compressibility of the simulated and experimental datasets, how they compare, and how they are correlated with reconstruction accuracy.

3.4.1 Tests on simulated tomographic data

We first applied the CS-ET reconstruction technique to a tilt series of simulated dark-field STEM images of circular and crescent-shaped objects, as from a cross-section of inorganic nanoparticles. This model of simple, high-contrast structures reflects the image statistics present in certain materials science applications of STEM tomography (Figure 3.5a). As shown in Figure 3.5b-j, the resulting CS-ET reconstructions are significantly clearer than their SIRT and WBP counterparts, in agreement with previous work of [73]. This type of simple nanoparticle specimen is ideal for application of the CS reconstruction approach; the results show that our novel real-space CS algorithm gives the expected result.

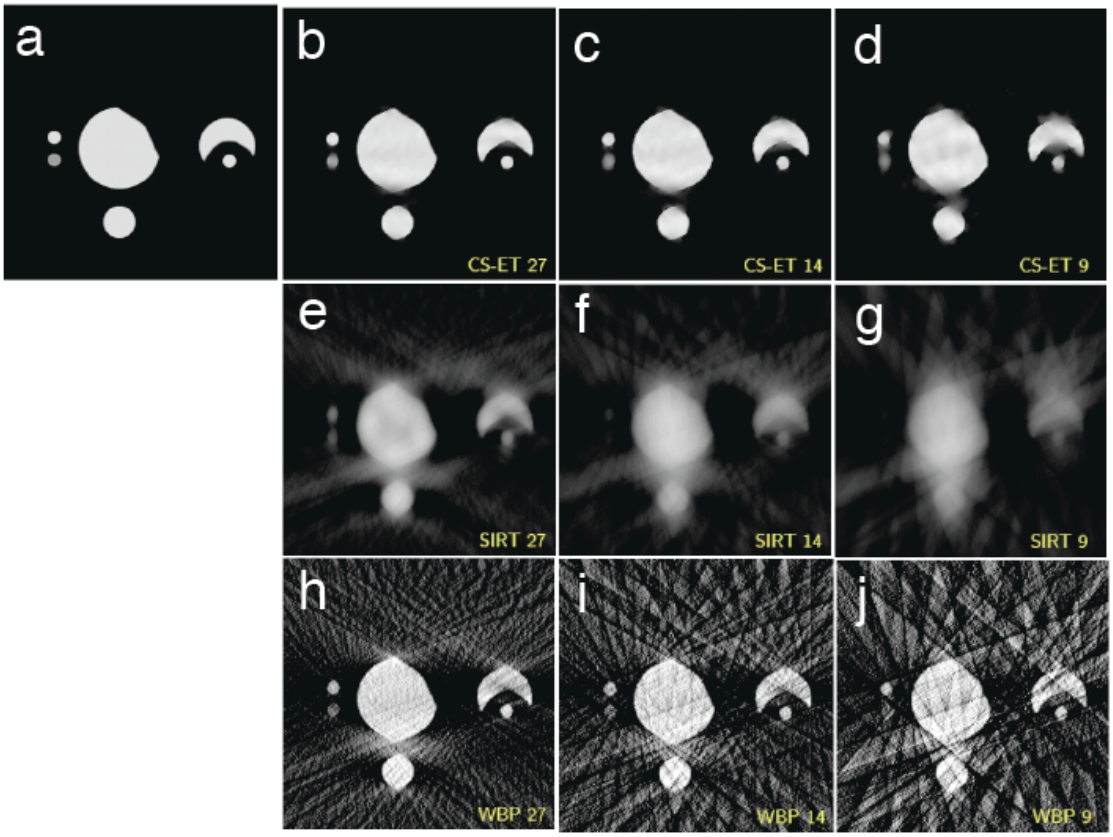

Figure 3.5: **Comparison of reconstruction techniques for undersampled recovery of an inorganic nanoparticle phantom**. Simulated projections at 5° increments between −70° and 60° are 2x and 3x undersampled to produce reconstructions from 27, 15, and 9 projections respectively. **(a)** simulated mass distribution in $x - z$ plane of single-phase nanoparticle phantom showing piecewise constant particles, **(b-d)** CS-ET reconstructions, **(e-d)** SIRT reconstructions, and **(h-j)** WBP reconstructions. Note the higher performance of CS-ET reconstruction for this simple, sparse model, in agreement with the earlier work of Leary et al. (2013).

We next tested the CS-ET technique on a simulated tomographic STEM bright-field tilt series generated from a 3D phantom imitating a collection of stained membrane-bound cellular compartments.

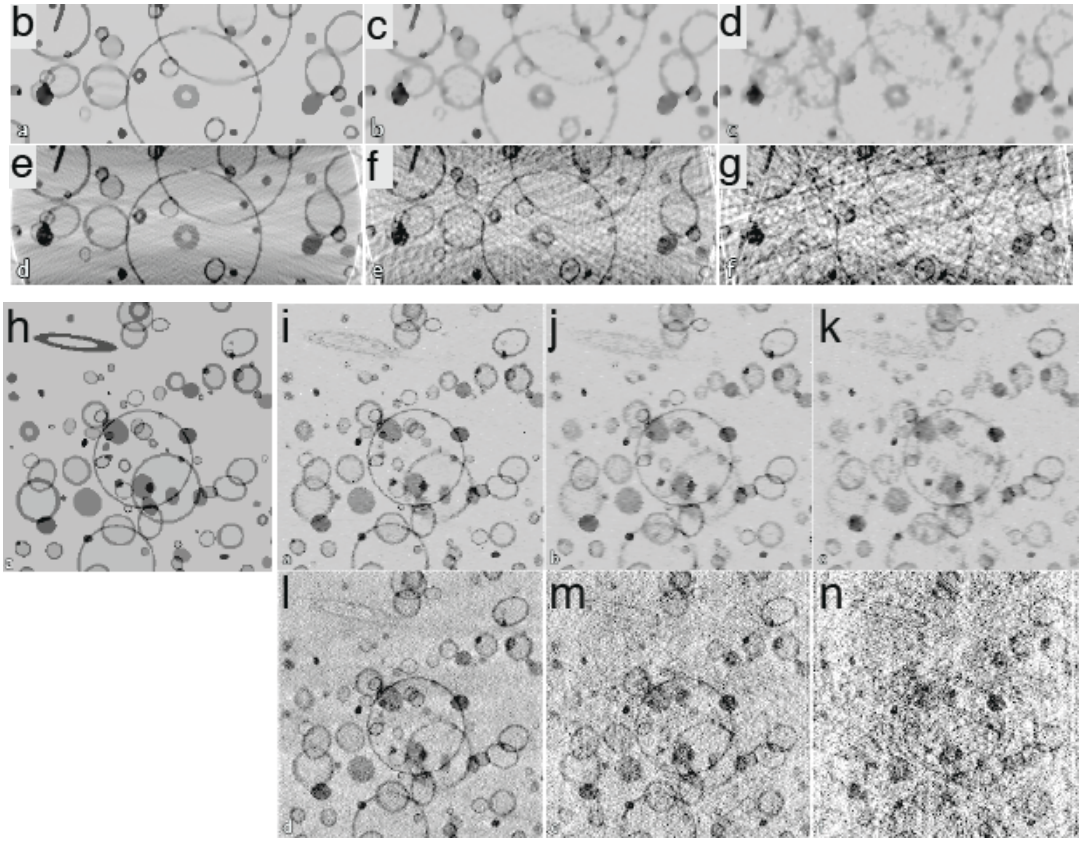

Figure 3.6: **(b)** CS-ET reconstruction of $x - z$ slice from simulated projections at $\pm 70°$ with $2°$ angular increment, **(c)** CS-ET reconstruction of $x - z$ slice with 3x undersampling of tilt angles, **(d)** CS-ET reconstruction of $x - z$ slice with 6x undersampling of tilt angles, **(e)** WBP reconstruction of $x - z$ slice with all tilt angles, **(f)** WBP reconstruction of $x - z$ slice with 3x undersampling of tilt angles, **(g)** WBP reconstruction of $x - z$ slice with 6x undersampling of tilt angles, **(h)** Slice through phantom in $x - y$ plane, **(i)** CS-ET reconstruction $x - y$ orthoslice with all tilt angles, **(j)** CS-ET reconstruction $x - y$ orthoslice with 3x undersampling of tilt angles, **(k)** CS-ET reconstruction $x - y$ orthoslice with 6x undersampling of tilt angles, **(l)** WBP reconstruction $x - y$ orthoslice with all tilt angles, **(m)** WBP reconstruction $x - y$ orthoslice with 3x undersampling of tilt angles, **(n)** WBP reconstruction $x - y$ orthoslice with 6x undersampling of tilt angles.

2D slices of this phantom are shown along the $x - z$ axis in Figure 3.1 and along the $x - y$ axis in Figure 3.6(h). Simulated projections were acquired independently for each $x - z$ slice and consisted of Radon transform data over a tilt range of $\pm 70°$ with an angular increment of $2°$. As described in the Methods, Poisson

noise simulating typical high-dose STEM measurement conditions was added to the projections. The projections were then further corrupted with Gaussian noise. The membrane phantom $x-z$ reconstructions were performed with noisy 1x, 3x, and 6x undersampling of the tilt angles, using CS-ET (Figure 3.6(i-k)) and WBP (Figure 3.6(l-n)). See Figure 3.7 for results from noiseless membrane phantom reconstructions.

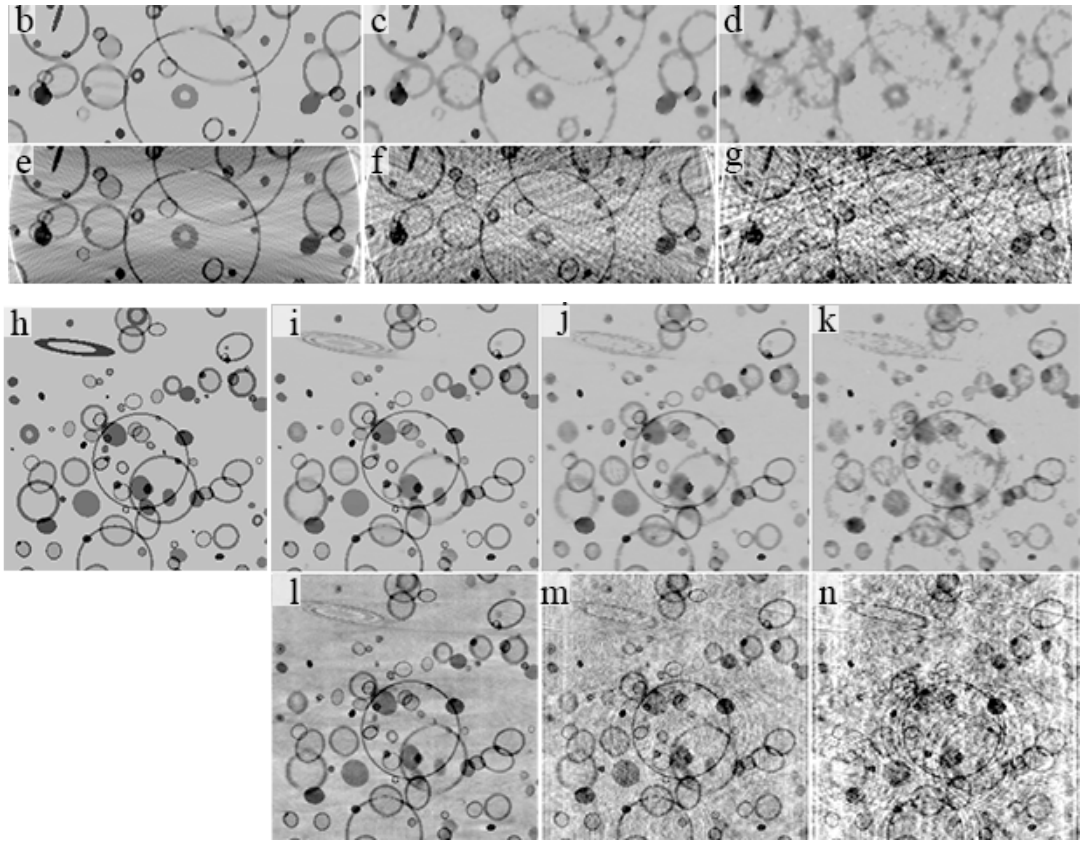

Figure 3.7: **CS-ET and WBP reconstructions from noiseless membrane phantom projections**. The figure layout is identical to Figure 3.6, but all reconstructions are shown from noiseless projection data. As expected, CS-ET substantially outperforms WBP and reconstruction quality is higher than that from noisy data.

It is evident that the performance of CS-ET exceeds that of WBP for the fully-sampled and undersampled tilt series, but the advantage of CS is greatest for the

6x undersampled tilt series. Figure 3.8 corroborates this observation with RMSE calculations for both reconstruction methods.

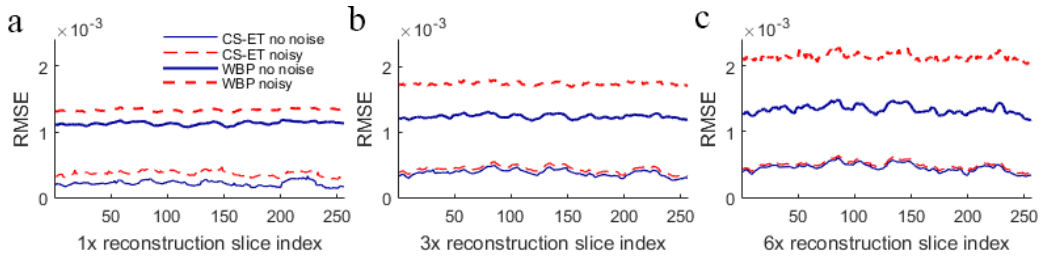

Figure 3.8: **Membrane phantom reconstruction $x-z$ slice RMSEs**. RMSE values are displayed for CS-ET and WBP reconstructions from noiseless and noisy **(a)** fully sampled tilt series **(b)** 3x undersampled tilt series, and **(c)** 6x undersampled tilt series. Note that CS-ET reconstructions consistently outperform their WBP counterparts in RMSE value. Moreover, CS-ET reconstruction volumes are more robust to the addition of Gaussian noise, with smaller relative and absolute increases in error values than those observed with WBP.

These results are in agreement with known results for CS recovery of sparse signals from undersampled data. However, the advantage of CS relative to WBP is not as significant for the membrane phantoms as it is for the nanoparticle phantom (Figure 3.5). We attempt to improve these results by creating projections at randomly-sampled angles over a tilt range of $\pm 70°$, but our numerical results shown in Figure 3.9 indicate that this sampling strategy performs worse than traditional uniform sampling with high probability.

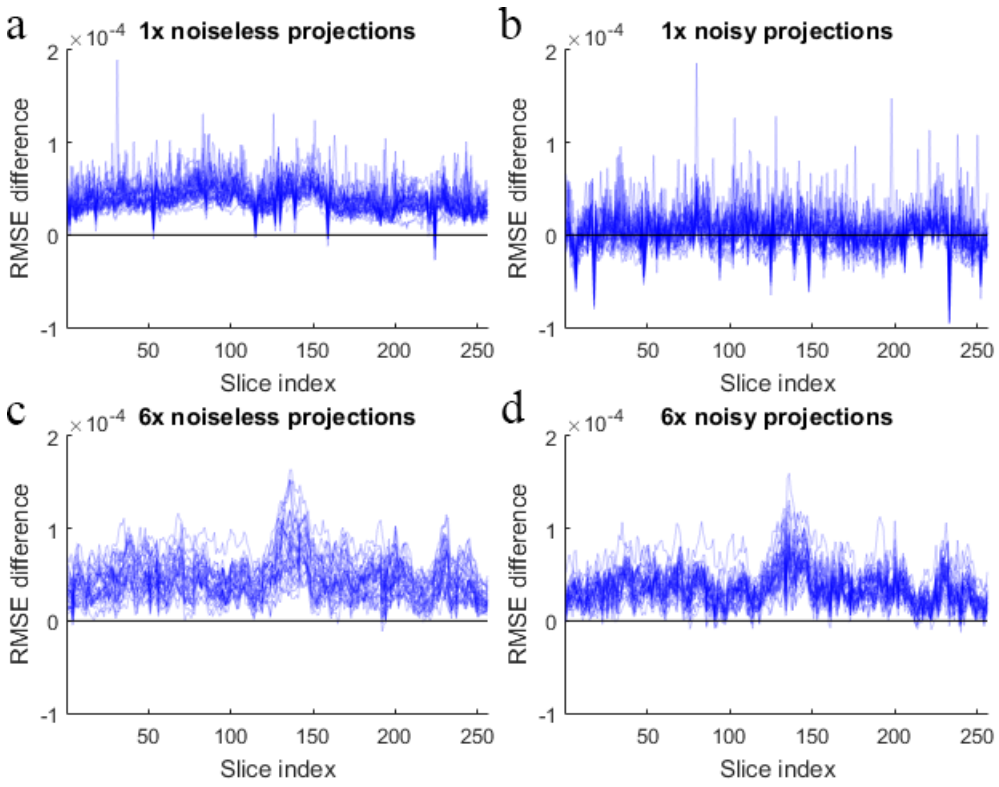

Figure 3.9: **Comparison of reconstruction errors from simulated STEM tomographic data based on uniform versus random sampling of tilt angles**. A simple variant of traditional, uniformly-sampled tomographic measurement is to choose tilt angles at random within the mechanically-feasible range. Numerically, this approach performs worse than uniform sampling for the membrane phantom. For the three membrane phantom reconstructions from fully sampled (1x), 3x undersampled, and 6x undersampled tilt series, 30 random reconstructions were performed using the same number of tilts at angles chosen uniformly at random between $-70°$ and $70°$. This procedure was repeated using both noiseless and noisy projection data. The difference between 30 random reconstruction RMSEs and uniform reconstruction RMSEs for the 1x and 6x trials are plotted above for each $x-z$ reconstruction slice, indexed along the x axis. A positive RMSE difference indicates that random sampling performed worse than uniform sampling.

Since CS performance is dependent on signal sparsity, we quantify this phenomenon by comparing the sparsity of the nanoparticle and membrane phantoms in the TV, identity, and wavelet domains in Figure 3.10, which shows that the membrane phantom is substantially less sparse than the nanoparticle phantom.

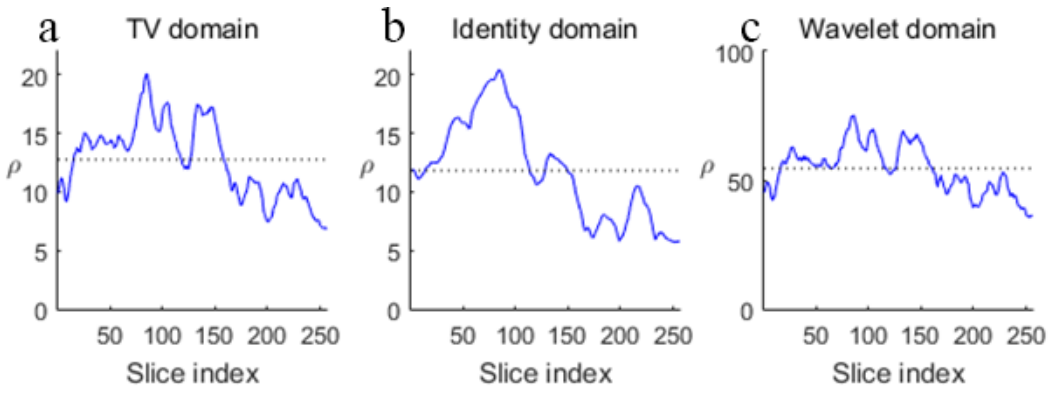

Figure 3.10: **Sparsity comparison for the membrane and nanoparticle phantoms**. The 1%-compressibility ratios of each of the 256 $x-z$ slices of the membrane phantom are calculated in each transform domain and plotted as a multiple ρ of the nanoparticle phantom's 1%-compressibility ratio. **(a)** ρ in the TV domain, with average equal to 12.7, **(b)** ρ in the identity domain, with average equal to 11.8, **(c)** ρ in the DB8 wavelet domain with average equal to 54.4. Note that the membrane phantom, significantly more simple in structure than experimental biological datasets, is markedly less compressible than the nanoparticle phantom.

3.4.2 CS-ET reconstructions of experimental tomographic tilt series

To test the CS-ET reconstruction algorithm on experimental data, we recorded STEM tomographic tilt series from insulin-secreting beta cells in a specimen of isolated mouse pancreatic islets of Langerhans, which was conventionally prepared by fixation with osmium tetroxide, dehydration, embedding and staining with uranium and lead [14, 100]. Cellular ultrastructure prepared in this way is rich in stained membranes, which are visible as dark lines in bright-field STEM images or bright lines in dark-field STEM images. Islet beta cells contain many internal structures of this type, including outer membranes of secretory granules, nuclear inner and outer membranes, mitochondria, and endoplasmic reticulum. In addition, there are many other structures such as ribosomes, cytoskeletal filaments, and other macromolecu-

lar assemblies. The sample represents a complex biological object on which different reconstruction techniques can be tested. Single-axis tilt series were reconstructed as 1024 independent $x - z$ slices, perpendicular to the y-axis, each containing images of size 1024×100 pixels, using CS-ET, SIRT, and WBP for fully sampled tilt series as well as 3x and 6x undersampled tilt series.

The reconstructions of bright-field STEM and dark-field STEM tomographic tilt series in Figure 3.11(a) and Figure 3.11(b), respectively, indicate that there is a significant reduction of noise in CS-ET and SIRT reconstructions relative to WBP. The distinction between CS-ET and SIRT reconstructions is less pronounced, though SIRT structures overall exhibit some blurring that is less apparent in CS-ET results.

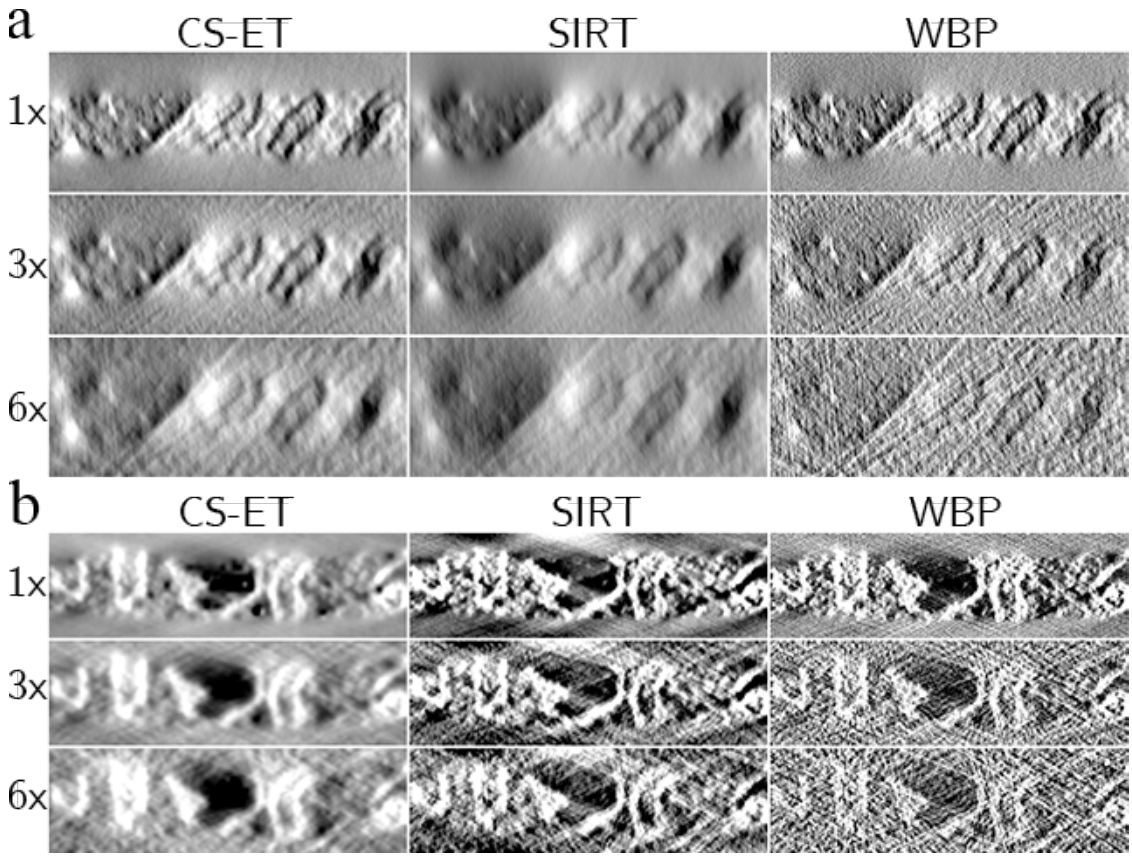

Figure 3.11: **Comparison of CS-ET, SIRT, and WBP reconstruction $x-z$ slices through mouse beta cell tissue**. **a** Reconstructions from high-dose BF STEM tilt series, **(b)** reconstructions from high-dose DF STEM tilt series.

Next, we compare CS-ET, SIRT, and WBP reconstructions from dark-field STEM tilt series recorded at low electron dose. The eight $x-z$ slices in Figure 3.12 indicate that both CS-ET and SIRT perform substantially better than WBP, and again CS-ET structural boundaries appear less diffuse than their SIRT counterparts.

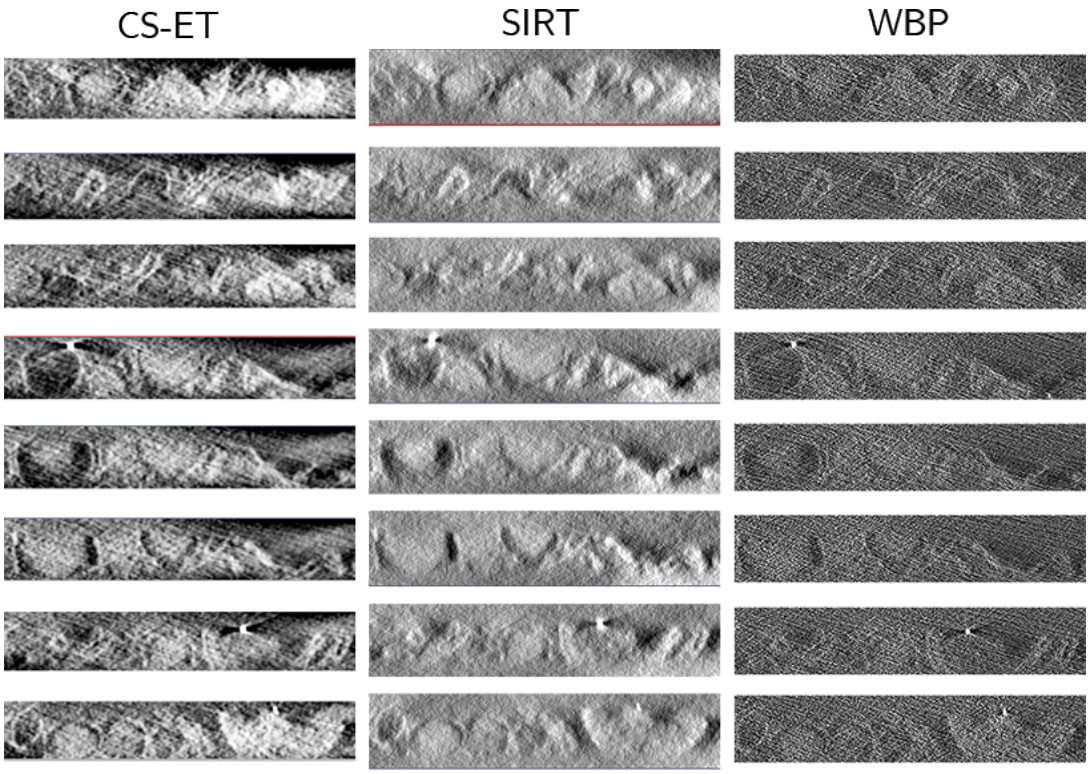

Figure 3.12: **Comparison of multiple CS-ET, SIRT and WBP reconstruction $x-z$ slices through mouse beta cell tissue**. This figure demonstrates multiple $x-z$ reconstruction slices from the high-dose DF dataset, to give a better sense of the differences in reconstruction quality between the three reconstruction methods.

In Figure 3.13, we compare CS-ET, SIRT, and WBP reconstruction $x-y$ slices from fully sampled, 3x, and 6x dark-field STEM tilt series recorded at low electron dose.

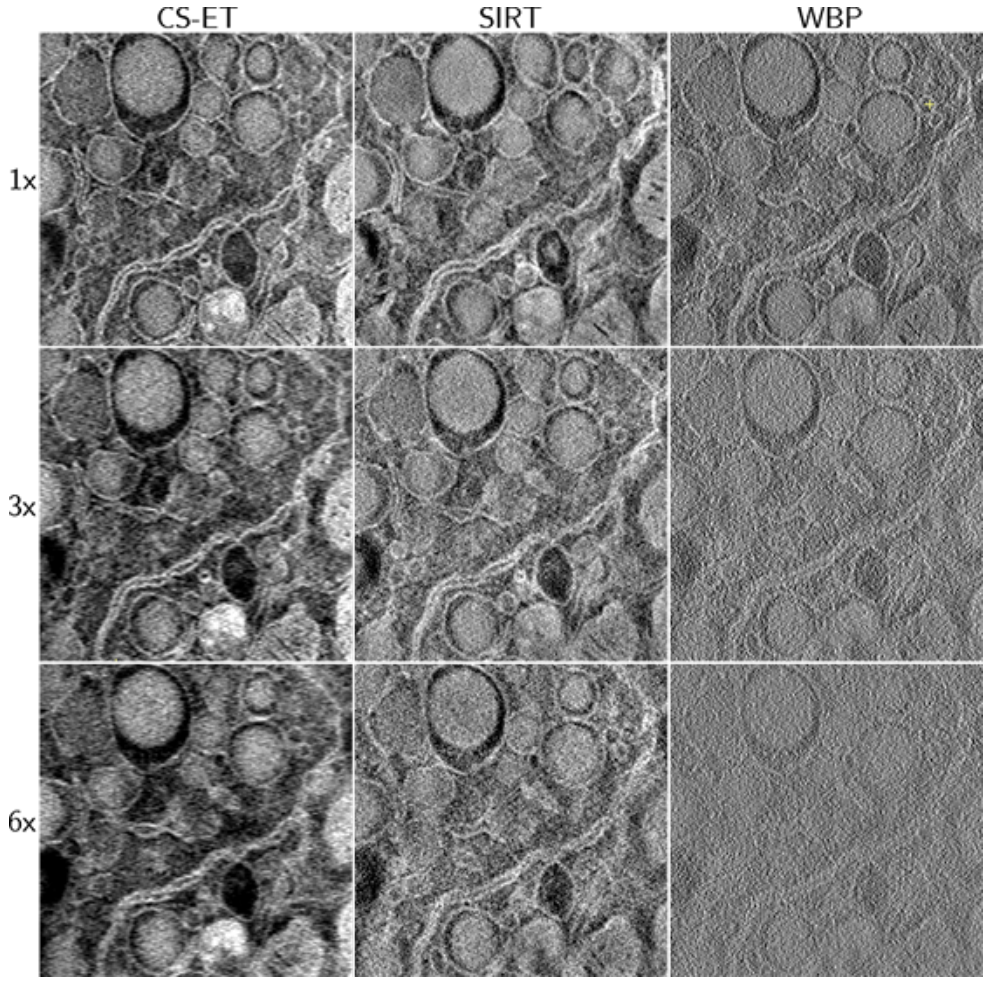

Figure 3.13: **Comparison of CS-ET, SIRT, and WBP reconstruction algorithms from dark-field STEM tilt series of mouse beta cells recorded at low electron dose**. **(a)** Reconstructed $x - z$ slices, **(b)** reconstructed $x - y$ slices from fully sampled, 3x undersampled, and 6x undersampled tilt series. Note that due to differing reconstruction procedures, a small discrepancy exists in the z coordinates of the CS-ET and SIRT $x - y$ slices, manifesting as small structural differences between the images.

Additional $x-y$ reconstruction slices from the high-dose BF and DF tilt series can be seen in Figures 3.14 and 3.15.

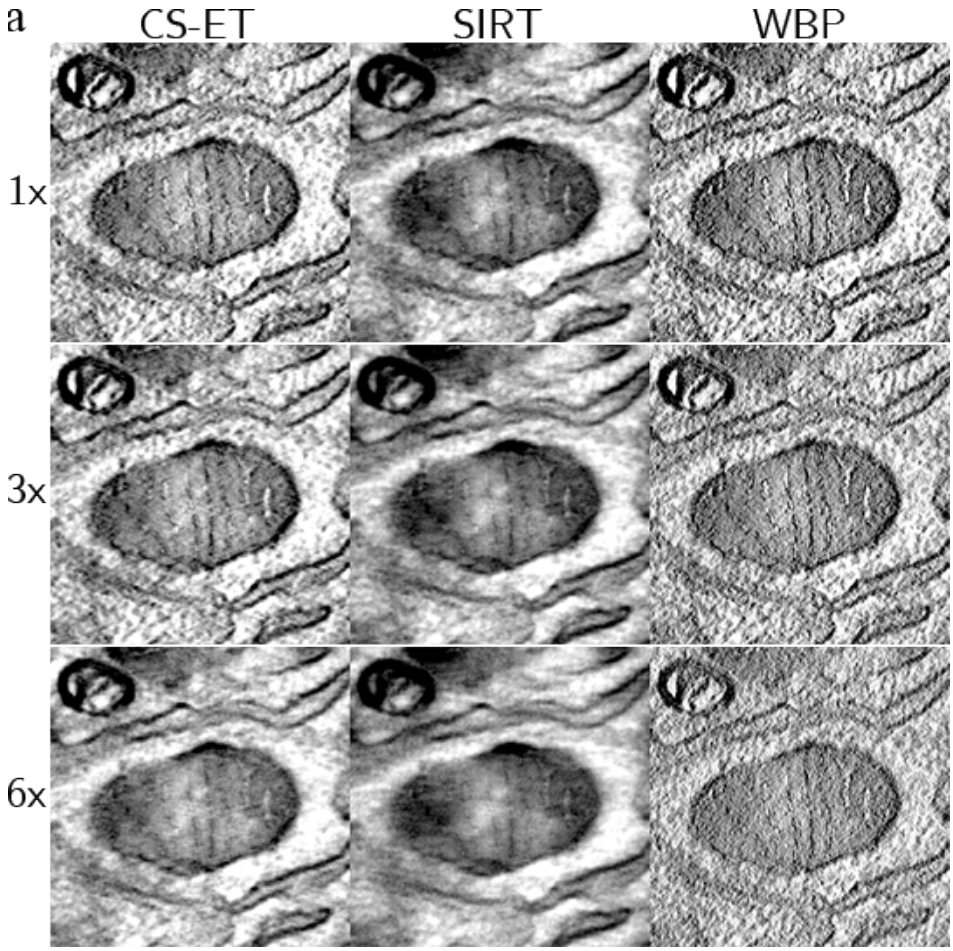

Figure 3.14: **Reconstruction $x - y$ slices from high-dose bright field tilt series**. For all three algorithms, the 3D tomogram volume is assembled as a series of independent 2D reconstructions in the $x - z$ plane. The orthogonal $x - y$ planes provide the highest resolution and are the most common human vantage point for the analysis of 3D ET data. Each row of the $x - y$ reconstruction view comes from a separate image reconstruction. This figure demonstrates $x - y$ reconstruction slice from the high-dose BF tilt series from 1x, 3x, and 6x undersampled data.

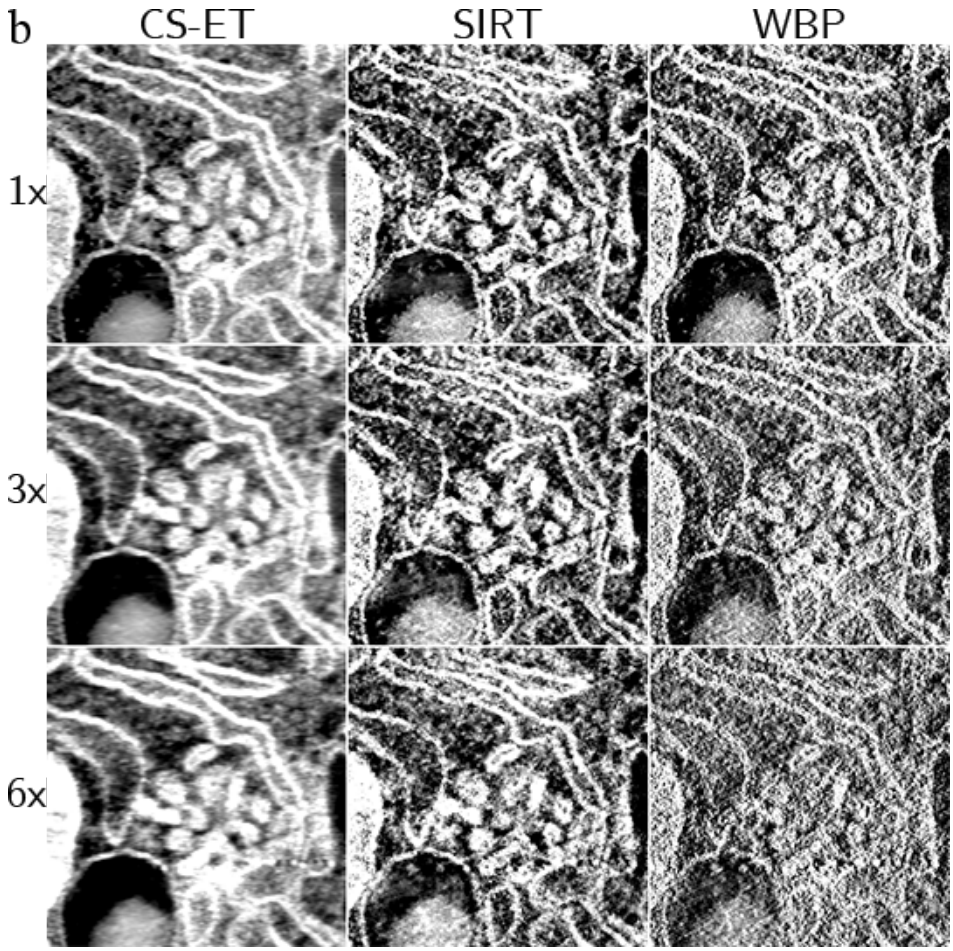

Figure 3.15: **Reconstruction $x-y$ slices from high-dose dark field tilt series**. This figure has the same layout as the previous figure. Visualized are $x-y$ reconstruction slices from the high-dose DF tilt series from 1x, 3x, and 6x sampled data.

From the $x-y$ perspective, the advantages of the CS-ET algorithm are more apparent. The WBP reconstruction becomes heavily degraded by noise with increased undersampling. A significant portion of this noise is absent from the SIRT reconstructions, but CS-ET exhibits still lower noise levels in the 3x and 6x undersampled reconstructions.

3.4.3 Comparisons between experimental and phantom sparsity

We expect that the advantage of CS methods for reconstructing sparse objects in electron tomography should manifest in a correlation between the root mean square error (RMSE) and the compressibility of each of the different $x - z$ planes in the reconstruction. We measure compressibility as $p\%$ compressibility ratios as defined in Section 2.2. The more compressible an image is, the smaller the compressibility ratio. This relationship is shown in Figure 3.16 for the membrane phantom, calculating a 5% compressibility ratio in the three sparsity domains analyzed in this study (TV, identity, wavelet). In all three transform domains, there is a clear positive correlation between reconstruction error and compressibility ratio for fully-sampled and 6x undersampled tilt series. This demonstrates the importance of signal sparsity for CS reconstruction methods.

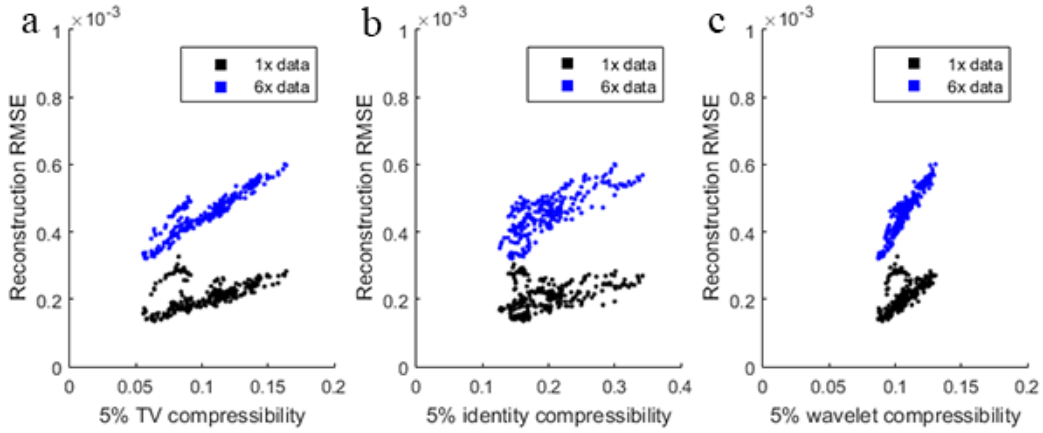

Figure 3.16: **Correlation between RMS error and compressibility ratio for CS-ET reconstructions of the membrane phantom**. The n%-compressibility ratio of an image is defined as the proportion of pixels with magnitude larger than n% of the maximum-magnitude pixel's magnitude. This can be calculated in any transform domain by applying the same procedure to the transformed image. The higher an image's compressibility ratio, the less compressible the image is. Plots show root mean square error (RMSE) versus 5%-compressibility ratio for **(a)** *TV* transform, **(b)** identity transform, and **(c)** DB8 wavelet transform.

Our lack of prior knowledge about the experimental data structure complicates any analysis of the dataset's compressibility in any of the transform domains. Nevertheless, we can estimate this by calculating the compressibility of a high-quality reconstruction obtained from a fully sampled tilt series across several relative compressibility thersholds. This is shown in Figure 3.17, which displays $\boldsymbol{\rho}$: the 2.5% compressibility ratios of the high-dose bright-field and dark-field datasets as a multiple of the 2.5% compressibility ratios of the membrane phantom.

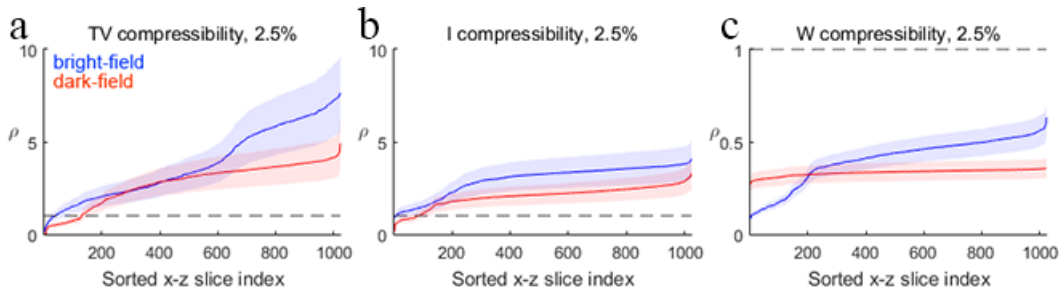

Figure 3.17: **Compressibility of CS-ET reconstructions relative to the membrane phantom.** 2.5% compressibility ratios of the high-dose bright-field and dark-field reconstructions' $x-z$ slices are expressed as a multiple of the 2.5% compressibility ratios of the membrane phantom; we denote this multiple as ρ. For each experimental $x-z$ slice, a distribution of ρ values are obtained for each of the 256 membrane phantom $x-z$ slices. The mean of ρ plus or minus one sample standard deviation is then plotted in each transform domain. Larger ρ values indicate a less sparse experimental $x-z$ slice, and the minimum number of measurements required to reconstruct an image scales linearly with ρ. A dashed line at $\rho = 1$ is included in each figure as a point of reference. **(a)** ρ values in the TV domain, **(b)** ρ values in the identity domain, **(c)** ρ values in the DB8 wavelet domain.

For each of the 1024 $x-z$ slices of each experimental reconstruction, a distribution of ρ values is calculated using each of the 256 membrane phantom $x-z$ slices; for each transform domain the mean of these distributions plus or minus one standard deviation are plotted in Figure 5a-c. Both experimental datasets are substantially less compressible than the membrane phantom in the TV and identity domains, while in the wavelet domain the opposite is true. This is likely due to the presence of sharp boundaries in the membrane phantom, leading to the slow decay of its wavelet coefficients. Regardless, while a wavelet regularization term does improve the experimental data reconstructions, even a large value is insufficient for regaining the CS-ET performance improvements exhibited in the phantom studies.

3.5 Discussion

Previous work has established that, at least in some imaging applications, CS techniques allow for high-quality recovery of tomograms from highly-undersampled data [73, 94]. Our work corroborates these results, and provides evidence that the quality of CS-ET matches or exceeds that of other recovery techniques for undersampled recovery of more complex images. However, the results also challenge the notion that undersampling is necessarily an attractive goal for CS techniques applied to STEM tomography of biological specimens. Figures 3 and 4 show that even moderate undersampling creates structurally-significant degradation of reconstruction resolution, likely a result of the significantly decreased compressibility of the biological data. A microscopist not constrained by electron dosage therefore has no incentive to perform undersampled measurements. When evaluating each algorithm's performance on fully sampled high-dose biological tilt series, the resolutions of all three reconstruction volumes are closely-matched. See Figures 3.4 and 3.14 for additional $x - y$ views of fully sampled reconstruction volumes.

Instead the CS-ET algorithm as a regularized reconstruction method appears to be more useful for improving reconstructions of noisy images in lower dosage conditions, a context in which the closely-matched CS-ET and SIRT reconstruction volumes provide marked improvement over WBP. The performance of CS-ET in undersampled noisy conditions suggests future utility in dosage-constrained imaging conditions.

Moreover, randomized sampling of projections in the feasible tilt angle range

does not appear to provide a superior randomized alternative to the traditional tilt series acquisition process. Open questions remain in the mathematical underpinnings of the recovery of sparse images from tomographic tilt series of either type, but this finding is contrary to the prevailing trends in existing CS application results, wherein random measurement ensembles are found to recovery sparse images with high probability. Therefore, a more careful exploration of this randomized measurement variant by acquiring new randomized STEM tilt series does not seem warranted, especially when weighed against investigating other developments. The application of the CS-ET algorithm to fully sampled, uniform tilt series is thus not compressed sensing in the strictest definition of the term, and is better understood as a $TV-$ and ℓ^1-regularized least-squares recovery algorithm.

Regardless of the terminology used, the results of this chapter show that optimization-based regularized image recovery algorithms taking advantage of prior structural knowledge may be feasibly applied to experimental STEM datasets from biological specimens, and offer performance competitive with or exceeding other common reconstruction algorithms.

This chapter focused on the recovery of TV-, identity-, and wavelet-sparse tomograms, but there are many natural variations based on alternative sparsity models which might be relevant to STEM imaging. Wavelet bases and frames have long been known to yield sparse representations of natural images, and more recent developments in sparse imaging via anisotropic frame representations (e.g., curvelets [16, 38], shearlets [69, 29]) may prove well-suited to CS-ET. Data-driven representation methods are another possible avenue of development, tailored to dataset struc-

ture and learned during the reconstruction procedure. Xu et al. (2012) demonstrate a similar idea for CT, while Gopinath et al. (2012) discuss a shape-based regularization procedure tailored to 3D tomogram structure.

Rather than decomposing the 3D reconstruction problem into 2D slices, the CS reconstruction theory is equally applicable to the recovery of 3D regions considered as a single domain. Currently, the memory requirements for large volumes prevent the creation of efficient fully-3D reconstruction algorithms, but that may change with the advance of computing technology. In terms of reconstruction quality, fully-3D sparse reconstructions are likely to be superior to their 2D equivalents, due to the added structural regularity information supplied by the spatial relationships between the $x - z$ slices.

Beyond these implementation details, the analysis in this chapter confronts interesting and challenging open questions in the theory of compressed sensing. The recovery of sparse images from non-RIP measurement matrices, such as those created by tomographic sampling, is still poorly-understood. Further research into understanding its empirically-observed efficacy would shed valuable light on interpreting experimental measurement procedures from many imaging domains in the CS framework.

Chapter 4: Sparse Olfactory Coding in the Locust

4.1 Overview

The contributions to signal processing and machine learning heralded by compressed sensing and the wider body of sparse signal representation research are remarkable. These concepts play an important role in a growing wave of technological advancement which promises to make our creations more intelligent, more responsive, and more productive. But our progress in creating a digital understanding of the world around us is still dwarfed by nature's achievement of intelligence in animal life. Even simple insects possess a brain which is capable of extracting complex information from its environment, and integrating that information with instinct and memory to execute the control processes governing movement and behavior. For most of history, the small size and complexity of the constituent elements of nervous systems have prevented us from understanding the physical basis by which natural intelligence is effected, and much remains unknown today. Still, we now possess the capability to measure neural activity at single-neuron resolution, across ever-growing spatial extents. *Tetrode recorders*, specialized microelectrodes inserted into neural tissue, can record the electrical activity of neurons within a radius of roughly 100-200μm [13]. Simultaneous recording of this electrical activity

from multiple locations allows the signal to be sorted into contributions from distinct neurons [13]. These recordings provide a window into how neural information is processed within interconnected clusters of neurons, and are vital to discovering the fundamental governing principles for biological neural networks.

The **neural sparse coding hypothesis** aims to describe one such governing principle. This hypothesis concerns itself with the networks of neurons devoted to the processing of sensory information, and states that populations of these neurons use an information encoding strategy whereby only a small percentage of the neurons are active at any point in time [87]. Support for this hypothesis as a general organizational strategy for at least certain portions of sensory processing networks has emerged from observations that, as one moves through the hierarchy of sensory processing centers, from the input receptors deeper into the brain, neural activity levels decrease [5]. To date, this pattern has been observed in vision, audition, and olfaction across animals at every level of complexity, from the simple nematode *Caenorhabditis elegans* [113] to insects and mammals [87]. In addition to lowering energy expenditure by the brain, sparse sensory codes create fewer collisions between different information states, a boon to associative memory processes. By decomposing environmental information into a small number of features, sparse sensory codes may also learn features which explicitly take into account environmental statistics [87]. It is also a predictive hypothesis; a highlight of this line of inquiry was the work of Olshausen and Field in the mid 90's, wherein they demonstrated that a population of simulated neurons trained to produce a sparse encoding of natural images learned receptive fields highly similar to the ones found in neurons in V1, a

visual processing area of the mammalian brain [85, 86].

The "codes" created by neural activity, carried by the membrane potential of participating neurons, differs in important ways from the signal decompositions via linear algebra performed in signal processing as in Chapter 2. Neural membrane potentials may be modelled computationally in multiple ways, depending on desired spatial resolution and population size. In this chapter I make use of Hodgkin-Huxley-type ODE models, which consider each neuron and synapse in a network as a single-point entity with a collection of time-varying states. The differential equations describing the membrane potential are highly nonlinear. Still, both encodings offer similar benefits to the systems which employ them - increase in coding capacity, better extraction of features of interest from signal data - but in order to understand how far this analogy can be carried, it is crucial to better understand the nature of neural activity dynamics in biology.

One important model biological system for understanding neural sparse coding is insect olfactory processing. Many features of olfactory processing networks are highly similar across insects, and even between invertebrates and vertebrates [42], but the olfactory system of the locust has emerged as a commonly-used model. This system is the focus of the work done in this chapter. After a brief introduction to neurons and neural networks in general, I explain a derivation of the Hodgkin-Huxley model of neuron activity, and give an overview of the aspects of locust neuroanatomy most relevant to my own research. Working with Dr. Mark Stopfer's Section on Sensory Coding and Neural Ensembles in the NICHD, I implemented and adapted a computational model of the Kenyon cells, the neurons responsible for

sparse olfactory coding in locusts. This model was used as part of a larger project to better understand the network structures responsible for the temporal patterning of Kenyon cell activity. In order for the computational model to be useful for that project, it is important that the simulated cells mimic their biological counterparts as closely as possible. To that end, we have worked to establish biological benchmarks by which the model's output may be judged. In addition to this detailed computational modeling, I created a simplified combinatorial probabilistic model of the activity of the networks responsible for producing sparse olfactory codes, guided by the lessons learned from the more granular computational modeling. In this chapter, I demonstrate that this model creates a framework which can explain important properties of the sparse olfactory codes observed in locusts.

4.2 Neurons and biological neural networks

4.2.1 The neuron, in brief

The research in this chapter requires a basic understanding of neuroanatomy and electrophysiology, and this is the goal of Sections 4.2-4.5. Much of this material is standard in the field, and my exposition is derived primarily from the excellent textbook *Theoretical Neuroscience*, by Dayan and Abbott [36], unless noted otherwise.

The neuron is a specialized cell found within the nervous system of most animals, which generates electrical pulses to rapidly communicate information throughout the body. These signals are generated in response to input from the external

environment or other neurons, and take the form of time-varying changes in the electrical potential of the neuron's interior relative to the exterior cellular environment, also called the **membrane potential**. Clusters of interconnected neurons govern the behavior of animals, and the processes of learning and memory are assumed to be governed by the changing of interconnection strengths over time.

The structure of a neuron can vary widely, within one organism and across multiple species. The locust neurons of interest in this chapter are all termed *pseudounipolar neurons*, see Figure 4.1 for a diagram.

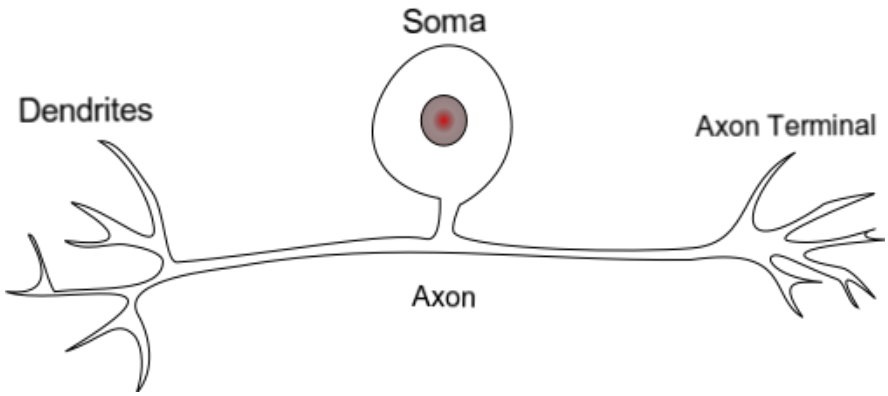

Figure 4.1: A diagram of a pseudounipolar neuron. Neural input arrives at the dendrites, while any generated output travels along the axon to the axon terminals. All three structures exist as part of a single process extending from the cell body (soma), which contains the cell nucleus.
Adapted from original work by Juoj8, CC BY-SA 3.0 licensed, https://commons.wikimedia.org/w/index.php?curid=17204559

For these neurons, the dendrites and axon terminals are connected by an axon which forms from a single process extending from the cell body. **Dendrites** receive input from upstream neurons across junctions called **synapses**. The activity of the synapses creates small changes in the membrane potential at the dendrites of the neuron, and these *postsynaptic potentials* sum and propagate towards the axon.

The potentials may be **inhibitory** - hyperpolarizing (decreasing) the membrane potential of the postsynaptic neuron, or **excitatory** - depolarizing (increasing) the postsynaptic neuron's membrane potential. If the summed synaptic inputs depolarize the membrane sufficiently, a feedback process generates an **action potential**, or spike. An action potential is a brief but rapid increase in the cell's membrane potential that persists for roughly 1 ms. The action potential travels along the **axon** to the axon terminals. At the **axon terminals**, the neuron forms synapses with dendrites of downstream neurons. Action potentials arriving at axon terminals most commonly trigger the release of chemicals across synaptic junctions, but may also form direct electrical connections with postsynaptic neurons.

Not all neurons receive input exclusively from other neurons. Some, the **sense receptors**, receive input from bodily structures outside the nervous system, or interactions between an organism and its exterior environment. These receptor cells are ultimately responsible for the senses of touch, taste, sight, hearing, and olfaction, in addition to several others. Each type of receptor cell possesses special structures for transducing environmental stimuli into the electrical signals used by neurons.

4.2.2 Neuron electrochemistry

The membrane potential of a neuron is governed by the relative concentration of ions between the cell's interior and the extracellular medium, most commonly Na^+, K^+, Ca^{2+}, and Cl^-. Specialized sodium-potassium *ion pumps* across the cell membrane actively transport Na^+ ions out of the cell and K^+ ions into the

cell. Without additional modulation, these pumps maintain an *equilibrium (resting) membrane potential* of approximately -70 mV relative to the extracellular medium, indicating an excess of negative charge within the cell and positive charge external to the cell. **Ion channels** across the cell membrane, when open, selectively permit different ion types to passively enter and exit the cell as dictated by the voltage and concentration gradients these pumps create. The control of the opening and closing of ion channels forms the basis for changing the neuron's membrane potential over time. Whether a class of ion channel is open or closed depends primarily on the cell's membrane voltage, but can also respond to other cellular signals (e.g. the presence of Ca^{2+} ions). Ion channel activity may also be controlled within the dendrites by the release of specialized chemicals called *neurotransmitters* from upstream synapses, as well as the action of special environmentally-sensitive cellular structures in receptor neurons.

4.2.3 Adaptation and learning in biological neural networks

A **biological neural network** is a collection of neurons which influence each other via synaptic interconnections. Taken as a whole, the nervous system of an animal may be considered as a neural network. However this network exhibits a modular structure, in which clusters of interconnected neurons share a functional role and share input and output targets within the brain at large. Therefore, efforts to understand the structure and function of the brain largely focus on delineating its regions into these network clusters, investigating the structure and function of each

such network, and understanding how each network's processing operations relate to the activity of the its upstream input network and downstream output network.

A fundamental feature of neural networks is their ability to alter their operation in response to environmental information. Many forms of **learning** are rooted in structural changes to synaptic connection strengths, which may increase or decrease the efficacy of a presynaptic neuron's action potential in eliciting changes in postsynaptic neuron membrane potentials. These changes may persist across a variety of time scales, from seconds to days or even longer, depending on the underlying synaptic mechanisms at hand. A classical example of this nature is *Hebb's rule*, which states that the strength of a connection between two neurons increases when both neurons fire within a short time interval. An extension of this rule proposes that synaptic strengths weaken when the firing of the neurons is not contiguous in time. Evidence from experiments exists for synaptic processes obeying this hypothesis in many parts of the nervous system.

In addition to dynamic processes governing synaptic strength, individual neurons themselves undergo adaptive processes in response to spiking activity. An example of this is given in Section 4.6.1 of this chapter, in which locust olfactory receptor cell firing rates decrease during prolonged repetitive stimulation. The combination of adaptive processes occurring within cells and adaptive and facilitatory processes occurring within synapses gives the brain its remarkable capacity for understanding and reacting to a changing environment over time.

4.3 Modeling neurons

As a biological system, neurons exhibit complex structure from the atomic level up to macroscopic scales across the entire nervous system. At the finest scales, one may model each neuron in a network as a collection of spatially-distributed compartments or a spatial continuum. At the coarsest scales, entire groups of neurons may be abstracted into single functional units. Each model scale comes with its own advantages and drawbacks. Fine-scale models accurately reproduce a wide range of individual neuron functions, but are computationally inefficient for modeling large collections of neurons and their large numbers of model parameters may be difficult to fit to limited data from experiments. Coarse-scale models allow for the modeling of large brain regions or entire brains, but may fail to capture the full range of computational activities performed by its constituent neural clusters. A common class of models for investigating the single-neuron-level activity for collections of thousands or millions of neurons are the **single-compartment models**, which treat each neuron and each synapse as single spatial points, each with associated collections of state variables. My research makes use of the *Hodgkin-Huxley-type* model, a single-compartment model offering an effective compromise between biophysical fidelity and computational tractability. A derivation of this model is the first objective of this section, followed by a derivation of a synaptic model used in tandem.

4.3.1 Modeling ion flow

A neuron whose membrane potential does not vary significantly across its spatial extent is termed *electrotonically compact*. Neurons approximating this characteristic are most amenable to accurate single-compartment modelling. In such models, the cell membrane acts as a single capacitor, separating the internal negative charge from the external positive charge and creating a *capacitance* constant C_m for the cell. The membrane potential V, capacitance C_m, and excess charge within the cell Q are then related by the equation $Q = C_m V$, and therefore the change in voltage over time is related to the change in charge over time as

$$C_m \frac{dV}{dt} = \frac{dQ}{dt}. \tag{4.1}$$

The $\frac{dQ}{dt}$ term is equal to the current passing into the cell, and the Hodgkin-Huxley model (among others) is created by constructing an appropriate model of those currents.

Two forces govern the flow of ions across neural membranes, and therefore govern the ionic currents a neuron experiences. Electrical forces drive ions away from regions of like charge and towards regions with an excess of the opposite charge, while diffusion drives molecules (charged or otherwise) from regions of high concentration to low concentration. Therefore for each type of ion channel, indexed by j, that a neuron possesses, there is a membrane **reversal potential** E_j at which the opposing diffusive and electrical forces are balanced, and the net ionic flow into the cell is 0.

When the membrane potential differs from this value, a nonzero flow of ions into or out from the cell then creates an ionic current I_j, equal to -1 times the voltage difference (*driving force*) $(V - E_j)$ multiplied by a conductance constant g_j controlled by the distribution of type-j ion channels across the membrane: $I_j = g_i(V - E_j)$. The basic single-compartment model form is thus

$$C_m \frac{dV}{dt} = -\sum_j g_j(V - E_j). \tag{4.2}$$

Additionally, an external current I_e is typically included in simulations, allowing the user to stimulate the neuron as with an electrode.

4.3.2 Modeling channel conductance

Individual ion channels toggle between open and closed states in a stochastic manner which may be dependent on membrane potential, or the presence of Ca^{2+} ions or other ligands. The following exposition develops conductance models under the assumption of voltage dependence, but Ca^{2+}-dependent models may be developed in a similar way with Ca^{2+} concentration replacing membrane potential as an independent variable. Due to the large number of each type of ion channel across a neuron membrane and the essentially-independent operation of the channels, their aggregate activity may be accurately modeled deterministically as $g_j(t) = \bar{g}_j P_j(t)$, where $\bar{g}_j$ is a constant *maximal conductance* for the channel type and $P_j(t)$ is the proportion of open channels at time t. There are several mechanisms which may govern the evolution of these $P_j(t)$ functions; the ones relevant to my work here are

constant, persistent, and transient conductances.

4.3.2.1 Constant conductance

For the simplest conductance models, $P_j(t) = 1$ for all t, and $g_j = \bar{g}_j$. This type of conductance accurately models the activity of sodium-potassium ion pumps, for instance. The net effect of constant conductance flows across a neuron membrane are typically aggregated into a single **leakage current**,

$$I_L = g_L(V - E_L). \tag{4.3}$$

In practice, the leakage current term also compensates for small model errors, and g_L and E_L are treated as free parameters which may be tuned to match model behavior to biological data after other, more principled parameter choices are made.

4.3.2.2 Persistent conductance

An ion channel exhibiting persistent conductance behaves as if its state is controlled by a single type of gate, with its probability of being open increasing as membrane potential increases. More generally, it is often convenient to understand such conductances as being gated by k identical independent subunits, so that the activation probability takes the form $P_j = n_j^k$ for some activation variable n_j denoting the probability of a single gate being open. While this k value may be determined in principle from the ion channel structure, in practice it is fit to data from experiments.

The single-gate transition probability n_j is modeled such that the closed $\rightarrow$ open transition occurs at a voltage-dependent rate $\alpha_n(V)$ and the open $\rightarrow$ closed transition occurs at a rate $\beta_n(V)$:

$$\frac{dn}{dt} = \alpha_n(V)(1 - n(t)) - \beta_n(V)n(t). \tag{4.4}$$

A useful equivalent formulation divides both sides by $\alpha_n(V) + \beta_n(V)$ and results in

$$\frac{dn}{dt} = \frac{n_\infty(V) - n(t)}{\tau_n(V)}, \tag{4.5}$$

where $n_\infty(V)$ is a voltage-dependent steady-state activation probability and $\tau_n(V)$ is a voltage-dependent time constant. The forms of $\alpha_n(V)$ and $\beta_n(V)$ may be estimated from first principles using a thermodynamic argument, but in practice these functions are fit parametrically from biological data.

4.3.2.3 Transient conductance

For persistent conductances, channel activation probabilities are monotonically increasing with depolarizing membrane potentials and independent of time. Transient conductances on the other hand remain open only transiently as the membrane potential is depolarized. This is described in terms of competing effects of an activation variable m_j and an inactivation variable h_j, which have opposite behaviors with regard to depolarization. For such a conductance, $P_j = m_j^k h_j^\ell$ for some integers k and ℓ. The values of m_j and h_j are governed by rate functions α_m, β_m, α_h, β_h

as in (4.4) and may also be written in terms of voltage-dependent steady state and time constant functions as in (4.5).

4.4 The Hodgkin-Huxley model

Most specifically, the **Hodgkin-Huxley** model of a single neuron refers to a particular system of nonlinear ODEs developed by Alan Hodgkin and Andrew Huxley in 1952 for modelling the membrane potential of the squid giant axon, a particularly large neuron well-suited to experimentation [60]. In this model, the neuron's ionic currents are modeled as a sum of a leakage current, a delayed-rectified K^+ current, and a transient Na^+ current. Additionally, capacitance and conductance variables were expressed as ratios per unit area.

$$C_m \frac{dV}{dt} = -\left(\bar{g}_L(V - E_L) + \bar{g}_K n^4 (V - E_K) + \bar{g}_{Na} m^3 h (V - E_{Na})\right), \tag{4.6}$$

where constants have the following values:

Constant	Value	Constant	Value
C_m	0.01 nF/mm^2		
$\bar{g}_L$	0.003 mS/mm^2	E_L	-52.402 mV
$\bar{g}_K$	0.036 mS/mm^2	E_K	-77 mV
$\bar{g}_{Na}$	1.2 mS/mm^2	E_{Na}	50 mV

and the rate functions take the following forms:

Function	Form	Function	Form
$\alpha_n(V)$	$\frac{0.01(V+55)}{1-\exp(-0.1(V+55))}$	$\beta_n(V)$	$0.125\exp(-0.0125(V+65))$
$\alpha_m(V)$	$\frac{0.1(V+40)}{1-\exp(-0.1(V+40))}$	$\beta_m(V)$	$4\exp(-0.0556(V+65))$
$\alpha_h(V)$	$0.07\exp(-0.05(V+65))$	$\beta_h(V)$	$\frac{1}{1+\exp(-0.1(V+35))}$

More generally, **Hodgkin-Huxley-like** models refer to systems of ODEs, similar to Equation (4.6), which take into account the activity of additional ion channels. The system used here in Section 4.7.2 is a model of this form.

4.5 Modeling synaptic activity

An additional term which must be added to Hodgkin-Huxley-like equations for modeling networks of neurons are synaptic currents. As with neurons there are many types of synapses throughout the nervous system which perform slightly different functions, but for the purposes of this chapter it mostly suffices to consider *spike-mediated chemical synapses*. See Figure 4.2 for an illustrative diagram of their action.

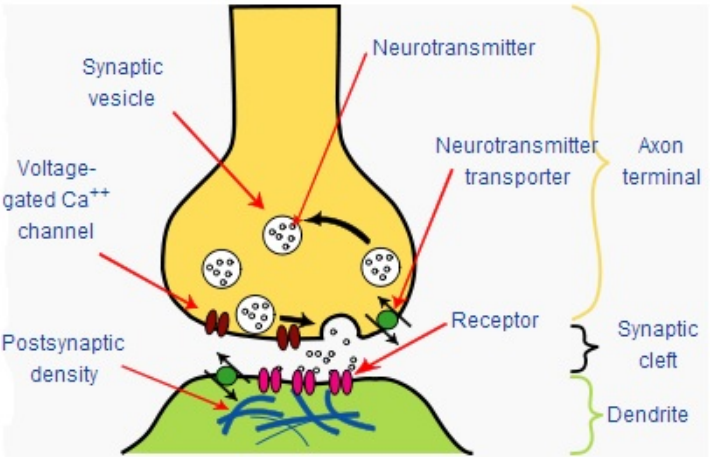

Figure 4.2: An illustrative diagram of a spike-mediated chemical synapse. A presynaptic action potential triggers the release of Ca^{2+} ions from voltage-dependent Ca^{2+} channels, which in turn causes the release of neurotransmitter molecules by vesicles into the synaptic cleft. These neurotransmitters diffuse across the cleft and bind to receptors on the postsynaptic neuron's cell membrane. This binding then alters the conductance of ion channels in the postsynaptic neuron, eliciting a change in membrane potential and effecting a transmission of the presynaptic action potential. Synapse_diagram_picture.jpg by Lhunter2099 is licensed under the CC BY-SA 4.0.

The activation of these synapses begins when an action potential in the presynaptic neuron enters the axon terminal connected to the synapse. This wave of depolarization causes the release of Ca^{2+} ions through voltage-dependent Ca^{2+} channels. The increase in Ca^{2+} concentration then causes presynaptic vesicles within the terminal to release molecules known as *neurotransmitters* into the synaptic cleft - the junction between the presynaptic neuron's axon terminal and postsynaptic neuron's dendrite. These neurotransmitters cross the cleft via diffusion, and bind to receptor sites on the postsynaptic neuron. This binding alters the conductance of postsynaptic ion channels which in turn creates an ionic current which alters the postsynaptic neuron's membrane potential. Depending on the neurotransmitter and synaptic receptors involved, the effect may be to depolarize or hyperpolarize the postsynaptic

neuron. If the postsynaptic neuron is depolarized, the synapse is termed **excitatory**. If the postsynaptic neuron is hyperpolarized, the synapse is **inhibitory**. An important property of the nervous system is that each neuron creates only inhibitory or excitatory connections with downstream neurons, never both.

As with neurons, synaptic activity may be modeled with varying levels of fidelity or abstraction. My research makes use of a common model as described in [89]. The ionic current generated by a synapse is modeled as

$$I_s = \bar{g}_s[O](V - E_s), \tag{4.7}$$

where $\bar{g}_s$ is the maximal conductance of the synapse, $[O]$ is the proportion of open synaptic channels, and E_s is the reversal potential of the synapse. The evolution of $[O](t)$ is governed by an ODE akin to (4.4):

$$\frac{d[O]}{dt} = \alpha(1 - [O])[T] - \beta[O], \tag{4.8}$$

where α and β are activation and inactivation constants respectively, and $[T]$ is a time-varying neurotransmitter concentration. Here, the form of $[T]$ is defined phenomenologically as a square pulse of fixed duration that begins when the presynaptic neuron generates an action potential:

$$[T](t) = AH(t_0 + t_{max} - t)H(t - t_0), \tag{4.9}$$

where A is a constant, $H(t)$ is the Heaviside function, t_{max} is the neurotransmit-

ter pulse width, and t_0 is the most recent time at which the presynaptic neuron generated a spike.

4.6 Olfaction in locusts

The details of sensory processing are studied across many animals and all sensory modalities, but certain animal/modality pairs are particularly fruitful. One such pair commonly studied is olfaction in the locust (*Schistocerca americana*). The insect's relatively large size makes it amenable to electrophysiological experimentation, and its olfactory system contains relatively few neurons and is well-characterized compared with that of mammals. An excellent overview of this system is presented in Chapter 11 of *Spike timing: mechanisms and function*, and much of the overview in this section is derived from there.

Research into locust olfaction to date focuses on the first three processing layers, as illustrated in Figure 4.3.

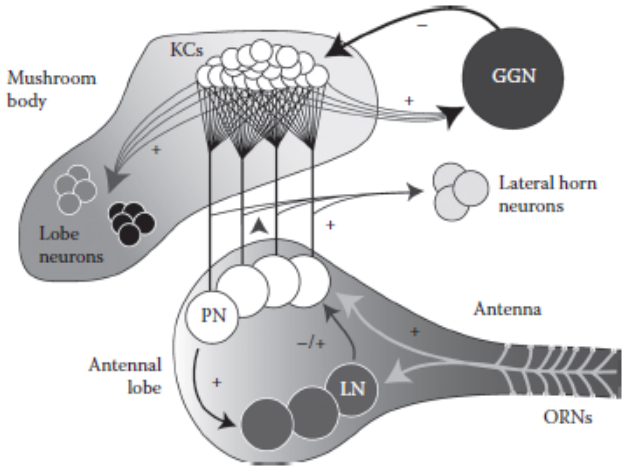

Figure 4.3: An illustrative diagram of the first three layers of olfactory processing in the locust. Olfactory receptor neurons (ORNs) transduce environmental information about chemical identity and concentration into neural firing rate activity. The output of the ORNs enters the antennal lobe (AL), where ORNs synapse onto excitatory projection neurons (PNs). In the AL, stimulus responses are formatted into a spatiotemporal code, distributed across the PNs and also segmented into temporal bins via oscillatory activity induced by feedback inhibition from reciprocally-connected local interneurons (LNs). The output of the PN activity is received by Kenyon cells (KC) in the mushroom body (MB), which produce output that is highly sparse both collectively in space across the population, as well as individually in time for any active neuron. The KCs are reciprocally connected with a single inhibitory neuron, the giant GABA-ergic neuron (GGN), which provides feedback inhibition to the KC population. The activity of the KCs also travels further downstream into the brain, to areas whose structure and function are less well-characterized.
Image adapted from Figure 11.1 in [42].

In brief, the $\sim 50,000$ excitatory olfactory receptor neurons (ORNs) in the antennae respond to the concentration and identity of chemicals in the environment, forming the first layer of the olfactory processing hierarchy. The ORNs synapse onto ~ 830 excitatory projection neurons (PNs) in the antennal lobe (AL). The PNs serve to average the noisy ORN inputs, but also perform additional, complex reformatting with the help of feedback connections with ~ 300 local neurons (LNs), most of which are inhibitory. The PNs synapse onto $\sim 50,000$ excitatory Kenyon cells (KCs) in the mushroom body (MB). The MB is a center of associative learning in the olfactory

system, and the KCs are noted for their highly-sparse activity. In response to a given odorant, only a handful of the roughly 50,000 KCs will respond, emitting only a few spikes at odor onset and, under certain conditions, at odor offset. Thus, the KCs exhibit both *population sparsity* (few responses across the population) and *lifetime sparsity* (responding KCs emit few spikes), forming a useful model system for investigating sparse sensory coding in general.

4.6.1 Olfactory receptor neurons

As with all receptor neurons, **olfactory receptor neurons** (ORNs) possess specialized structures for transducing environmental information into neural activity. In this case, these structures are *olfactory receptors* (ORs) - proteins embedded in the membrane of ORN dendrites. Unlike vision and audition, where receptor neurons must be sensitive to two dimensions of stimulus structure (frequency and intensity), the stimulus space for olfaction is high-dimensional, reflecting the variety and complexity of biologically-relevant volatile molecules. Correspondingly, there are a large number of different OR types, approximately 60 in insects. Each ORN expresses a single OR type, or two working in concert [66]. In the absence of stimulation, ORNs will fire at a base rate of $\sim 1-30$ Hz. When stimulated, depending on the concentration and identity of the stimulus odorant and the type of OR expressed, the ORN firing rate may increase, decrease, or sometimes combinations of the two. The time scale, latency, and amplitude of these effects depend on the combination of odorant and ORN, and in both cases prolonged exposure to a simulus can result

in a **sensory adaptation** effect, whereby the firing rate of excited ORNs decreases and the firing rate of inhibited ORNs increases. The magnitude of this effect is also varies with stimulus and OR types.

There are additional important variants between ORN types that serve to increase the encoding capacity of the olfactory system. Some ORNs may be broadly tuned, responding to many types of odorants. Others are narrowly tuned, and respond only to a single or a few types of stimuli. In general, higher concentrations of an odorant elicit activity from more ORN types. These properties serve to create a *combinatorial code* of odorant concentration and identity among the ORNs - by distributing a stimulus representation across multiple neurons, each receptive to different properties of the stimulus, N neurons can encode up to 2^N different stimuli. This ORN activity is then passed to the second layer of olfactory processing, in the antennal lobe.

4.6.2 The antennal lobe

The *antennal lobe* (AL) consists of two neuron populations. Both the excitatory **projection neurons** (PNs) and the predominantly-inhibitory **local neurons** (LNs) receive output from the ORNs, but only the activity of the PNs is relayed downstream outside of the AL. The tuning profiles of PNs are broader than those of ORNs, but their spiking activity is more reliable - each PN receives direct synaptic input from multiple ORNs expressing the same set of ORs, and so signal averaging is one function of the PNs. PNs also get input from other ORNs through excita-

tory and inhibitory lateral connections through local neurons. However, PNs serve a more complex role than just averaging their ORN inputs - the AL gives rise to oscillatory activity which creates a spatiotemporal encoding of stimulus information, distributed both across space among the PN population and in time across successive oscillation windows.

The oscillatory activity of the AL stems from the reciprocal connections between the PN and LN populations. The action potentials generated by stimulated PNs elicit increased electrical activity in the nonspiking inhibitory LNs, which in turn stimulates a release of inhibitory neurotransmitters that hyperpolarize targeted PNs. This cycle generates oscillations in PN activity at roughly 20 Hz. This oscillation pushes stimulated PNs to spike closely-together in time, creating synchrony in PN responses within an oscillation window. Each active PN reliably phase-locks to these oscillatory windows for only a subset of the full AL response to an odor.

The nature of the AL's encoding of odor information can be interpreted by examining the trajectory of large PN populations over time via dimension reduction techniques [79]. Given repeated recordings from N PNs, the results can be averaged per-PN and binned across time to give average firing rates in each bin for each PN. This procedure results in a trajectory in an N-dimensional population state space. These high-dimensional trajectories lie within low-dimensional manifolds in the ambient population space, and may be visualized in 3D using nonlinear dimension reduction techniques such as locally-linear embedding (LLE) [92].

The results of this analysis are striking - population responses reliably fall along odorant-specific manifolds, while odorant concentration determines the position of

the trajectory within the manifold. During prolonged stimulation, the population trajectory moves from an initial point along a stereotyped path to another fixed point in the state space, and then returns to the initial point via a different stereotyped path [79]. These results illustrate the efficacy of the odor encoding within the AL - the intrinsically-high-dimensional odor stimuli are reliably encoded in a higher-dimensional representation with both spatial and temporal components, allowing different stimuli to be well-separated in their representations for further parsing downstream in the Kenyon cells.

4.6.3 The mushroom body

The *mushroom body* (MB) contains the third layer of the olfactory processing hierarchy. Here, the roughly 830 PNs of the AL synapse onto some 50,000 **Kenyon cells** (KCs), whose activity greatly differs from the PNs' output. The KCs are tuned more narrowly than PNs and ORNs, responding to far fewer odors each. Moreover, the KC population has a much more sparse response to stimuli, with only a handful of the neurons responding to any given odorant. KCs that do respond to a stimulus emit only a small number of spikes following odor onset, and following odor offset under certain conditions as well. These properties, as well as the amenability of the locust MB to physiological recordings, make the KCs an ideal model system for studying the sparse sensory coding phenomenon.

The sparsity of the KC responses derives from cellular as well as network factors. Intrinsic cellular properties conspire to give KCs high firing thresholds -

each KC in the locust receives synaptic input from approximately half of the 830 PNs, and roughly 100 of these PNs must spike nearly simultaneously to elicit a spike from the KC. The synchronized oscillatory behavior of PNs allows for this to occur during stimulation, but the decrease in firing rates following sensory adaptation of the ORNs drops PN activity levels below the threshold required for KCs to fire.

Network mechanisms assist in achieving this behavior within the mushroom body. All KCs in the mushroom body synapse onto a single inhibitory neuron, the **giant GABAergic neuron** (GGN), which in turn synapses onto each KC in the population. This provides feedback inhibition to the KCs - the first KCs to spike stimulate the GGN, which in turn creates a powerful wave of inhibition to silence the KCs. This process creates a window within each larger oscillation cycle in which sufficiently-stimulated KCs may fire. The idea of winner-take-all competition in neural networks via lateral or feedback inhibition has been studied for some time [31], but differs in important ways from the process taking place in the MB. A generalization of the winner-take-all model, the *E%-max winner-take-all* model [37], is more relevant. In this model, the time delay between initial KC activation and the return of inhibitory signals from the GGN to the KC population creates a small window in which those KCs whose input is at least some percentage of the input to the maximally-driven KC ($E\%$ of the maximum), fire as well. The value of E is determined by the travel time of the action potentials to and from the GGN. This can be likened to an adaptive filter in signal processing, which instead of preserving a fixed number of the largest coefficients of a signal's representation in some dictionary, preserves a varying number of coefficients with sufficiently large

energy.

Moving further downstream, the KCs synapse onto other targets within the MB, the so-called α- and β-lobe neurons. Current knowledge of the function of these networks and those further downstream is tenuous compared to the levels of the processing hierarchy described here. Further investigations of these areas, as well as their inevitable pathways into the more generalized multimodal sensory integration regions which must underlie the decision-making and behavior of even relatively-simple creatures such as locusts, will doubtlessly further illuminate the physiological basis for sensory perception. Until then, the earlier layers continue to provide challenging problems in understanding the efficient representation and analysis of the complex, high-dimensional signals which olfaction creates.

4.7 Offset spiking in the locust mushroom body

Olfaction serves many purposes for insects - finding food, finding mates, and sending messages among them. In a turbulent medium, e.g., wind in air or waves in water, odors emitted from a point source do not diffuse evenly into the surrounding space. Instead they exhibit complex plume patterning, in which packets of odorant molecules extend through the environment. Within a plume, odorant concentration may vary significantly over millimeter and smaller spatial scales [81]. Insects must navigate these plumes, and there is evidence that they do so by zig-zagging across them, seeking out the transition boundaries between plume and ambient space [81]. How does an insect parse this information from the patterns of activity created in

the ORNs? KCs are active predominantly at the beginning of a prolonged stimulus, and are also observed to fire following the cessation of a stimulus under certain conditions, as demonstrated in Figure 4.4.

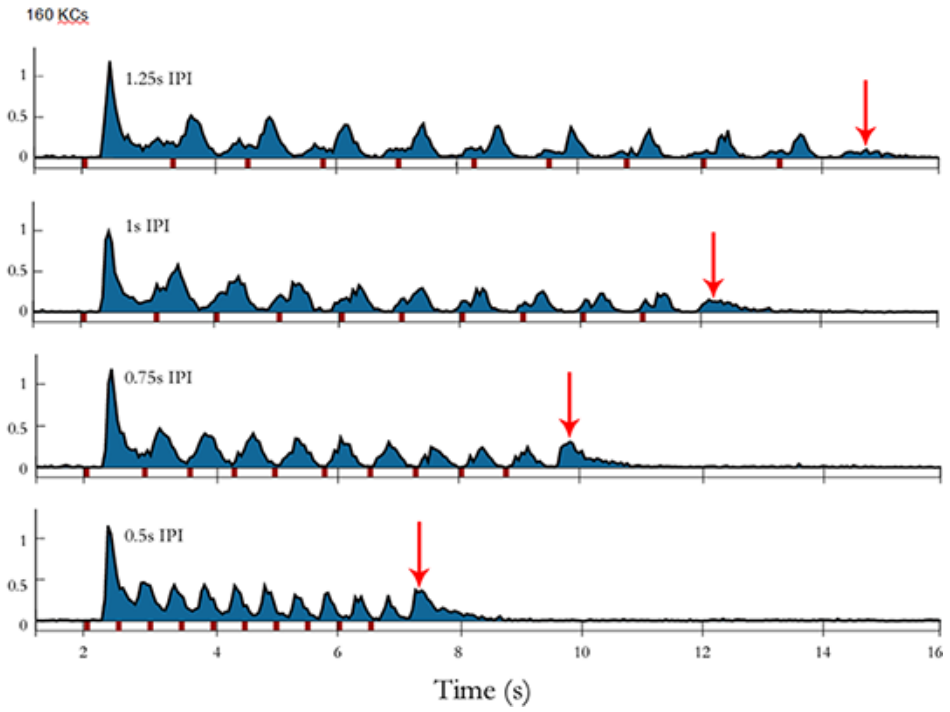

Figure 4.4: An illustration of the offset spiking phenomenon in Kenyon cells, using **peristimulus time histograms** (PSTHs), which plot neuron firing rates over time, averaged across a population of neurons. When a brief odor pulse, indicated in red on the x axis, is repeated with a long inter-pulse interval (IPI), little activity occurs after the final odor offset. When the odor pulses are presented in more rapid succession, offset spiking activity increases.
Image taken with permission from unpublished work in a presentation by Joby Joseph, Stacey Daffron, and Mark Stopfer to the Stopfer lab in NICHD.

This indicates that one function of the KC odor encoding is a sort of temporal feature extraction, revealing the concentration boundaries relevant to the insect's plume navigation. What is the origin of this functionality within the olfactory processing networks? The case for onset spiking is clear-cut - sensory adaptation in the ORNs drives down activity levels after the first second or so of a stimulus

response, in turn driving down PN activity levels below the threshold required for KCs to fire. What is the analogous explanation for offset spiking? That is the question we seek to address with this project.

4.7.1 Relevant network properties

Understanding the genesis of this offset spiking phenomenon is an important aspect of understanding the network and cellular factors that cause KCs to respond to changes in a stimulus, as opposed to the ongoing state of the olfactory stimulus itself. For this reason, I set out along with the Stopfer Lab to investigate the causes of this phenomenon and to model it in simulation. A computational model with adequate fidelity to the biological behavior of KCs allows us to test hypotheses about network structure and function relevant to this behavior.

The simplest explanation for this behavior is that it is a consequence of stimulus offset activity in the PNs. Both PN and LN populations receive input from ORNs, but it is known that LN activity tracks stimulus intensity closely, while ORN $\rightarrow$ PN activity persists for some time, or even begins only at offset. See Figure 4.5 for an example of this discrepancy.

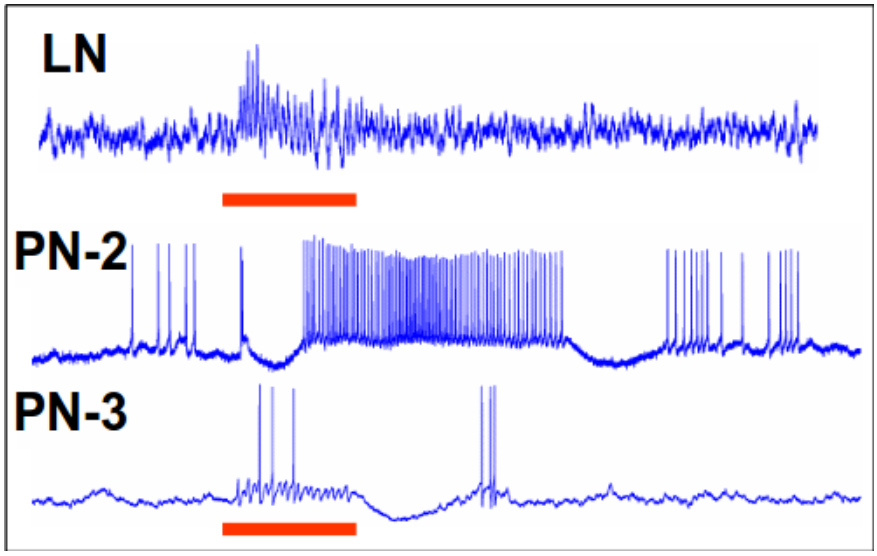

Figure 4.5: A demonstration of the difference between LN responses and PN responses to a stimulus. The stimulus window is indicated by a horizontal red bar beneath the LN and PN responses.
Image taken with permission from unpublished work in a presentation by Joby Joseph, Stacey Daffron, and Mark Stopfer to the Stopfer lab in NICHD.

Indeed, PN activity can be divided into two classes - one in-phase with stimulus presentation and one in anti-phase with the stimulus. The latter category underlies this PN offset activity. ORNs whose firing rates are inhibited by a stimulus presentation will briefly overshoot their baseline firing rates when the stimulus is withdrawn, providing a potential foundation in the antennae for the anti-phase category of PNs.

This burst of PN offset activity is a plausible origin for KC offset spiking, but without inhibitory input from the LNs, which ends promptly with the stimulus, it lacks the precise temporal coordination of the oscillatory PN activity during a stimulus presentation. This must be addressed on two levels to fully explore its importance to KC offset spiking. First, is the bump in PN offset activity sufficient to make up for the decreased temporal coherence, in order to generate KC action

potentials? Second, how do adaptive processes in the AL shape this offset PN activity, both in terms of its intensity and temporal coherence?

There is less data from experiments to draw upon to answer this second question. It is known that an adaptive process known as **fast learning** shapes the coherence of PN activity during a stimulus presentation [6]. The first time a new odor is presented to a locust, the LNs' and PNs' responses do not display oscillatory behavior; this only manifests after repeated presentations of the odor. The biological mechanisms which cause this adaptation are not fully known, but it is hypothesized to involve the facilitation (increase in strength) of PN $\leftrightarrow$ LN synaptic interconnections, and potentially the existence and facilitation of PN $\leftrightarrow$ PN interconnections as well. This is supported by computational simulations in [6], but biological evidence is lacking. If fast learning does indeed work in this manner, offset activity following adaptation should increase in intensity, conditions more appropriate for downstream coincidence detection. More information from experiments is needed to fully understand the effects of this adaptation, but it provides a plausible mechanism for the appearance of KC offset spiking under a more narrow range of stimulus conditions which bears further study. In the meantime, the first question of the tradeoff between increased activity and decreased temporal coherence can be explored via computational modeling.

4.7.2 Computational modeling of KC offset spiking

The computational modeling done in this section seeks to answer one question: if N PNs spike while connected to a single KC, how close together in time must they do so to elicit a spike from the KC? Solving this problem would help in determining the relationship between offset PN activity and offset KC activity. Beyond the model implementation, this arrangement implies another difficult problem: How can we determine if the model's output is similar to the behavior of the corresponding biological system? I propose several benchmarks to better establish comparisons between simulated and biological behavior. Ultimately, the model appears to still exhibit discrepancies with its biological counterpart, but there is value in this result, too. Inferring biological behavior from simulation is complicated, and too often there is a divide between experimentalists uninterested in simulation and theoreticians unconcerned with the full consequences of modeling abstractions. The model exhibited here is an iteration of a work in progress. It is a focal point for devising better procedures for effectively matching neural simulation with reality.

The modeling of KC offset spiking starts simple: we can ignore the effect of the GGN's feedback inhibition, which cannot factor into the initiation of KC offset spikes (though may curtail them afterward). Then, an absence of KC $\leftrightarrow$ KC synapses means KC responses are only correlated through their shared PN afferents. So, we only need to model one KC and its upstream PN synapses. In accordance with estimates obtained from experiments, we model 415 PN $\rightarrow$ KC synapses as in Section 4.5, connected to a single KC using a Hodgkin-Huxley-like model derived

primarily from [89]. At a high level, the KC membrane potential V is modeled as

$$C_m \frac{dV}{dt} = -(I_{leak} + I_{ionic} + I_{syn} + I_{noise}). \tag{4.10}$$

I_{leak} consists of two components: a "general" leakage current I_L and a potassium leakage current I_{KL}. The intrinsic ionic current I_{ionic} has five components: a fast sodium current I_{Na}; a delayed rectified potassium current I_K; a Ca^{2+}-dependent potassium current $I_{K(Ca)}$; a transient, so-called A-type potassium current [56] $I_{K,A}$; and a transient calcium current I_{Ca}. The synaptic current I_{syn} is a sum of 415 individual synaptic currents, one per upstream PN, each of the form described in Equation (4.7). I_{noise} is a random variable, with $I_{noise}(t)$ being sampled from a normal distribution with mean 0 and standard deviation of 1.5 (nA). In all, V is governed by the equation

$$\begin{aligned} -C_m \frac{dV}{dt} = {} & g_L(V - E_L) + g_{KL}(V - E_{KL}) + g_{Na} m_{Na}^3 h_{Na}(V - E_{Na}) \\ & + g_K m_K^4 h_K (V - E_K) + g_{K(Ca)} n_{K(Ca)}^2 (V - E_K) + g_{K,A} n_{K,A} (V - E_K) \\ & + g_{Ca} m_{Ca}^2 h_{Ca} (V - E_{Ca}) + \sum_{j=1}^{415} g_{syn,j} [O]_j (V - E_{syn}) + I_{noise}(t), \end{aligned} \tag{4.11}$$

where $C_m = 2.9 \times 10^{-4}$ μF, and the leakage and ionic conductance variables take the values

Variable	Value (μS)	Variable	Value (μS)
g_L	2.9×10^{-2}	g_{KL}	1.16×10^{-3}
g_{Na}	26.1	g_K	2.9
$g_{K(Ca)}$	2.9×10^{-1}	$g_{K,A}$	1.45×10^{-2}
g_{Ca}	2.9×10^{-2}		

and the reversal potentials take the values

Variable	Value (mV)	Variable	Value (mV)
E_L	-65	E_{Na}	50
E_K	-95	E_{Ca}	$12.8 \log(2/[\mathrm{Ca}^{2+}])$

where $[\mathrm{Ca}^{2+}]$ is the concentration of Ca^{2+} ions (in $m$mol) and is updated via the first-order ODE

$$\frac{d[Ca]}{dt} = -1.7862 I_{Ca} - \frac{([Ca] - 2.4 \times 10^{-4})}{100}. \tag{4.12}$$

The gating variables $n.$, $m.$, and $h.$ are updated via ODEs of the forms described in

(4.4) and (4.5).

$$\frac{dm_{Na}}{dt} = \alpha_{m_{Na}}(V)(1 - m_{Na}) - \beta_{m_{Na}}(V)m_{Na} \tag{4.13}$$

$$\frac{dh_{Na}}{dt} = \alpha_{h_{Na}}(V)(1 - h_{Na}) - \beta_{h_{Na}}(V)h_{Na} \tag{4.14}$$

$$\frac{dm_K}{dt} = \alpha_{m_K}(V)(1 - m_K) - \beta_{m_K}(V)m_K \tag{4.15}$$

$$\frac{dh_K}{dt} = \alpha_{h_K}(V)(1 - h_K) - \beta_{h_K}(V)h_K \tag{4.16}$$

$$\frac{dn_{K(Ca)}}{dt} = \frac{n_{\infty,K(Ca)}([\mathrm{Ca}^{2+}]) - n_{K(Ca)}}{\tau_{n_{K(Ca)}}([\mathrm{Ca}^{2+}])} \tag{4.17}$$

$$\frac{dn_{K,A}}{dt} = \frac{n_{\infty,K,A}(V) - n_{K,A}}{\tau_{n_{K,A}}(V)} \tag{4.18}$$

$$\frac{dm_{Ca}}{dt} = \frac{m_{\infty,Ca}(V) - m_{Ca}}{\tau_{m_{Ca}}(V)} \tag{4.19}$$

$$\frac{dh_{Ca}}{dt} = \frac{h_{\infty,Ca}(V) - h_{Ca}}{\tau_{h_{Ca}}(V)} \tag{4.20}$$

Each of these ODEs defines two additional functions - activation and inactivation functions for I_{Na} and I_K, steady-state and time-constant functions for the other currents. The form of those functions is described below.

Function	Form
$\alpha_{m_{Na}}(V)$	$\frac{0.32(13-V-65)}{e^{(13-V-65)/4}-1}$
$\beta_{m_{Na}}(V)$	$\frac{0.28(V-105)}{e^{(V-105)/5}-1}$
$\alpha_{h_{Na}}(V)$	$0.128e^{(17-V-65)/18}$
$\beta_{h_{Na}}(V)$	$\frac{4}{e^{(40-V-65)/5}+1}$
$\alpha_{m_K}(V)$	$\frac{0.032(15-V-65)}{e^{(15-V-65)/5}-1}$
$\beta_{m_K}(V)$	$0.5e^{(10-V-65)/40}$
$\alpha_{h_K}(V)$	$0.028e^{(15-V-65)/15} + \frac{2}{e^{(85-V-65)/10}+1}$
$\beta_{h_K}(V)$	$\frac{0.4}{e^{(40-V-65)/10}+1}$
$n_{\infty,K(Ca)}([\mathrm{Ca}^{2+}])$	$\frac{3333[\mathrm{Ca}^{2+}]^2}{3333[\mathrm{Ca}^{2+}]^2+1}$
$\tau_{n_{K(Ca)}}([\mathrm{Ca}^{2+}])$	$\max\left(0.1, \frac{0.896}{100[\mathrm{Ca}^{2+}]^2+0.03}\right)$
$n_{\infty,K,A}(V)$	$\frac{1}{1+e^{-(V+60)/8.5}}$
$\tau_{n_{K,A}}(V)$	$\left(\frac{1}{e^{(V+35.82)/19.69}+e^{-(V+79.69)/12.7}} + 0.37\right)/3.74$
$m_{\infty,Ca}(V)$	$\frac{1}{1+e^{-(V+52)/7.4}}$
$\tau_{m_{Ca}}(V)$	$\left(3 + \frac{1}{e^{(V+27)/10}+e^{-(V+102)/15}}\right)/9.9$
$h_{\infty,Ca}(V)$	$\frac{1}{1+e^{(V+80)/5}}$
$\tau_{h_{Ca}}(V)$	$\left(85 + \frac{1}{e^{(V+48)/4}+e^{-(V+407)/50}}\right)/3.74$

The remaining variables to be accounted for are the synaptic conductance variables, $\{g_{syn,j}\}_{j=1}^{415}$. The strengths of the PN $\rightarrow$ KC synaptic connections are not uniform.

Experiments done by Jortner et al. as described in [64] Figure 6A have provided an empirical characterization of the distribution of synaptic strengths, as measured by a distribution of peak EPSP amplitudes. We were able to recreate this distribution, linearly interpolating between the normalized histogram values to create an empirical probability distribution. It is not obvious *a priori* that using this same distribution for synaptic conductances will result in simulated peak EPSPs matching the distribution described in Jortner et al. This simulated EPSP distribution can be calculated, and indeed it closely matches the results from [64]. See Figure 4.6 for a comparison of the distributions.

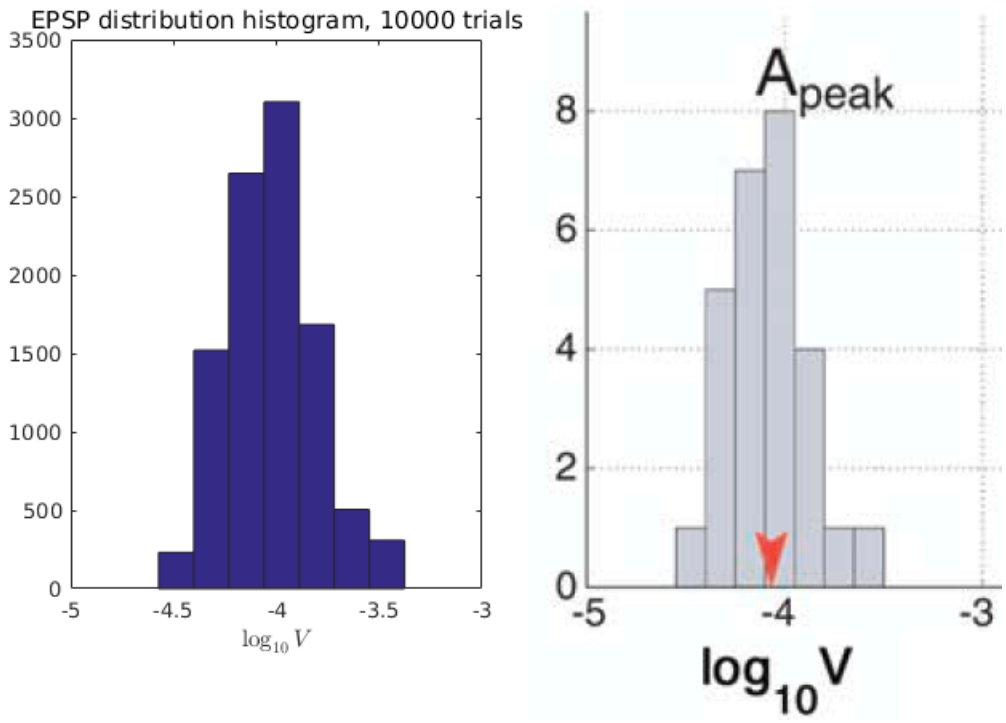

Figure 4.6: A comparison of the empirical peak EPSP distribution generated by PN $\rightarrow$ KC synapses and its simulated counterpart. Simulated synaptic conductances were drawn from an empirical distribution as described in [64] Figure 6A, reproduced on the right side of this figure, but simulated peak EPSP measurements were measured as in the Jortner et al. experiment - activating a single synapse and measuring the peak depolarization of the simulated KC. This was repeated for 10000 trials to generate the histogram on the left of this figure. The close agreement between the two indicates that this distribution is a suitable characterization of synaptic conductances.

While [89] provided the majority of the parameters and functions used in this model, several alterations had to be made, due to errors we discovered in the paper as well as disagreements between the model's initial behavior and data from experiment. The I_{Ca} term derives from unpublished code written by Maxim Bazhenov, generously shared with us. The parameters used in Equation (4.12) stem from the values used in that code, as well. Additionally, the leakage conductance g_L was changed from 2.9×10^{-3} to 2.9×10^{-2}. While the previous value is a valid model

parameter, this change improved the behavior of the KC model. The KC resting potential increased from ~ -90mV to ~ -68mV, and the KC became less quiescent.

fixed the simulated KC resting potential and brought its responses to stimulation in line with experiment.

4.7.3 Simulation procedure

Using this model, we devised a protocol for testing the relationship between the temporal coherence of PN $\rightarrow$ KC activity, the number of active PN $\rightarrow$ KC synapses, and spiking activity in the KC. A variable number N of the 415 synapses were designated as "active", spiking at a time drawn uniformly at random from an interval $[t_0, t_0 + \Delta t]$ for some choice of interval width Δt. In addition to these active spikes, it is known that PNs have a baseline firing rate of approximately 2.5 Hz [79]. This was modeled as a Poisson process in the remaining $415 - N$ synapses by generating spike times with inter-spike arrival times, measured in ms, drawn from an exponential distribution with rate parameter $\lambda = 0.0025$. With the spike times generated, the synapses and KC are simulated for $t_0 + \Delta t$ ms, plus a short additional amount of time.

This procedure was repeated for $K = 100$ trials for each $N \in \{50 + 10i\}_{i=0}^{15}$ and each $\Delta t \in \{10 + 10i\}_{i=0}^{5} \cup \{100 + 15i\}_{i=0}^{20}$. For each combination of N and Δt, a **peristimulus time histogram** (PSTH) was constructed to quantify the average KC activity over the K trials. To construct a PSTH, the time course of a simulation is divided into bins of width δ ms. For each bin i, the average spike count s_i equals

the sum of the number of spikes on the interval $[i\delta, (i+1)\delta)$ across all K trials, divided by K. The PSTH then plots each s_i, and this information may be used to determine which combinations of N and Δt are sufficient for generating consistent KC spike activity.

4.7.3.1 Numerical implementation

To perform the KC offset spiking simulations, an explicit **four-step Runge-Kutta** (RK4) method [2] is used, as described in Section 2.3.2, with a timestep of $h = 0.03$ ms. For this model, the components of $\boldsymbol{y}$ are $V(t)$, $[\mathrm{Ca}^{2+}](t)$, the gating variables $m_{Na}(t)$, $h_{Na}(t)$, $m_K(t)$, $h_K(t)$, $n_{K(Ca)}(t)$, $n_{K,A}(t)$, $m_{Ca}(t)$, $h_{Ca}(t)$, and the synaptic activation variables $\{[O]_j(t)\}_{j=1}^{415}$. $f(t, \boldsymbol{y})$ is defined component-wise in Equations 5.11-5.20, as well as Equation (4.8) with $\alpha = 0.94$ and $\beta = 0.18$ for the synapses. An example of a $V(t)$ generated by this model can be seen in Figure 4.7.

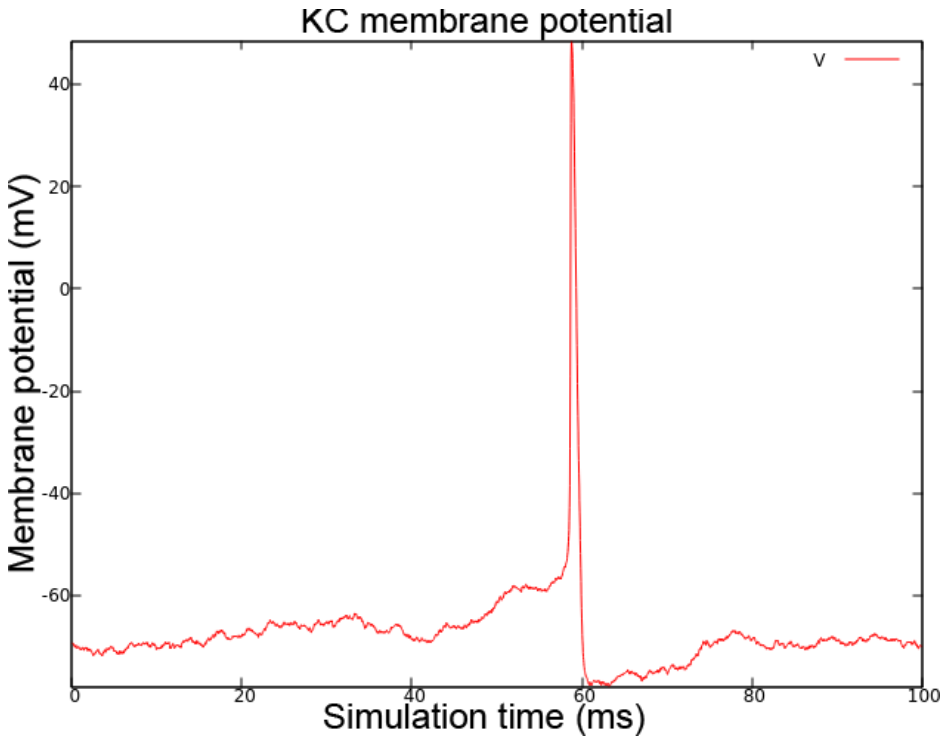

Figure 4.7: An example of a numerical solution to Equation 5.11, produced using an RK4 integration scheme with a C++ implementation. This demonstrates the time course of a typical KC action potential.

To perform the simulation procedure, along with the biological benchmarking described in Section 4.7.5, the RK4 method detailed in Section 2.3.2 was implemented in C++. The relevant biological structures were modeled as objects which implemented RK4 updates of their dynamic parameters. A Network object composed of a KenyonCell object and CholinergicSynapse objects forms the basis of the code, allowing for functions to be written to create and manipulate Networks and save their state information. In this way, quantitative analyses of the simulated Network state may be performed, leading to both the conclusions of this section and the possible biological validation tests. Creating the simulated KC PSTHs in Section 4.7.4 required running many instances of these Networks to generate trial

data. To improve performance, the simulation protocol ran these trials in parallel, making use of OpenMP [33] to distribute the trials across multiple threads.

4.7.4 Results

The results of this procedure can be seen in Figures 4.8 - 4.10.

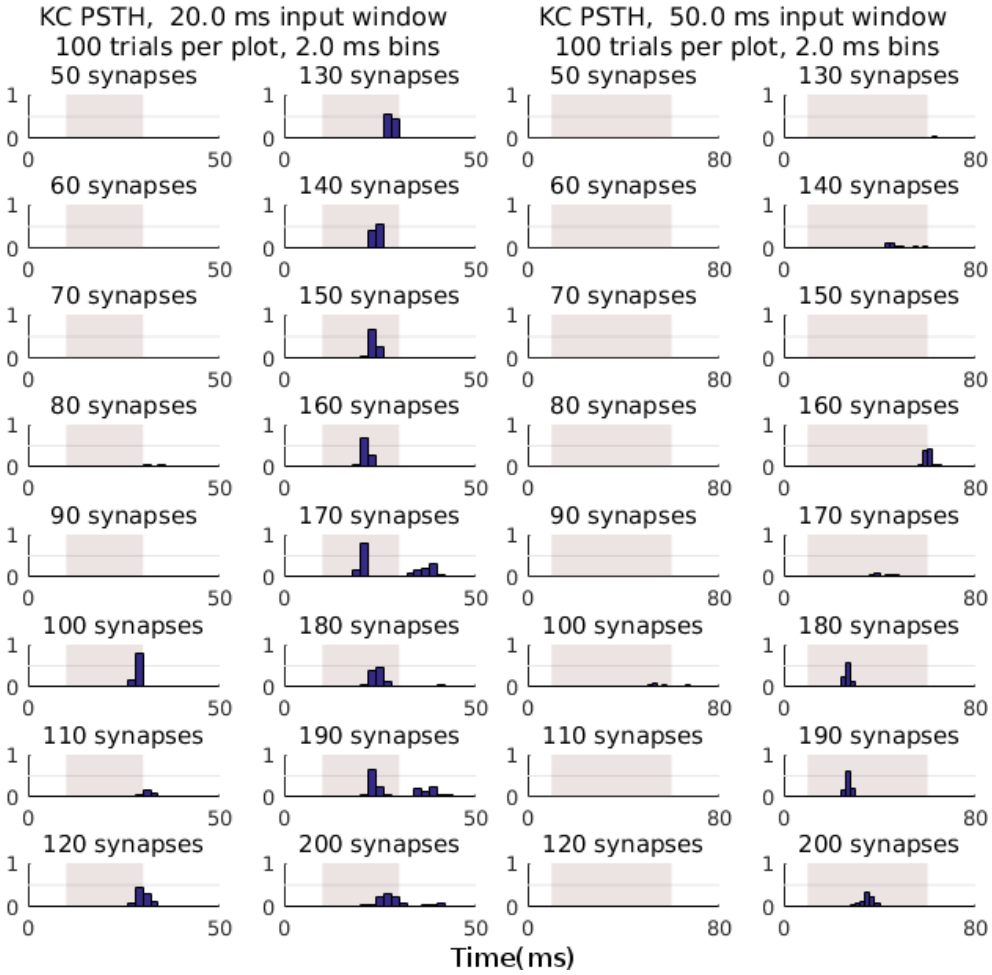

Figure 4.8: PSTHs of KC activity for variable numbers of coordinated PN → KC synapse activations within 20 ms (left) and 50 ms (right) activation windows.

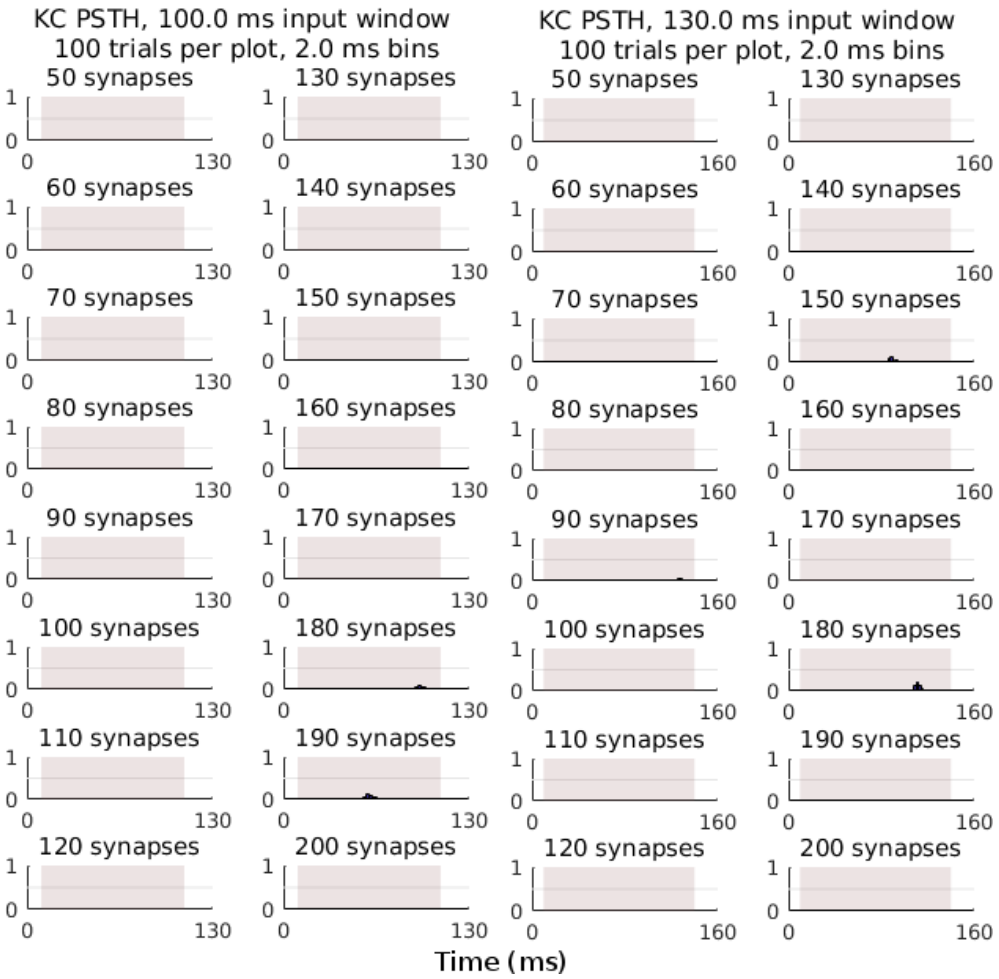

Figure 4.9: PSTHs of KC activity for variable numbers of coordinated PN → KC synapse activations within 100 ms (left) and 130 ms (right) activation windows.

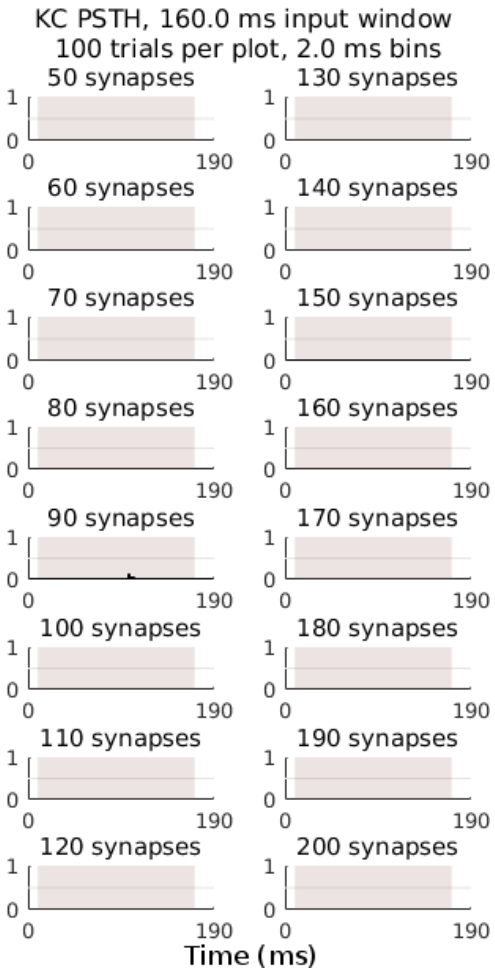

Figure 4.10: PSTHs of KC activity for variable numbers of coordinated PN → KC synapse activations within a 160 ms activation window.

The combination of N and Δt values most relevant to understanding the KC input conditions during a stimulus offset can be estimated from [79], which shows that 40% of recorded PNs were not silent during the roughly 300ms spike in activity at stimulus offsets, and approximately 10% of recorded PNs were responsive in any single 50ms window. Assuming, as suggested by experiments, 50% PN → KC connectivity with synaptic connections distributed uniformly across the PN population, we are interested in the outcome of 40-50 neurons spiking in a 50ms window, and 180-200 neurons spiking in a 300ms window. According to the tests

performed, these activity levels are insufficient to consistently elicit spikes in the simulated KC - with a 160 ms activation window very little activity at all is elicited, and with a 205 ms window (not shown), no KC spikes are generated at all.

The implications of this result on our understanding of offset activity in biological KCs hinges on the fidelity of this computational model to the corresponding biology, as well as our assumptions about the statistics of PN $\rightarrow$ KC activity derived from [79]. Does this result imply that we are missing some important factor governing KC offset spiking, or must the model be improved? We conclude Section 4.7 by discussing the measures taken to ensure biological fidelity so far, and the future additions that can be made as more biological data is obtained.

4.7.5 Biological fidelity

Comparing our computational model with experiments, and using that comparison to update our model, is an ongoing process. The easiest comparison to draw between the two is the match for the KCs' equilibrium potentials. According to [65], a KC receiving no input activity has an equilibrium potential of approximately -68 mV. Our model KC does as well, though this is unsurprising since one aspect of model tuning explicitly addressed this. Similarly, the use of the empirical synaptic conductance distribution allows the inputs to the KC to more closely emulate their biological counterpart, and we confirmed in Figure 4.6 that the resulting peak EPSP amplitude distributions match.

We can also use [64] to compare our KC's firing threshold with those observed

biologically. The firing threshold of a neuron is defined to be the difference between its resting potential and the highest potential the cell may attain before a spike is generated. In Figure 4.11, we show that our KC's observed firing threshold of approximately 12mV falls squarely within the distribution observed in [64].

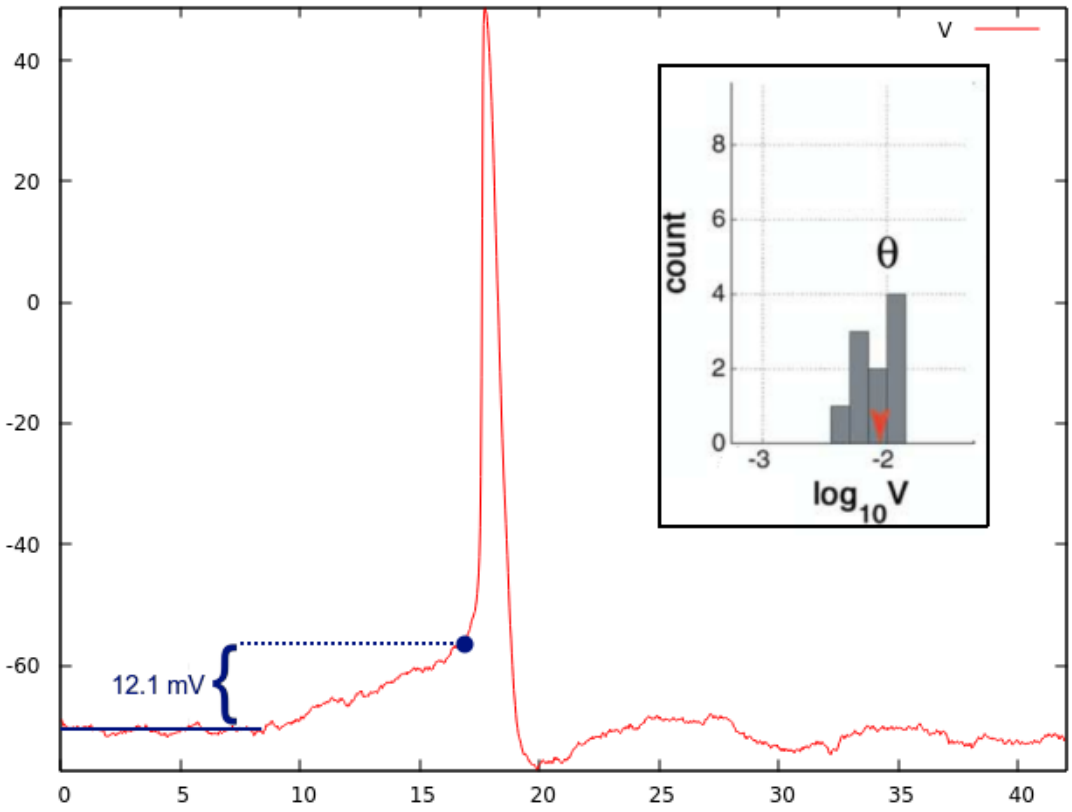

Figure 4.11: A comparison of our simulated KC's firing threshold and the distribution of this quantity observed in [64]. Note that the simulated value falls within the support of the empirical distribution. Inset figure is adapted from [64].

These measures provide some benchmarks for assessing our model, but further feedback from new experiments will provide additional, more powerful measures to aid with validation. One example of this will be measuring the *membrane time constant* of the locust KC. This constant is defined as the amount of time required for a neuron to transition $(1 - 1/e) \approx 63.2\%$ of the distance from a membrane

potential depolarized by a square current pulse, back to equilibrium. This can be done easily in simulation, and the results of this process can be seen in the following table, for various current pulse amplitudes.

Current (nA)	0.025	0.05	0.1	0.2	0.25
STC (ms)	6.89	6.83	6.77	6.89	7.37

Comparing these values with those derived from experiment will tell us whether our simulated KC responds to input current on the same time scales as biological KCs.

The previous benchmarks provide a series of tests for assessing how much our KC membrane potential behaves like biological KCs, but membrane potential is only one of the several states associated with the simulated KC. This cell also has multiple ion channels whose behavior is responsible for the time course of membrane potential activity, and it would be greatly beneficial to assess the similarity of our simulated ion channels to the biological ones. In a *voltage clamp* experiment, a control loop feedback mechanism is used to inject varying amounts of current into a neuron in order to clamp its membrane potential at a desired constant level. Using this setup, we can quantify the magnitude of various ionic currents through a KC at several voltages, and compare them with its simulated counterpart. I have added a simulated voltage clamp procedure to our project's codebase; once a suitable voltage clamp experiment has been performed, I will be able to replicate its conditions to obtain a comparison of ionic current flows. Using this, in combination with the other benchmarks established for determining biological fidelity, will make our model more informative for understanding the KC offset spiking phenomenon.

4.8 Understanding sparse olfactory codes in Kenyon cells

In this section, we account for the ways in which the network activity among the PNs and KCs following an odor stimulus generates KC activity that exhibits population sparsity and is an encoding of odor state information. Remarkably, both of these properties can be accounted for using a simple combinatorial model of the network activity, relying on two assumptions about the working of locust olfaction:

(1) Each KC is connected to approximately 50% of the PN population [42].

(2) The oscillatory activity of the PN population in response to an odor stimulus segments PN activity into temporal integration windows. A KC spikes if approximately 100 of its connected PNs spike within one of these integration windows.

We show that, in this model, PN stimulus response activity levels within the range observed experimentally generate information-preserving activity patterns in the KCs using only a small number of the 50000 KCs.

4.8.1 Network model specification

PN population: The collection of PNs is a set of points X_{PN}, with $|X_{PN}| = N_{PN} = 830$.

KC population: The collection of KCs is a set of points X_{KC}, with $|X_{KC}| = N_{KC} = 50000$.

PN→KC synapses: The PN→KC synapses are modeled as a set of E_{PN}^{KC} of

pairs (x, y), where $x \in X_{PN}$ and $y \in Y_{PN}$. We assume each of the KCs has $K = 415$ (50% of N_{PN}) such synapses, for a total of KN_{KC} synapses.

Synapse distribution: For each KC y_j, the collection of K synapses targeting y_j has as sources K PNs chosen uniformly at random from X_{PN}, independently for all KCs y_j. Each KC is therefore associated with a subset $I_j \subseteq X_{PN}$ of source PNs it is connected to.

Stimulus responses: Following the introduction of olfactory stimulus i into the locust headspace, the PNs and KCs emit a series of spikes at recorded times. These responses are aggregated into T disjoint time windows of length δt ms to form response activity vectors $\boldsymbol{A}_i(t) \in \{0,1\}^{N_{PN}}$ and $\boldsymbol{B}_i(t) \in \{0,1\}^{N_{KC}}$, where $A_{j,i}(t) = 1$ if PN j emitted a spike in time bin t and $B_{j,i}(t) = 1$ if KC j emitted a spike in time bin t. In the following analysis, we assume that the PN response activities are independently and identically distributed across all time bins, and it is therefore convenient to fix t and analyze

$$\boldsymbol{a}_i = \boldsymbol{A}_i(t) \tag{4.21}$$

and

$$\boldsymbol{b}_i = \boldsymbol{B}_i(t). \tag{4.22}$$

PN active sets: Since $\boldsymbol{a}_i$ is a binary vector, it partitions X_{PN} into two sets, an **active set** S_i and its complement, so that for all $x_j \in X_{PN}$, $a_{j,i} = 1$ if $x_j \in S_i$ and $a_{j,i} = 0$ otherwise. Define $M_i = |S_i|$ to be the size of this active set.

KC activity: In this model, KC j spikes if τ or more of its upstream PNs are active in the same time bin. In the development to follow, we set $\tau = 100$.

Therefore,

$$b_{j,i} = \begin{cases} 1, & |S_i \cap I_j| \geq \tau \\ 0, & |S_i \cap I_j| < \tau \end{cases} \tag{4.23}$$

4.8.2 Explaining KC population sparsity

The activity of the KCs exhibits **population sparsity** in its response to stimuli as well. On the order of 10 KCs out of a population of approximately 50,000 will respond to a given stimulus presentation, yet the identity and concentration of a stimulus' odorants can be inferred from KC population activity at least as well as from PN population activity [96]. The population of KCs in the MB is also the center of associative olfactory learning in the locust [70], lending weight to the theoretical results concerning the efficacy of sparse coding for associative learning. The ability to construct such a low-dimensional representation of olfactory state is remarkable, given the intrinsically-high-dimensional nature of olfactory signals [76]. What is the relationship between PN activity and number of active KCs in the model we have introduced? As we demonstrate below, an interesting pattern emerges when $K = 415$ and $\tau = 100$. When $M_i < 150$, the probability of observing one or more spikes across the N_{KC} KCs is very low; when $M_i = 150$ this probability is 0.2. As M_i increases from 150 to 160, the KC spike count distribution changes, so that when $M_i = 160$ a mean of 13 spiking KCs are expected. This relationship between active PN count and the number of responding KCs fits closely with observations from experiments [79, 42]. Figure 4.12 provides a visualization of the evolution of the KC spike count distribution, Figure 4.13 shows the change in the expected KC spike

count, and Figure 4.14 shows how the probability of observing more than 0 spikes increases with M_i.

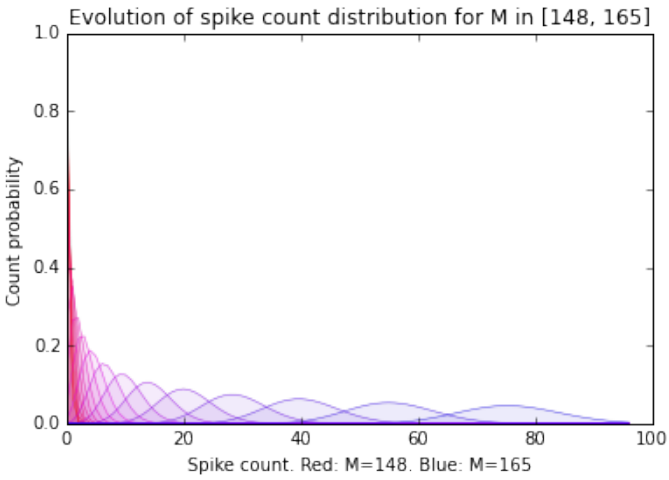

Figure 4.12: Examples of single time bin KC spike count distributions for varying values of M, the number of active PNs. Illustrated is an overlay of spike count distributions with M varying from 148 to 165. Small M values are in red, transitioning to blue as M increases. $M = 150$ marks an important transition between no KC activity and biologically-plausible levels of KC activity.

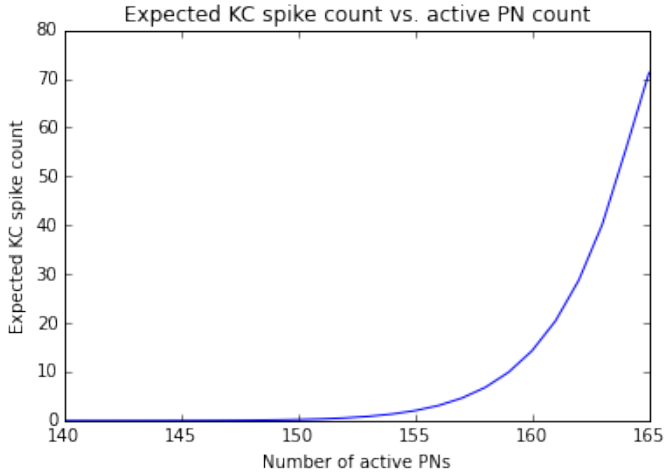

Figure 4.13: A plot of the expected number of spiking KCs as a function of the number of active PNs.

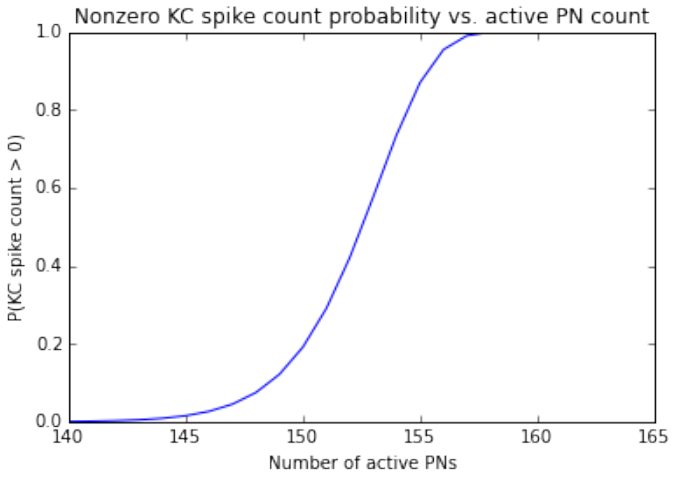

Figure 4.14: A plot of the probability of observing more than 0 spiking KCs as a function of the number of active PNs.

The derivation of this family of distributions follows. If a KC receives K synaptic inputs, there are $\binom{N_{PN}}{K}$ distinct input configurations. If in a given time bin there are M active PNs, we calculate that the probability of a KC having exactly k active inputs is is

$$\rho_k(M, N_{PN}, K) = P(k \text{ active inputs}) = \frac{\binom{M}{k}\binom{N_{PN}-M}{K-k}}{\binom{N_{PN}}{K}}, \tag{4.24}$$

and refer to this probability as ρ_k when there is no danger of ambiguity due to the omission of parameters. Therefore, the probability $\rho(M, N_{PN}, K)$ of a KC spiking is the probability of receiving τ or more of the M active inputs.

$$\rho(M, N_{PN}, K) = P(\text{KC spikes}) = \sum_{k=\tau}^{M} \rho_k(M) = \sum_{k=\tau}^{M} \frac{\binom{M}{k}\binom{N_{PN}-M}{K-k}}{\binom{N_{PN}}{K}}. \tag{4.25}$$

The value of ρ is highly sensitive to the value of M; $\rho(140) \approx 1.5 \times 10^{-8}$, while $\rho(160) \approx 2.9 \times 10^{-4}$. Figure 4.15 illustrates the relationship between ρ and M. Note

that $\rho(M)$ scales at least exponentially with M across this range.

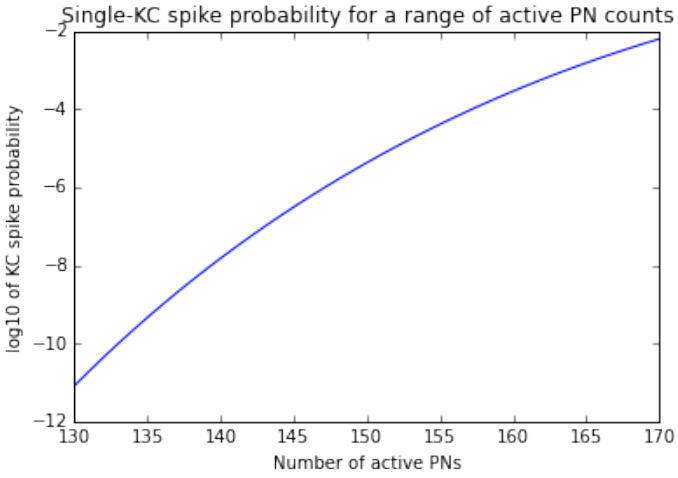

Figure 4.15: Single Kenyon cell spike probabilities ρ as a function of active projection neuron counts M, displayed as a log plot.

If each of N_{KC} KCs receives synaptic input according to this same random model, we can compute the distribution of population-wide spike counts in a single time bin. For a fixed M and $\rho = \rho(M)$,

$$P(s \text{ total KC spikes}) = \binom{N_{KC}}{s} \rho^s (1-\rho)^{N_{KC}-s}. \tag{4.26}$$

This is the family of distributions, parameterized by M, which is described in Figures 4.12, 4.13, and 4.14. Again, note that 150 active PNs marks the transition from little-to-no KC activity to KC activity levels in line with observations from experiments.

The small fluctuations of M between successive time bins then elicit substantial changes in KC population activity during the PNs' stimulus response. Deriving an exact estimate using Equation (4.26) depends on the exact time course of M over a response, but across the ten 50ms time bins that comprise the first 500ms of

a stimulus response, this method predicts an observation of approximately 10-100 active KCs across the PN response trajectory. Therefore, this random PN→KC synaptic model agrees well with biological data about the firing threshold of each KC and the response statistics of the PN population, and we may understand this combination of factors as the origin of the sparsity in the KCs' sparse population code.

4.8.3 Explaining KC olfactory encoding

The fact that the KC population is able to preserve the olfactory information encoded in the PNs is perhaps not surprising - the 60-fold expansion in population size provides greatly increased coding capacity, even given the temporal aspects of PN stimulus responses. Here, we provide one demonstration of this network model's ability to explain KC olfactory coding capabilities by bounding the distribution of $\mathcal{D}_{KC} = ||\boldsymbol{B}_1 - \boldsymbol{B}_2||_1$ for two KC stimulus responses following distinct PN stimulus responses, conditioned on the size of the overlap of the stimulus responses. In this way, we characterize how easily discriminated two KC population responses are, depending on the similarity of the PN upstream activity.

Given two single-time PN activity states $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$ with active sets S_1 and S_2, $|S_1| = M_1$ and $|S_2| = M_2$, we are interested in the probability that a single downstream KC will spike in response to one of $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$ but not both, as this event allows us to use that KC to distinguish between the two.

The sets S_1 and S_2 share $S_{12} = S_1 \cap S_2$, an overlap representing PNs which are

active in both $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$. Assume $|S_{12}| = L$. For convenience, define $\sigma_1 = S_1 - S_{12}$, $|\sigma_1| = m_1 = M_1 - L$, and $\sigma_2 = S_2 - S_{12}$, $|\sigma_2| = m_2 = M_2 - L$. Define also $\nu = N_{PN} - L$, the number of PNs outside of S_{12}. A diagram of these sets and their labels can be seen in Figure 4.16.

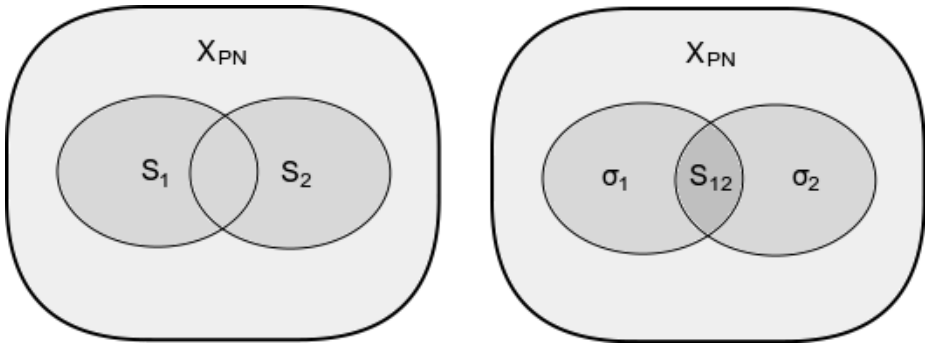

Figure 4.16: A diagram of the subsets of X_{PN} associated with a pair of single-time PN response states $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$. For either $\boldsymbol{a}_i$ there is an active set S_i such that $a_{j,i} = 1$ if and only if PN $x_j \in S_i$. Their intersection is labeled S_{12}, and $\sigma_i = S_i - S_{12}$. Set sizes are $|X_{PN}| = N_{PN}$, $|S_i| = M_i$, $|S_{12}| = L$, $|\sigma_i| = m_i$, $|X_{PN} - S_{12}| = \nu$.

To analyze the activity of KC y_j in response to the PN states $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$, consider the intersection of I_j, the set of K PNs that form the sources of y_j's incoming synapses, with S_1 and S_2. The members of I_j are chosen uniformly at random from X_{PN} in the same way as S_1 and S_2, and we can compute the distribution of sizes of the overlap sets $\lambda_j = I_j \cap S_{12}$, $\iota_{j,1} = I_j \cap \sigma_1$, and $\iota_{j,2} = I_j \cap \sigma_2$. A diagram of these sets can be seen in Figure 4.17.

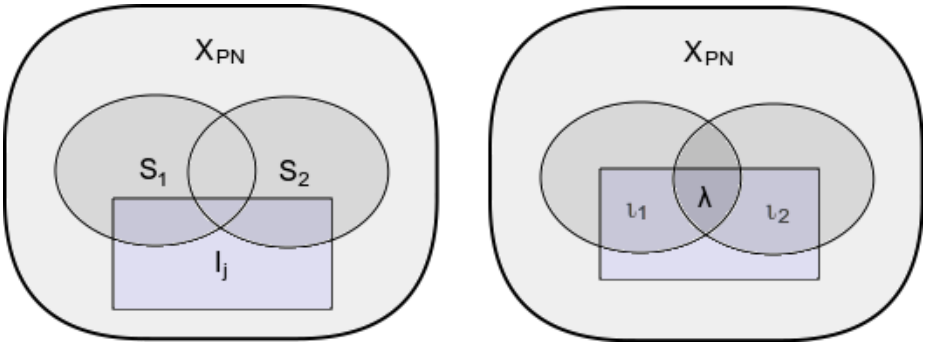

Figure 4.17: A diagram of the subsets of X_{PN} associated with a pair of active PN sets S_1 and S_2 and KC y_j's source PN set I_j. Subscript j's for λ_j, $\iota_{j,1}$, and $\iota_{j,2}$ are omitted for visual clarity.

To understand the relationship between PN response overlap S_{12} and the distribution of $\mathcal{D}_{KC}$, we seek to calculate the KC response difference probabilities

$$\Pi_{L,\ell,j} = P_{L,\ell}(b_{j,1} \neq b_{j,2}) = P\Big(b_{j,1} \neq b_{j,2} \;\Big|\; |S_{12}| = L,\, |\lambda_j| = \ell\Big) \tag{4.27}$$

for any j and a range of values of L and ℓ. Similar conditional probabilities show up in further calculations, so for any event E_j depending on KC j define

$$P_{L,\ell}(E_j) = P\Big(E_j \;\Big|\; |S_{12}| = L,\, |\lambda_j| = \ell\Big). \tag{4.28}$$

To compute the quantity in Equation (4.27), note that if the firing threshold for any KC is τ active upstream PNs, PN state $\boldsymbol{a}_i$ will generate a spike in KC y_j if at least $\tau_j = \tau - \ell_j$ elements of I_j are contained within σ_1, i.e., $|\iota_i| \geq \tau_j$. Conversely, there

will be no spike if $|\iota_i| < \tau_j$. Since $b_{j,i} \in \{0, 1\}$ for all j and i, we compute that

$$\Pi_{L,\ell,j} = P_{L,\ell}(b_{j,1} = 1 \wedge b_{j,2} = 0) + P_{L,\ell}(b_{j,1} = 0 \wedge b_{j,2} = 1) \tag{4.29}$$

$$= P_{L,\ell}(|\iota_1| \geq \tau_j \wedge |\iota_2| < \tau_j) + P_{L,\ell}(|\iota_2| \geq \tau_j \wedge |\iota_1| < \tau_j). \tag{4.30}$$

Each term in this sum can be analysed in the same way, so we now focus on the first of them.

$$P_{L,\ell}(|\iota_1| \geq \tau_j \wedge |\iota_2| < \tau_j) = \sum_{k_1=\tau_j}^{m_1} \sum_{k_2=1}^{\tau_j - 1} P_{L,\ell}(|\iota_1| = k_1 \wedge |\iota_2| = k_2). \tag{4.31}$$

Each term of the sum in the RHS of Equation (4.31) can be computed combinatorially. There are $\binom{m_1}{k_1}$ ways to choose k_1 elements from σ_1, $\binom{m_2}{k_2}$ ways to choose k_2 elements from σ_2 and $\binom{\nu - m_1 - m_2}{\kappa - k_1 - k_2}$ ways to distribute the remaining PN→KC synapses to the elements of $X_{PN} - S_{12}$. There are $\binom{\nu}{\kappa}$ ways to distribute the κ synapses amongst all ν elements of $X_{PN} - S_{12}$. Therefore,

$$P_{L,\ell}(|\iota_1| = k_1 \wedge |\iota_2| = k_2) = \frac{\binom{m_1}{k_1}\binom{m_2}{k_2}\binom{\nu - m_1 - m_2}{\kappa - k_1 - k_2}}{\binom{\nu}{\kappa}}, \tag{4.32}$$

so

$$P_{L,\ell}(b_{j,1} = 1 \wedge b_{j,2} = 0) = \sum_{k_1=\tau_j}^{m_1} \sum_{k_2=1}^{\tau_j - 1} \frac{\binom{m_1}{k_1}\binom{m_2}{k_2}\binom{\nu - m_1 - m_2}{\kappa - k_1 - k_2}}{\binom{\nu}{\kappa}}. \tag{4.33}$$

Combining Equations (4.29) and (4.33), we get that

$$\Pi_{L,\ell,j} = \sum_{k_1=\tau_j}^{m_1} \sum_{k_2=1}^{\tau_j - 1} \frac{\binom{m_1}{k_1}\binom{m_2}{k_2}\binom{\nu - m_1 - m_2}{\kappa - k_1 - k_2}}{\binom{\nu}{\kappa}} + \sum_{k_1=\tau_j}^{m_2} \sum_{k_2=1}^{\tau_j - 1} \frac{\binom{m_2}{k_1}\binom{m_1}{k_2}\binom{\nu - m_1 - m_2}{\kappa - k_1 - k_2}}{\binom{\nu}{\kappa}}. \tag{4.34}$$

An example plot of this quantity, computed with $L = 30$, $M_1 = M_2 = 160$, and $\ell \in [10, 20]$, can be found in Figure 4.18.

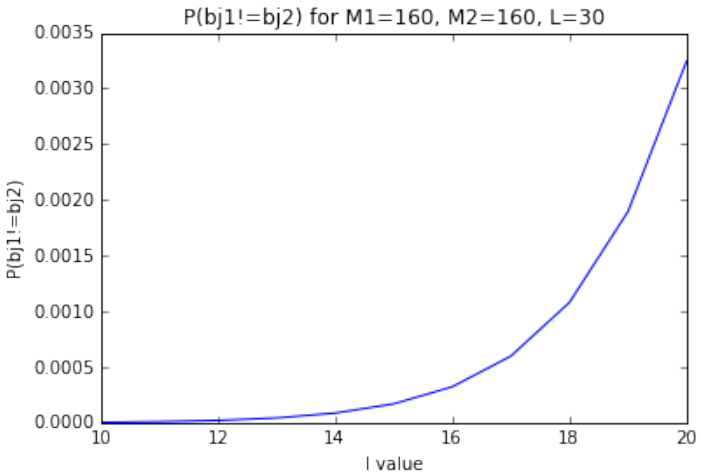

Figure 4.18: A plot of $\Pi_{L,\ell,j} = P_{L,\ell}(b_{j,1} \neq b_{j,2})$ for $L = 30$, $M_1 = M_2 = 160$, and $\ell \in [10, 20]$.

Note that when ℓ is constant across a KC population, $\Pi_{L,\ell,j}$ is constant across its j index, and the mean population KC response difference will be $\Pi_{L,\ell,j} \cdot N_{KC}$. While it is unlikely that ℓ would be constant across a population, this figure gives a way to estimate the relationship between $\Pi_{L,\ell,j}$ and the difference in activity levels between two KC population responses.

We can compute this for ranges of M_1, M_2, L, and ℓ matching the ranges observed or predicted from experiments. Given S_1 and S_2 distributed uniformly at random in X_{PN}, we have

$$P(|S_{12}| = k) = \frac{\binom{M_1}{k}\binom{N_{PN}-M_1}{M_2-k}}{\binom{N_{PN}}{M_2}}. \tag{4.35}$$

An example of this distribution with $M_1 = M_2 = 150$ can be seen in Figure 4.19.

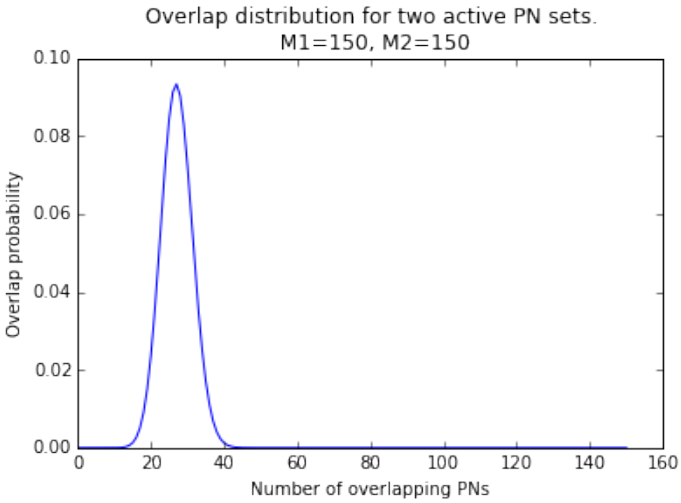

Figure 4.19: An overlap probability distribution for two populations of 150 PNs chosen uniformly at random from a population of 830. Approximately 98.6% of this distribution is contained in the interval $[17, 37]$.

Slightly more than 98.6% of this distribution is supported on the interval $L \in [17, 37]$. We therefore focus on L of this size in the following analysis, though there may be reason to consider other ranges since true PN response distributions are unlikely to be well-described by this model.

For a choice of L, one next considers the distribution of $\lambda_j = |S_{12} \cap I_j|$. Since $K = 0.5N_{PN}$ and the elements of I_j are selected uniformly at random from X_{PN}, we expect that $\ell \approx 0.5L$ when L is large. For smaller L, deviations from this mean may be more substantial. To understand the values of $\Pi_{L,\ell,j}$ produced by our model in the parameter ranges suggested by experiment, we plot $\Pi_{L,\ell,j}$ for $M_1 = M_2 = M$ in the range $[140, 160]$, for $L \in \{20, 30, 40\}$ and ranges of ℓ distributed around $0.5L$. The results can be seen in Figures 4.20, 4.21, and 4.22.

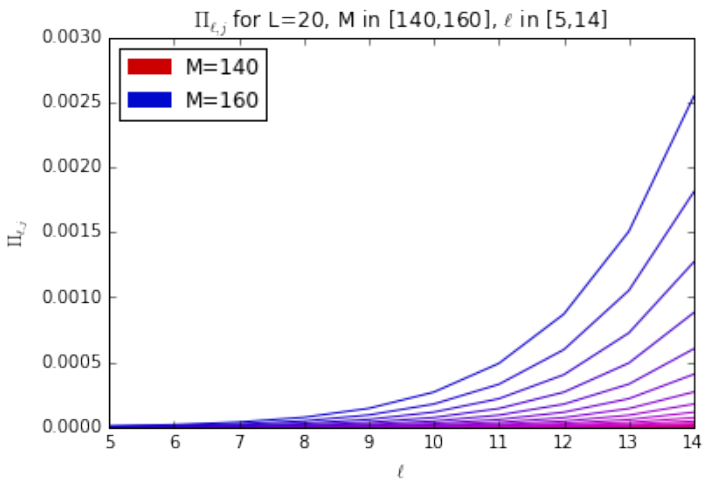

Figure 4.20: A family of $\Pi_{L,\ell,j}$ distributions over $\ell \in [5, 14]$. For reference, if $\Pi_{\ell,j} = 0.0015$ for each $y_j \in X_{KC}$, the mean of $\mathcal{D}_{KC} = 75$.

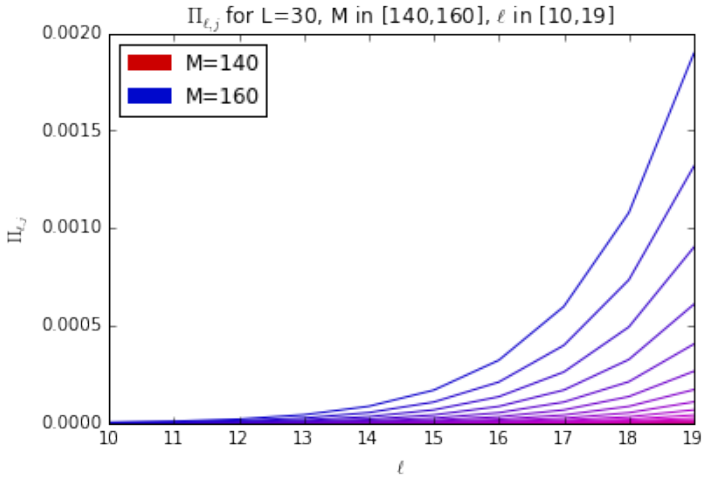

Figure 4.21: A family of $\Pi_{L,\ell,j}$ distributions over $\ell \in [10, 19]$. For reference, if $\Pi_{\ell,j} = 0.001$ for each $y_j \in X_{KC}$, the mean of $\mathcal{D}_{KC} = 50$.

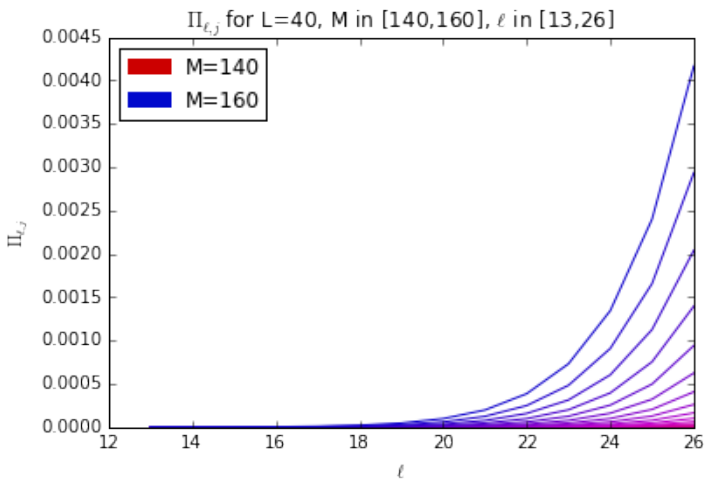

Figure 4.22: A family of $\Pi_{L,\ell,j}$ distributions over $\ell \in [13, 26]$. For reference, if $\Pi_{\ell,j} = 0.002$ for each $y_j \in X_{KC}$, the mean of $\mathcal{D}_{KC} = 100$.

Having computed $\Pi_{\ell,j}$, one would like to use this quantity to calculate $\mathcal{D}_{KC}$. Keeping in mind the choice of notation that $\boldsymbol{b}_i = \boldsymbol{B}_i(t)$ for some fixed t, and defining $\Pi_{\ell,j}(t)$ to be $\Pi_{\ell,j}$ in time bin t, we note that

$$\mathcal{D}_{KC} = \sum_{t=1}^{T} ||\boldsymbol{B}_1(t) - \boldsymbol{B}_2(t)||_1 \tag{4.36}$$

$$= \sum_{t=1}^{T} \sum_{j=1}^{N_{KC}} |\boldsymbol{B}_{j,1}(t) - \boldsymbol{B}_{j,2}(t)| \tag{4.37}$$

and $|\boldsymbol{B}_{j,1}(t) - \boldsymbol{B}_{j,2}(t)|$ Bernoulli$(\Pi_{\ell,j}(t))$. Given the assumptions of independent and identical distributions of each KC's synaptic connections, and each time bin's PN responses, it follows that $\mathcal{D}_{KC}$ is a sum of these Bernoulli distributions. Since each $\Pi_{\ell,j}(t)$ is small, one may try and use Le Cam's theorem to approximate $\mathcal{D}_{KC}$ as a Poisson distribution with parameter $\mu = \sum_{t=1}^{T} \sum_{j=1}^{N_{KC}} \Pi_{\ell,j}(t)$ [72]. However, since the goodness of fit for this approximation depends on the particular ranges chosen for M_1, M_2, L, and ℓ, a useful but simpler approach is to compare $\mathcal{D}_{KC}$ with the

binomial distributions $B(N_{KC}T, \Pi)$, where $\Pi < \Pi_{\ell,j}(t)$ for each j and t. Then for any $k > 0$,

$$P\left(\mathcal{D}_{KC} < k\right) \leq P\left(B(N_{KC}T, \Pi) < k\right). \tag{4.38}$$

Equation (4.38) lets us perform the final computations to characterize the capacity of the KC responses to distinguish PN population states. We will use the following definition in that process:

Definition 12. *Two binary random vectors* $\boldsymbol{v}_1$ *and* $\boldsymbol{v}_2$ *are* $\boldsymbol{(d, \epsilon)}$***-distinguishable*** *if*

$$P\left(||\boldsymbol{v}_1 - \boldsymbol{v}_2||_1 < d\right) < \epsilon. \tag{4.39}$$

In effect, one sets a threshold d which determines when two vectors are "different enough" to be distinguished, and then calculates the probability that the vectors will be indistinguishable by that criterion. There is limited evidence from experiment for the values of d most relevant to the decoding of KC activity by downstream processing centers, but here we can present the relationship between d and the maximum values of ϵ for which the (d, ϵ)-distinguishability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$ holds. To use Equation (4.38) to get a lower bound on the distinguishability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$, we must choose a value for Π. Appropriate choices may be made using the information from Figures 4.20, 4.21, and 4.22. For values of ℓ close to $0.5L$, it appears that $\Pi_{\ell,j}$ ranges from 0.00001 to 0.0001. Additionally, since $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$ are time series of KC states, we must choose a value for T, the number of time bins to use. According to [79], KC activity rates significantly diminish after the first

500ms of an odor response, so given that $\delta t = 50$ms we choose $T = 10$ to model the distinguishability of this initial response trajectory. Figures 4.23, 4.24, and 4.25 show binomially-distributed lower bounds for $P(\mathcal{D}_{KC} < d)$ for $T = 10$ and a range of d values using $\Pi = 0.00001$, $\Pi = 0.00005$, and $\Pi = 0.0001$ respectively. Note the different scales on the horizontal axes.

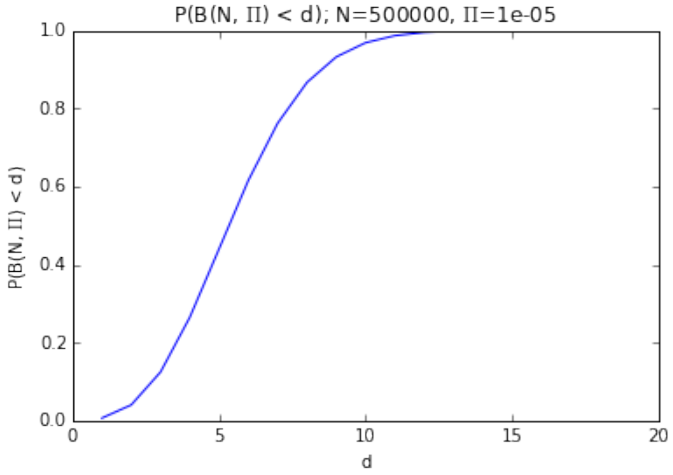

Figure 4.23: A plot of $B(N_{KC}T, \Pi)$ for $T = 10$ and $\Pi = 0.00001$. This quantity provides a lower bound for the (d, ϵ)-distinguishability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$.

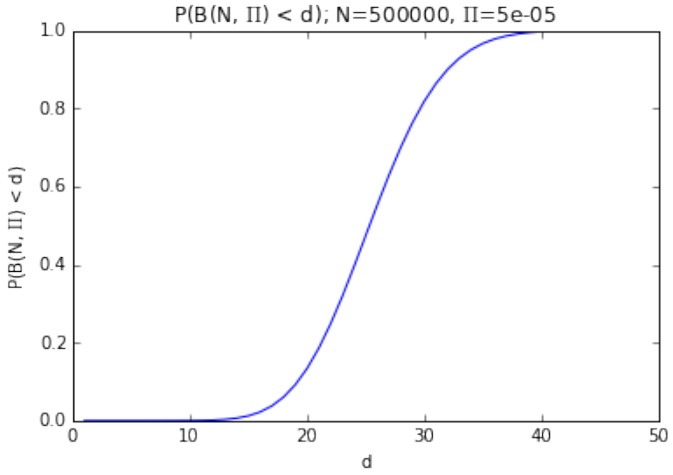

Figure 4.24: A plot of $B(N_{KC}T, \Pi)$ for $T = 10$ and $\Pi = 0.00005$. This quantity provides a lower bound for the (d, ϵ)-distinguishability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$.

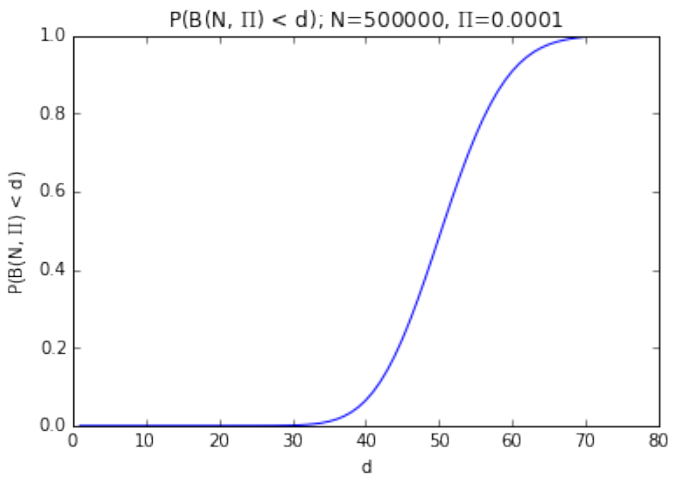

Figure 4.25: A plot of $B(N_{KC}T, \Pi)$ for $T = 10$ and $\Pi = 0.0001$. This quantity provides a lower bound for the (d, ϵ)-distinguishability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$.

Thus, if we can only bound $\Pi_{\ell,j}(t)$ from below by 0.00001, there is a high probability that $\mathcal{D}_{KC}$ will be concentrated near 0 and therefore $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$ will be difficult to distinguish. A bound of $\Pi = 0.00005$ assures us that at least $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$ will differ by at least 10 with very high probability. For a bound of $\Pi = 0.0001$, there is a high probability of $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$ differing in at least 30 locations.

This analysis demonstrates that the KC activity trajectories created by the combinatorial network model introduced in this chapter are capable of distinguishing between dissimilar PN activity trajectories.

4.9 Discussion

The Kenyon cells in the locust exhibit two patterns of activity deemed "sparse". As far as they are responsive primarily to changes in odor state, spiking mainly at onset and offset of a stimulus, the KCs exhibit lifetime sparsity. This behavior allows

for the detection of environmental concentration gradients across time, effectively providing a temporal counterpart to the spatial edge detectors common to models of visual processing. This selectivity is evidently rooted in the sensory adaptation process in ORNs, which reduces ORN and PN activity shortly after a stimulus change to drive PN activity below the KCs' firing thresholds. However, the absence of synchrony in PN activity upon offset calls into question contemporary understanding of the role of oscillations in eliciting a KC spike. We sought to supplement an investigation of this phenomenon via electrophysiological experimentation by using a computational model of these networks as dynamical systems. Model adjustments are still underway, but the validation process has developed into an effective collaboration for creating neuron population models with descriptive value for biological systems.

As far as olfactory information is carried in only a small percentage of the KC population at any point in time, the KCs also serve as a hallmark example of a network which exhibits population sparsity. While the processes governing biological neural network evolution bear little resemblance to the computational processes employing linear algebra to produce sparse signal decompositions, the advantages to each type of system are similar. Using the information learned during my collaboration with Dr Stopfer's lab, I devised a simplified binary activity model for the PN and KC populations in the locust. The behavior of this model is dependent on several parameters, but by investigating ranges of parameters most closely matching current biological data, I demonstrated that the model system produces KC output with activity levels in line with experiment, when PN activity levels are in the range

observed in experiment as well.

I also computed dissimilarity distributions for KC population responses to different PN trajectories, and showed that KC population activity in this model distinguishes between different stimuli with high probability. The discrimination analysis relied on additional assumptions which do not fully capture the range of conditions that the locust olfactory system may operate under, but given time and interest this model could be easily extended to accommodate more elaborate PN response and population-wide synaptic distribution models. Ultimately though, this model provides an increased understanding of how the nonlinear dynamical systems describing neural activity produce activity patterns which are sparse, and which are also representations of sensory signal information.

4.10 Bibliography

Bibliography

[1] I Aganj, A Bartesaghi, Ma Borgnia, HY Liao, G Sapiro, and S Subramaniam. Regularization for inverting the radon transform with wedge consideration. In *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on*, pages 217–220. IEEE, 2007.

[2] Uri M Ascher and Linda R Petzold. *Computer methods for ordinary differential equations and differential-algebraic equations*, volume 61. SIAM, 1998.

[3] C Avinash and M Slaney. *Principles of computerized tomographic imaging*. Society for Industrial and Applied Mathematics, 2001.

[4] R Bai, A. Cloninger, W Czaja, and P Basser. Efficient 2D MRI relaxometry using compressed sensing. *Journal of Magnetic Resonance*, 255:88–99, 2015.

[5] HB Barlow. Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1:371–394, 1972.

[6] Maxim Bazhenov, Mark Stopfer, Terrence J Sejnowski, and Gilles Laurent. Fast odor learning improves reliability of odor responses in the locust antennal lobe. *Neuron*, 46(3):483–492, 2005.

[7] P Binev, W Dahmen, R DeVore, P Lamby, D Savu, and R Sharpley. Compressed sensing and electron microscopy. In *Modeling Nanoscale Imaging in Electron Microscopy*, pages 73–126. Springer, 2012.

[8] Ake Björck. *Numerical methods for least squares problems*. SIAM, 1996.

[9] RN Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.

[10] Michael Brin and Garrett Stuck. *Introduction to dynamical systems*. Cambridge University Press, 2002.

[11] A Bruckstein, D Donoho, and M Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM review*, 51(1):34–81, 2009.

[12] Malcolm Burrows. *The neurobiology of an insect brain*. Oxford University Press Oxford, 1996.

[13] György Buzsáki. Large-scale recording of neuronal ensembles. *Nature neuroscience*, 7(5):446–451, 2004.

[14] T Cai, H Hirai, G Zhang, M Zhang, N Takahashi, H Kasai, LS Satin, RD Leapman, and AL Notkins. Deletion of ia-2 and/or ia-2β in mice decreases insulin secretion by reducing the number of dense core vesicles. *Diabetologia*, 54(9):2347–2357, 2011.

[15] R Calderbank, S Howard, and S Jafarpour. Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property. *Selected Topics in Signal Processing, IEEE Journal of*, 4(2):358–374, 2010.

[16] E Candes, L Demanet, D Donoho, and L Ying. Fast discrete curvelet transforms. *Multiscale Modeling & Simulation*, 5(3):861–899, 2006.

[17] EJ Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9):589–592, 2008.

[18] EJ Candes and J Romberg. l1-magic : Recovery of sparse signals via convex programming. Technical report, Caltech, 2005.

[19] EJ Candes and J Romberg. Practical signal recovery from random projections. In *Proc. SPIE Computational Imaging*, volume 5674, pages 76–86, 2005.

[20] EJ Candes and J Romberg. Sparsity and incoherence in compressive sampling. *Inverse problems*, 23(3):969, 2007.

[21] EJ Candès, 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.

[22] Emmanuel J Candès et al. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.

[23] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.

[24] S Chen, D Donoho, and M Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.

[25] X. Chen and A Powell. Almost sure convergence of the kaczmarz algorithm with random measurements. *Journal of Fourier Analysis and Applications*, 18(6):1195–1214, 2012.

[26] O Christensen. *An introduction to frames and Riesz bases*. Springer, 2002.

[27] A. Cloninger, W Czaja, R Bai, and P Basser. Solving 2d fredholm integral from incomplete measurements using compressive sensing. *SIAM Journal on Imaging Sciences*, 7(3):1775–1798, 2014.

[28] Alexander Cloninger. *Exploiting Data-Dependent Structure for Improving Sensor Acquisition and Integration*. PhD thesis, University of Maryland, College Park, 2014.

[29] F Colonna, K Easley, Gand Guo, and D Labate. Radon transform inversion using the shearlet representation. *Applied and Computational Harmonic Analysis*, 29(2):232–250, 2010.

[30] JA Connor and CF Stevens. Voltage clamp studies of a transient outward membrane current in gastropod neural somata. *The Journal of Physiology*, 213(1):21, 1971.

[31] Robert Coultrip, Richard Granger, and Gary Lynch. A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, 5(1):47–54, 1992.

[32] W Czaja and J Tanis. Kaczmarz algorithm and frames. *International Journal of Wavelets, Multiresolution and Information Processing*, 11(05):1350036, 2013.

[33] Leonardo Dagum and Rameshm Enon. Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.

[34] I Daubechies. Orthonormal bases of compactly supported wavelets. *Communications on pure and applied mathematics*, 41(7):909–996, 1988.

[35] M Davenport, M Duarte, Y Eldar, and G Kutyniok. Introduction to compressed sensing. *http://www.dfg-spp1324.de/download/preprints/preprint093.pdf*, 2011.

[36] Peter Dayan and Laurence F Abbott. *Theoretical Neuroscience*, volume 806. Cambridge, MA: MIT Press, 2001.

[37] Licurgo de Almeida, Marco Idiart, and John E Lisman. A second function of gamma frequency oscillations: an e%-max winner-take-all mechanism selects which cells fire. *The Journal of Neuroscience*, 29(23):7497–7503, 2009.

[38] MV De Hoop, H Smith, G Uhlmann, and RD Van der Hilst. Seismic imaging with the generalized radon transform: a curvelet transform perspective. *Inverse Problems*, 25(2):025005, 2009.

[39] S Deans. *The Radon Transform and some of its Applications*. Courier Dover Publications, 2007.

[40] A Delaney and Y Bresler. Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography. *Image Processing, IEEE Transactions on*, 7(2):204–221, 1998.

[41] R DeVore. Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4):918–925, 2007.

[42] Patricia M DiLorenzo and Jonathan D Victor. *Spike timing: mechanisms and function*. CRC Press, 2013.

[43] D Donoho and X Huo. Uncertainty principles and ideal atomic decomposition. *Information Theory, IEEE Transactions on*, 47(7):2845–2862, 2001.

[44] D Donoho, Y Tsaig, I Drori, and JL Starck. Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 58(2):1094–1121, 2012.

[45] J A Fessler and B Sutton. Nonuniform fast fourier transforms using min-max interpolation. *Signal Processing, IEEE Transactions on*, 51(2):560–574, 2003.

[46] Michel Fortin and Roland Glowinski. *Augmented Lagrangian methods: applications to the numerical solution of boundary-value problems*. Elsevier, 2000.

[47] J Frank. *Electron tomography: methods for three-dimensional visualization of structures in the cell*. Springer, 2006.

[48] D. Ge, DX JiGe, and Y Ye. A note on the complexity of lp minimization. *Mathematical programming*, 129(2):285–299, 2011.

[49] T Goldstein and S Osher. The split bregman method for l1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.

[50] A Gopinath, G Xu, D Ress, O Oktem, and S Subramaniam. Shape-based regularization of electron tomographic reconstruction. *Medical Imaging, IEEE Transactions on*, 2012.

[51] B Goris, T Roelandts, KJ Batenburg, H Heidari Mezerji, and S Bals. Advanced reconstruction algorithms for electron tomography: From comparison to combination. *Ultramicroscopy*, 2012.

[52] B Goris, W Van den Broek, KJ Batenburg, H Heidari Mezerji, and S Bals. Electron tomography based on a total variation minimization reconstruction technique. *Ultramicroscopy*, 113:120–130, 2012.

[53] D Gottlieb, B Gustafsson, and P Forssen. On the direct fourier method for computer tomography. *Medical Imaging, IEEE Transactions on*, 19(3):223–232, 2000.

[54] C Graydon, J Zhang, N Oesch, AA Sousa, R Leapman, and J Diamond. Passive diffusion as a mechanism underlying ribbon synapse vesicle release and resupply. *The Journal of Neuroscience*, 34(27):8948–8962, 2014.

[55] T Groves. Thick specimens in the cem and stem. resolution and image formation. *Ultramicroscopy*, 1(1):15–31, 1975.

[56] Bernd Grünewald. Differential expression of voltage-sensitive K+ and Ca2+ currents in neurons of the honeybee olfactory pathway. *Journal of experimental biology*, 206(1):117–129, 2003.

[57] K Grünewald, P Desai, Dennis C Winkler, J Heymann, D Belnap, W Baumeister, and A Steven. Three-dimensional structure of herpes simplex virus from cryo-electron tomography. *Science*, 302(5649):1396–1398, 2003.

[58] A Hafftka, H Celik, A Cloninger, W Czaja, and R Spencer. 2d sparse sampling algorithm for nd fredholm equations with applications to nmr relaxometry. *SampTA 2015 – Sampling Theory and Applications*, American University, Washington, D.C., 2015.

[59] Ariel Hafftka. *Tensor Completion for Multidimensional Inverse Problems with Applications to Magnetic Resonance Relaxometry*. PhD thesis, University of Maryland, College Park, 2016.

[60] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500, 1952.

[61] M Hohmann-Marriott, AA Sousa, A Azari, S Glushakova, G Zhang, J Zimmerberg, and R Leapman. Nanoscale 3D cellular imaging by axial scanning transmission electron tomography. *Nature Methods*, 6(10):729–731, 2009.

[62] Jiang Hsieh. Computed tomography: principles, design, artifacts, and recent advances. SPIE Bellingham, WA, 2009.

[63] M Hügel, H Rauhut, and T Strohmer. Remote sensing via l1-minimization. *Foundations of Computational Mathematics*, 14(1):115–150, 2014.

[64] Ron A Jortner, S Sarah Farivar, and Gilles Laurent. A simple connectivity scheme for sparse coding in an olfactory system. *The Journal of Neuroscience*, 27(7):1659–1669, 2007.

[65] Joby Joseph, Felice A Dunn, and Mark Stopfer. Spontaneous olfactory receptor neuron activity determines follower cell response properties. *The Journal of Neuroscience*, 32(8):2900–2910, 2012.

[66] U Benjamin Kaupp. Olfactory signalling in vertebrates and insects: differences and commonalities. *Nature Reviews Neuroscience*, 11(3):188–200, 2010.

[67] B Kolman and D Hill. *Introductory linear algebra*. Prentice-Hall, 2005.

[68] J Kremer, D Mastronarde, and J McIntosh. Computer visualization of three-dimensional image data using IMOD. *Journal of Structural Biology*, 116(1):71–76, 1996.

[69] Gitta Kutyniok et al. *Shearlets: Multiscale analysis for multivariate data*. Springer Science & Business Media, 2012.

[70] Gilles Laurent and Mohammad Naraghi. Odorant-induced oscillations in the mushroom bodies of the locust. *the Journal of Neuroscience*, 14(5):2993–3004, 1994.

[71] Gilles Laurent, Michael Wehr, and Hananel Davidowitz. Temporal representations of odors in an olfactory network. *The Journal of Neuroscience*, 16(12):3837–3847, 1996.

[72] Lucien Le Cam et al. An approximation theorem for the poisson binomial distribution. *Pacific J. Math*, 10(4):1181–1197, 1960.

[73] Rowan Leary, Zineb Saghi, Paul A Midgley, and Daniel J Holland. Compressed sensing electron tomography. *Ultramicroscopy*, 131:70–91, 2013.

[74] G Lesaja. Introducing interior-point methods for introductory operations research courses and/or linear programming courses. *Open Operational Research Journal*, 3:1–12, 2009.

[75] M Lustig, D Donoho, and J Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic resonance in medicine*, 58(6):1182–1195, 2007.

[76] Marcelo O Magnasco, Andreas Keller, and Leslie B Vosshall. On the dimensionality of olfactory space. *bioRxiv*, page 022103, 2015.

[77] D Mastronarde. Dual-axis tomography: an approach with alignment methods that preserve resolution. *Journal of Sructural Biology*, 120(3):343–352, 1997.

[78] S Matej, J Fessler, and I Kazantsev. Iterative tomographic image reconstruction using fourier-based forward and back-projectors. *Medical Imaging, IEEE Transactions on*, 23(4):401–412, 2004.

[79] Ofer Mazor and Gilles Laurent. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48(4):661–673, 2005.

[80] J Milne and S Subramaniam. Cryo-electron tomography of bacteria: progress, challenges and future prospects. *Nature Reviews Microbiology*, 7(9):666–675, 2009.

[81] J Murlis and CD Jones. Fine-scale structure of odour plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiological Entomology*, 6(1):71–86, 1981.

[82] R Narasimha, I Aganj, A Bennett, M Borgnia, D Zabransky, G Sapiro, S McLaughlin, J Milne, and S Subramaniam. Evaluation of denoising algorithms for biological electron tomography. *Journal of Sructural Biology*, 164(1):7, 2008.

[83] Peter D Nellist. Scanning transmission electron microscopy. In *Science of Microscopy*, pages 65–132. Springer, 2007.

[84] A Neumaier. Solving ill-conditioned and singular linear systems: A tutorial on regularization. *Siam Review*, 40(3):636–666, 1998.

[85] Bruno A Olshausen et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.

[86] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37(23):3311–3325, 1997.

[87] Bruno A Olshausen and David J Field. Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14(4):481–487, 2004.

[88] WJ Palenstijn, KJ Batenburg, and J Sijbers. Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs). *Journal of Structural Biology*, 176(2):250–253, 2011.

[89] Javier Perez-Orive, Maxim Bazhenov, and Gilles Laurent. Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *The Journal of Neuroscience*, 24(26):6037–6047, 2004.

[90] M Radermacher. Weighted back-projection methods. In *Electron tomography*, pages 245–273. Springer, 2006.

[91] M Raginsky, R Willett, Z Harmany, and R Marcia. Compressed sensing performance bounds under poisson noise. *Signal Processing, IEEE Transactions on*, 58(8):3990–4002, 2010.

[92] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[93] L Rudin, S Osher, and E Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.

[94] Z Saghi, D Holland, R Leary, A Falqui, G Bertoni, A Sederman, L Gladden, and P Midgley. Three-dimensional morphology of iron oxide nanoparticles with reactive concave surfaces. a compressed sensing-electron tomography (cs-et) approach. *Nano letters*, 11(11):4666–4673, 2011.

[95] Claude E Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[96] Kai Shen, Sina Tootoonian, and Gilles Laurent. Encoding of mixtures in a simple olfactory system. *Neuron*, 80(5):1246–1262, 2013.

[97] Jonathan Richard Shewchuk. *An introduction to the conjugate gradient method without the agonizing pain*. Carnegie-Mellon University. Department of Computer Science, 1994.

[98] D Sima. *Regularization techniques in model fitting and parameter estimation*. PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, 2006.

[99] K Song, L Comolli, and M Horowitz. Removing high contrast artifacts via digital inpainting in cryo-electron tomography: An application of compressed sensing. *Journal of Structural Biology*, 178(2):108–120, 2012.

[100] AA Sousa, A Azari, G Zhang, and R Leapman. Dual-axis electron tomography of biological specimens: Extending the limits of specimen thickness with bright-field stem imaging. *Journal of structural biology*, 174(1):107–114, 2011.

[101] Mark Stopfer, Vivek Jayaraman, and Gilles Laurent. Intensity versus identity coding in an olfactory system. *Neuron*, 39(6):991–1004, 2003.

[102] Gilbert Strang and Wellesley-Cambridge Press. *Introduction to linear algebra*, volume 3. Wellesley-Cambridge Press Wellesley, MA, 1993.

[103] T Strohmer and R. Vershynin. A randomized kaczmarz algorithm with exponential convergence. *Journal of Fourier Analysis and Applications*, 15(2):262–278, 2009.

[104] J Tropp and A Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 53(12):4655–4666, 2007.

[105] Joel Tropp. Greed is good: Algorithmic results for sparse approximation. *Information Theory, IEEE Transactions on*, 50(10):2231–2242, 2004.

[106] E Van Den Berg and M Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.

[107] A Veeraraghavan, A Genkin, S Vitaladevuni, L Scheffer, S Xu, H Hess, R Fetter, M Cantoni, G Knott, and D Chklovskii. Increasing depth resolution of electron microscopy of neural circuits using sparse tomographic reconstruction. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1767–1774. IEEE, 2010.

[108] C Vogel and M Oman. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17(1):227–238, 1996.

[109] D Walnut. *An introduction to wavelet analysis*. Springer, 2002.

[110] Rongrong Wang. *Global Geometric Conditions on Sensing Matrices for the Success of L1 Minimization Algorithm*. PhD thesis, University of Maryland, College Park, 2013.

[111] David B Williams and C Barry Carter. *The transmission electron microscope*. Springer, 1996.

[112] Q Xu, HY Yu, XQ Mou, L Zhang, J Hsieh, and G Wang. Low-dose x-ray ct reconstruction via dictionary learning. *Medical Imaging, IEEE Transactions on*, 31(9):1682–1697, 2012.

[113] Alon Zaslaver, Idan Liani, Oshrat Shtangel, Shira Ginzburg, Lisa Yee, and Paul W Sternberg. Hierarchical sparse coding in the sensory system of caenorhabditis elegans. *Proceedings of the National Academy of Sciences USA*, 112(4):1185–1189, 2015.