Understanding the nervous system as an information processing machine:

Dense, nonspecific, canonical microcircuit architecture of inhibition in neocortex

and

A neural circuit for angular velocity computation

Adam Max Packer

Submitted in partial fulfillment of the
Requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2011

ABSTRACT

Understanding the nervous system as an information processing machine:

Dense, nonspecific, canonical microcircuit architecture of inhibition in neocortex

and

A neural circuit for angular velocity computation

Adam Max Packer

This thesis is the combination of two separate lines of work linked by one common goal: understanding the nervous system as an information-processing machine. David Marr (1982) put forth the idea that in order to fully understand an information-processing machine one must understand it at three separate levels. The computational *goal* of the system must be understood separately from the *algorithm* by which it is computed and the *hardware* in which it is computed. During my time as a graduate student I have been fortunate enough to work on two different levels in two very different systems. Chapter 1 focuses on the hardware of neural circuitry, specifically on how inhibitory interneurons connect to excitatory neurons. Chapter 2 focuses on the algorithmic problem of how flies could use gyroscopic sensors to calculate angular velocity.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I am grateful to my advisor, Rafael Yuste, who initially offered me a job as a technician after Hurricane Katrina. I found a mentor whose excitement for science is unfettered. I appreciate that he encouraged my curiosity in varied scientific endeavors—often far from the core of my thesis!

Working in the Yuste laboratory is what convinced me to pursue a PhD. It is a fantastic environment for doing science and I thank my collaborators in the lab over the years for their support, encouragement, and friendship.

A special thanks goes to Samuel Snider whose energy and intelligence made Saturdays in the lab, thinking about how flies fly, most enjoyable.

I thank my committee members, Larry Abbott, Dmitri Chklovskii, Wesley Grueber, and Darcy Kelley for the fruitful discussions and critical insights they shared in reviewing this dissertation.

I acknowledge my parents and sisters for their unwavering support in all that I choose to do. I thank my wife, Lelia Packer, whose friendship, honesty, sensitivity, and love have enriched my life in ways which simply cannot be quantified.

DEDICATION

To my best friend, my wife, Lelia Packer.

PREFACE

Chapter 1 is a manuscript currently in preparation. Chapter 2 has been previously published as *A neural circuit for angular velocity computation* in *Frontiers in Neural Circuits* **4**: 123 (Snider SB, Yuste R, and Packer AM, 2010).

CHAPTER 1

Summary

GABAergic interneurons play a major role in the function of the mammalian neocortex but their circuit connectivity is still poorly understood. We use two-photon RuBi-Glutamate uncaging to optically map how the largest population of cortical interneurons, the parvalbumin-positive cells (PV+), are connected to pyramidal cells (PCs). We find locally dense connectivity from PV+ interneurons onto PCs across cortical areas and layers. In many experiments, all nearby PV+ cells are connected to every local PC sampled. In agreement with this, we find no evidence for connection specificity, as PV+ interneurons contact PC pairs similarly regardless of whether they are synaptically connected or not. Moreover, using anatomical reconstructions of connected neurons, we find that the overlap of PV+ axons with PC dendrites (Peters' Rule) accurately predicts the spatial profiles of connectivity from PV+ to PCs, without any additional mechanisms. We conclude that the canonical microcircuit architecture for PV+ interneurons, and probably inhibition in general, is an unspecific, densely homogenous matrix covering all postsynaptic targets nearby.

Introduction

The mammalian neocortex is a marvel of biological engineering capable of impressive computational feats. This structure appears to be organized in a stereotypical and hierarchical manner, presumably to achieve a high level of parallel processing power (Mountcastle, 1982). Unfortunately, the structure of cortical microcircuits appears to be an "impenetrable jungle" (Ramon y Cajal, 1923) because of the dense mixing of different types of neurons, which makes deducing the way it operates difficult. Indeed, although cortical preparations have been studied anatomically and physiologically for over a century, remarkably little has been revealed about the fine-scale organization of these neurons into circuits. It is clear that the cortex is divisible into layers and areas, each with distinct connectivity patterns with other layers and areas, presumably resulting in specific functions for each element (Douglas and Martin, 2004). It is also known that individual neurons belong to particular classes, although there is still no valid classification for many neurons (Ascoli et al., 2008), or theoretical framework describing how they function in a concerted way to process information.

Part of the reason for this ignorance is the experimental difficulty in revealing how these neurons are connected to each other to form functional circuits (Crick, 1979). In the last decade several new techniques have been developed to reveal synaptically connected neurons, including paired whole-cell recordings (Thomson and Lamy, 2007), serial EM reconstructions (Bock et al 2010), viral tracing (Wickersham et al., 2007), calcium imaging probing (Peterlin et al., 2000), reverse correlation imaging (Aaron and Yuste, 2006) and photostimulation with caged glutamate (Callaway and Katz, 1993). We have recently introduced a variant of photostimulation that uses two-photon uncaging of

glutamate to map connected neurons with single cell resolution (Nikolenko et al., 2007). When combined with uncaging of RuBi-Glutamate, a recently developed caged glutamate compound (Fino et al., 2009), two-photon photostimulation can be applied to map inhibitory connections. Using Rubi-Glutamate, we recently mapped the synaptic circuits of somatostatin-positive interneurons in the upper layers of mouse frontal cortex, finding a very dense and unspecific connectivity from these interneurons to neighboring pyramidal cells (Fino and Yuste, 2011). This dense connectivity has not been reported before and we wondered if it was a peculiar feature of the somatostatin-positive interneurons or of the frontal cortex.

To answer these questions, we have now mapped the connectivity from fast-spiking parvalbumin-positive (PV+) basket cells — the largest population of interneurons in the neocortex — to pyramidal neurons, and compared the synaptic circuits in which they are embedded in different cortical layers and areas. We find locally dense and unspecific connections from PV+ interneurons on to pyramidal cells (PCs) in two different cortical areas (somatosensory and frontal) as well as in two different layers (2/3 and 5). Although we use brain slices, in which many connections are severed, we often find that every local PV+ cell is connected to every PC sampled. We also find that Peters' Rule (i.e., the spatial overlap of axons and dendrites) is sufficient to predict the spatial pattern of connectivity of PV+ interneurons, supporting the idea that these circuits are not specific. Our results support the view that the dense architecture of interneuron connectivity, whereby a “blanket of inhibition" is extended on local pyramidal cells, is a universal feature in the design of neocortical microcircuits.

Results

Two-photon photostimulation of parvalbumin interneurons

To study the connectivity from a well-defined subset of interneurons onto nearby pyramidal cells (PCs), we used the G42 transgenic mouse line in which parvalbumin-expressing interneurons are labeled with GFP (Chattopadhyaya et al., 2004). In coronal brain slices, we targeted and recorded from 91 GFP-positive cells, and performed anatomical and electrophysiological characterizations of these neurons. Anatomically, all targeted PV+ interneurons had basket cells features with promiscuous and densely branching axons (n=15 reconstructions; Figure 1-1A). Electrophysiologically, GFP-positive neurons were all fast-spiking cells, firing at frequencies greater than 50 Hz (97 ± 6 Hz, n=42) (Ascoli et al., 2008). They also had high rheobases (312 ± 102 pA, n=42), strong fAHP currents, and rectifying I/V curves (Woodruff et al., 2009)(Figure 1-1B-C).

To map synaptic connections from PV+ interneurons onto postsynaptic PCs, we followed the protocol initially developed by (Nikolenko et al., 2007) and recently refined for inhibitory connections by (Fino and Yuste, 2011). This protocol relies on the ability to optically stimulate individual neurons with two-photon uncaging of RuBi-Glutamate (Fino et al., 2009). We achieved single-cell resolution by multiplexing the two-photon laser beam across time, focusing it to eight points around the target cell soma (Figure 1-1D), and across space, by using a diffractive optical element (Nikolenko et al., 2007) to split the laser beam into five beamlets (Figure 1-1E). This strategy preserved the integrity of the point-spread function of the two-photon beam (Figure S1-1), thus enabling single-cell precision. When photostimulating a recorded PV+ interneuron with

this method, one to several action potentials (APs) were elicited at intermediate laser powers (150 mW on sample, Figure 1-1F) while bursts of APs were elicited at greater powers (300 mW on sample, Figure 1-1G). Despite the high rheobase of these cells, these laser powers (>150mW) never failed to produce APs (APs in 21 out of 21 cells tested). APs induced by uncaging displayed the fAHP waveform and high frequency trains characteristic of fast-spiking cells. In addition, when the stimulation of the PV+ interneuron was repeated every five seconds, identical responses were produced each time (14.22 ± 0.08 millisecond delay to first spike, 11.6 ± 0.2 action potentials per stimulation, n=10 photostimulations). Thus, the photostimulation strategy we employed was precise and reliable, activating neurons with single-cell resolution.

Mapping monosynaptic connections from PV+ interneurons to pyramidal cells

We then mapped inputs from interneurons onto PCs, following the strategy developed by (Fino and Yuste, 2011). Briefly, a PC was patched while nearby PV+ interneurons were sequentially photostimulated (Figure 1-2A). Each PV+ interneuron was activated with a range of at least five laser powers, including the maximum of 300 mW on the sample. Postsynaptic responses in the PC that were time-locked to the stimulus were measured in voltage clamp (Figure 1-2B). The PC was initially held at +40mV to increase the driving force on inhibitory currents and then at -40mV to test that these postsynaptic currents were indeed inhibitory.

If the photostimulation of a PV+ cell generated a synaptic response in the PC with a GABAergic reversal potential (set to -80mV), we identified them as “true positive”, i.e., putatively connected neuron, since it was time-locked to the optical stimulation of an

inhibitory cell. To test whether interneurons generating these “true positive” responses were indeed connected to the postsynaptic PC cell, we patched them with a second electrode (Figure 1-2A-D, cell #4-red arrow), stimulated them with intracellular pulses of current to produce action potentials, and confirmed the generation of postsynaptic IPSCS (Figure 1-2E). Overall, we found that, in 9 out of 10 cases, true-positive PV+ cells generated short latencies IPSCs in the PC, indicative of monosynaptic transmission (Figure 1-2G-J; 1.12 ± 0.19 milliseconds, n=9 pairs). These presynaptic PV+ interneurons always fired APs when photostimulated (Figure 1-1G) and generated inhibitory responses with large outward currents (Figure 1-2D and 2H).

To explore if interneurons that did not generate a response in the PC could be potentially also connected to it, we also patched 8 photostimulated PV+ interneurons for which no response was observed ("negatives"). We confirmed that there was no synaptic connection in any of these cases (Figure 1-2K-N), concluding that negatives were "true negatives". In other words, we never observed any "false negatives" in which a PV+ interneuron which did not appear to be connected optically was in fact connected. Even PV+ interneurons that were very close to each other could be disambiguated, as evidenced by the complete lack of response (Figure 1-2B, cell #5, inset 2F) from cells directly adjacent to true positive, connected cells (Figure 1-2A, grey arrow vs. red arrow).

Besides true positive and negative responses, we occasionally observed unexpected postsynaptic currents that reversed direction at the glutamate reversal potential of 0mV (“false positives” responses). Most of these “false positives” occurred when directly stimulating in the vicinity of the recorded PC, just outside its soma (Figure 1-2O). In addition to a different reversal potential, these responses had slower kinetics

(Figure 1-2P) than the true positive responses from connected PV+ interneurons (Figure 1-2H). We also observed occasional false positive responses that arose from stimulating sites far away from the dendritic tree of the recorded neuron (Figure 1-2Q, black circle). We presume that these excitatory responses (Figure 1-2R) are due to PCs located nearby to the targeted PV+ interneuron, that were accidentally stimulated (Figure S1-2). These two types of false positive responses were rare and easily distinguished from true positive responses because of their glutamatergic reversal potentials.

In summary, we classified all stimulated interneurons into three types, based on their responses in the PC cell: (i) true positive cells, connected to the PC with monosynaptic IPSCs, (ii) negative cells, likely not connected and (iii) false positive cells, which generated contaminating EPSCs. For the remaining of the study, we used this classification to analyze the maps of input connectivity.

Dense local connectivity from PV+ interneurons onto pyramidal cells in different cortical layers and areas

We assembled a total of 82 input maps from layer 2/3 and 5 of somatosensory cortex as well as from layer 2/3 of frontal cortex, arising from recordings from 13 single PCs, 12 PC pairs, 11 PC triplets and 3 PC quadruplets. In optimal cases, we could map the inputs onto 4 PCs in different focal planes (Figure 1-3). Using this three dimensional strategy with simultaneous whole-cell recordings from multiple PCs, we were able to sample the connectivity from 2002 PV+ interneuron to PC pairs. On average, 24 ± 1 PV+ interneurons were tested for each map from 5 ± 0.2 focal planes.

We found that PCs received connections from most local PV+ interneurons, but not from more distantly located ones (Figures 3 and 4A). Overall, we tested a total of 21 ± 1, 29 ± 3, and 25 ± 2 PV+ interneurons in S2/3, S5, and F2/3 respectively ($p > 0.05$, Kruskal-Wallis), observing an average of 3 ± 0.4, 10 ± 1.4, and 7 ± 0.8 connections in S2/3, S5, and zF2/3 ($p<0.0001$ one-way ANOVA; $p<0.001$ S2/3 vs. S5, $p<0.01$ S2/3 vs. F2/3, $p<0.05$ S5 vs. F2/3 Tukey multiple comparison test; Figure 1-4B).

To analyze the dependence of the connectivity versus intersomatic distance, we defined the connection probability from PV+ interneurons to a PC as a percentage (connected interneurons / connected and unconnected interneurons) and measured it as a function of intersomatic distance between interneurons and PCs. We assembled input maps from rectangular slabs of cortex measuring 800 μm by 600 μm by 52 ± 3 μm (Figure 1-4C) so we were able to calculate intersomatic distance in three dimensions (Figure 1-4D, green bars.) False positives were rare and excluded from subsequent analyses (Figure 1-4D, black bars).

In all areas, the probability of connection was remarkably high in the local range and the probability of connection was inversely proportional to intersomatic distance between the PV+ interneuron and the PC (Figure 1-4E). An exponential fit of the connection probability versus intersomatic distance revealed spatial decay constants of 124, 183, and 180 μm for S2/3, S5, and F2/3 respectively ($R^2 = 0.92$, 0.91, and 0.86).

For intersomatic distances of less than 200 μm, the probability of connection from a PV+ interneuron to a PC was high, ranging from 43 ± 6% in somatosensory layer 2/3, 67 ± 6% in somatosensory layer 5, and up to 76 ± 4% in frontal layer 2/3 ($p<0.0001$ one-way ANOVA, $p<0.05$ S2/3 vs. S5, $p<0.001$ S2/3 vs. F2/3, $p>0.05$ S5 vs. F2/3 Tukey

multiple comparisons post test; n=38, 20, and 23 maps in S2/3, F2/3, and S5; Figure 1-4F). In fact, at these close distances we found 13 cases of completely connected local maps (n=5 for S2/3; n=4 for S5, n=4 for F2/3; ranging from 1-7 interneurons, average 3.4 ± 0.6 interneurons, n=13 maps).

Taking into account all distances, the average connection probability (connected out of connected and unconnected neurons) was lower in S2/3, 18 ± 3%, than in S5, 36 ± 4%, or F2/3, 33 ± 3 % ($p = 0.001$, Kruskal-Wallis, $p < 0.01$ for S2/3 vs. S5 and S2/3 vs. F2/3, $p > 0.05$ for S5 vs. F2/3, Dunn's multiple comparison test; Figure 1-4G). A similar result was found when comparing the histograms of connectivity versus distance (Figure 1-4E) in a bin-by-bin fashion ($p < 0.0001$, Friedman test). This difference was not due to a difference in the number of targets tested in the different areas and layers, variables which were not significantly different (Figure 1-4B).

Confirmation of dense local connectivity with dual whole-cell recordings

The previous results indicated that PV+ interneurons can be very densely connected to their local PCs, in some case this dense connectivity can exist at the physical limit, where every PV+ interneuron is connected to every PC sampled.

Given that this result was obtained with a new mapping method, we sought to confirm the high local connectivity with dual whole cell recordings in non-optical experiments, by recording from randomly selected pairs of PV+ interneurons and PCs whose cell bodies were within 160 μm of each other (Figure 1-5A). In 27 of 33 pairs recorded in S2/3, the PV+ interneuron was connected to the PC (Figure 1-5C and D). These monosynaptic connections had low failure rates and large amplitudes which

depressed over the course of an eight action potential train delivered at 50 Hz, but recovered one second later (Table 1-1, Figure 1-5B). Restricting the analysis to intersomatic distances of less than 100 μm showed that 24/27 of pairs were synaptically connected.

We also confirmed the dense local connectivity in S5 and L/23 with dual patch-clamp recordings which revealed that 86% (6 out of 7) of pairs in S5 and 80% (4 out of 5) of pairs in F2/3 were connected. In two instances in S5, one PV+ interneuron was connected to two nearby PCs (at distances of 60 and 84 μm in one experiment, and 77 and 92 μm in the other.)

These dual whole-cell recordings confirmed the locally dense connectivity of PV+ interneurons to PCs observed with optical stimulation. For the optical mapping, we observed high probabilities of connection within 100 μm from the PC: 71 ± 9% in S2/3 (n=21 maps), 92 ± 8 % in S5 (n=12 maps), and 80 ± 7% in F2/3 (n=14 maps, Figure 1-5E). Results obtained from optical mapping and patch-clamp recordings were not significantly different (S2/3, p=0.28; S5, p=0.38; F2/3, p=0.13 Mann-Whitney; Figure 1-5E). The probability of connection within 100 μm also did not differ between areas and layers (p = 0.15 for mapping, p = 0.86 for patching, Kruskal-Wallis). This implies that the locally dense inhibitory connectivity is similar across the cortical areas and layers tested.

Lack of specificity in interneuron- pyramidal cell connectivity

A corollary of locally dense connectivity from PV+ interneurons to PCs is that these interneurons must contact all, or most, PCs nearby, regardless of their identity or

adherence to a particular subcircuit. We tested this prediction by comparing the input maps, obtained at the same time, of several PCs from the same slice (Figure 1-3 and 1-6). We used these simultaneously acquired maps to examine whether pairs of PCs that were connected among themselves had a preferential innervation from nearby interneurons, as compared to pairs of PCs that were unconnected between themselves (Figure 1-6A; (Yoshimura et al., 2005).

We found that the common connection probability for pairs of PCs was similar, regardless of whether these PC pairs were connected or not (0.34 ± 0.09 and 0.28 ± 0.07 for connected or unconnected PC pairs; $p = 0.33$, n=8 connected PCs, n=26 unconnected PCs, S2/3, Mann-Whitney; Figure 1-6B). The common connection probability for all pairs of PCs, connected or not, fell as the distance between PCs increased ($p < 0.05$, n = 34 S2/3 PC pairs, Figure 1-6B). Given that connected pairs often occur at close distances, we were concerned that the common connection probabilities for connected PCs could be biased if the distances between those PCs were systematically smaller. In fact, the average distance of connected PCs was significantly closer than that for unconnected PCs (44 ± 5 μm vs 65 ± 6 μm $p < 0.05$, unpaired t-test with Welch correction). To avoid this distance effect, we restricted our analysis to PC pairs within 70 μm of each other. This reduced the average distance of unconnected PCs to 45 ± 5 μm, which was not significantly different from the average distance of connected PCs ($p = 0.9284$, n=8 connected PCs, n=16 unconnected PCs, Mann-Whitney). In this reduced dataset, the common connection probability for unconnected or connected PCs was statistically indistinguishable (0.35 ± 0.11 for unconnected PCs vs. 0.34 ± 0.09 for connected ones; p

= 0.7286, Mann-Whitney). This analysis indicated that PV+ interneurons do not discriminate between PCs that are forming subcircuits or not.

Spatial patterns of connected PV+ interneurons varies between layers

The previous analysis of connectivity was performed by taking into account the distance of the interneurons to the pyramidal cell. To further examine the spatial structure of connections for the three cortical locations examined, we plotted the position of connected and unconnected PV+ interneurons within the coordinate plane of the cortical circuit and used these plots to explore whether the angular position of connected PV+ interneurons relative to PCs displayed any particular geometry (Figure 1-7). To visualize this angular distribution, we computed polar plots in which the PCs were at the center and a sector was plotted with a radius proportional to the average connection probability of PV+ interneurons from that angular region (Figure 1-7B). At nearby distances (0-200 μm), the average connection probability from any given angle was homogenously high across areas and layers (Figure 1-7B, top row). At further distances (200-600 μm), a spatial pattern appeared, whereby PCs preferentially received connections from interneurons located in vertically oriented positions (Figure 1-7B, middle row). Averaging all distances together also revealed a preference for connections from vertical orientations and also a difference between layer 2/3 and layer 5 of somatosensory and frontal cortices (Figure 1-7B, bottom row). While layer 2/3 PCs received connections from locations closer to the pial surface, layer 5 PCs received connections from locations in deeper layers ($p = 0.0027$, circular ANOVA). Under the convention of 90 degrees being towards the pial surface, the average direction from which PCs received

connections from PV+ interneurons was 86 degrees in S2/3, 258 degrees in S5 and 105 degrees in F2/3, and these angles differed significantly ($p < 0.01$ between S2/3 and S5, $p < 0.01$ between F2/3 and S5, $p = 0.6494$ between S2/3 and F2/3, Watson-Williams test).

We therefore concluded that there were significant differences in the spatial pattern of connectivity between upper and lower layer PV+ cells, although in both cases, and both cortical areas, there was a preferential vertical arrangement of presynaptic neurons.

Overlap of axons and dendrites predicts spatial connectivity profiles

We sought to understand the mechanisms underlying the dense synaptic connectivity maps we observed, and for the differences in the spatial connectivity patterns observed between layers. For this purpose, we next considered the representative anatomy of axons of PV+ interneurons and dendrites of PCs and wondered whether the morphologies of those cells alone could have predicted the spatial connectivity profiles we measured. For example, the vertical location of interneuron inputs across layers (Figure 1-7B) could be explained by the fact that the axons of PV+ interneurons in layer 2/3 of both somatosensory and frontal cortex extend vertically down from the cell body, whereas axons of PV+ interneurons in layer 5 of somatosensory cortex extend vertically up from the cell body (Figure 1-7C). One could invoke Peter' s rule, which states that a neuron connects with target structures in the same proportion as they exist in the tissue (i.e., without any selectivity; Peters and Feldman, 1976), to argue that the axonal patterns can explain the vertical connectivity. But are those arborizations sufficient to account for the connectivity we observed?

To address this question quantitatively, and test whether Peters' Rule could explain our findings, we calculated the average axodendritic overlap between axons of PV+ interneurons and dendrites of PCs in S2/3 and S5. We lacked enough morphological reconstructions from frontal L2/3 for a meaningful analysis across regions. We used the cell bodies and proximal dendritic tree of the PCs (within 100 μm of the soma) because PV+ interneurons contact the perisomatic region of their postsynaptic targets (Karube et al., 2004; Wang et al., 2002).

First, we computed the average morphological structures present in our dataset by projecting together the morphologies of eight S2/3 PV+ interneuron axons, 4 S5 PV+ interneuron axons, and five PC cell bodies and dendrites within 100 μm (Figure 1-8A). Next, we multiplied these calculated morphological densities among themselves to generate the density of spatial overlap of axons and dendrites. We then repeated this calculation, systematically offsetting the positions of the cell bodies of the PV+ interneurons and PCs, by using the coordinates of every interneuron tested from the input maps. This resulted in a series of overlap maps which were then summed to yield a amount of axodendritic overlap for each offset of the cell bodies.

We found that the calculated axodendritic overlap was remarkably similar to the observed connectivity patterns of PV+ interneurons and PCs (Figure 1-8 B-E). Specifically, the axodendritic overlap calculated at each intersomatic distance was not significantly different from the connection probability we had previously measured ($R=0.95$, $p<0.0001$ for S2/3; $R=0.78$, $p<0.0001$ for S5; Figure 1-8B-C). The result was algebraically identical if we calculated the axodendritic overlap for each combination of PV+ interneuron axon and PC dendrite and cell body and then averaged the results

together. Apparently, details such as position of axonal boutons or the fine targeting on postsynaptic processes are not necessary to accurately recapitulate the correct connectivity profile, highlighting the idea that these axons connect in an unspecific manner with whatever postsynaptic structure they encounter.

In addition, the vertical orientation of the connectivity profile was also remarkably similar to the calculated axodendritic overlap profile (Figure 1-8D-E). Indeed, not only the vertical orientation but also the bias for S2/3 PCs to receive a higher probability of connection from PV+ interneurons more towards the pial surface while S5 PCs had the opposite tendency (Figure 1-8E) was recapitulated in the axodendritic overlap (Figure 1-8D). Moreover, while the correlation between the predicted connections and the observed connections was very high within the same layer ($R=0.80$, $p<0.0001$ for S2/3; $R=0.83$, $p<0.0001$ for S5; Figure 1-8F), the correlation across layers were both below zero (between S2/3 predictions and S5 observations $R= -0.5059$; between S5 predictions and S2/3 observations, $R= -0.1170$). This indicates that differences in the spatial patterns of connectivity observed between different cortical layers can be explained solely by taking into account their differences in the morphologies of the PV+ axon and PC dendrites.

We conclude that anatomical overlap between axonal arborizations of PV+ cells and dendritic arborizations of PCs was enough to predict the spatial patterns of connectivity, and even the difference observed across different layers, without the need of any additional mechanisms.

Estimation of convergence and divergence in PV- PCs circuits

The high density of connection we observe from PV+ interneurons to nearby PCs has two implications: 1) Each PC must receive convergent connections from many PV+ interneurons and 2) Each PV+ interneuron must contact many PCs. Using the probability function that describes connectivity, we calculated these numbers and explore whether they are self-consistent. To do this, we constructed a cube of neocortex, 500 μm per side, with one PC located at the origin (Figure 1-9, black triangle), and an average density of 40,000 neurons/mm^3 (Knox, 1980) (see Figure 1-9). We assume that 80% of neurons are excitatory and 20% are inhibitory (Beaulieu et al., 1992) and that one-quarter of the inhibitory neurons are parvalbumin-positive in S2/3 (Xu et al., 2009). We then used the experimentally determined function describing the probability of connection of the PV+ interneurons to the centrally located PC, which is an exponentially decaying function of distance with decay constant of 124 μm with no other adjustable parameters (Figure 1-4E). We assumed that the probability of connection decayed similarly in every direction and chose the central PC as representative of any S2/3 PC in the neocortex.

When restricted to a 50 μm thick slab centered on the PC, this model accurately replicated our results predicting that there would be 6 PV+ interneurons connected to the PC. Although in the maps we only observed 3 connected PV+ interneurons on average in S2/3, in this transgenic mouse only half the parvalbumin-positive interneurons are labeled (Chattopadhyaya et al., 2004).

We then extended the model to answer the two questions posed at the beginning of this section, analyzing the full 500 μm cube. Our model predicted there to be 46 PV+ interneurons connected to the PC. With the estimation of how many PV+ interneurons contact each PC, and the ratio of the numbers of PCs to PV+ interneurons, we then

calculated how many PCs each PV+ interneuron contacts on average. For this calculation, we define:

C	number of PCs
V	number of PV+ interneurons
R= C/V	ratio of PCs to PV+ interneurons
N	number of PV+ interneurons that connect to each PC
N·C	total number of PV+ interneuron to PC connections
N·C/V = N·R	number of PCs contacted by one PV+ interneuron

Given that the number of PV+ interneurons that connect to each PC (variable N) is 46 and that the ratio of PCs to PV+ interneurons (variable R) is 80% to 5% (see above), or 16, the number of PCs contacted by one PV+ interneuron is N·R=46*16=736. It is known that PV+ interneurons contact other PV+ interneurons as well. Therefore, the total number of outgoing connections made by PV+ interneurons is likely even higher. If we assume the same convergence of 46 PV+ interneurons onto one PV+ as we found for PV+ to PC connections, the number of outgoing PV+ interneuron connections changes only from N·C to N(C+V), such that the number of PCs and PV+ interneurons contacted by one PV+ interneuron algebraically reduces to N(R+1), which is 46*17=782. Therefore, we estimate that the typical PV+ interneuron contacts approximately 800 postsynaptic neurons, both PCs and other PV+ interneurons.

Do PV+ interneurons in somatosensory layer 2/3 have enough boutons to contact 800 neurons? FS basket cells in layer 2/3 have approximately 4,000 boutons (Karube et

al. 2004; Tamas et al., 1997; Wang et al. 2002). There are on average 4 boutons from PV+ interneurons onto each postsynaptic soma at the age tested here (Chattopadhyaya et al., 2003). Therefore, we might expect approximately 1,000 outgoing connections from PV+ interneurons.

A similar calculation was performed for the S5 and F2/3 datasets. We constructed a similar cube of neocortex for somatosensory layer 5, in which the proportion of interneurons which are PV+ is around 60% (Xu et al., 2009) and our experimentally determined spatial decay constant for connectivity was 183 μm, resulted in 184 PV+ interneurons contacting each PC. As there are more PV+ interneurons overall, the number of neurons contacted by each PV+ interneuron was still around 1,000: $N(R+1)= 184 * (80/12 + 1) = 1,410$. For frontal cortex layer 2/3, the proportion of interneurons which are PV+ is around 40% (Xu et al., 2009) and the decay constant was 180 μm, resulting in 121 PV+ interneurons contacting each PC. The number of neurons contacted by each PV+ interneuron was $121 * (80/8 +1) = 1,331$.

In summary, using a model we estimate that the typical PV+ interneuron in neocortex contacts several hundred to one thousand postsynaptic targets locally. At the same time, each PC is contacted by a large number of PV+ interneurons, ranging from 50 to almost 200. Therefore, this PV+ interneuron—PC circuit has a high degree of local divergence and convergence.

Discussion

In summary, we find locally dense connectivity from PV+ interneurons onto PCs across cortical areas and layers, approaching in some experiments a completely connected circuit. We find no evidence for target specificity since PV+ interneurons equally contact PC pairs which are synaptically connected or not. Although the location of presynaptic PV+ interneurons varies between layer 2/3 and layer 5, Peters' Rule is sufficient to predict the spatial features of these inputs maps and account for differences observed across layers. Thus, the canonical microcircuit architecture for PV+ interneurons, and perhaps inhibition in general (Fino and Yuste, 2011), may be a densely homogenous matrix innervating all PCs nearby.

Methodological consideration

We are using a novel method to map synaptic connections, so it is pertinent to discuss its advantages and potential limitations. This method relies on two-photon uncaging of glutamate to activate individual neurons, while recording synaptic currents in other neurons, whose inputs we are attempting to map. Two-photon activation is critical to the spatial resolution of the technique as the small point spread function produced by the two-photon process (Figure S1-1) allows for an accuracy down to the level of a single cell.

Two potential problems with this optical technique could arise: 1) failure to activate a targeted neuron and 2) inadvertent activation of a neuron which was not optically targeted. To ensure that the targeted neuron actually fired, we always used different laser power levels when photostimulating neurons, including a high power

(300mW on sample), which, when calibrating the technique, always fired PV+ cells. At the same time, we cannot guarantee that every optically targeted neuron actually fired. In addition, these experiments were performed in brain slices that have many severed inputs, so we think that it is reasonable to assume that our input mapping strategy yields an underestimate of the true connectivity, which would imply the density of connections is even higher than we observe.

To address the issue of inadvertently activating neurons we did not intend to stimulate, we tested whether a neuron fired when the photostimulation was not centered on the soma. We rarely observe activation of a neuron when it was not photostimulated or when we targeted the laser outside its soma (Figure S1-2), even if we purposefully choose a PC with a lower rheobase than the PV+ interneurons for this calibration experiment. In addition, we used the "switching" test (Figure 1-2B) to confirm the inhibitory identity of the photostimulated cell. Given the lower incidence of inhibitory interneurons relative to excitatory neurons in neocortex, it seems unlikely that another interneuron would be close enough to give rise to the inhibitory event observed.

But the direct confirmation that the maps are correct arises from two independent sets of experiments with dual whole-cell recordings, which not only validate the method but also confirm the main results. First, predicted connections were confirmed by patching the targeted presynaptic interneuron during the experiment. We found that 9 out of 10 putatively connected pairs were indeed connected, while 8 out of 8 putatively unconnected pairs were unconnected. Second, the results we observe with the optical mapping matched those observed with dual whole-cell recordings from randomly selected PV+ and PC pairs (Figure 1-5).

This technique therefore provides a robust and accurate measure of the connectivity from a defined set of neurons onto any recorded neuron of interest. One key advantage of two-photon photostimulation is its single-cell resolution. Other techniques, such as one photon laser scanning photostimulation (Callaway and Katz, 1993) or similar approaches relying on optogenetics (Katzel et al., 2011), do not yet have sufficient resolution to distinguish among inputs from different cells, since they activate groups of neurons, or axons from many cells together. At the same time, inhibitory inputs coming from all interneuron subtypes measured by these other techniques has been shown to come from predominantly local sources (Dantzker and Callaway, 2000; Yoshimura et al., 2005) and have similar intralaminar patterns across areas (Katzel et al., 2011); these results are consistent with our analysis yet they also fall short of providing the actual detailed account of the cell-to-cell connectivity diagram.

A second advantage of our two-photon technique is that it provides a measurement of what proportion of cells are connected out of the total population of presynaptic neurons, a critical number to assess the specificity of the connectivity. In contrast to whole-cell recordings, in which the entire population of cells cannot be easily sampled, our method allows for an estimate of the true connectivity in each experiment. This not only avoids the possible introduction of statistical artifacts as the results from averaging different experiments together, but also provides an important absolute count of the nearby connected neurons.

One caveat with our experiments, however, is that there could be different subtypes of parvalbumin-positive interneurons labeled in the transgenic mouse we used (Chattopadhyaya et al., 2004). In the neurons sampled in this study, however, we have

only observed basket –like morphologies in our biocytin-filled reconstructions, and have not observed other morphologies, such as chandelier cells. Lacking any evidence for a heterogeneity in the population of PV cells tested, we treat these PV+ cells as equivalent for the purposes of this study.

Dense connectivity of PV+ interneurons

Our main result is the finding of a strikingly high density of connectivity from PV+ interneurons onto nearby pyramidal cells in somatosensory cortex layer 2/3 and 5 and frontal cortex layer 2/3 (Figure 1-4). This result is confirmed by dual whole-cell recordings (Figure 1-5) and can approach in some experiments the limit of having every interneuron being connected to every local PC sampled.

The reported connectivity rates between interneurons identified as PV+, fast-spiking (FS), or basket cells, and nearby PCs vary greatly (review by Thomson and Lamy, 2007; Holmgren et al., 2003; Kapfer et al., 2007; Oswald et al., 2009; Reyes et al., 1998; Tamas et al., 1997; Thomson et al., 1996; Thomson et al., 2002; Wang et al., 2002; Yoshimura and Callaway, 2005). Consistent with the high density that we observed, connectivity rates of 67% have been reported for fast-spiking cells to PCs within 50 μm in layer 2/3 of rat somatosensory cortex (Kapfer et al., 2007). However, rates as low as 16% have been reported in layer 3 of rat cortex (Thomson et al. 2002). Intermediate values of 50-60% have also been observed (Holmgren et al. 2003). In addition, a small decrease in FS to PC connectivity with intersomatic distance has been reported for layer 2/3 of visual or somatosensory (Holmgren et al., 2003) and auditory cortex (Oswald et al., 2009)).

In the present study, we found 468 PV+ interneuron to PC connected pairs among 1,747 pairs tested, a significant increase in the number of connections probed, as compared with past studies. The increase in the numbers of pairs we could quickly test allowed us not only to draw conclusions about the basic rates of connectivity, but also to directly assess the true rates for each slice. As this study shows, knowledge of the number of cells tested and the exact distances over which a given probability of connection is observed are important factors to understanding connectivity among neurons. Moreover, with a more significant statistical sample, we can also begin to probe the mechanisms by which that connectivity is achieved.

Using our large dataset, we calculated the function describing how the probability of connection varies with intersomatic distance. This allowed us to estimate the degree of divergence from PV+ interneurons and the degree of convergence onto PCs, finding a high divergence from PV+ interneurons outputs, which goes hand in hand with the convergence of many PV+ interneurons onto one PC we observed in our highly dense maps. We estimate a higher density of converging connections onto PCs than previously calculated, even though our data may in fact underestimate the true probability of connection. Previous studies in neocortex (Tamas et al. 1997) and hippocampus (Buhl et al. 1994) estimated approximately 25 PV+ interneurons converge onto one PC. This estimate is lower than what we found but it also predicts only $N(R+1) = 25*17 = 425$ connections from PV+ interneurons to PCs, which is a low number when one considers the large number of boutons present in a single PV+ axon (~4,000, (Karube et al., 2004; Wang et al., 2002). Therefore, the high density of connections we estimate from PV+ interneurons to PCs appears more consistent with the observed high number of boutons.

Finally, we have recently described how somatostatin-positive (SOM+) interneurons also make locally dense connections onto nearby PCs (Fino and Yuste, 2011) and our current data on PV+ interneurons are in good agreement with the SOM+ interneuron data. The average probability of connection from a SOM+ interneuron onto a PC in layer 2/3 of frontal cortex of P11-P16 mice was 43%, somewhat higher than the average probability of connection we observed for PV+ interneurons in any layer or area. But the probability of connection for SOM+ interneurons within 200 μm of a PC was 71%, similar to the probabilities of connection we observed within 100 μm. The similarly dense matrix of local connections from these two distinct subpopulations of interneurons onto PCs implies that inhibitory connections may be built according to the same basic plan.

Unspecific innervation of PC cells

Our second main result is the finding that the probability of 2 PCs receiving connections from the same PV+ interneurons does not depend on whether the 2 PCs are synaptically connected (Figure 1-6). This result is consistent with the high density of connections we observe as there is no room for specificity within such a high probability — in other words, there are no neurons to specifically avoid contacting.

In a past study, considering a pair of PCs which are synaptically connected as part of the same 'subcircuit', revealed no difference in the inhibition received by the pair of PCs (Yoshimura et al., 2005). However, FS cells which were reciprocally connected to PCs shared more common input from excitatory sources than those which were not reciprocally connected (Yoshimura and Callaway, 2005). In addition, FS cells

preferentially targeted PCs which provided reciprocal excitatory connections (Yoshimura and Callaway, 2005). Our method, with single cell resolution, allowed us to directly determine whether PV+ interneuron connections to nearby PCs were specific. We found no evidence that PV+ interneurons preferentially contact PC pairs in the same subcircuit (Figure 1-6A). Our results therefore differ from those of (Yoshimura and Callaway, 2005) which showed preferential connections from FS cells to PCs providing reciprocal excitation. But, at the same time, we did not test the same hypothesis as our optical strategy did not allow us to detect the response of the PV+ interneurons to PC stimulation. In addition, the previous study assessed connectivity in the visual cortex of rats aged P21-P26, i.e. a different cortical area, age, and species than we used.

It is still possible that there could be specificity in the PV+ interneuron-PC connections which we were unable to test. For example, the postsynaptic target of the PCs may influence from which PV+ interneurons they receive connections, as has been shown for some excitatory circuits (Brown and Hestrin, 2009; Morishima and Kawaguchi, 2006). It is also possible that this locally dense architecture is pruned during development, although this was not found in the similarly dense connections made by somatostatin interneurons (Fino and Yuste, 2011), and fast-spiking cell connections onto pyramids have been reported to be dense in mature animals (Oswald et al., 2009).

Mechanisms of interneuronal connectivity

The third goal of our study focused on understanding how the lack of specific connectivity emerges and we conclude that the spatial pattern of connectivity between PV+ interneurons and PC can be explained by Peters' Rule (Peters and Feldman, 1976).

This rule states that neurons connect to each other based solely on the overlap of their axonal and dendritic arborizations, without any additional specificity. We find that a simple application of this rule is sufficient to predict the spatial connectivity profiles we observe and their differences between layers. This is surprising because Peters' Rule is often, but not always (Shepherd et al., 2005), violated for excitatory circuits (Brown and Hestrin, 2009; Petreanu et al., 2009). Also, ultrastructural analyses often show failures of Peters' Rule (Mishchenko et al., 2010; Stepanyants et al., 2002).

Our finding that Peters' Rule is sufficient to explain the spatial distribution of connectivity from PV+ interneurons does not mean that other aspects of their connectivity are unspecific. For example, it is well know that different interneurons subtypes have subcellular specificy in their targeting (Somogyi et al., 1998). Indeed, PV+ interneurons preferentially contact perisomatic regions of PCs, and this violates Peter's rule because PV+ interneuron axons must course through areas in which non-perisomatic dendrites of PCs reside. Also, at the macroscopic level, there are clear specific patterns of interlaminar projections in the axons of PV interneurons, as can be observed in some of our reconstructions, where certain layers are specifically targeted (Figure 1-8). We therefore view the interneuron innervation pattern as resulting from several layers of mechanisms, some of which could be specific, like the interlaminar and subcellular targeting, but also unspecific ones, such as the intralaminar spatial connectivity pattern with PCs that we observe.

A canonical microcircuit for inhibition

Our results demonsatrate that inhibitory connections lack specificity in terms of contacting PCs. Consistent with this idea, interneurons in layer 2/3 of visual cortex are broadly tuned to the orientation of a stimulus (Kerlin et al., 2010; Ma et al., 2010; Niell et al., 2008; Sohya et al., 2007; Zariwala et al., 2011), although some PV+ interneurons in the more superficial part of layer 2/3 appear to be sharply tuned (Runyan et al., 2010). The difference between these studies may be due to methodological reasons or differences in PV+ interneuron subtypes (Hasenstaub and Callaway, 2010). But interneurons also receive input from neurons tuned to different orientations (Bock et al., 2010), providing an anatomical basis for the observation of broad tuning in interneurons. If PV+ interneurons are indeed broadly tuned functionally, it would make sense that their anatomical outputs are also "broadly tuned" in the sense that they do not discriminate between postsynaptic targets.

In addition, the broad inhibition provided by PV+ interneurons (present study) and SOM+ interneurons (Fino and Yuste, 2011) is consistent with previous findings showing inhibition is very strong in a local region. Indeed, individual neurogliaform cells, for instance, release enough GABA to mediate volume transmission to nearby neurons, in a completely unspecific fashion (Olah et al., 2009). Given that interneurons act as “foreign invaders” during development (Anderson et al., 1997), migrating tangentially into the layers which develop in a radial fashion, it is parsimonious that their wiring program may be relatively simple, without the need for specific mechanisms to connect the neural circuit in the appropriate functional manner. Extending this "blanket of inhibition" unspecifically over all local PCs could represent the key canonical circuit design for all inhibitory connections in the neocortex.

Experimental Procedures

Slice preparation and electrophysiology

All animal handling and experimentation was done according to the National Institutes of Health and local Institutional Animal Care and Use Committee guidelines. Somatosensory or frontal coronal slices 350 μm thick were prepared from P12-17 (14.4±0.18) day-old G42 mice (Chattopadhyaya et al., 2004) using a Leica VT1000S or VT1200S vibratome with ice-cold sucrose solution containing (in mM): 27 $NaHCO_3$, 1.5 NaH_2PO_4, 222 sucrose, 2.5 KCl, 3 $MgSO_4$, and 1 $CaCl_2$. Slices were incubated at 36°C for 30 minutes in ACSF containing (in mM): 126 NaCl, 26 $NaHCO_3$, 1.1 NaH_2PO_4, 10 glucose, 3 KCl, 3 $MgSO_4$, and 1 $CaCl_2$. During recordings made at room temperature or 32°C, ACSF was similar except for the following (in mM): 2 $MgSO_4$ and 2 $CaCl_2$. All sucrose and ACSF solutions were saturated with 95% O_2 and 5% CO_2.

Whole-cell recordings were made using the Multiclamp 700B amplifier (Molecular Devices, Sunnyvale, CA), digitized with National Instruments 6259 multichannel cards, and recorded using custom software written using the LabView platform (National Instruments, Austin, TX). Current-clamp recordings were performed with intracellular solution (pH 7.2) containing (in mM): 135 K-methylsulfate, 8 NaCl, 10 HEPES, 2 Mg-ATP, 0.3 Na-GTP, 7 phosphocreatine, 0.02 Alexa Fluor 594, and 10.7 biocytin. Voltage-clamp recordings were performed with intracellular solution (pH 7.3) containing (in mM): 128 CH_3CsO_3S, 10 HEPES, 10 Na-phosphocreatine, 2 $MgCl_2$, 3 $MgSO_4$, 4 Mg-ATP, 0.4 Na-GTP, 0.02 Alexa Fluor 594, and 10.7 biocytin.

Imaging and RuBi-Glutamate uncaging

Imaging and uncaging were performed as described in (Fino and Yuste, 2011). Briefly, images were acquired using a custom-made two-photon laser scanning microscope based on the Olympus FV-200 system (side-mounted to a BX50WI microscope with a 40X 0.8NA or 20X 0.5NA water immersion objective) and a Ti:sapphire laser (Chameleon Ultra II, Coherent, >3 W, 140 fs pulses, 80 MHz repetition rate). Images were acquired at 850 or 900nm for GFP and 800 nm for Alexa 594 with minimal power to avoid uncaging Rubi-Glutamate.

300 μM Rubi-Glutamate (Tocris, Ellisville, MO) was added to the oxygenated ACSF during mapping experiments and was recirculated with a peristaltic pump (RP-1, Rainin Instruments, Oakland, CA). All mapping experiments were conducted using the 20X 0.5NA objective.

We used custom software (Nikolenko et al., 2007) to create the complex targets around cell bodies shown in Figure 1-1D. Each complex target, which consisted of five beamlets due to the diffractive optical element we used (Figure S1-1, (Nikolenko et al., 2007)), was illuminated for eight milliseconds, resulting in a stimulation duration of 72 milliseconds given an additional millisecond to move between each target. A Pockels cell (Conoptics, Danbury, CT) allowed us to control power precisely over these short durations. The power at the sample ranged from 150 to 330 milliwatts for mapping experiments. We used multiple different power levels in each mapping experiment, always testing the maximum level to be sure not to miss connections from any PV+ interneurons with particularly high rheobases.

In order to confirm that our photostimulation did not inadvertently activate nearby neurons, we patched PCs with a low rheobase (40pA) and photostimulated the entire area

surrounding the cell in a grid fashion using the same parameters as we did to stimulate interneurons (Figure S1-2). First, we confirmed that this cell could be photostimulated by eliciting action potentials via uncaging over the soma directly ten out of ten times (Figure S1-2A). Next, we checked 180 locations in the field of view (Figure S1-2B) three times each and found that the cell only produced action potentials twice: one out of three times when one of the targets was just slightly offset from the soma (Figure S1-2D) and two out of three times in one other location (Figure S1-2C). Thus, it is very rare that action potentials are elicited unless stimulating a cell soma directly. This makes it extraordinarily unlikely that neurons could be led to produce action potentials by inadvertent photostimulation, especially for interneurons which have higher rheobases, again confirming the high resolution of this technique.

We were concerned that the depth of the patched PC may have affected the connection probability we observed. However, the correlation between the connection probability and the depth of the patched PC is small ($R^2 = 0.08$, n=82 maps, Figure S1-3).

Biocytin histochemistry and reconstructions

At the end of an experiment, slices were fixed and kept overnight in 4% paraformaldehyde in 0.1M phosphate buffer (PB) at 4ºC. The slices were then rinsed three times for five minutes per rinse on a shaker in 0.1M PB. They were placed in 30% sucrose mixture (30g sucrose dissolved in 50ml ddH_20 and 50 ml 0.24M PB) for 2 hours and then frozen on dry ice in tissue freezing medium. The slices were kept overnight in a -80ºC freezer. The slices were defrosted and the tissue freezing medium was removed by three 20 minute rinses in 0.1M PB while on a shaker. The slices were kept in 1%

hydrogen peroxide in 0.1M PB for thirty minutes on the shaker to pretreat the tissue, then were rinsed twice in 0.02M potassium phosphate saline (KPBS) for twenty minutes on the shaker. The slices were then kept overnight on the shaker in Avidin-Biotin-Peroxidase Complex. The slices were then rinsed three times in 0.02M KPBS for 20 minutes each on the shaker. Each slice was then placed in DAB (0.7 mg/ml 3,3"-diaminobenzidine, 0.2 mg/ml urea hydrogen peroxide, 0.06M Tris buffer in 0.02M KPBS) until the slice turned light brown, then immediately transferred to 0.02M KPBS, and finally transferred again to fresh 0.02M KPBS after a few minutes. The stained slices were rinsed a final time in 0.02M KPBS for 20 minutes on a shaker. Each slice was observed under a light microscope and then mounted onto a slide using crystal mount.

Successfully filled and stained neurons were then reconstructed using Neurolucida software (MicroBrightField, Williston, VT). The neurons were viewed with a 100x oil objective on an Olympus IX71 inverted light microscope or an Olympus BX51 upright light microscope. The Neurolucida program projected the microscope image onto a computer drawing tablet. The neuron's processes were traced manually while the program recorded the coordinates of the tracing to create a digital, three-dimensional reconstruction. The x- and y-axes form the horizontal plane of the slice, while the z-axis is the depth. The user defined an initial reference point for each tracing. The z-coordinate was then determined by adjustment of the focus. In addition to the neuron, the pia and white matter were drawn. Axon and dendrite densities were calculated from the Neurolucida reconstruction using the TREES toolbox (Cuntz et al., 2010). The densities were calculated with voxels 5 μm on each side.

The number of neurons contacted within 400 μm was estimated with the following MATLAB commands:

```
dx=.0000001;
integral=cumtrapz(exp(-(0:dx:0.4)/0.124).*4.*pi.*100000.*(0:dx:0.4).^2);
integral(end)*dx
```

Data analysis

Off-line analysis was conducted using MATLAB (Mathworks, Natick, MA), InStat (GraphPad, La Jolla, CA), MiniAnalysis (Synaptosoft, Decatur, GA), and Oriana (Kovach Computing Services, Wales, U.K.). Additional circular statistics tests were performed with the MATLAB CircStat Toolbox (Berens, 2009). All results are expressed as mean ± SEM.

Monosynaptic connections were characterized by evoking single action potentials in the PV+ interneuron and recording the IPSC evoked in the PC to measure the latency, amplitude, rise time, decay time, rate of rise, and failure rate. These values were recorded at +40mV with a chloride reversal potential of -80mV. Trains of eight action potentials delivered at 50 Hz followed by a single action potential one second after the last action potential in the train were evoked in the PV+ interneuron to measure the paired pulse ratio, summation, train depression, and recovery (Markram et al., 1998). Paired pulse ratio is the amplitude of the second IPSC divided by the amplitude of the first IPSC. Summation is the amplitude of the second IPSC relative to the baseline of the first IPSC divided by the amplitude of the first IPSC. Train depression is the amplitude of the last IPSC in the train divided by the amplitude of the first IPSC in the train.

Recovery is the amplitude of the IPSC evoked by the action potential one second after the train divided by the amplitude of the first IPSC in the train.

References

Aaron, G., and Yuste, R. (2006). Reverse optical probing (ROPING) of neocortical circuits. Synapse *60*, 437-440.

Anderson S.A., Eisenstat D.D., Shi .L, Rubenstein J.L. (1997). Interneuron migration from basal forebrain to neocortex: dependence on Dlx genes. Science *278*, 474-6.

Ascoli, G.A., Alonso-Nanclares, L., Anderson, S.A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsaki, G., Cauli, B., Defelipe, J., Fairen, A., *et al.* (2008). Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex. Nat Rev Neurosci *9*, 557-568.

Beaulieu, C., and Colonnier, M. (1983) The number of neurons in the different laminae of the binocular and monocular regions of area 17 in the cat. J Comp Neuro *217*, 337-344.

Beaulieu, C., Kisvarday, Z., Somogyi, P., Cynader, M., Cowey, A. (1992). Quantitative distribution of GABA-immunopositive and -immunonegative neurons and synapses in the monkey striate cortex (area 17). Cerebral Cortex *1992*, 295-309.

Berens, P. (2009). CircStat: A MATLAB Toolbox for Circular Statistics. J Stat Softw *31*, 1-21.

Binzegger, T., Douglas, R.J., and Martin, K.A. (2004). A quantitative map of the circuit of cat primary visual cortex. J Neurosci *24*, 8441-8453.

Blatow, M., Rozov, A., Katona, I., Hormuzdi, S.G., Meyer, A.H., Whittington, M.A., Caputi, A., and Monyer, H. (2003). A novel network of multipolar bursting interneurons generates theta frequency oscillations in neocortex. Neuron *38*, 805-817.

Bock, D., Lee, W.C., Kerlin, A., Andermann, M.L., Soucy, E., Yurgenson, S., Wetzel, A.W., Hood, G., and Reid, R.C. (2010). Local anatomical connectivity of a cluster of physiologically characterized cells in mouse visual cortex. In Society for Neuroscience (San Diego, CA).

Brown, S.P., and Hestrin, S. (2009). Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. Nature *457*, 1133-1136.

Bruno, R.M., and Simons, D.J. (2002). Feedforward mechanisms of excitatory and inhibitory cortical receptive fields. J Neurosci *22*, 10966-10975.

Buhl, E.H., Halasy, K., Somogyi, P. (1994). Diverse sources of hippocampal unitary inhibitory postsynaptic potentials and the number of synaptic release sites. Nature *368*, 823-828.

Callaway, E.M., and Katz, L.C. (1993). Photostimulation using caged glutamate reveals functional circuitry in living brain slices. Proc Natl Acad Sci U S A *90*, 7661-7665.

Chattopadhyaya, B., Di Cristo, G., Higashiyama, H., Knott, G.W., Kuhlman, S.J., Welker, E., and Huang, Z.J. (2004). Experience and activity-dependent maturation of perisomatic GABAergic innervation in primary visual cortex during a postnatal critical period. J Neurosci *24*, 9598-9611.

Crick, F.H. (1979). Thinking about the brain. Sci. Am. *241*, 219-232.

Cuntz, H., Forstner, F., Borst, A., and Hausser, M. (2010). One rule to grow them all: a general theory of neuronal branching and its practical application. PLoS Comput Biol *6*.

Dantzker, J.L., and Callaway, E.M. (2000). Laminar sources of synaptic input to cortical inhibitory interneurons and pyramidal neurons. Nat Neurosci *3*, 701-707.

Douglas, R.J., and Martin, K.A.C. (2004). Neuronal circuits of the neocortex. Annu Rev Neurosci *27*, 419-451.

Fino, E., Araya, R., Peterka, D.S., Salierno, M., Etchenique, R., and Yuste, R. (2009). RuBi-Glutamate: Two-Photon and Visible-Light Photoactivation of Neurons and Dendritic spines. Front Neural Circuits *3*, 2.

Fino, E., and Yuste, R. (2011). Dense inhibitory connectivity in neocortex. Neuron.

Galarreta, M., and Hestrin, S. (2002). Electrical and chemical synapses among parvalbumin fast-spiking GABAergic interneurons in adult mouse neocortex. Proc Natl Acad Sci U S A *99*, 12438-12443.

Hasenstaub, A.R., and Callaway, E.M. (2010). Paint it black (or red, or green): optical and genetic tools illuminate inhibitory contributions to cortical circuit function. Neuron *67*, 681-684.

Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. J Physiol *551*, 139-153.

Kapfer, C., Glickfeld, L.L., Atallah, B.V., and Scanziani, M. (2007). Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex. Nat Neurosci *10*, 743-753.

Karube, F., Kubota, Y., and Kawaguchi, Y. (2004). Axon branching and synaptic bouton phenotypes in GABAergic nonpyramidal cell subtypes. J Neurosci *24*, 2853-2865.

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

Kerlin, A.M., Andermann, M.L., Berezovskii, V.K., and Reid, R.C. (2010). Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. Neuron *67*, 858-871.

Knox, C.A. (1980). Effects of aging and chronic arterial hypertension on the cell populations in the neocortex and archicortex of the rat. Acta Neuropathol *56*, 139-34.

Kraushaar, U., and Jonas, P. (2000). Efficacy and stability of quantal GABA release at a hippocampal interneuron-principal neuron synapse. J Neurosci *20*, 5594-5607.

Ma, W.P., Liu, B.H., Li, Y.T., Huang, Z.J., Zhang, L.I., Tao, H.W. (2010). Visual representations by cortical somatostatin inhibitory neurons--selective but with weak and delayed responses. J Neurosci *30*, 14371-9.

Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. Proc Natl Acad Sci U S A *95*, 5323-5328.

Marshel, J.H., Mori, T., Nielsen, K.J., and Callaway, E.M. (2010). Targeting single neuronal networks for gene expression and cell labeling in vivo. Neuron *67*, 562-574.

Mayhew, T.M., and Momoh, C.K. (1973). Contribution to the quantitative analysis of neuronal parameters: the effects of biased sampling procedures of estimates of neuronal volume, surface area and packing density. J Comp Neurol *148*, 217-228.

McGarry, L.M., Packer, A.M., Fino, E., Nikolenko, V., Sippy, T., and Yuste, R. (2010). Quantitative classification of somatostatin-positive neocortical interneurons identifies three interneuron subtypes. Front Neural Circuits *4*, 12.

Miki, T., Fukui, Y., Takeuchi, Y., and Itoh, M. (1995). A quantitative study of the effects of prenatal X-irradiation on the development of cerebral cortex in rats. Neurosci Res *23*, 241-247.

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

Morishima, M., and Kawaguchi, Y. (2006). Recurrent connection patterns of corticostriatal pyramidal cells in frontal cortex. J Neurosci *26*, 4394-4405.

Mountcastle, V.B. (1982). An organization principle for cerebral function: the unit module and the distributed system. In Mindful Brain, F.O. Schmitt, ed. (Cambridge, MA: The MIT Press), pp. 1-50.

Niell, C.M., and Stryker, M.P. (2008) Highly selective receptive fields in mouse visual cortex. J Neurosci *28*, 7520-7536.

Nikolenko, V., Poskanzer, K.E., and Yuste, R. (2007). Two-photon photostimulation and imaging of neural circuits. Nat Methods *4*, 943-950.

Olah, S., Fule, M., Komlosi, G., Varga, C., Baldi, R., Barzo, P., and Tamas, G. (2009). Regulation of cortical microcircuits by unitary GABA-mediated volume transmission. Nature *461*, 1278-1281.

Oswald, A.M., Doiron, B., Rinzel, J., and Reyes, A.D. (2009). Spatial profile and differential recruitment of GABAB modulate oscillatory activity in auditory cortex. J Neurosci *29*, 10321-10334.

Peterlin, Z.A., Kozloski, J., Mao, B.Q., Tsiola, A., and Yuste, R. (2000). Optical probing of neuronal circuits with calcium indicators. Proc Natl Acad Sci U S A *97*, 3619-3624.

Peters, A., and Feldman, M.L. (1976). The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. General description. J Neurocytol *5*, 63-84.

Peters, A., Kara, D.A., and Harriman, K.M. (1985). The neuronal composition of area 17 of rat visual cortex. III. Numerical considerations. J Comp Neurol *238*, 263-274.

Petreanu, L., Mao, T., Sternson, S.M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. Nature *457*, 1142-1145.

Philp, J.R., and Buchanan, T.J. (1971). Quantitative measurement on finite tissue sections. J Anat *108*, 89-97.

Pouille, F., Marin-Burgin, A., Adesnik, H., Atallah, B.V., and Scanziani, M. (2009). Input normalization by global feedforward inhibition expands cortical dynamic range. Nat Neurosci *12*, 1577-1585.

Ramon y Cajal, S. (1923). Recuerdos de Mi Vida: Historia de Mi Labor Cientıfica.

Reyes, A., Lujan, R., Rozov, A., Burnashev, N., Somogyi, P., and Sakmann, B. (1998). Target-cell-specific facilitation and depression in neocortical circuits. Nat Neurosci *1*, 279-285.

Runyan, C.A., Schummers, J., Van Wart, A., Kuhlman, S.J., Wilson, N.R., Huang, Z.J., and Sur, M. (2010). Response features of parvalbumin-expressing interneurons suggest precise roles for subtypes of inhibition in visual cortex. Neuron *67*, 847-857.

Schuz, A., and Palm, G. (1989). Density of neurons and synapses in the cerebral cortex of the mouse. J Comp Neurol *286*, 442-455.

Shepherd, G.M., Stepanyants, A., Bureau, I., Chklovskii, D., and Svoboda, K. (2005). Geometric and functional organization of cortical circuits. Nat Neurosci *8*, 782-790.

Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K., Tsumoto, T. (2007) GABAergic neurons are less selective to stimulus orientation than excitatory neurons in layer II/III of visual cortex, as revealed by in vivo functional Ca2+ imaging in transgenic mice. J Neurosci *27*, 2145-2149.

Stepanyants, A., Hof, P.R., and Chklovskii, D.B. (2002). Geometry and structural plasticity of synaptic connectivity. Neuron *34*, 275-288.

Somogyi, P., Tamas, G., Lujan, R., Buhl, E.H. (1998). Salient features of synaptic organisation in the cerebral cortex. Brain Research Reviews *26*, 113-135.

Tamas, G., Buhl, E.H., and Somogyi, P. (1997). Fast IPSPs elicited via multiple synaptic release sites by different types of GABAergic neurone in the cat visual cortex. J Physiol *500 (Pt 3)*, 715-738.

Thomson, A.M., and Lamy, C. (2007). Functional maps of neocortical local circuitry. Front Neurosci *1*, 19-42.

Thomson, A.M., West, D.C., Hahn, J., and Deuchars, J. (1996). Single axon IPSPs elicited in pyramidal cells by three classes of interneurones in slices of rat neocortex. J Physiol *496 (Pt 1)*, 81-102.

Thomson, A.M., West, D.C., Wang, Y., and Bannister, A.P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2-5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. Cereb Cortex *12*, 936-953.

Wang, Y., Gupta, A., Toledo-Rodriguez, M., Wu, C.Z., and Markram, H. (2002). Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex. Cereb Cortex *12*, 395-410.

White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1986). The Structure of the Nervous System of the Nematode Caenorhabditis elegans. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences *314*, 1-340.

Wickersham, I.R., Lyon, D.C., Barnard, R.J., Mori, T., Finke, S., Conzelmann, K.K., Young, J.A., and Callaway, E.M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. Neuron *53*, 639-647.

Woodruff, A., Xu, Q., Anderson, S.A., and Yuste, R. (2009). Depolarizing effect of neocortical chandelier neurons. Front Neural Circuits *3*, 15.

Xu, X., Roby, K.D., and Callaway, E.M. (2009). Immunochemical characterization of inhibitory mouse cortical neurons: three chemically distinct classes of inhibitory cells. J Comp Neuro *518*, 389-404.

Yoshimura, Y., and Callaway, E.M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. Nat Neurosci *8*, 1552-1559.

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

Zariwala, H.A., Madisen, L., Ahrens, K.F., Bernard, A., Lein, E.S., Jones, A.R., and Zeng, H. (2011). Visual tuning properties of genetically identified layer 2/3 neuronal types in the primary visual cortex of cre-transgenic mice. Front. Syst. Neuroscience *4*, 162.

Figure Legends

Figure 1-1: Characterization and photostimulation of parvalbumin-positive fast-spiking basket cells

A: Reconstruction of a GFP-positive interneuron in somatosensory cortex layer 2/3 from the G42 transgenic mouse line (axons in blue, dendrites in red). B: Electrophysiological recordings revealed the classic high rheobase and fAHP current (top panel). A fast-spiking response was elicited upon injecting the interneuron with 1nA of current (bottom panel). C: This subtype of interneuron is characterized by a rectifying IV curve (Woodruff et al, 2010). D, E: A two-photon laser beam was multiplexed across both time, by sequentially stimulating multiple targets around a cell body (D), and space, by engaging a diffractive optical element to split the laser beam into five beamlets (E), to uncage Rubi-Glutamate around a cell body (laser spots are drawn at actual size, see Figure S1-1). F: Single action potentials were elicited via illumination with 150mW of 800nm light on the sample using a 20X 0.5NA objective (black bar shows duration of illumination). G: At a higher power, 300 mW, multiple action potentials were elicited with the characteristic fast-spiking waveform.

Figure 1-2: Mapping inhibitory inputs of pyramidal cells

A: Mapping potential connections from 5 interneurons (numbered 1 to 5) to one PC (blue arrow) in somatosensory layer 2/3. Unconnected interneurons in gray, connected in red circles.. B: Electrophysiological recordings obtained from the PC at holding potentials of +40mV (top trace) and -40mV (bottom trace) during the photostimulation of the

interneurons shown in A. Note the synaptic response from one interneuron (red arrow in A, inset in D) while an interneuron directly nearby (gray arrow in A, inset in F) showed no response. C, E: The red PV+ interneuron neuron was patched and confirmed to be synaptically connected while no response was recorded for the gray neuron (F), which was determined to be unconnected by photostimulation. G: True positive response. A pyramidal cell (PC) patched in a field of view showing many GFP-positive PV+ interneurons nearby, one of which (circled in red) was determined to be putatively connected by optical stimulation. H: While holding the PC at +40mV (top trace) or -40mV (bottom trace), we photostimulated the cell circled in red in G, eliciting IPSCs. The reversal potential for GABA was set at -80mV in the internal solution so that EPSCs, but not IPSCs, would change directions between +40 and -40mV. I, J: The photostimulated cell was patched, confirming electrophysiologically that it was indeed synaptically connected. K: Negative responses. A PV+ interneuron (circled in gray) was targeted for photostimulation during whole-cell recording of two nearby PCs. L: The recordings from the PCs in K show no evidence of a synaptic connection during photostimulation of the PV+ interneuron. M, N: The photostimulated cell was patched, confirming electrophysiologically the lack of synaptic connections from the PV+ interneuron onto either PC. O: False positive response. A PV+ interneuron (circled in black) directly next to the recorded PC was targeted for photostimulation. P: A slow response which was outward at +40mV but flipped at -40mV was distinguished from the synaptic event shown in H as a direct stimulation of the patched neuron. Q: A PV+ interneuron (circled in black), located distal to the recorded PC (blue arrow), was targeted for photostimulation. R: An excitatory cell connected to the recorded PC was stimulated

resulting in a false positive distinguished by the presence of excitatory postsynaptic currents.

Figure 1-3: Simultaneous mapping of PV+ inputs to four PCs.

Four PCs in somatosensory cortex layer 2/3 were patched and the surrounding PV+ interneurons in four different focal planes were optically stimulated. Each column shows the result of mapping the GFP-labeled PV+ interneurons for each of the four PCs. Each row shows the result from a different focal plane, with the projection of all the focal planes in the top row. Connected PV+ interneurons are circled in red, unconnected PV+ interneurons are circled in gray, and PV+ interneurons at which a direct stimulation of the patched PC occurred are circled in black.

Figure 1-4: Dense PV+ interneuron inputs to PCs

A: Representative map from S2/3 showing inputs from PV+ interneurons to PCs (red - connected, unconnected - gray, false positive - black). B: Number of PV+ interneurons tested (no significant differences) and how many were connected in each cortical area and layer sampled. C: Schematic of the cube of tissue we sampled in our maps. D: Histogram showing all the intersomatic distances of the PV+ interneurons from the PCs whose connectivity was tested (green bars) in the different cortical layers and areas (left, middle, right panels). False positive responses (black bars) are few in number. E: Histogram of the connection probability depending on the intersomatic distance between interneurons and PCs. The connection probability of PV+ interneurons to PCs in all areas and layers was very high when the PV+ interneurons were close to the PC, and fell off

with distance. F: The probability of observing a connection from a PV+ interneuron to a PC within 200 μm in S2/3, S5, and F2/3. G: Connection probability for different cortical areas and layers. The probability was lower in somatosensory layer 2/3 than in somatosensory layer 5 and frontal layer 2/3.

Figure 1-5: Confirmation of dense PV+ innervation of PCs

A: Anatomical reconstruction of a connected pair with a PV+ interneuron connected to a PC. PV+ axon in blue, PV+ dendrites in red, PC axon in green, PC dendrite in black. B: Electrophysiological recording of this PV+ interneuron to PC connection showing IPSCs in the PC in response to a train of 8 action potentials evoked at 50Hz in the PV+ interneuron. The IPSCs depress over the course of the action potential train. C: Intersomatic distances of the tested pairs from dual whole-cell recording experiments (red - connected, gray - unconnected). D: Probability of connection observed at different intersomatic distances. E: Connection probability within 100 μm did not differ whether observed with mapping or patching in all tested cortical areas and layers. Within this nearby range, there was no difference in connection probability between the layers and areas.

Figure 1-6: Lack of specificity in PV+ interneuron connectivity

A: The common connection probability, calculated as the number of PV+ interneurons which connect to two PCs out of the number of PV+ interneurons which connect to either PC, does not differ depending on whether the postsynaptic PCs are connected to each

other. B: The common connection probability is inversely proportional to the distance between the PCs.

Figure 1-7: Spatial pattern of connected PV+ interneurons varies between layers

A: Position of the connected (red) and unconnected (gray) PV+ interneurons plotted relative to the recorded PC (center of each plot). B: Polar plots show the probability of connection from a given angular region with the pial surface at the top of the circle. The probability of connection at the outer edge of a circle is 1 and 0.5 for the middle dotted circle. Polar plots in each column represent different cortical areas and layers. Rows show the angular distribution of connection counting only interneurons at certain distances (within 200 μm, from 200-600 μm, or all distances). C: Anatomical reconstructions of PV+ interneurons from each cortical area/layer (axons - blue, dendrites - red). Note the striking vertical organization at greater distances and the fact that layer 2/3 receives from interneurons more towards the pial surface while layer 5 is the opposite.

Figure 1-8: Axodendritic overlap can predict spatial patterns of PV+ connectivity

A: Average density of eight somatosensory layer 2/3 PV+ interneuron axons (left), 5 PC cell bodies and dendrites within 100 μm (middle), and the average density of five somatosensory layer 5 PV+ interneuron axons (right) were multiplied to produce a measure of axodendritic overlap between axons of PV+ interneurons and soma and dendrites of PCs. This value was calculated using the intersomatic distances of the PV+ interneurons and PCs tested in the maps. B-E: The amount of overlap between the axons

of the PV+ interneurons and the dendrites and cell bodies of the PCs predict the observed connectivity profiles with regard to both distance (B,C) and angular distribution (D,E) in both S2/3 and S5 (left and right columns).

Figure 1-9: Model showing convergence of PV+ interneurons onto one PC

A cube of neocortex 500 μm on each side from somatosensory layer 2/3 with 5,000 neurons. The cube has been constructed with PV+ interneurons and PCs using the connectivity profile we observed experimentally in S2/3, with one representative PC in black at the center. Other PCs are in blue, connected PV+ interneurons in red, and unconnected PV+ interneurons in gray. The density of connected PV+ interneurons converging on this one PC can be visualized in three dimensions.

Supplemental Figure 1-1: Point spread function of DOE beamlets

The point spread function of the DOE beamlets was measured with the 20X 0.5 NA objective we used for mapping and 0.17 μm fluorescent beads. A: The lateral size of each DOE beamlet is approximately 1 μm (inset, raw image). B: The axial resolution of each beamlet is 6 μm (FWHM, full width at half maximum).

Supplemental Figure 1-2:

A: The PC indicated by the yellow arrow in B was photostimulated directly to produce action potentials 10 times. B: The locations of the 180 stimulation targets which were stimulated during the recording of the pyramidal cell (PC) indicated by the yellow arrow. C: Photostimulation of this target (lower red circle in B) caused the patched PC to fire an

action potential two out of three times. D: This target (upper red circle in A), which was just slightly offset from the center of the patched PC, caused the cell to fire only once.

Supplemental Figure 1-3:

Correlation between the connection probability and the depth of the patched PC is small ($R^2 = 0.08$, n = 82 maps).

Latency (ms)	0.87±0.05
Amplitude (pA)	222±88
Rise time (ms)	3.33±0.42
Decay time (ms)	12.21±1.45
Rate of rise (pA/ms)	76.44±27.66
Failure rate	0.10±0.06
Paired pulse ratio (%)	0.91±0.12
Summation (%)	1.34±0.17
Train depression (%)	0.64±0.11
Recovery (%)	0.90±0.09

Table 1-1: Characteristics of the synaptic connections from PV+ interneurons to PCs (n=10 from S2/3, S5, and F2/3). Note that during the recordings the postsynaptic PC was held at +40mV with the chloride reversal set to -80mV.

Figure 1-1

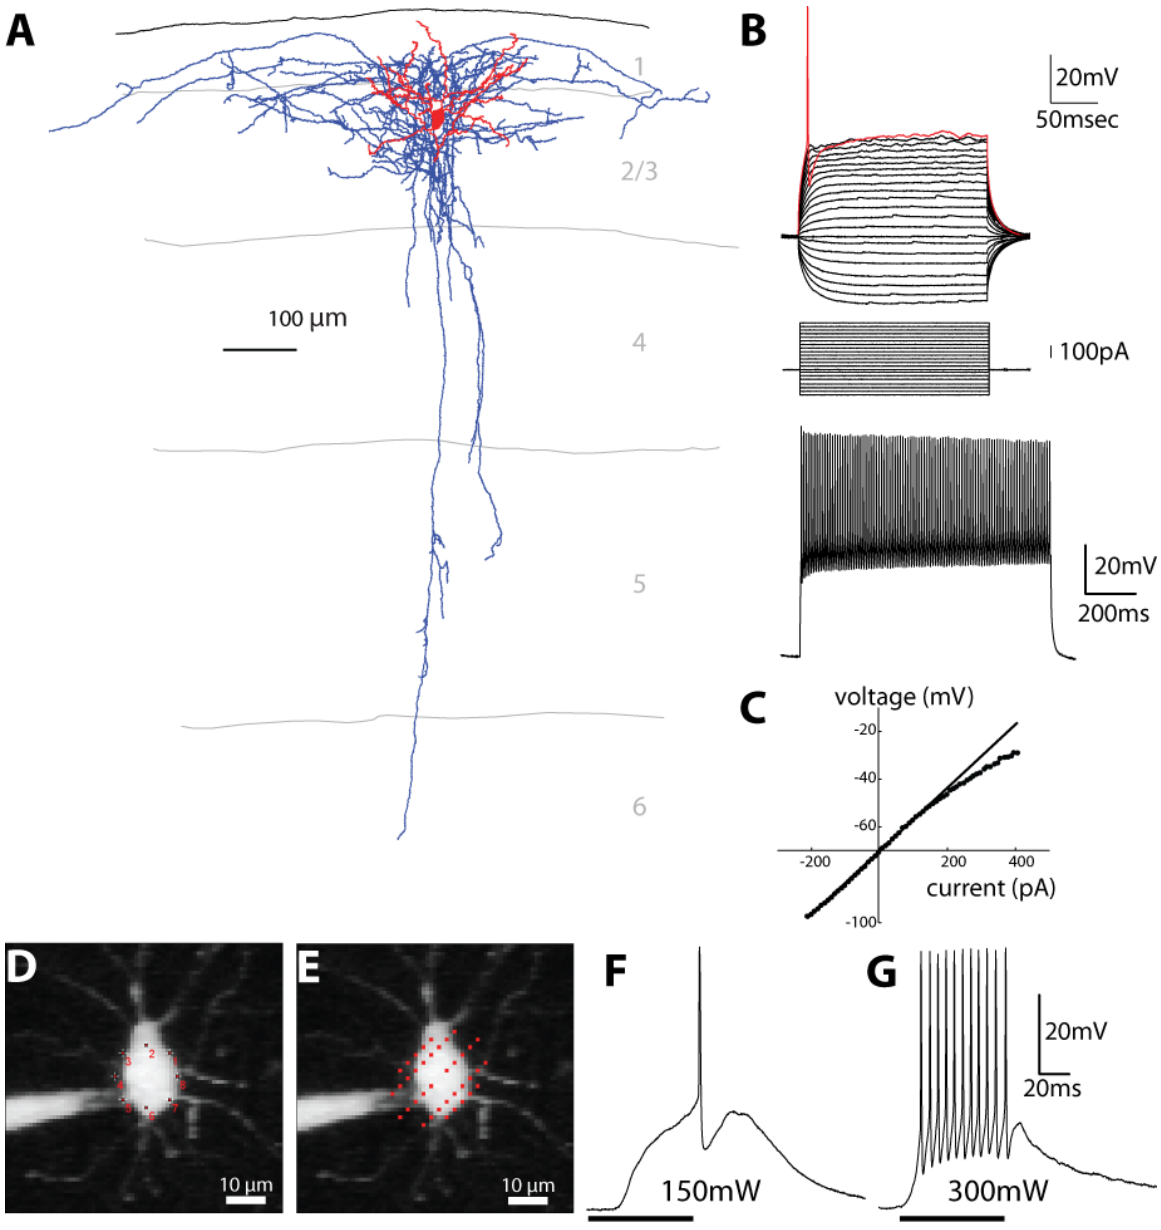
A
1
2/3
4
5
6
100 µm
B
20mV
50msec
100pA
20mV
200ms
C
voltage (mV)
-20
-40
-60
-200
200
400
current (pA)
-100
D
10 µm
E
10 µm
F
150mW
G
20mV
20ms
300mW

Figure 1-2

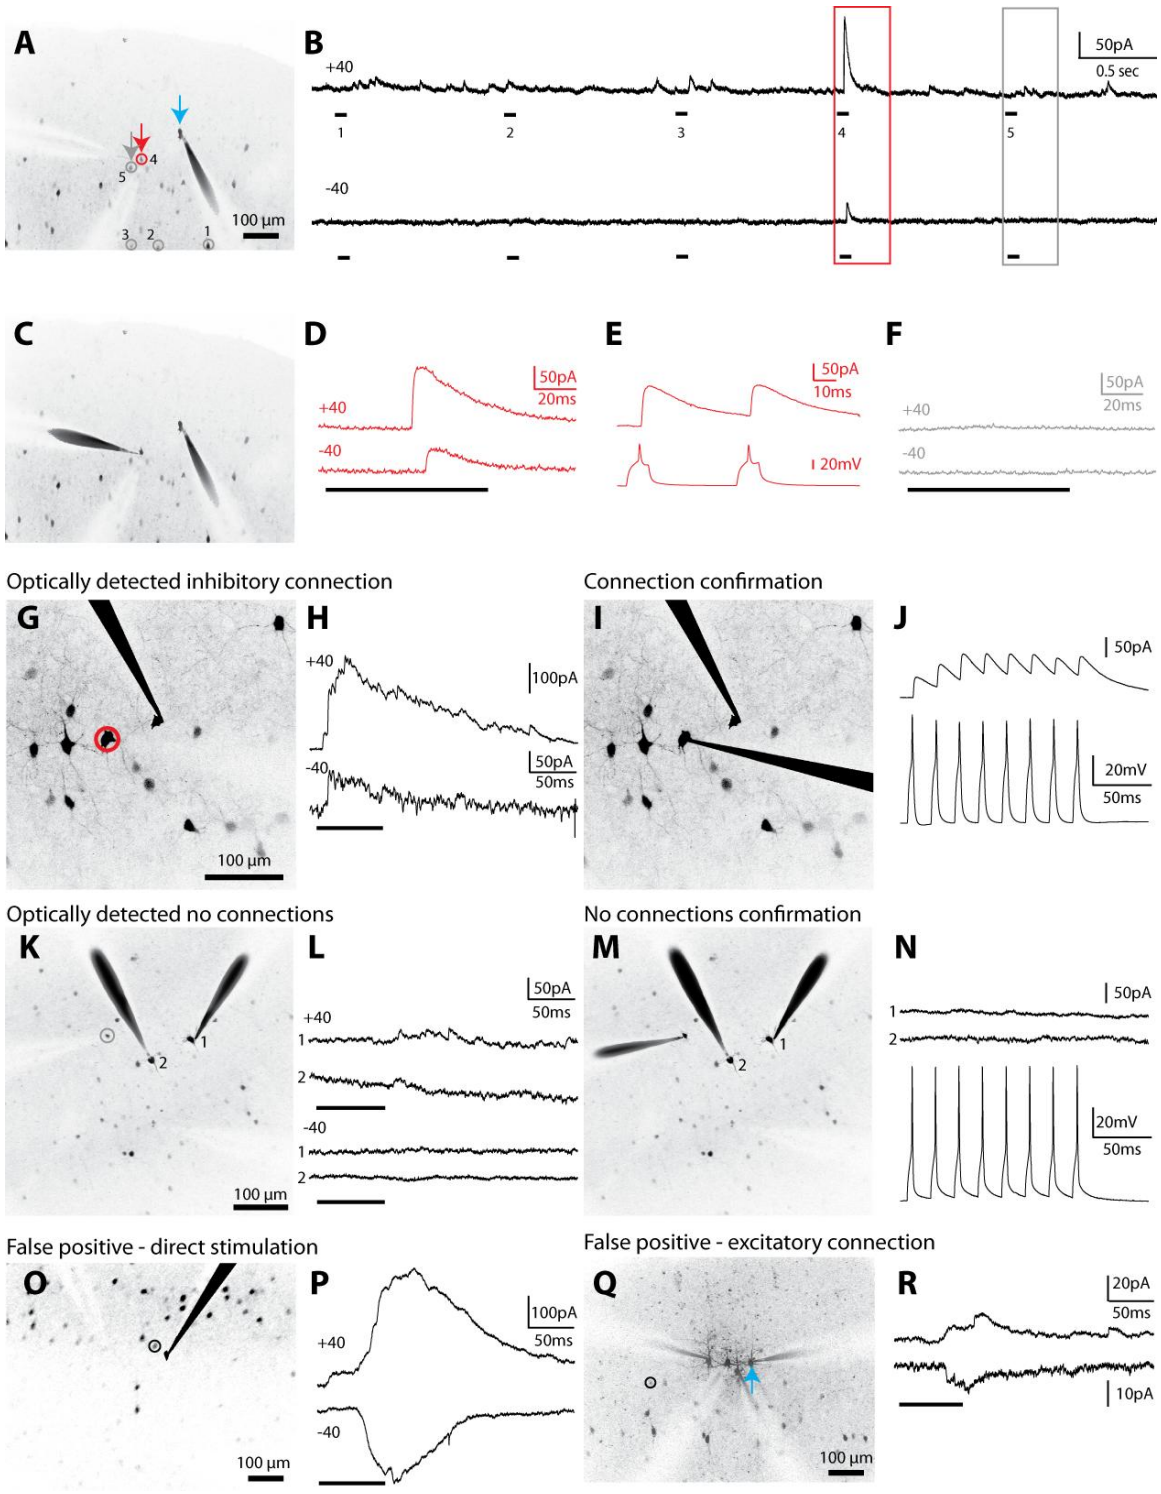
A
B
+40
-40
50pA
0.5 sec
1
2
3
4
5
100 µm
C
D
50pA
20ms
+40
-40
E
50pA
10ms
20mV
F
50pA
20ms
+40
-40
Optically detected inhibitory connection
Connection confirmation
G
H
100pA
50pA
50ms
100 µm
I
J
50pA
20mV
50ms
Optically detected no connections
No connections confirmation
K
L
50pA
50ms
M
N
20mV
50ms
False positive - direct stimulation
False positive - excitatory connection
O
P
100pA
50ms
Q
R
20pA
50ms
10pA

Figure 1-3

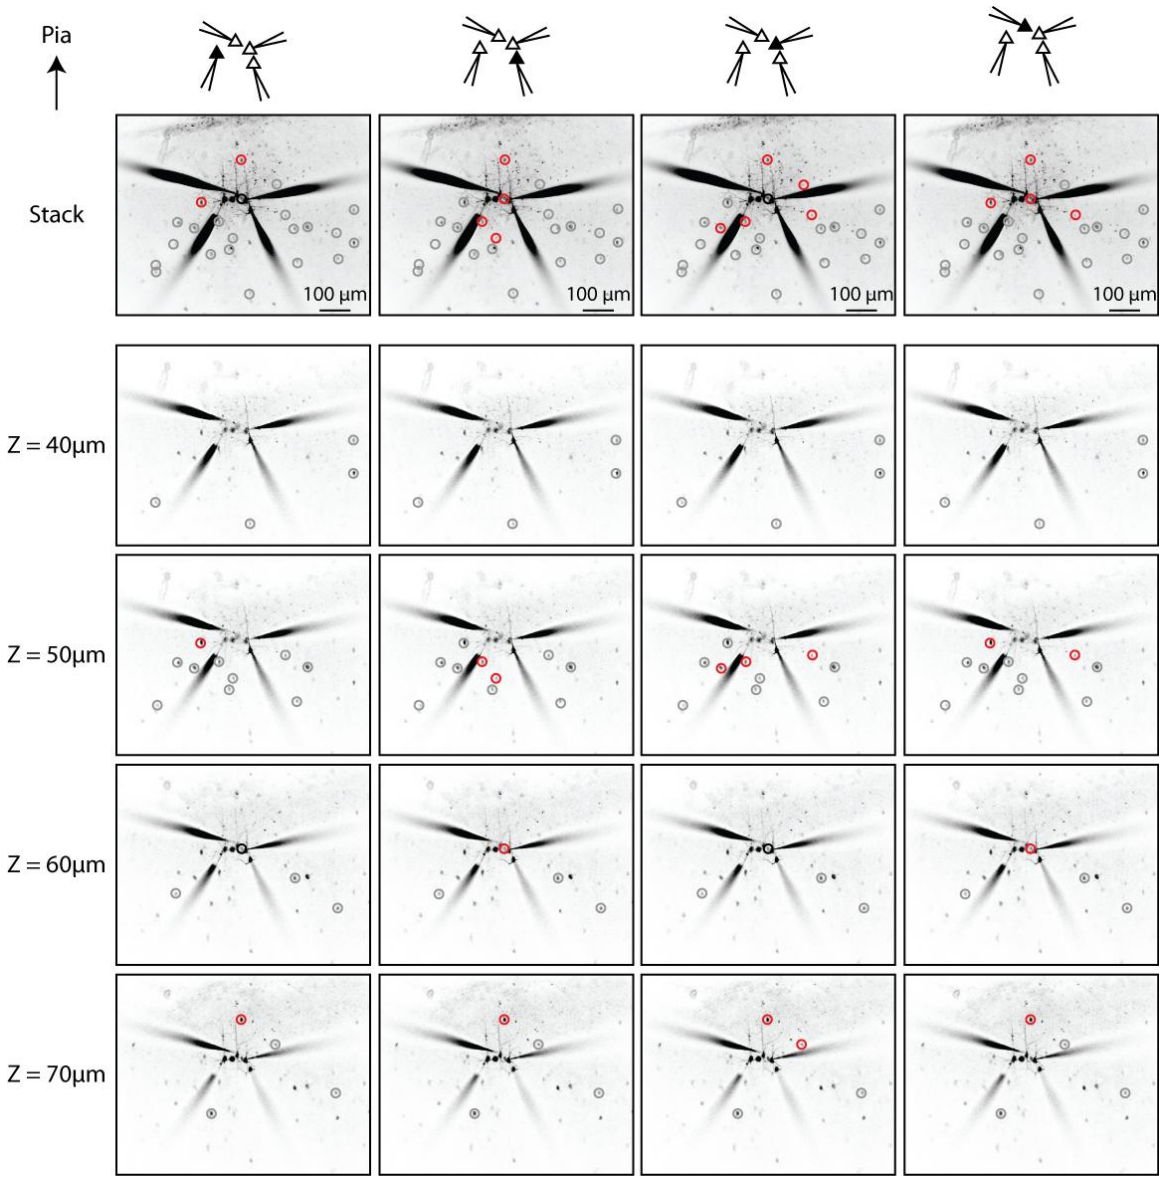

Figure 1-4

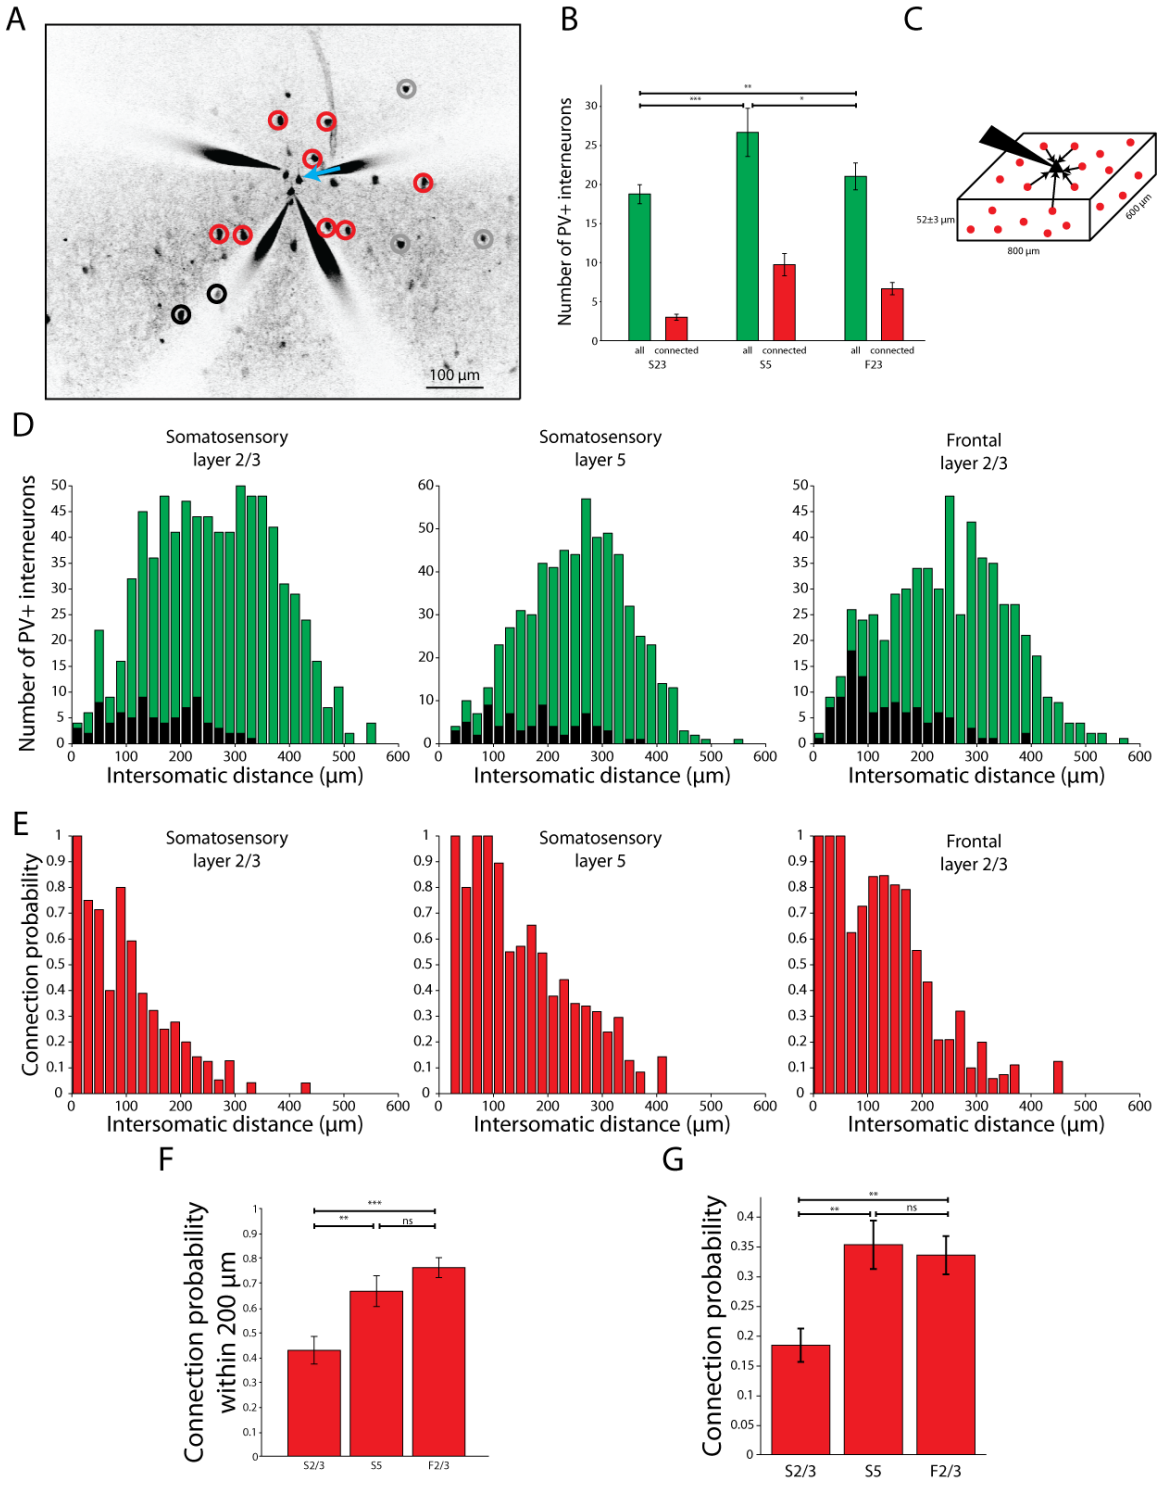
A
100 µm
B
Number of PV+ interneurons
all connected
S23
all connected
S5
all connected
F23
C
52±3 µm
800 µm
600 µm
D
Somatosensory layer 2/3
Somatosensory layer 5
Frontal layer 2/3
Number of PV+ interneurons
Intersomatic distance (µm)
E
Connection probability
Intersomatic distance (µm)
F
Connection probability within 200 µm
S2/3
S5
F2/3
G
Connection probability
S2/3
S5
F2/3

Figure 1-5

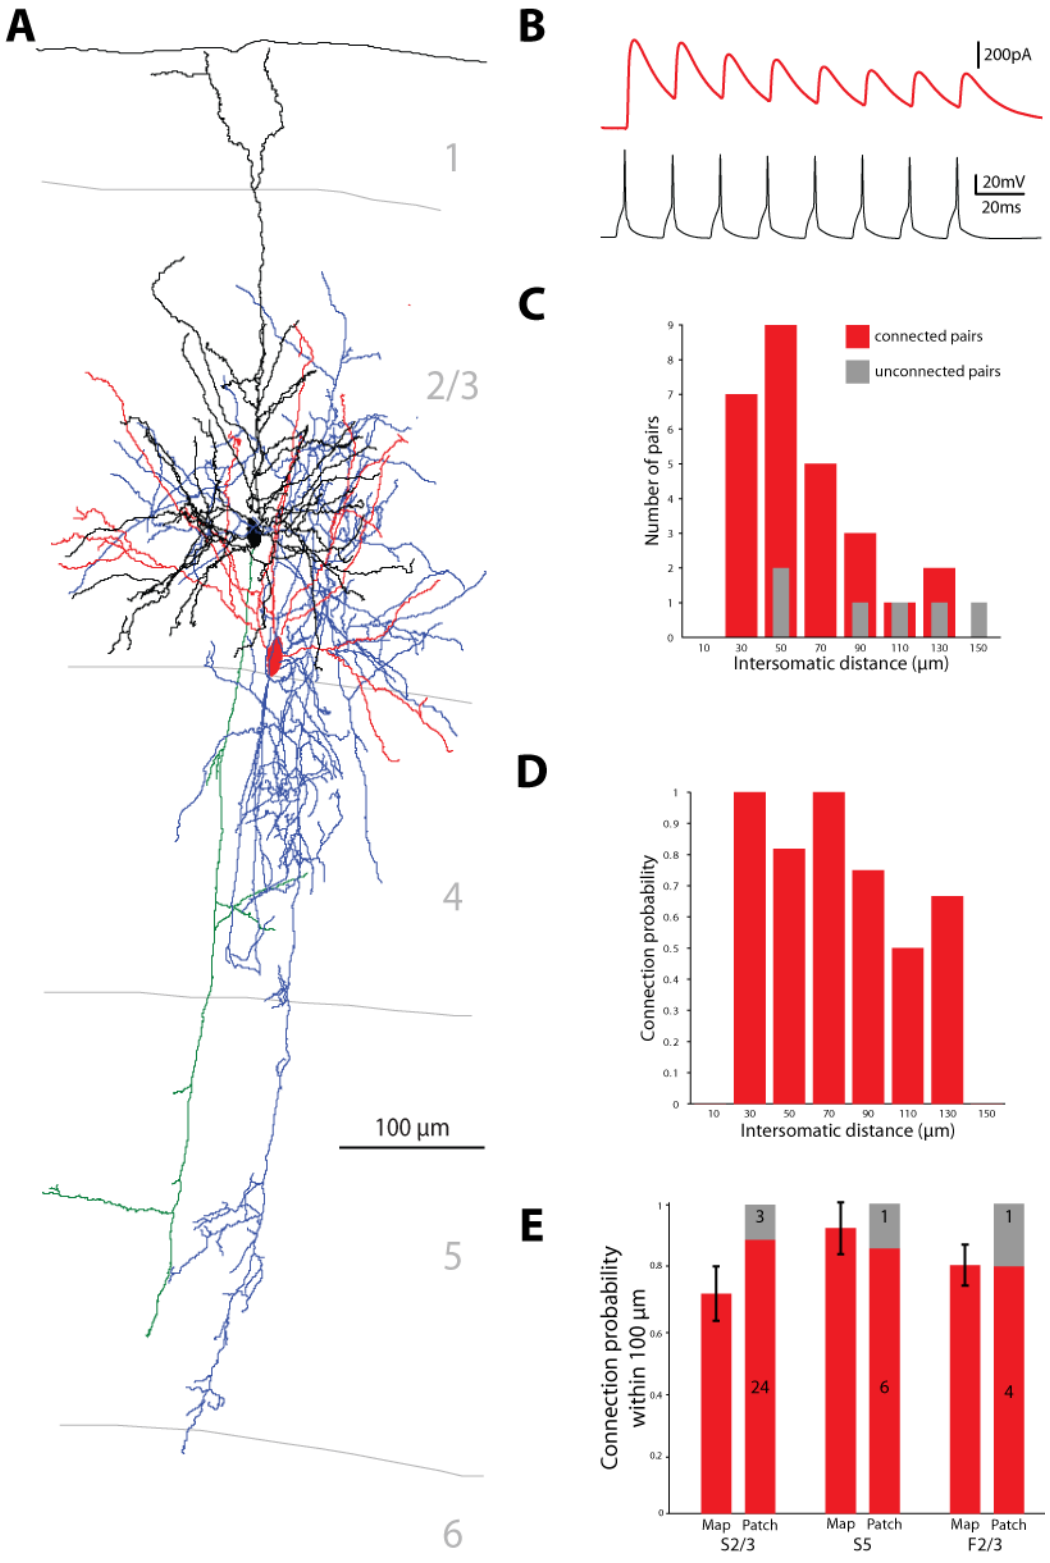
A
1
2/3
4
5
6
100 µm
B
200pA
20mV
20ms
C
connected pairs
unconnected pairs
Number of pairs
Intersomatic distance (µm)
D
Connection probability
Intersomatic distance (µm)
E
Connection probability within 100 µm
3
1
1
24
6
4
Map Patch
S2/3
Map Patch
S5
Map Patch
F2/3

Figure 1-6

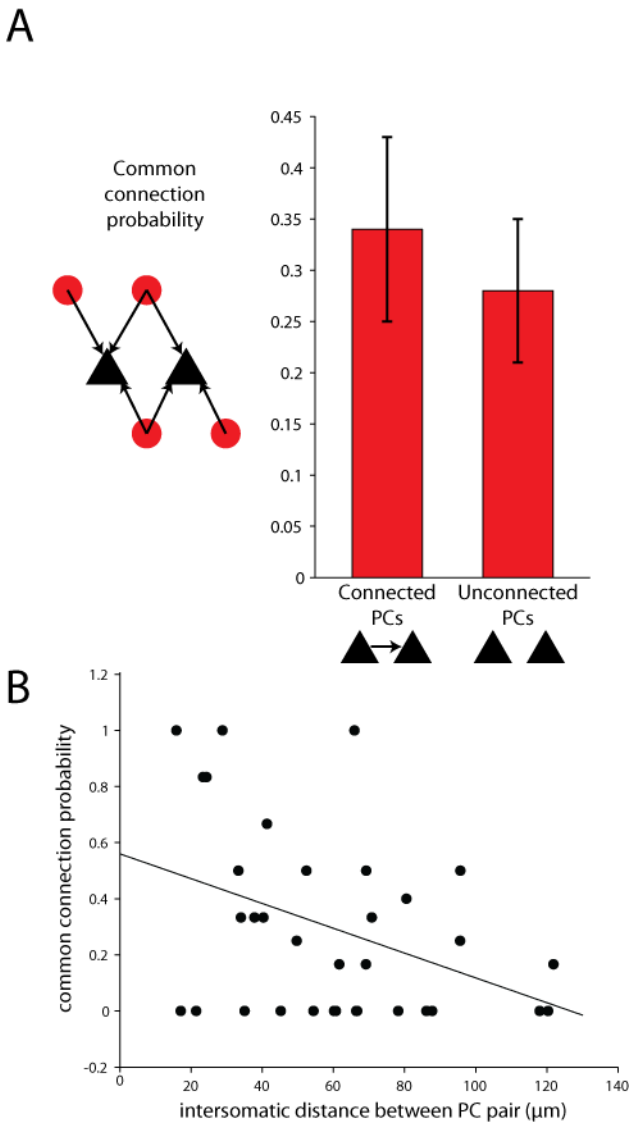
A
Common connection probability
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Connected PCs
Unconnected PCs
B
common connection probability
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120
140
intersomatic distance between PC pair (µm)

Figure 1-7

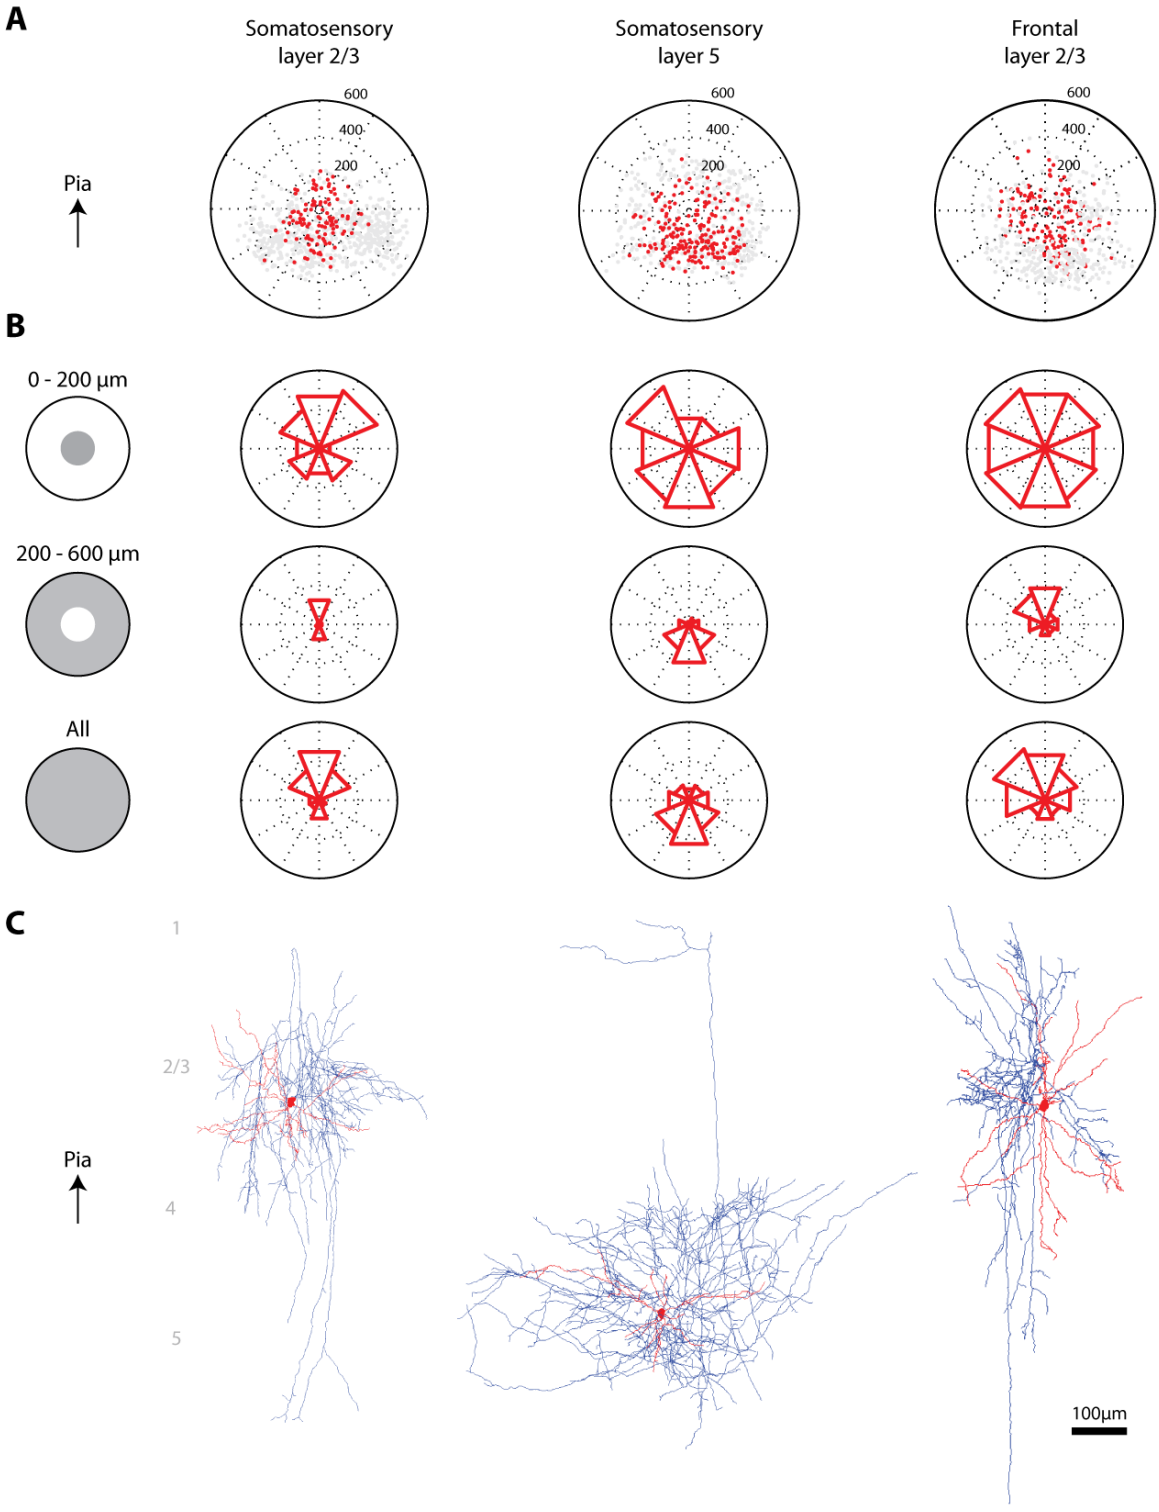
A
Somatosensory layer 2/3
Somatosensory layer 5
Frontal layer 2/3
600
400
200
Pia
B
0 - 200 µm
200 - 600 µm
All
C
1
2/3
4
5
Pia
100µm

Figure 1-8

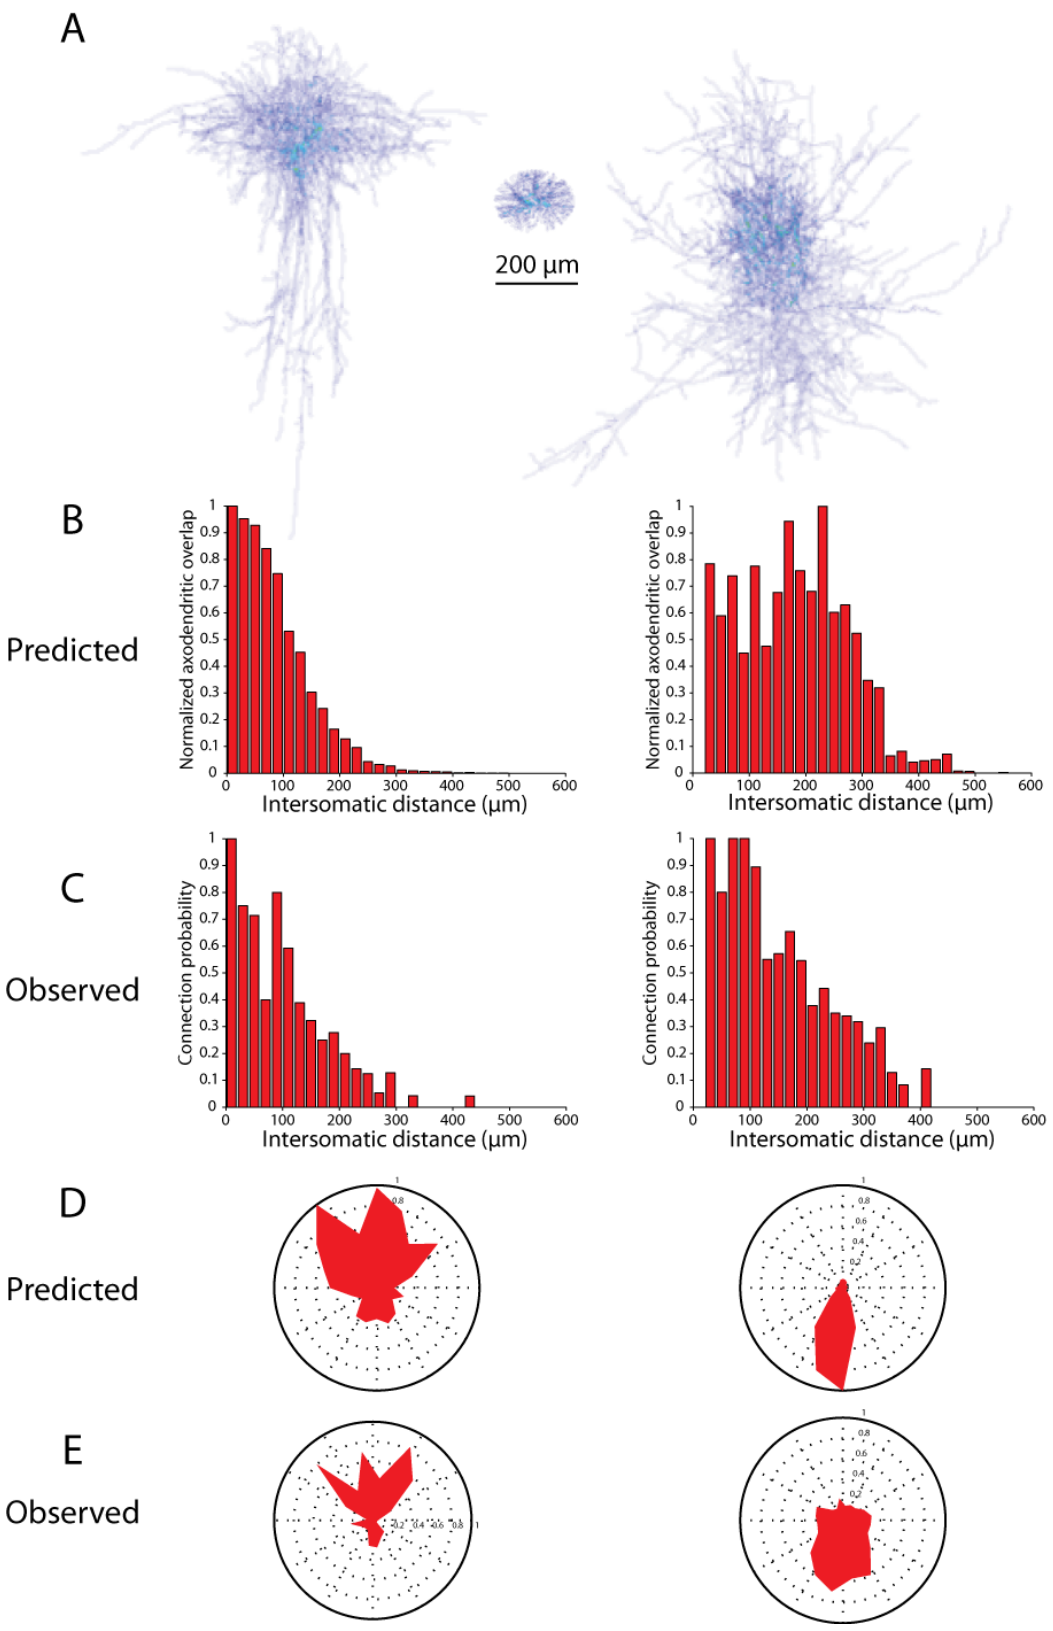
A
200 µm
B
Predicted
Normalized axodendritic overlap
Intersomatic distance (µm)
C
Observed
Connection probability
Intersomatic distance (µm)
D
Predicted
E
Observed

Figure 1-9

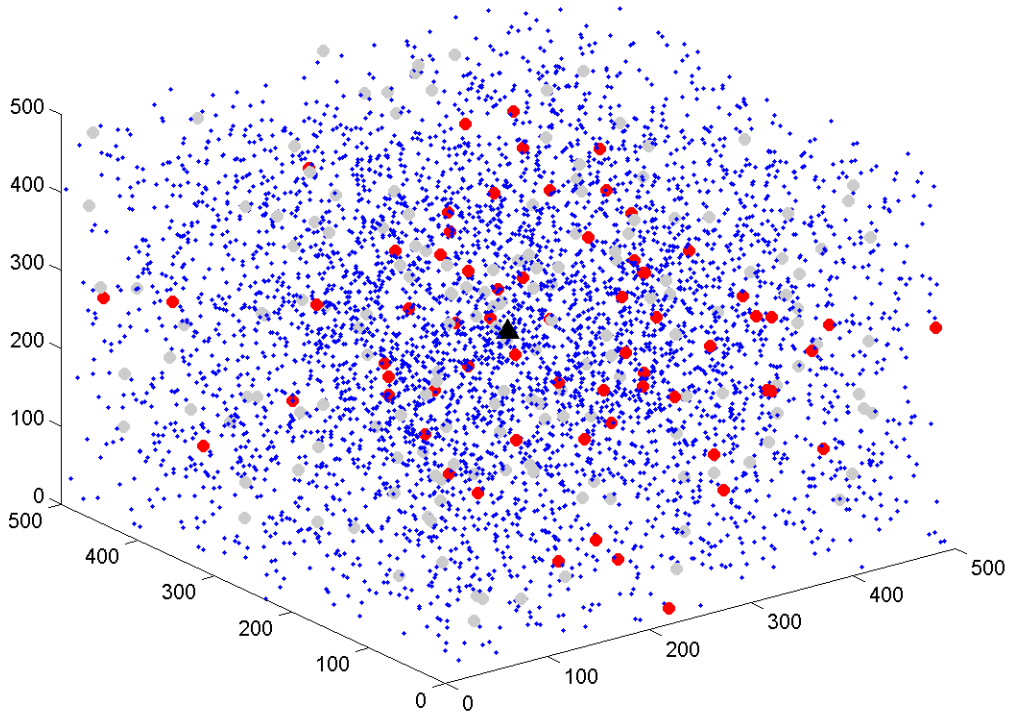

Supplemental Figure 1-1

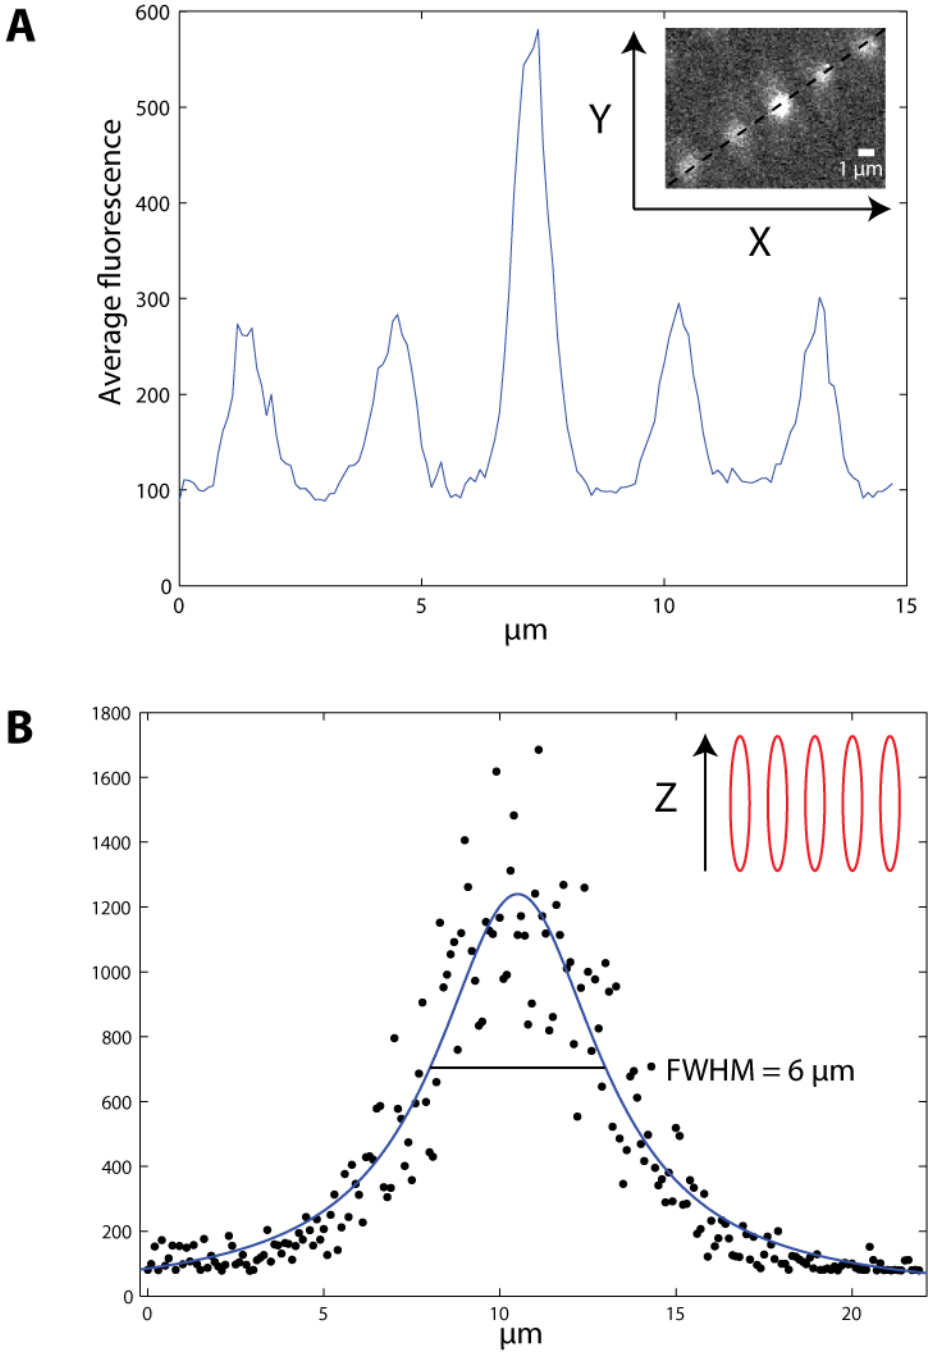

Supplemental Figure 1-2

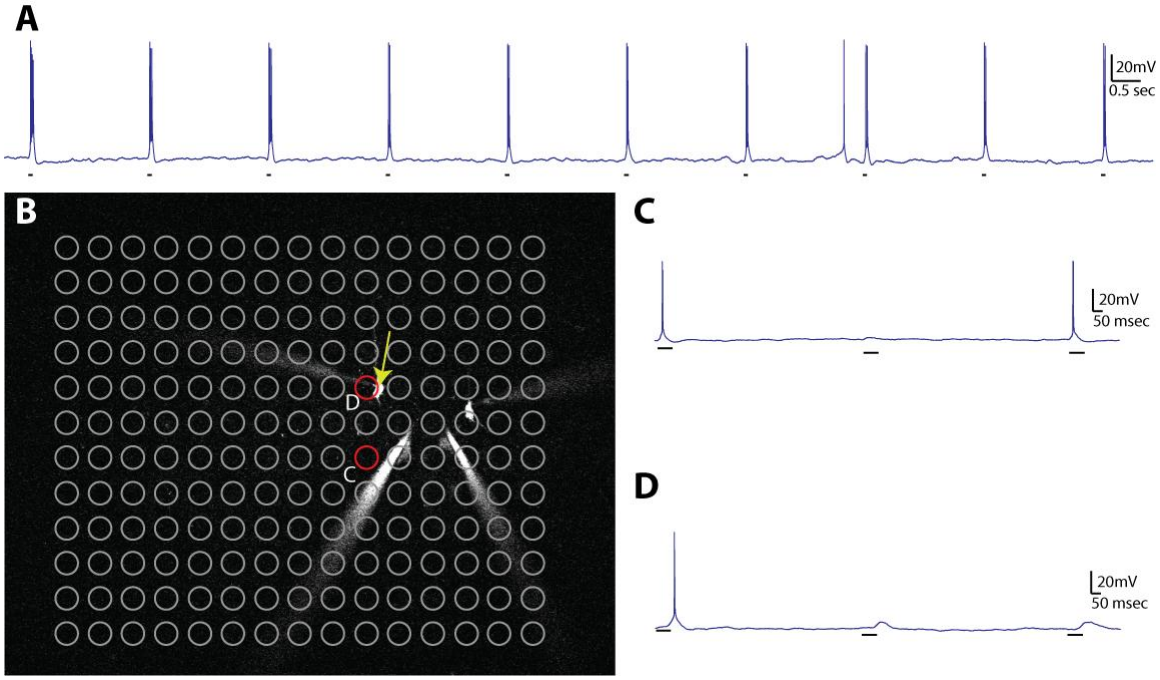

Supplemental Figure 1-3

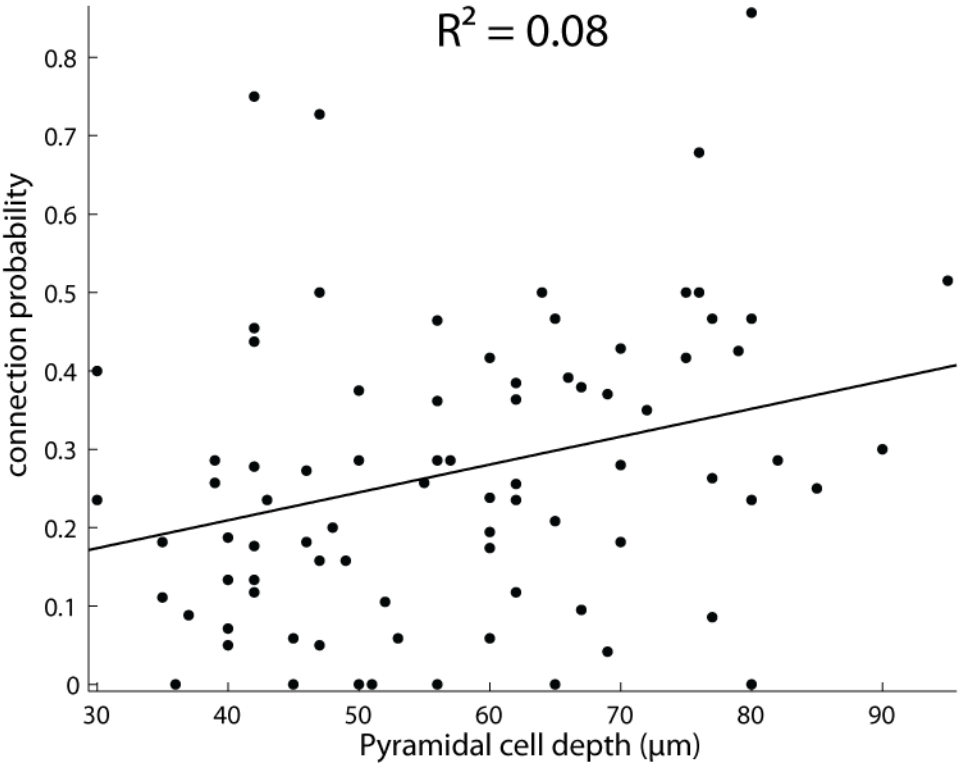

CHAPTER 2

Summary

In one of the most remarkable feats of motor control in the animal world, some Diptera, such as the housefly, can accurately execute corrective flight maneuvers in tens of milliseconds. These reflexive movements are achieved by the halteres, gyroscopic force sensors, in conjunction with rapidly-tunable wing-steering muscles. Specifically, the mechanosensory campaniform sensilla located at the base of the halteres transduce and transform rotation-induced gyroscopic forces into information about the angular velocity of the fly's body. But how exactly does the fly's neural architecture generate the angular velocity from the lateral strain forces on the left and right halteres? To explore potential algorithms, we built a neuro-mechanical model of the rotation detection circuit. We propose a neurobiologically plausible method by which the fly could accurately separate and measure the three-dimensional components of an imposed angular velocity. Our model assumes a single sign-inverting synapse and formally resembles some models of directional selectivity by the retina. Using multidimensional error analysis, we demonstrate the robustness of our model under a variety of input conditions. Our analysis reveals the maximum information available to the fly given its physical architecture and the mathematics governing the rotation-induced forces at the haltere's end knob.

Introduction

The true flies, insects of the order Diptera, are capable of astounding aerobatic feats of agility with only a single pair of wings. A male housefly can make corrective course changes to track a female fly within 30 msec (Land, 1974). This complex task requires substantial information processing but is performed in an organism whose brain contains only 100,000 neurons. As such, flies need an accurate circuit capable of detecting changes in body rotation. A highly specialized mechanical apparatus is used to provide appropriate information to these local circuits and enable flight stabilization with incredibly short response latencies. Mechanosensation of rotational stimuli is generated by the halteres, barbell-shaped appendages that evolved from hind wings. They function as gyroscopic sensors of rotational force by virtue of their constant oscillatory motion anti-phase to the wing (Nalbach, 1993; Pringle, 1948).

The halteres are positioned behind the forewings and are sheltered from many aerodynamic forces during flight (Dickinson, 1999; Nalbach, 1993). During a rotation, however, the halteres' motion occurs in a non-inertial reference frame and the halteres are accordingly subjected to various forces and pseudoforces. Acceleratory rotations about the axes of the fly's body generate centrifugal, Euler's (due to angular acceleration), and Coriolis forces, among others, which put lateral strain on the moving haltere (with the lateral direction defined as perpendicular to the beating plane of the haltere.) The haltere oscillates approximately harmonically, generating forces in a sinusoidal fashion (see Nalbach 1993). During a rotation, the forces described above are thought to be detected by mechanoreceptor strain gauges called campaniform sensilla in the base of the halteres. Detection of these forces generates the input necessary to calculate the angular velocity

of the fly's body. The output of the halteres impinges on the steering muscles of the wing, causing them to stiffen or relax, changing the course of the wing and thus a change in the trajectory of the fly.

The flies only sense the lateral components of force acting on the halteres, as has been shown in experiments in which flies do not show any corrective responses to rotations around the axes of the halteres. Such rotations only generate forces in the radial direction, or along the axis of the haltere, with apparently no component projecting laterally from the haltere beat plane (Dickinson, 1999; Nalbach, 1993). In addition, haltere afferents likely evolved to detect only the lateral forces due to the massive, continuous presence of primary inertial forces in the radial and tangential directions, which outweigh the Coriolis forces in those directions by orders of magnitude. There are several anatomically distinct groups of campaniform sensilla; it is not clear whether all groups sense lateral force or only a subset (Fayyazuddin and Dickinson, 1996). These sensors are located at the haltere's hinge, endowing them with maximal strain detection (Wu, 2002).

In this study, we sought to better understand the computations underlying these aspects of the fly’s motor control and modeled how the halteres could be transducing Coriolis forces into accurate measurements of angular velocity. The only angular variable correlated with wing beat amplitude that directly modulates the flight trajectory is angular velocity, implying that flies have access to this information (Dickinson, 1999). The complete aerodynamic system of the fly has been modeled previously (Dickson, 2008), but the neural architecture and computation that could produce reliable estimates of angular velocity from the lateral forces on the haltere during a body rotation have yet

to be explored sufficiently. More recently, Thompson and colleagues also modeled how such a calculation could be performed (Thompson et al., 2009), but it is unclear how their algorithm could be implemented neurobiologically.

The understanding of a complex information processing system such as this one requires multiple layers of analysis. In this system, the layer of analysis that has received the least attention is the algorithmic one (Marr, 1982): how does the hardware of the fly allow for the software to perform the necessary calculations? The question in this specific context is diagrammed in Figure 2-1. How does the fly's neural architecture (Figure 2-1C) algorithmically produce the desired output, angular velocity, from the given input, i.e. the waveform of total lateral force on the left and right halteres (Figure 2-1B middle)? The neuromechanical model we propose in this paper not only explains observations of haltere structure and function but also yields novel and readily testable predictions. Our analytical approach underscores the power of incorporating the biomechanics of a physical system into a model of its function (Briggman and Kristan, 2008).

Materials and Methods

Haltere motion

Parameterization of the equation for simple harmonic motion[1], $\Phi(t) = a\sin(\omega t + \varphi)$, leads to the following expressions for the position of the left and right halteres. We define the x (pitch) axis to be the transverse axis of the fly, the y (roll) axis to be the longitudinal axis of the fly, and the z (yaw) axis to point orthogonally away from the other two axes. The following equations describe the position of the haltere in this coordinate system:

$$X_L = A\cos\big(\Phi(t)\big)\cos\theta$$

$$X_R = -A\cos\big(\Phi(t)\big)\cos\theta$$

$$Y_L = A\cos\big(\Phi(t)\big)\sin\theta$$

$$Y_R = A\cos\big(\Phi(t)\big)\sin\theta$$

$$Z_L = A\sin\big(\Phi(t)\big)$$

$$Z_R = A\sin\big(\Phi(t)\big)$$

where Φ(t) is the function describing the angle the haltere makes with the xy axis, a is the amplitude of the angular oscillation of the haltere, ω is the frequency of the haltere's angular oscillation, φ is the phase shift of the oscillation, A is the length of the haltere, and θ is the angle the haltere makes with the transverse axis, in the xy plane.

---

1 The oscillation of the haltere is more like a square wave than a sinusoid (Nalbach, 1993). However, for the purposes of this model simple harmonic motion produces sufficient detail.

We now proceed to describe the forces which impinge on the haltere in a lateral direction. (Note: In this section we discuss only lateral forces, where the lateral direction is defined as orthogonal to the beat plane of the haltere. Positive and negative lateral directions are arbitrarily determined to be pointing towards and away from the fly's thorax, respectively. The radial direction is along the stalk of the haltere, and the tangential direction is the direction of the tangential velocity of the haltere.)

The Coriolis force on the haltere in the lateral direction

Coriolis forces act on the mass of the moving haltere in a rotating reference frame. The Coriolis forces are the only forces which directly contain information about the angular velocity, so we focus on them in our algorithm to recover this value so critical to flight stabilization.

Notation: **bold** denotes a three dimensional vector, e.g. $\mathbf{X} = \begin{cases} \mathrm{x} \\ \mathrm{y} \\ \mathrm{z} \end{cases}$

a dot on top of a variable denotes a time derivate, e.g. $\dot{\Omega} = \frac{d\Omega}{dt}$

The velocities of the left and right halteres are:

$$\boldsymbol{V}_L = \begin{cases} -A\dot{\Phi}(t)\cos\theta\sin\Phi(t) \\ -A\dot{\Phi}(t)\sin\theta\sin\Phi(t) \\ A\dot{\Phi}(t)\cos\Phi(t) \end{cases}$$

$$\boldsymbol{V}_R = \begin{cases} A\dot{\Phi}(t)\cos\theta\sin\Phi(t) \\ -A\dot{\Phi}(t)\sin\theta\sin\Phi(t) \\ A\dot{\Phi}(t)\cos\Phi(t) \end{cases}$$

Given an angular rotation, $\boldsymbol{\Omega}$, defined as:

$$\boldsymbol{\Omega} = \begin{cases} \Omega_{pitch} \\ \Omega_{roll} \\ \Omega_{yaw} \end{cases}$$

And the equation for the Coriolis force:

$$\boldsymbol{F}_{coriolis} = -2m(\boldsymbol{\Omega} \times \boldsymbol{V})$$

We can solve for the Coriolis force on the left haltere:

$$\frac{\boldsymbol{F}_{L,coriolis}}{-2m} = \begin{cases} \Omega_{roll} A\dot{\Phi}(t)\cos\Phi(t) + \Omega_{yaw} A\dot{\Phi}(t)\sin\theta\sin\Phi(t) \\ -\Omega_{yaw} A\dot{\Phi}(t)\cos\theta\sin\Phi(t) - \Omega_{pitch} A\dot{\Phi}(t)\cos\Phi(t) \\ -\Omega_{pitch} A\dot{\Phi}(t)\sin\theta\sin\Phi(t) + \Omega_{roll} A\dot{\Phi}(t)\cos\theta\sin\Phi(t) \end{cases}$$

To obtain the lateral component of the Coriolis force on the left haltere, project the force onto the vector defining the lateral projection of the haltere (-sinθ, cosθ, 0):

$$F_{L,coriolis,lateral} = \frac{\boldsymbol{F}_{L,coriolis}}{-2m} \cdot \begin{cases} -\sin\theta \\ \cos\theta \\ 0 \end{cases}$$

$$= -\sin\theta\left(\Omega_{roll} A\dot{\Phi}(t)\cos\Phi(t) + \Omega_{yaw} A\dot{\Phi}(t)\sin\theta\sin\Phi(t)\right)$$

$$+ \cos\theta\left(-\Omega_{yaw} A\dot{\Phi}(t)\cos\theta\sin\Phi(t) - \Omega_{pitch} A\dot{\Phi}(t)\cos\Phi(t)\right)$$

Grouping like terms:

$$\frac{F_{L,coriolis,lateral}}{-2m} =$$

$$A\dot{\Phi}(t)\big[-\Omega_{roll}\cos\Phi(t)\sin\theta - \Omega_{yaw}(\sin^2\theta + \cos^2\theta)\sin\Phi(t)$$

$$- \Omega_{pitch}\cos\Phi(t)\cos\theta\big]$$

Note that the yaw term Ω_{yaw} reduces due to the trigonometric identity: $\sin^2\theta + \cos^2\theta = 1$

Leaving the following expression for the Coriolis force on the left haltere *in the lateral direction only*:

$$F_{L,coriolis,lateral} = 2mA\dot{\Phi}\big[\Omega_{roll}\cos\Phi(t)\sin\theta + \Omega_{yaw}\sin\Phi(t) + \Omega_{pitch}A\dot{\Phi}(t)\cos\Phi(t)\cos\theta\big]$$

The calculation for the right haltere is similar except for minor changes in sign:

$$\frac{\boldsymbol{F}_{R,coriolis}}{-2m} = \begin{cases} \Omega_{roll}A\dot{\Phi}(t)\cos\Phi(t) + \Omega_{yaw}A\dot{\Phi}(t)\sin\theta\sin\Phi(t) \\ \Omega_{yaw}A\dot{\Phi}(t)\cos\theta\sin\Phi(t) - \Omega_{pitch}A\dot{\Phi}(t)\cos\Phi(t) \\ -\Omega_{pitch}A\dot{\Phi}(t)\sin\theta\sin\Phi(t) - \Omega_{roll}A\dot{\Phi}(t)\cos\theta\sin\Phi(t) \end{cases}$$

$$F_{R,coriolis,lateral} = \frac{\boldsymbol{F}_{R,coriolis}}{-2m}\cdot\begin{cases}\sin\theta\\ \cos\theta\\ 0\end{cases}$$

$$= \sin\theta\left(\Omega_{roll}A\dot{\Phi}(t)\cos\Phi(t) + \Omega_{yaw}A\dot{\Phi}(t)\sin\theta\sin\Phi(t)\right) + \cos\theta\left(\Omega_{yaw}A\dot{\Phi}(t)\cos\theta\sin\Phi(t) - \Omega_{pitch}A\dot{\Phi}(t)\cos\Phi(t)\right)$$

$$\frac{F_{R,coriolis,lateral}}{-2m} =$$

$$A\dot{\Phi}(t)\big[\Omega_{roll}\cos\Phi(t)\sin\theta + \Omega_{yaw}(\sin^2\theta + \cos^2\theta)\sin\Phi(t) - \Omega_{pitch}\cos\Phi(t)\cos\theta\big]$$

$$F_{R,coriolis,lateral}$$

$$= 2mA\dot{\Phi}\left[-\Omega_{roll}\cos\Phi(t)\sin\theta - \Omega_{yaw}\sin\Phi(t)\right.$$

$$\left. + \Omega_{pitch}\cos\Phi(t)\cos\theta\right]$$

Other forces on the haltere in the lateral direction

During a rotation, three other forces — angular acceleration, centrifugal, and gravitational — have components which project laterally. These forces can be derived using the same trigonometric identity as in the derivation for the Coriolis force.

Angular Acceleration: $F_{aa} = -m\left(\dot{\Omega} \times r\right)$

The lateral components of the angular acceleration force generated by pitch, roll, and yaw rotations are equal to:

$F_{aa,L,pitch} = F_{aa,R,pitch} = mA\dot{\Omega}_{pitch}\sin\Phi(t)\cos\theta$

$F_{aa,L,roll} = mA\dot{\Omega}_{roll}\sin\Phi(t)\sin\theta$

$F_{aa,R,roll} = -mA\dot{\Omega}_{roll}\sin\Phi(t)\sin\theta$

$F_{aa,L,yaw} = -mA\dot{\Omega}_{yaw}\cos\Phi(t)$

$F_{aa,R,yaw} = mA\dot{\Omega}_{yaw}\cos\Phi(t)$

where m is the mass of the haltere, $\dot{\Omega}$ is the angular acceleration of the fly's body, and r is the position of the haltere within the rotating reference frame.

Centrifugal: $F_{cf} = -m\big(\Omega \times (\Omega \times \mathrm{r})\big)$

The lateral components of the centrifugal force generated from roll and pitch rotations are equal to:

$$F_{cf,L,pitch} = mA\Omega_{pitch}^2 \cos\Phi(t) \cos\theta \sin\theta$$

$$F_{cf,R,pitch} = mA\Omega_{pitch}^2 \cos\Phi(t) \cos\theta \sin\theta$$

$$F_{cf,L,roll} = -mA\Omega_{roll}^2 \cos\Phi(t) \cos\theta \sin\theta$$

$$F_{cf,R,roll} = -mA\Omega_{roll}^2 \cos\Phi(t) \cos\theta \sin\theta$$

Gravitational: $F_g = mg \sin\Psi$

The gravitational force depends on ψ, the angle the lateral direction with respect to the haltere makes with the gravitational force. g is the acceleration due to gravity. Because ψ depends on the lateral direction of the haltere, it is affected only by the angular position of the fly's body with respect to the ground, not by the position of the haltere with respect to the fly's body. We use quaternions to calculate the lateral direction of the haltere. From the initial angular velocities and angular accelerations, an axis in space and an angle Θ are computed through which the fly rotates about during the time until the required measurement is taken. The vector indicating the lateral direction with respect to the haltere's beat plane is derived with quaternion multiplication:

$$q_1 = \left(\cos\frac{\Theta}{2}, \langle\mu\rangle * \sin\frac{\Theta}{2}\right)$$

$$p_1 = (0, \langle v_0\rangle)$$

$$q_1' = \left(\cos\frac{\Theta}{2}, -\langle\mu\rangle * \sin\frac{\Theta}{2}\right)$$

$$p_2 = (q_1) * (p_1) * (q_1') = (0, \langle p\rangle)$$

where $\langle\mu\rangle$ is a vector representing the three-dimensional Euclidian direction of the axis around which the fly rotates, Θ is the angle through which the fly rotates, $\langle v_0\rangle$ is the initial, three-dimensional Euclidian direction of the lateral direction with respect to the haltere beat plane, and p_2 is a quaternion equal to $(0, \langle p\rangle)$, a vector indicating the final 3 dimensional direction of the lateral direction with respect to the haltere beat plane. Note: During our simulations, the fly often begins from a pitched and rolled position in space, so the pitch, roll, and yaw axes must be recalculated with respect to the fly's body, not to the ground. The pitch, roll, and yaw axes must be rotated to maintain their definition with respect to the fly's body. This is calculated via quaternion multiplication as well.

At any point during a rotation, a haltere is exposed to a sum of these forces. Forces due to angular acceleration, or Euler's forces, are orders of magnitude smaller than Coriolis forces except at low angular velocities. The gravitational force, F_g, only has a lateral component when the fly's body is angled relative to flat ground. The lateral component of the gravitational force, however, is not constant during a rotation as has been claimed previously (Nalbach, 1993; Wu, 2002) but rather changes depending on the angular position of the fly's body in space.

The total lateral component of the forces on the left and right halteres during a rotation is the sum of the lateral components of each of the four forces described above:

$$F_{L,lateral} = mA\big[2\dot{\Phi}\big(\Omega_{roll}\cos\Phi(t)\sin\theta + \Omega_{yaw}\sin\Phi(t) + \Omega_{pitch}(t)\cos\Phi(t)\cos\theta\big) + \dot{\Omega}_{roll}\sin\Phi(t)\sin\theta + \dot{\Omega}_{pitch}\sin\Phi(t)\cos\theta - \dot{\Omega}_{yaw}\cos\Phi(t) + \left(\Omega_{pitch}^2 - \Omega_{roll}^2\right)\cos\Phi(t)\cos\theta\sin\theta\big] + F_{g,l}$$

$$F_{R,lateral} = mA\big[2\dot{\Phi}\big(-\Omega_{roll}\cos\Phi(t)\sin\theta - \Omega_{yaw}\sin\Phi(t) + \Omega_{pitch}(t)\cos\Phi(t)\cos\theta\big) - \dot{\Omega}_{roll}\sin\Phi(t)\sin\theta + \dot{\Omega}_{pitch}\sin\Phi(t)\cos\theta + \dot{\Omega}_{yaw}\cos\Phi(t) + \left(\Omega_{pitch}^2 - \Omega_{roll}^2\right)\cos\Phi(t)\cos\theta\sin\theta\big] + F_{g,R}$$

Results

Algorithm

If pitch, roll, and yaw rotations generate forces with lateral components, and the campaniform sensilla are sensitive only to forces in the lateral direction, how can the different rotations be distinguished from one another? In effect, how can the fly generate labeled lines for the pitch, roll, and yaw angular velocity components of a rotation when none are inherent to the haltere? The answer to these questions requires extracting meaningful measurements of the magnitude and direction of the body's rotation from the complex waveforms of the total lateral forces on each haltere, described by the equations above.

To recover the pitch component of the angular velocity, Ω_{pitch}, the total lateral forces on each haltere are added. This removes the contributions of the roll and yaw Coriolis and angular acceleration forces to the signal, as they are out of phase between the left and right halteres. This results in twice the lateral force due to pitch rotation, as well as other, smaller-magnitude terms, whose contribution may be removed via means mentioned in the Discussion.

$$\begin{aligned} F_L + F_R = {} & 4mA\dot{\Phi}\left[\Omega_{pitch}\cos\Phi(t)\cos\theta\right] \\ & +2mA\left[\dot{\Omega}_{pitch}\sin\Phi(t)\cos\theta + \left(\Omega_{pitch}^2 - \Omega_{roll}^2\right)\cos\Phi(t)\cos\theta\sin\theta\right] \\ & +F_{g,l} + F_{g,R} \end{aligned}$$

This expression has a sinusoid time dependence at the wing beat frequency, $\omega/2\pi$. The information about the pitch component of angular velocity is conveyed in the changing amplitude of this waveform, as can be seen in Figure 2-2a. Ω_{pitch} can be recovered by taking

amplitude measurements of this signal at specific times (equivalent to convolving the summed force signal with a delta function, see Chapter 2 appendix for details) and division by a constant, *k*:

$$\Omega_{pitch}(t_p) = \frac{F_{L(t_p)} + F_{R(t_p)}}{k}$$

$$k = 4mAa\omega \cos\theta$$

where $t_p = t_{p,0} + \frac{\omega}{2\pi} n$ for all positive integers n and $t_{p,0}$ is the first time at which the force due to pitch rotation reaches an extremum value. This time corresponds to a particular angular position of the haltere, Φ_p, for the pitch component (Figure 2-2a). *k* is the constant of proportionality between the angular velocity and the Coriolis force. See Chapter 2 appendix for details about Φ_p. This constant of proportionality must be built-in, arising through evolutionary selection (producing flies with increasingly accurate velocity readings) and/or tuned by some learning mechanism, not treated here. This constant may also include additional constants of proportionality relating to the strain detection by which the campaniform sensilla mechanically transduce the Coriolis forces.

The lateral component of the gravitational force changes during a pitch, roll, or yaw rotation, is sinusoidal with respect to the body's pitch angle, and remains a somewhat significant source of error in measuring the pitch angular velocity. Because the remaining pitch angular acceleration term is multiplied by $\sin\Phi(t)$, sampling at a Φ_p near zero radians effectively removes the angular acceleration pitch term. Though the lateral centrifugal force term is not removed, it is significantly smaller than the Coriolis force at most velocities. In addition, during a rotation with similar roll and pitch velocity components, this term will be effectively zero.

To recover the roll and yaw components of an applied angular velocity, the total lateral forces on the right haltere are subtracted from the left. Though it is more difficult to envision a neural manifestation of the subtraction operation, it is certainly not impossible. For example, ON/OFF cells in the visual system perform an algorithmically similar operation by employing inhibition and signed labeled lines (Ballard, 2008). Subtraction of total lateral force on the right haltere from that induced on the left haltere yields the following equation:

$$F_L - F_R = 4mA\dot{\Phi}\left[\Omega_{roll}\cos\Phi(t)\sin\theta + \Omega_{yaw}\sin\Phi(t)\right]$$
$$+2\mathrm{mA}\left[\dot{\Omega}_{roll}\sin\Phi(t)\sin\theta - \dot{\Omega}_{yaw}\cos\Phi(t)\right]$$
$$+F_{g,l} - F_{g,R}$$

The roll and yaw angular acceleration forces are not removed from the subtraction, as they are out of phase between the left and right halteres. They are, however, an order of magnitude smaller than the Coriolis force, and because they are not proportional to angular velocity, they only cause significant error in rotation detection at very low angular velocities.

This subtraction removes the Coriolis force due to pitch as well as the centrifugal force, as these forces are in phase between the halteres. This expression has a more complex time dependence, as is seen in Figure 2-2b and 2-2c. Though difficult to discern by eye, the resulting waveform has two frequency components, including the forces generated by both roll and yaw rotations. The roll component yields an oscillation at the wingbeat frequency, $\omega/(2\pi)$, while the yaw component yields an oscillation at twice the

wingbeat frequency, ω/π. Again, amplitude measurements of this signal at appropriate intervals, as shown in Figure 2-2b and 2-2c, followed by divisions with separate constants approximate Ω_{roll} and Ω_{yaw}. The sample points were chosen to minimize the effects of the angular velocity component not being detected.

$$\Omega_{roll}(t_r) = \frac{F_{L(t_r)} - F_{R(t_r)}}{k'}$$

$$k' = 4mAa\omega \sin\theta$$

$$\Omega_{yaw}(t_y) = \frac{F_{L(t_y)} - F_{R(t_y)}}{k''}$$

$$k'' = 4mAa\omega l$$

To detect roll, the force signal is sampled at $\Phi_r = 0$ radians, and to detect yaw the signal is sampled at $\Phi_y = {}^{5\pi}/_{11}$ radians, so for roll detection only yaw angular acceleration terms will be nonzero, and for yaw detection, only roll angular acceleration terms will be nonzero. These cause significant error in roll and yaw detection respectively only at very low angular velocities and nonzero angular accelerations. See Chapter 2 appendix for further description of the motivations for choosing these sample points and constants.

Neural Circuit

We can now posit a reasonable model for the circuitry in the black box in Figure 2-1. The algorithm we have proposed can be implemented in neural hardware conforming to the already known nervous system organization of the fly. The series of amplitude measurements can be taken assuming a neuron that fires in a phase locked pattern with respect to the haltere's angular position (Figure 2-3). This signal may be

simply an efferent copy from the central pattern generator (CPG) setting the rhythmic motion of the haltere or wing. The CPG may be driven or modulated by the opposing stretch-activated muscles that power the haltere. The signal from the CPG biases the intermediate cells performing the measurement, called N_{pitch}, N_{roll}, and N_{yaw}. These cells then drive the mnb1 motor neurons which act directly on the wing steering muscle, b1. At different stages in this diagram, inherent biophysical properties of neurons could easily provide the necessary subtractions and divisions required by the computational model. Certainly, the number of synapses here is small enough such that the required computation could be performed within the time restrictions imposed by the fly's quick reaction time.

Some of the anatomical connections between the haltere afferents and the steering muscles that would be necessary to implement this neural model have already been delineated in the fly. Contralaterally projecting haltere afferents, which would be necessary for addition and subtraction of the lateral force on each haltere, are provided by the contralaterally projecting haltere interneurons, or cHINs, identified by dye injections into haltere afferents (Strausfeld, 1985; Trimarchi and Murphey, 1997).

Error Analysis

This model mathematically produces accurate estimates of angular velocity for many types of physiologically relevant angular acceleratory and angular velocity inputs in all directions, but also yields systematic errors in conditions that have yet to be tested experimentally. We determined eight parameters which most affect the fly's ability to recover angular velocity. These parameters are the initial angular velocity and

acceleration in three dimensions as well as the initial pitch and roll angles of the fly with respect to ground. We systematically tested ten different values for each of these eight parameters. This required running our model with 10^8 different initial conditions. We tested initial angular velocities from -16.1 to 19.9 radians/sec as well as initial angular accelerations from -50 to 62.5 radians/sec^2 in the three orthogonal directions corresponding to pitch, roll, and yaw. We also tested initial pitch and roll angles of the fly's body relative to the ground from -50 to 60 degrees. After allowing the position and velocity of the fly to change from their initial values and calculating the relevant forces, we recovered angular velocities with our proposed model at the first available sample points (Φ_p, Φ_r, Φ_y).

For each combination of conditions, we calculated the absolute and relative errors of the recovered angular velocity. For each of the pitch, roll, and yaw components of the recovered angular velocity, we generated 10^8 absolute error values in radians and 10^8 relative error values in percentage of actual velocity at that particular sample point. Visualizing this data is inherently difficult, as the error values we generated depend on eight independent variables, which would require nine spatial dimensions to represent graphically. To solve this issue, we use the same graphing technique for visualizing data in high-dimensional spaces used by Taylor et. al (2006.) This method, based on clutter-based dimension reordering (Peng 2004; Peng 2005) and dimensional stacking, produces a figure in which the dimensions containing the least variability are compressed within the other dimensions. The dimensions are arranged such that salient trends in the data can be most clearly seen. The basics of this method are explained graphically in Figure 2-4. Akin to visualizing a three-dimensional stack as a set of two-dimensional slices, the

method yields a "montage of montages" such that each of the eight axes are embedded inside each other. The color of the pixel corresponds to the error at that particular set of conditions.

Figure 2-5a shows that the proposed model is worst at recovering roll velocities near zero, as indicated by the red horizontal stripe. Here, the center square of the figure is not entirely red, and further inspection shows that roll velocities near zero can be recovered more accurately if imposed alongside initial roll angles and pitch angles near zero (which are respectively defined by the center of the largest horizontal axis and second-largest vertical axis.) This makes sense: the fly's mechanical apparatus is best at encoding roll when it is flying straight and level, the condition in which it is most often found. The largest horizontal axis, (βr), indicates that a large absolute initial roll angle can cause significant error in roll velocity detection; this can be seen by noting the warmer colors of the left- and right-most columns of the figure. Additionally, all of the variables on the smallest axes – yaw velocity, roll acceleration, pitch acceleration, and pitch velocity -- do not contribute significantly to error in recovering roll velocity. As was shown by equation (13) and the sample phase delay, Φ_r , the angular acceleration force from the yaw component of an angular acceleration remains when recovering roll angular velocity. Looking at the grid defined by the set of second largest horizontal and vertical axes, it is clear the error resulting from the magnitude of the yaw component of the imposed acceleration and the initial pitch angle is derived from a non-linear relationship between the two terms, and that the polarity of this relationship is defined by the sign of the initial roll angle (the position to the left or the right of the largest vertical

axis). Figure 2-5a makes clear that the lowest relative errors (indicated by the cooler colors) occur for large absolute roll angular velocities and low absolute initial roll angles.

The relative error results for recovering pitch are displayed in Figure 2-5b. The red horizontal stripe in the center shows the conditions in which pitch detection is worst (greater than 36.94% error). This shows that the fly makes the worst assessment of its angular velocity in the pitch direction when the angular velocity in that direction is near zero, regardless of whatever other initial conditions may be present.

Continuing to analyze Figure 2-5b shows that as for large absolute values of pitch angle (βp), error in pitch rotation detection increases, due to the increased lateral gravitational force in these conditions. This can be deduced from the figure by noting the warmer colors in the left- and right-most columns of Figure 2-5b. Some of the boxes that checker the figure have a change in color that depends on the second largest horizontal and vertical axis, corresponding to initial roll angle, βr, and initial yaw velocity respectively. The error appears to depend on whether the values of initial roll angles and yaw velocities fall along a line of a certain slope. The magnitude of the slope of this line appears to increase with the magnitude of the initial pitch angle, and the slope switches from positive to negative as the initial pitch angle switches from negative to positive. One example of a testable prediction that can be gleaned from such a detailed investigation of this figure is that for negative initial pitch angles, the error is generally higher when initial roll angles and yaw velocities are positively correlated. As the magnitude of the initial pitch angle increases, the recovery error is generally lower as the magnitude of the initial roll angle increases.

Finally, Figure 2-5c, similar to 2-5a and 2-b, with the red horizontal stripe running across the middle, reveals that recovering yaw velocity is worst at initial yaw velocities near zero, regardless of the other conditions. The trend of warmer colors in the left-and right-most columns reveals that yaw recovery is more error prone as the magnitude of the initial roll angle increases. There is also a large dependence on the values of the second largest vertical axes, the initial imposed roll velocity. The polarity of the relationship of initial roll velocity to error is defined by the magnitude of the initial roll angle. The areas of lowest error in yaw recovery are low magnitude initial roll angles and high magnitude initial yaw velocities.

Predictions

There are a number of testable predictions resulting from our algorithm and error analysis. For one, the fly should also have different response latencies (on the order of milliseconds) to applied angular velocities depending on the angular position of the haltere when the velocity is imposed. Additionally, flies with increasing pitch angle, but lower pitch angular velocities, should significantly over-estimate pitch velocity. The same is true for roll, with increasing body roll angle, but lower roll angular velocities. With increasing body roll angle and low yaw angular velocities, yaw angular velocity will be significantly under-estimated by the fly. If the haltere input is not down-weighted at low angular velocities, these could lead to excessive or insufficient corrective rotations, because a larger lateral gravitational force component will partially obscure an accurate force measurement. A similar phenomenon is observed in the desert ant *Cataglyphis Fortis*, whose error in their ability to calculate their home-direction vector increases with

experimentally-imposed large-angle turns at their feeding site (Muller and Wehner, 1988).

One general trend in the error graphs is the poor performance of this model at angular velocities below four radians/sec. Current experimental data suggests that at low angular velocities, mechanosensory input is relied upon less than with visual input. At low angular velocities (around 200 degrees/sec, or ~3.5 radians/sec), providing mechanosensory stimulation alone triggers a minimal corrective response in the fly, while providing visual input alone yields an appropriate corrective response (Sherman and Dickinson, 2004). This trend is reversed at larger-magnitude angular velocities, for example above 500 degrees/sec or 8.7 radians/sec in the pitch direction. Sherman and Dickinson have experimentally determined that the haltere's mechanosensory apparatus only responds to angular velocities within a certain range. This analysis offers a potential reason for their observation. At low angular velocities of the fly's body, other lateral forces on the haltere overwhelm Coriolis forces, preventing accurate angular velocity measurements. Though our model predicts why haltere output may not be accurate for low angular velocities, it does not treat how the gain of this input is differentially reduced in favor of visual inputs in such conditions. The mechanism could incorporate two non-mutually exclusive means: (1) the biomechanical properties of the campaniform sensilla, which could set an appropriately high threshold for lateral strain detection, or (2) top-down neural architecture which combines visual and haltere input.

Discussion

We have proposed a neurobiologically plausible model of how the fly can have accurate readings of pitch, roll, and yaw angular velocities with only information about the total lateral forces on the left and right halteres. Following the work of Dickinson explaining how the fly can integrate different types of signals to generate a flight correction pattern, we now seek to identify and delineate the elements inside their "black box" that are used to compute the angular velocity. Specifically, our model assumes a simple scenario that can be implemented in agreement with the already known nervous system organization of the fly. One assumption is that a neuron fires in a phase locked pattern with respect to the haltere's angular position and that this signal biases the intermediate cells performing the measurement. These cells would then directly drive the motor neurons of the wing steering muscle.

Our model does not assume labeled lines for pitch, roll, and yaw at the level of the campaniform sensilla. Indeed, we feel that this is a gross over-simplification of the mechanism of the haltere's action. The directional components of the forces generated by the different rotations overlap significantly. There is no direction in which any physical sensor could be oriented to set up a line only sensitive to a particular rotational direction at this level of the architecture. As our circuit shows, these forces must be discriminated downstream of the campaniform sensilla. It is important to note that while we chose a three-dimensional orthogonal system to complete the analysis, it is possible the fly's axes are not orthogonal and thus somewhat redundant. In fact, previous work has shown that flies are least sensitive to yaw rotations, perhaps indicating that the reference system used by the flies encodes variability in yaw the least (Sherman and Dickinson, 2003). We

could have chosen any coordinate system -- for example, spherical or cylindrical -- to complete this physical analysis.

As we show using error analysis, the model appears robust. Moreover, our model not only fits a majority of the data about fly flight mechanosensation, but also points to future experiments that may further elucidate this highly complex information processing system. The general computational framework proposed in this paper has broader implications for any system that relies on accurate estimates of a complex time-varying signal. For example, the asymmetric crossed inhibition resembles the crossed interaction of the "correlation model" proposed for fly motion detection (Egelhaaf et al, 1989). In fact, one could abstract the actual input and appreciate the similarities between the computation of an angular velocity signal and the directional selectivity signals in the two systems.

Relation to past work

As mentioned in the introduction, Thompson and colleagues recently proposed a model that attempted to tackle this problem (Thompson 2008). This work represents a very salient contribution to the problem and nicely complements our approach, where we further elaborate on its potential solution. In their algorithm, the halteres could obtain accurate measurements of the pitch, roll, and yaw components of an imposed, constant angular velocity. In modeling haltere motion in response to an imposed angular velocity, they lifted the assumption of infinite lateral rigidity, and modeled the lateral motion of the haltere as a point mass "on a rigid massless structure with a torsional spring and damper at the base" (Thompson 2008). Though they model the out-of-plane

displacement of the haltere with a system of nonlinear (with respect to imposed angular velocity and acceleration) differential equations, it is assumed that the halteres' out-of-plane displacements are linearly related to strain generated in the campaniform sensilla field. We did not model strain directly, but if it is linearly related to the displacement, our model is sufficient. Thompson et. al. claim that accurate measurements of the pitch, roll, and yaw components of the angular velocity of the fly's body can be obtained by haltere position-specific average strain and average strain-rate measurements.

The authors do not discuss exactly how many haltere oscillations it takes for the out of plane displacement to reach a steady state. Out of 40 oscillations of the haltere, only the last 20 are used to produce the out-of-plane displacement versus haltere position plots. This suggests that the out-of-plane displacement of the halteres is not the same for the first 20 oscillations, and the motion has not reached steady state. Since their proposed method of velocity decoupling relies on averaged out-of-plane motion (strain) measurements, if the phase plot does not converge on a single loop within a few haltere oscillations, the fly will recover erroneous readings of angular velocity. How is the haltere capable of delaying strain measurement until the phase-plane trajectory has stabilized, meaning the out-of-plane motion has reached a repeating pattern with respect to in-plane oscillation? What would prevent the haltere from taking a measurement at the appropriate point in the stroke as soon as any out-of-plane motion is experienced?

In addition, it is difficult to determine a neurobiologically feasible way to measure average strain and average strain rate between two points in a stroke cycle. Taking an average displacement value is potentially feasible, if one measurement were stored until the next measurement were taken, or with an afferent campaniform sensillae acting as a

low-pass filter, as they suggest, but encoding an instantaneous derivative of displacement is more difficult to imagine in the physiology of the campaniform sensillae. They list several methods which are common in engineering but have no known neural correlates. Interestingly, they mention the "feasibility of encoding and reconstructing the full body rate vector using only discreet (sic) compressive strain magnitude measurements to describe the symmetric and asymmetric aspects of the haltere trajectory." This is indeed similar to our proposed method.

Finally, Thompson et al. test their model with constant angular velocities, but including only two different roll velocity components, zero and five radians per second. They also exclude gravitational acceleration in their equation for out-of-plane displacement. In contrast, we show that the gravitational force has a significant effect on the accuracy of angular velocity recovery, and test 10^8 different parameters.

Future Directions

Our model could be tested by further anatomical and functional measurements of neurons in the haltere circuitry, to determine if there are units that behave as we predict. While the exact neuronal components and circuits are still mysterious, paradoxically, there is an electrical synapse between certain haltere afferents and the b1 steering muscle (Fayyazuddin and Dickinson, 1996; Trimarchi and Murphey, 1997). However, it is difficult to fathom an algorithm by which the haltere afferents could be electrically coupled via these gap junctions directly to the steering muscles and produce accurate reflexes to imposed angular velocities. But both of the above studies also found a slower, chemical, poly-synaptic pathway to b1 from the haltere afferents. It is not clear which

single pathway, if not both, campaniform sensilla use to convey information. The gap junctions may simply function to set the phase of the wingbeat relative to the halteres. To alter the flight trajectory, a wingbeat phase advance or recession could come through the chemical, poly-synaptic pathway. The computation that would enable a correct alteration to the flight path does not seem possible without slightly more complex circuitry than direct electrical coupling.

Recently, the output of the haltere nerve has been recorded directly. It has been shown to be directionally selective and fire with sub-millisecond jitter (Fox and Daniel, 2008) Additionally, recordings from neck motor neurons have been shown to receive phase-locked output from the halteres at physiological frequencies up to 105 Hz (Huston and Krapp, 2009) These experimental data confirm that our model, which requires precise timing and phase-locking, is plausible. One important future step will be to add random noise modeled on the known jitter of the haltere output to our analysis, which has so far only revealed the systematic errors due to the physical nature of the system.

More complex analyses could make use of two genetic mutants that affect the sensors or the neural circuit directly. The *Ultrabithorax* mutant yields anomalous campaniform sensilla structures in the halteres (Roch and Akam, 2000). The *shaking-B*^{2} mutant disrupts electrical synapses from haltere afferents (Trimarchi and Murphey, 1997). One could test, for example, whether *Ultrabithorax* or *shaking-B*^{2} mutant flies show corrective responses to mechanically imposed angular velocities. While these flies have proven unable to fly, (not surprising given their absolute lack of electrical synapses), one wonders if their haltere-mediated reflex system is at least partially intact. This would implicate the haltere/b1 chemical-synapse pathway in angular velocity

detection. The ability to discriminate between sensory deficits versus neural circuitry deficits will certainly aid further elucidation of the engineering principles at work.

One last remaining question is how the fly can distinguish externally generated rotations from internally generated rotations. If the fly initiates a rotation, it must also inhibit the reflexive correction automatically calculated by the haltere circuitry. In addition, when reflexively responding to a stimulus, the fly should not respond to angular velocity information generated by the haltere during the corrective rotation. It seems unlikely inhibition of this counter-reflexive action would be a regulatory property of the flight control system given the speed of the reflex. If a top-down "stop steering" command were given (perhaps via inhibition of mnb1), it would have to be timed with extraordinary precision. Alternatively, the counter-reflexive action could simply be an inherent property of the reflex. Perhaps the biophysical properties of the b1 muscle driver neurons are simply unresponsive to haltere input during a short (few ms) refractory period after performing a corrective maneuver.

Our model assumes a train of delta functions at the specific phase-delayed times is convolved with the summed and subtracted force signals. One improvement to our model would be to actually deconvolve the summed and subtracted force signals and find the best filters for extracting the different signals.

Finally, it is known the visual system can impact the motion of the haltere directly in order to transform visual-mediated corrections into rotation corrections. It may also be possible that a horizontal indicator signal is sent from the visual system to the haltere in order to allow for a correction of gravity. Our scheme allows for both modifications -- inhibition of self-initiated movements and gravity correction -- by simply allowing for

another incoming signal to subtract from the force estimates before that are fed back to the wing.

References

Ballard, D., Jehee, J (2008). On the Coding of Negative Quantities in Cortical Circuits. Available from Nature Precedings.

Briggman, K.L., and Kristan, W.B. (2008). Multifunctional pattern-generating circuits. Annual review of neuroscience *31*, 271-294.

Dickinson, M.H. (1999). Haltere-mediated equilibrium reflexes of the fruit fly, Drosophila melanogaster. Philosophical transactions of the Royal Society of London *354*, 903-916.

Dickson, W., Straw, A., Dickinson, M. (2008). Integrative Model of Drosophila Flight AIAA Journal *46*, 2150-2164.

Egelhaaf M, Borst A, Reichardt W. (1989) Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system. J Opt Soc Am A. *6*, 1070-87.

Fayyazuddin, A., and Dickinson, M.H. (1996). Haltere afferents provide direct, electrotonic input to a steering motor neuron in the blowfly, Calliphora. J Neurosci *16*, 5225-5232.

Fox, J.L., and Daniel, D.L. (2008). A neural basis for gyroscopic force measurement in the halteres of *Holorusia*. J Comp Physiol A *194*, 887-897.

Hengstenberg, R. (1988). Mechanosensory control of compensatory head roll during flight in the blowfly Calliphora erythrocephala. J. Comp. Phys. A. *163*, 151-165.

Huston, S. J., and Krapp, H. G. (2009). Nonlinear integration of visual and haltere inputs in the fly neck motor neurons. J Neurosci *29*, 13097-13105.

Land, M.F., and Collett, T. S. (1974). Chasing behaviour of houseflies (Fannia canicularis). Journal of Comparative Physiology *89*, 331-357.

Marr, D. (1982). Vision : a computational investigation into the human representation and processing of visual information (San Francisco: W.H. Freeman).

Muller, M., and Wehner, R. (1988). Path integration in desert ants, Cataglyphis fortis. Proceedings of the National Academy of Sciences of the United States of America *85*, 5287-5290.

Nalbach, G. (1993). The halteres of the blowfly Calliphora. J Comp Physiol A *173*, 293-300.

Peng W. (2005). Clutter-Based Dimension Reordering in Multi-Dimensional Data Visualization (Master's thesis). Worcester MA: Worcester Polytechnic Institute.

Peng W, Ward MO, and Rundensteiner EA. (2004). Clutter reduction in multidimensional data visualization using dimensional reordering. Proceedings of the IEEE Symposium on Information Visualization 2004, edited by Keahey A, Austin TX, p. 89–96.

Pringle, J.W.S. (1948). The Gyroscopic Mechanism of the Halteres of Diptera. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences *233*, 347-384.

Roch, F., and Akam, M. (2000). Ultrabithorax and the control of cell morphology in Drosophila halteres. Development (Cambridge, England) *127*, 97-107.

Sherman, A. (2003). The integration of visual and haltere feedback in Drosophila flight control. Ph.D. dissertation, University of California, Berkeley with the University of California, San Francisco.

Sherman, A., and Dickinson, M.H. (2003). A comparison of visual and haltere-mediated equilibrium reflexes in the fruit fly Drosophila melanogaster. The Journal of experimental biology *206*, 295-302.

Sherman, A., and Dickinson, M.H. (2004). Summation of visual and mechanosensory feedback in Drosophila flight control. J Exp Biol, 207(Pt 1):133-42.

Strausfeld, N.J., and Seyan, H.S. (1985). Convergence of visual, haltere, and prosternal inputs at neck motor neurons of Calliphora erythrocephala. Cell Tissue Res *240*, 601-615.

Thompson, R.A., Wehling, M.F., Evers, J.H., and Dixon, W.E. (2008). Body rate decoupling using haltere mid-stroke measurements for inertial flight stabilization in Diptera. Journal of comparative physiology *195*, 99-112.

Trimarchi, J.R., and Murphey, R.K. (1997). The shaking-B2 mutation disrupts electrical synapses in a flight circuit in adult Drosophila. J Neurosci *17*, 4700-4710.

Wu, W.C., Wood, R.J., and Fearing, R.S. (2002). Halteres for the Micromechanical Flying Insect. IEEE Int. Conf. on Robotics and Automation, Washington, DC May 11-15, 2002.

CHAPTER 2 APPENDIX

The general form for extracting angular velocities from the complex waveforms shown in Figure 2-2 are a series of amplitude measurements at different times. A form of this algorithm has been presented for detecting wing deformations (Dickinson, 1990). The measurements are taken at a constant frequency, but they have differing time lags, or phase shifts, from when the first measurement is taken. The phase shift for extracting pitch, roll, and yaw, depends not on the time of force onset, but on the angular position, Φ, of the haltere. The pitch, roll, and yaw signals are each measured at a different value of Φ, henceforth known as Φ_p , Φ_r , Φ_y respectively. The F_{pitch} signal can be measured at the Φ_p which indicates the extremum of the $F_L + F_R$ waveform. The F_{roll} can also be discriminated from yaw by setting Φ_r to correspond to a point on the $F_L - F_R$ waveform at which the F_{roll} is at a local extremum, because the yaw signal, at twice the frequency, will be zero. (Note: this does not mean that the $F_L - F_R$ waveform will necessarily be at a maximum). To distinguish yaw from roll Φ_y must correspond not to an extremum, but to a point on the $F_L - F_R$ waveform where the F_{roll} contribution is approximately equal to zero, but the F_{yaw} contribution is nonzero. The value can then be divided by an additional constant constant, *l* (included in k"), representing the inverse of the ratio of the sample point's y position on the F_{yaw} signal to the y position of the extremum of the F_{yaw} signal.

Measurement begins in a pitch, roll, or yaw sampling neuron when the angular position of the haltere reaches Φ_p , Φ_r , or Φ_y respectively (See Figure 2-2.) Given the known wing beat frequency, the resulting time delay for sampling will always be between about one and eight milliseconds. The exact phase shift will depend on the initial angular

position of the haltere when the rotation is induced. These time bounds are strict, as they allow determination of the rotation's sign.

CHAPTER 2 Figure Legends

Figure 2-1: Force diagram and experimental goal.

(A) Force diagrams and (B) plots are shown for a situation in which the fly is rotated in the pitch direction. (C) The location of the circuit which this study seeks to uncover situated within the known anatomy of the fly's flight system. Modified from (Dickson, 2008).

Figure 2-2: Analysis of variables

(A) The sample points for pitch, roll, and yaw (Φ_p , Φ_r and Φ_y) are plotted against time on the graph of the haltere's position (Φ) given an initial haltere angle of zero degrees. (B-C) Φ_p, Φ_r , and Φ_y are plotted against time on the graphs of the equations for F_L - F_R and F_L + F_R given a constant applied angular velocity with pitch, roll, and yaw components of 30 radians/sec^2 each.

Figure 2-3: Neural circuit for angular velocity computation.

Squares - anatomical structures; circles - neurons; smaller open circles - excitatory connections; smaller filled circle - inhibitory connection. The haltere CPG drives the left and right halteres which sense lateral forces. These force signals are summated at the neuron labeled N+ and subtracted at the neuron labeled N-. The signals from these neurons are then transmitted to three intermediate neurons: N_{pitch}, N_{roll}, and N_{yaw}. These cells are biased by the time delays tau_{pitch}, tau_{roll}, and tau_{yaw} via delay lines from the haltere's phase-locked CPG. These delays must be set by the angular position of the halteres such that the three intermediate neurons take measurements at the appropriate

times. The results of this computation are fed to each wing via the mnb1 cells which control the b1 steering muscles.

Figure 2-4: Visualization of the dimensional stacking technique.

A standard rectangular coordinate system has two axes defining a plane on which a point is uniquely specified by two values equal to the distance of the point from each axis. A three dimensional coordinate system has three axes defining a cube but can only be drawn on a two dimensional surface by making use of perspective. An eight dimensional coordinate system is impossible to visualize in three dimensions but would function the same as any other coordinate system by specifying eight unique values for a given point in space. Here, in order to visualize such a coordinate system in two dimensions, we have "collapsed" the dimensions beyond the first two inside the other dimensions in a recursive fashion. This so-called "dimensional stacking" technique allows one to visualize trends in a high-dimensional space. The dimensions are collapsed in a specific order such that similar output values from the given multi-dimensional inputs are clustered together. In the example shown here, the red square defines a coordinate system over which the first through fourth dimensions vary for given values of dimensions five through eight. (Specifically, those values are the last value of dimension eight, the sixth value of dimension seven, the fourth value of dimension six, and the fifth value of dimension five.)

Figure 2-5: Error analysis

(A-C) Dimensional stacks and histograms of the relative error in pitch, roll, and yaw recovery. 10 values of each parameter were used: initial velocities (Vr, Vp, Vy) from -16.1 to 19.9 radians/sec, angular accelerations (Ar, Ap, Ay) from -50 to 62.5 radians/sec^2, initial body pitch angle (βp) and initial body roll angle (βr), from -50 to 60 degrees.

FIGURE 2-1

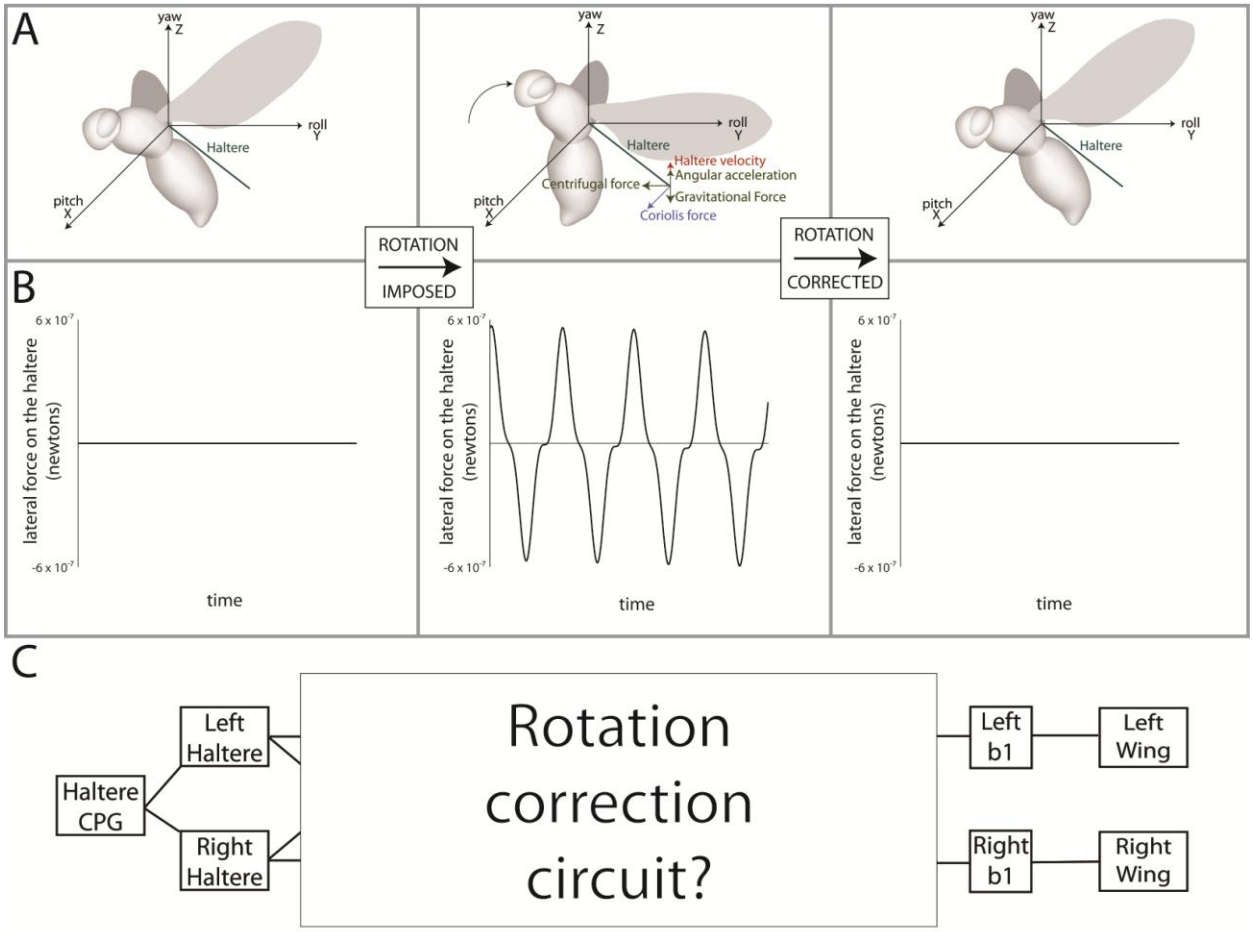
A
yaw
Z
roll
Y
pitch
X
Haltere
Haltere velocity
Angular acceleration
Centrifugal force
Gravitational Force
Coriolis force
ROTATION
IMPOSED
ROTATION
CORRECTED
B
6 x 10⁻⁷
-6 x 10⁻⁷
lateral force on the haltere
(newtons)
time
C
Left
Haltere
Haltere
CPG
Right
Haltere
Rotation
correction
circuit?
Left
b1
Left
Wing
Right
b1
Right
Wing

Figure 2-2

A

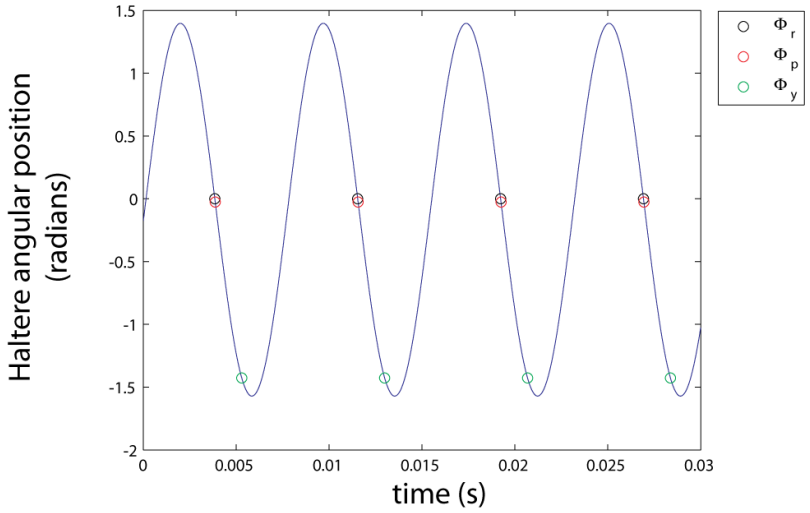

B

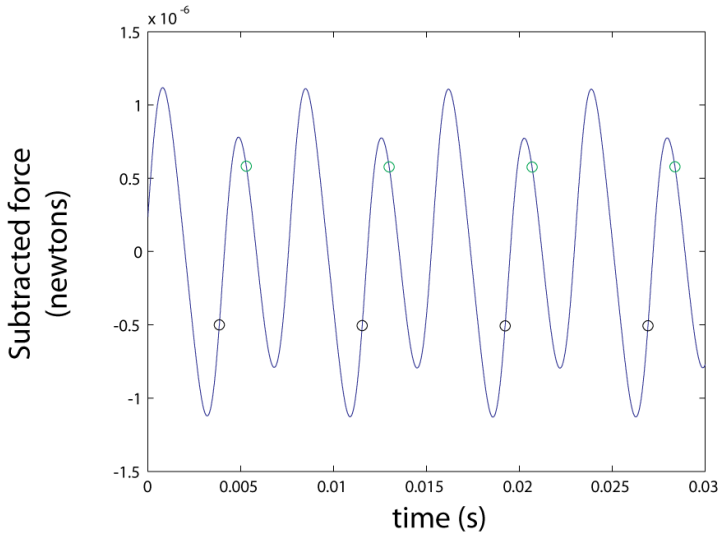

C

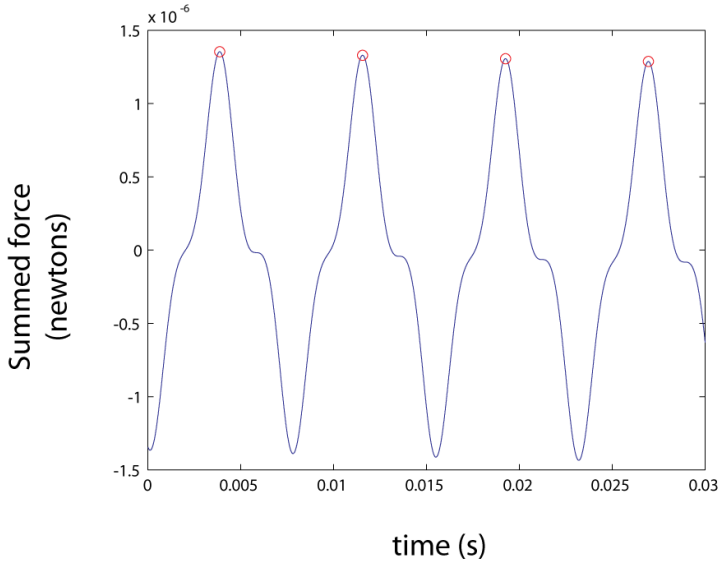

Figure 2-3

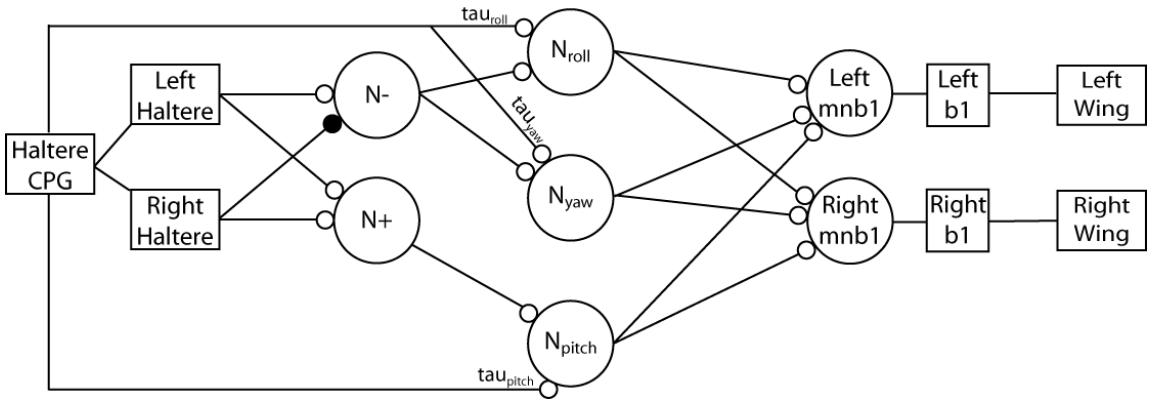
tau_roll
N_roll
Left Haltere
N-
tau_yaw
Haltere CPG
N_yaw
Right Haltere
N+
Left mnb1
Left b1
Left Wing
Right mnb1
Right b1
Right Wing
N_pitch
tau_pitch

Figure 2-4

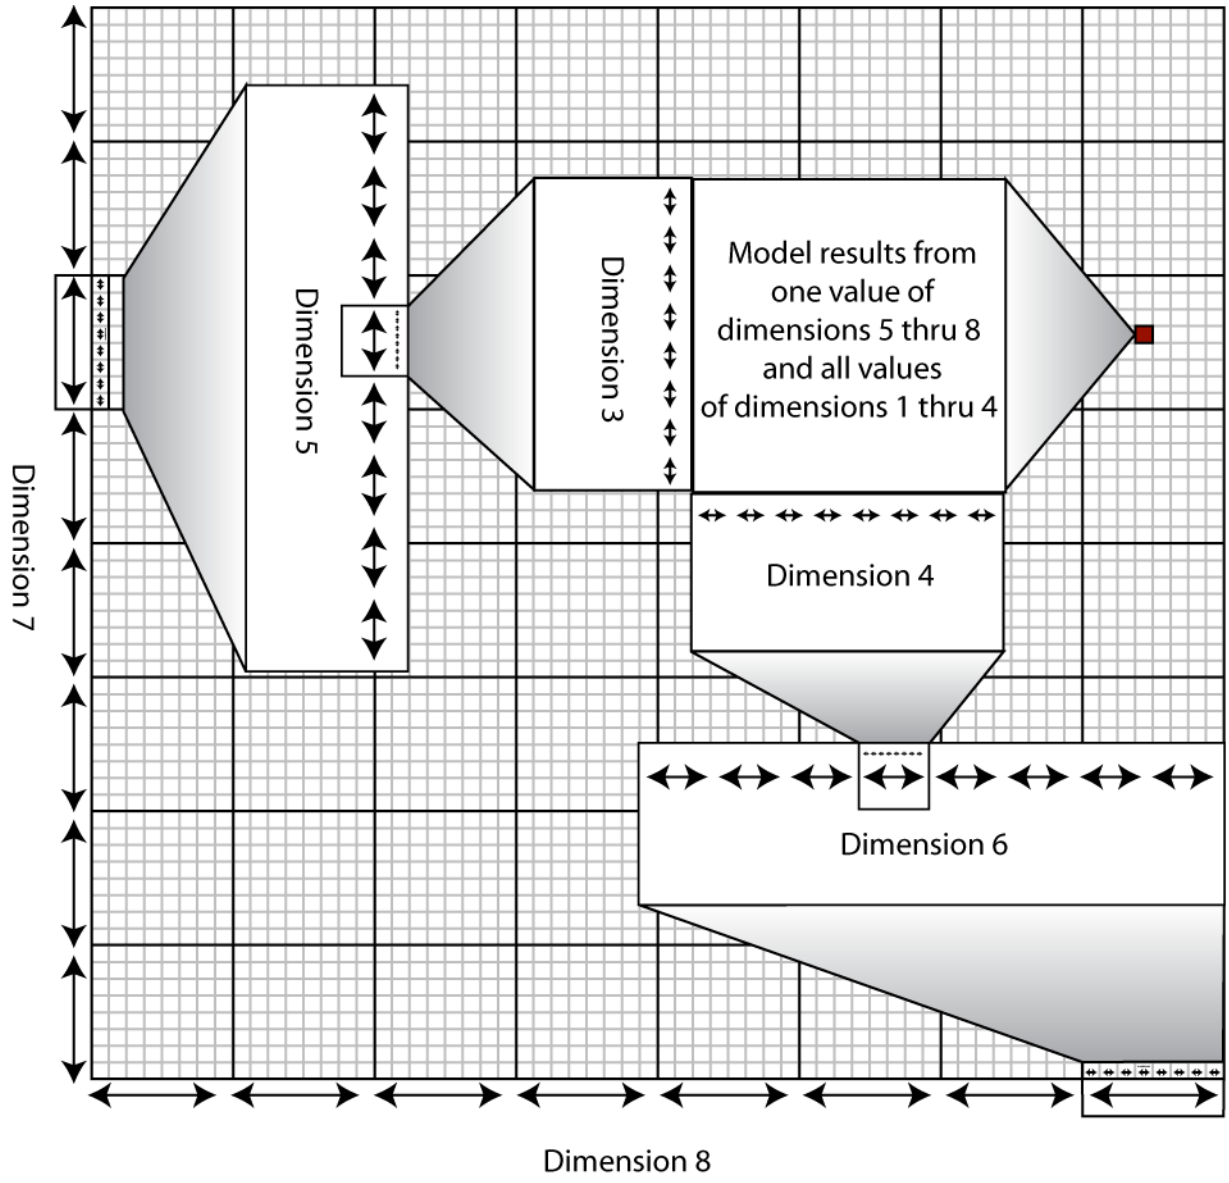
Dimension 7
Dimension 5
Dimension 3
Model results from one value of dimensions 5 thru 8 and all values of dimensions 1 thru 4
Dimension 4
Dimension 6
Dimension 8

Figure 2-5
A. ROLL

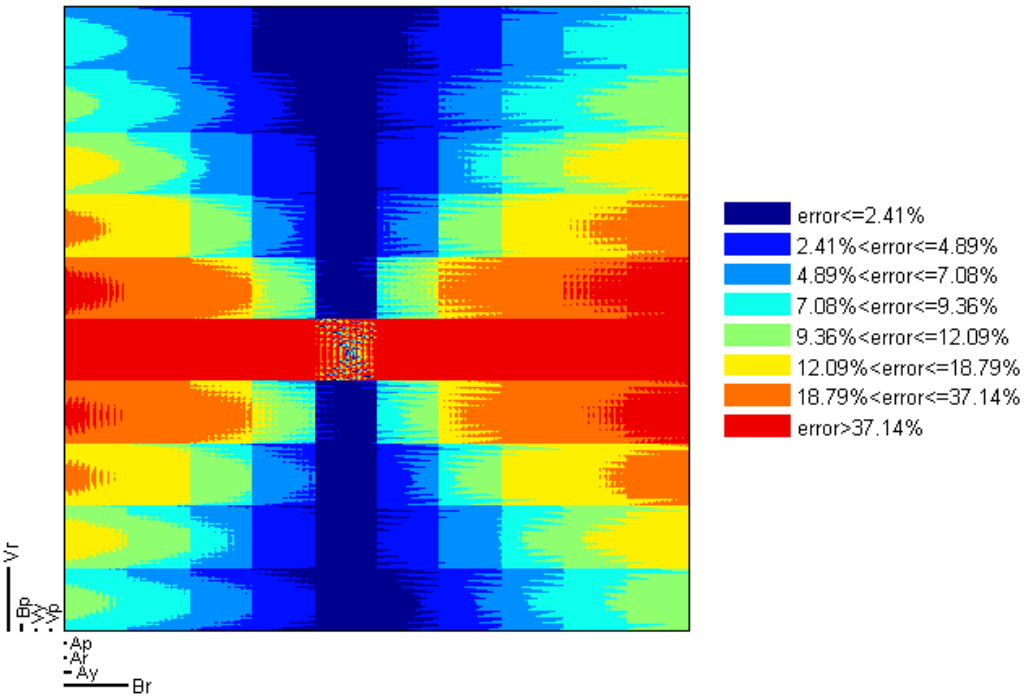

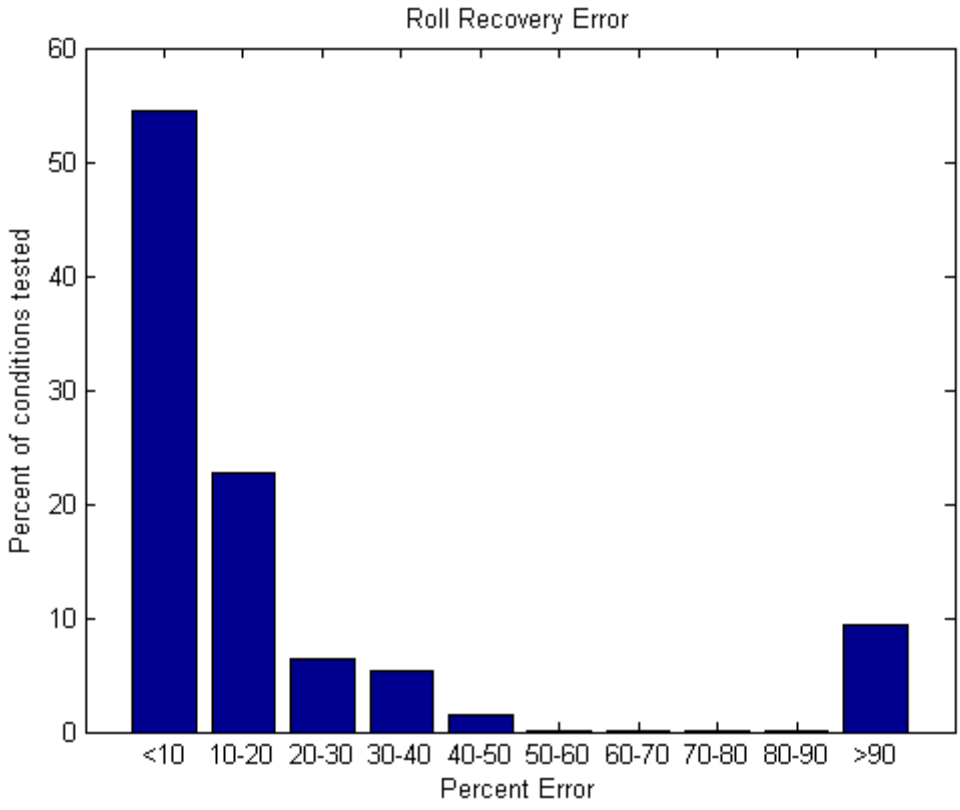

B. PITCH

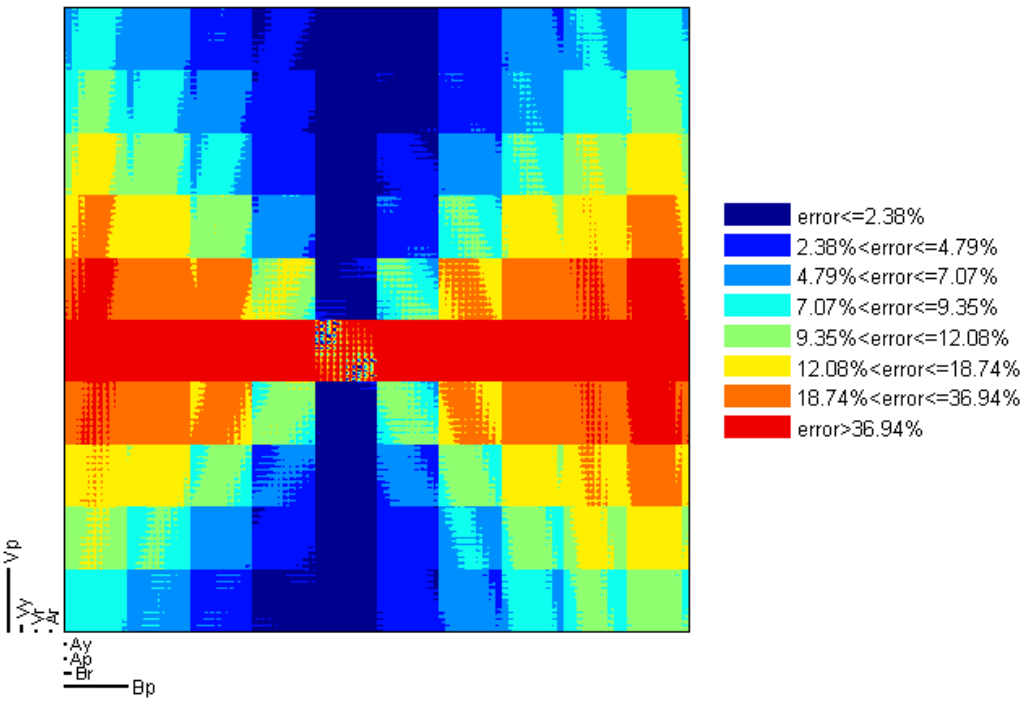
error<=2.38%
2.38%<error<=4.79%
4.79%<error<=7.07%
7.07%<error<=9.35%
9.35%<error<=12.08%
12.08%<error<=18.74%
18.74%<error<=36.94%
error>36.94%
Vp
Vy
Vr
Ar
Ay
Ap
Br
Bp

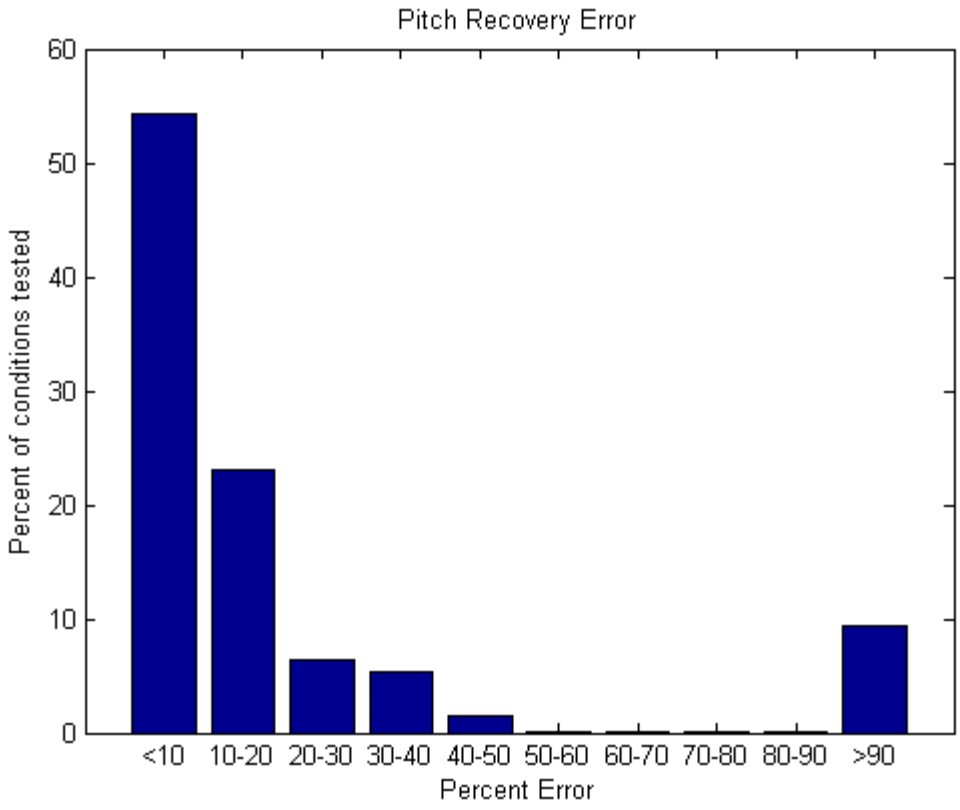
Pitch Recovery Error
Percent of conditions tested
0
10
20
30
40
50
60
<10
10-20
20-30
30-40
40-50
50-60
60-70
70-80
80-90
>90
Percent Error

C. YAW

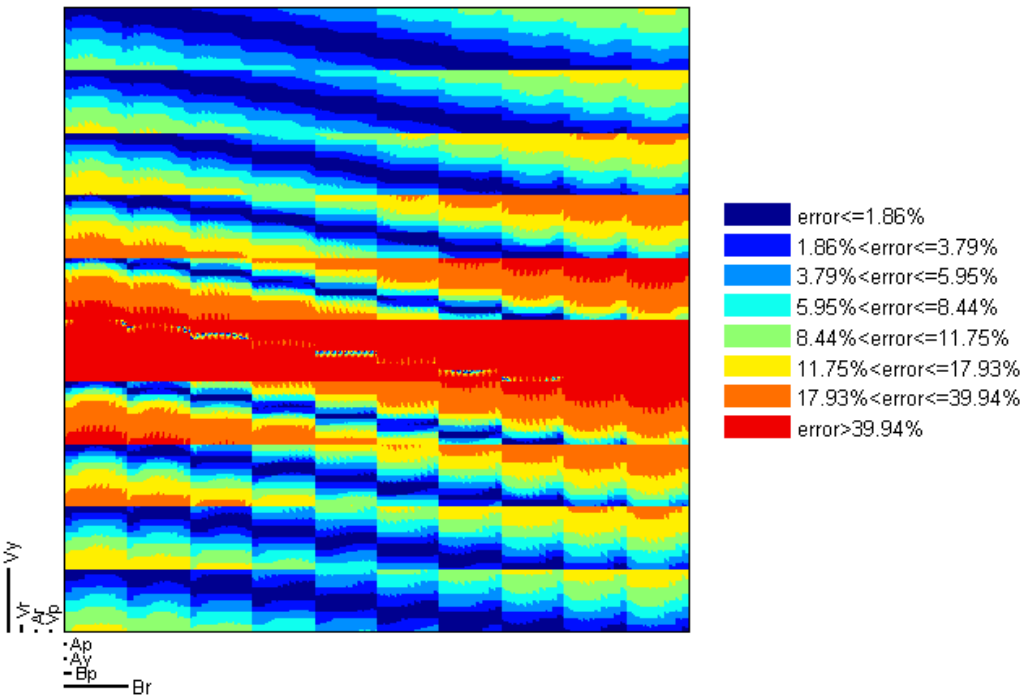
error<=1.86%
1.86%<error<=3.79%
3.79%<error<=5.95%
5.95%<error<=8.44%
8.44%<error<=11.75%
11.75%<error<=17.93%
17.93%<error<=39.94%
error>39.94%
Vy
Vr
Ar
Vp
Ap
Ay
Bp
Br

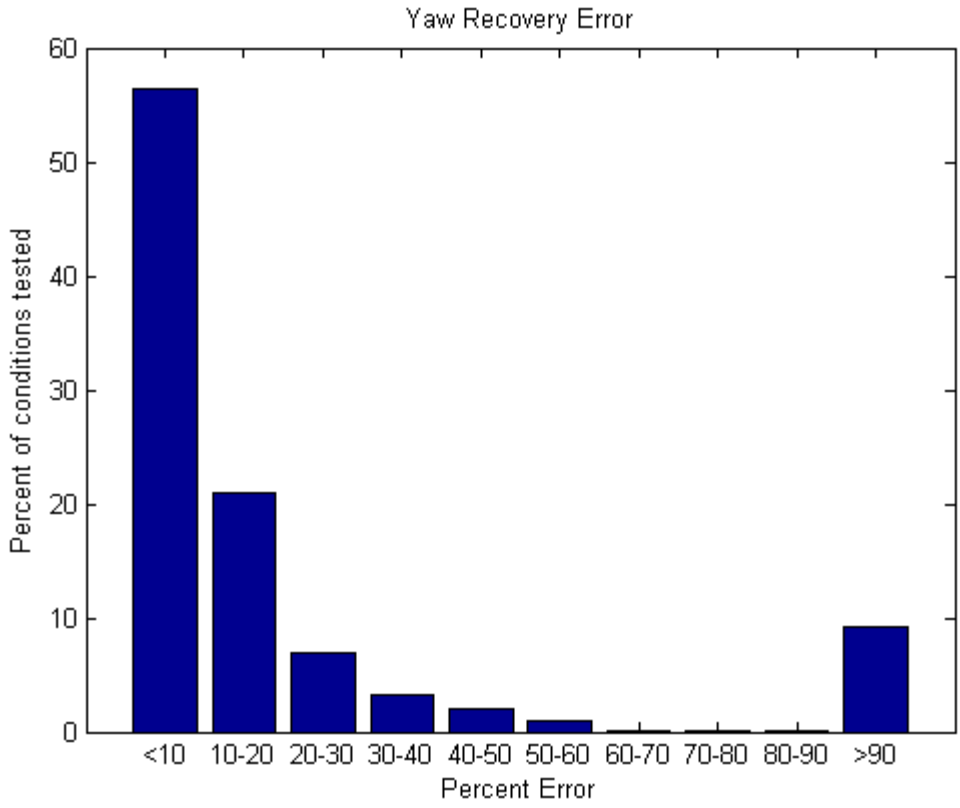
Yaw Recovery Error
Percent of conditions tested
0
10
20
30
40
50
60
<10
10-20
20-30
30-40
40-50
50-60
60-70
70-80
80-90
>90
Percent Error

APPENDIX A: PackIO electrophysiology suite

What follows is a description of the electrophysiology suite of computer programs I developed during my time as a graduate student. I present here a guide to its operation because I used it to make all of the recordings in this dissertation.

While making electrophysiological recordings and trying to synchronize various pieces of equipment (cameras, lasers, galvanometers, etc.), I found no software that was up to the task. So, I began a project to develop an uber-system capable of performing any data acquisition or generation operation that could be completed with the National Instruments DAQ hardware in use in the lab at the time. I originally referred (somewhat in jest) to the project as "PackIO", a combination of my name and IO, as in input/output, since the software is supposed to be extensible in its ability to record any input or generate any output in any synchronized fashion. Unfortunately, the name stuck.

Now, PackIO is in use by a majority of the members of the Yuste laboratory, and by Jason Maclean at the University of Chicago, whose laboratory has made improvements not treated here.

The current version number of the software is 264 and it is 1.6 MB in size. A Professional installation of LabView 2009 with the associated DAQmx drivers is required.

Recorded data files can be viewed with the EphysViewer software package written in MATLAB by Brendon Watson and modified by me. It is available from the Yuste CVS repository in the PaqTools module (http://yustecvs.biology.columbia.edu/cgi-bin/cvsweb/).

When first run, PackIO will create the following directories in the directory the main program library is in (if they do not already exist):

- TempData: where temporary data files produced by the Go back feature (see Figure A1-1) are stored.
- Configurations: contains four subdirectories
    - Epoch: individual epoch files should be stored here for use in the Epoch Loader
    - Input: input configurations should be stored here for use in the Master Setup input tab (see Figure A1-2)
    - Output: output configurations should be stored here for use in the Master Setup output tab (see Figure A1-3)
    - Seal: configurations for Seal Test (see Figure A1-5) should be stored here

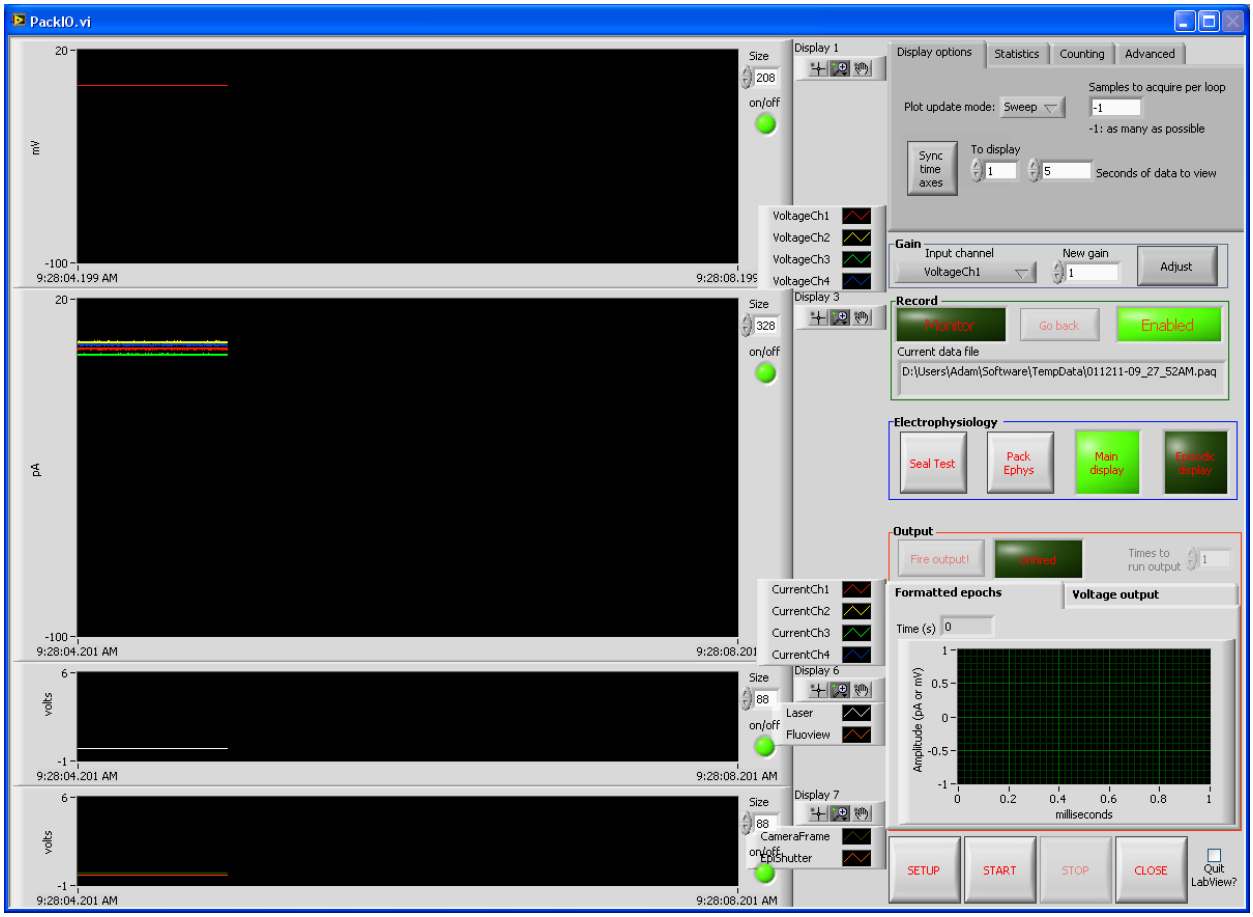

Figure A1-1: The main window of PackIO.

The main window of PackIO (Figure A1-1) is where all the major features of the program can be accessed. On the left side are the seven numbered main displays which show incoming data. The controls on the right are described below:

- Tabs on the upper right
    - "Display options"
        - "Plot update mode" controls the manner in which data is updated to the main displays at the left.
        - "Samples to acquire per loop" is the number of samples the card attempts to read from every channel of an acquisition during each

loop of the main program's operation. The default of -1 will acquire as many samples as are available.

        - "Sync time axes" locks the x-axes of the main display windows to the x-axis shown in the "To display" box.
        - "Seconds of data to view" indicates how much time to display in the main displays.
    - "Statistics" displays information regarding the condition of PackIO. These statistics are useful when diagnosing issues with slower computers.
    - "Counting" implements certain functions of the counters which most DAQ cards have. As of this writing, the count is merely displayed, and can be divided by a user-set value.
    - "Advanced" allows the user to set a value to which an output channel will be set after completion of a task. Normally, this value is zero, but there are instances, for example when controlling certain pieces of equipment such as laser controllers, in which a negative value may be desired.
- Gain controls
    - The gain of an "input channel" can be adjusted during a recording to the value in "New gain" by clicking the "Adjust" button.
- Recording controls
    - If the button on the upper left is not lit, it will read "Monitor" which indicates no data will be recorded to a file selected by the user (but data may still be recorded if the "Go back" feature is enabled, see below.) If the "Monitor" button is clicked, it will become lit and read "Record" in

which case the user will be prompted for a file name when START is clicked and recording commences.

    - The button on the far right enables the "Go back" feature. This feature records all data, even when the user does not explicitly specify to record. The user can click "Go Back" if the user decides he or she actually wants the data just displayed. The file name of where the data has been stored is then displayed in a dialog box.

- Electrophysiology controls
    - "Seal test" opens the seal test window, see Figure A1-5.
    - "Pack Ephys" opens the PackEphys software module, see Figure A1-6.
    - "Main display" controls whether the main displays are plotted or not. Turning off the main displays can be useful during episodic recordings, especially when using a slow computer, as it decreases the computational load of the program.
    - "Episodic display" opens the episodic window, see Figure A1-7.
- Output controls
    - "Fire output" commences the generation of the output sequence selected in Master Setup (see Figure A1-3) when that sequence's generation time is set by the user, not a trigger or acquisition event.
    - "Unfired" is lit and reads "Fired" when an output sequence was successfully generated.
    - "Times to run output" is enabled when the Master Setup Experimental type is set to "Acquire immediately and output specified number of times"

(See Figure A1-2). Then the output is generated the number of times specified. This option is useful when the same output needs to be run repeatedly, but note that this option is not hardware-timed, so the delay between repetitions is variable. (If precise timing is required, the output desired should be completely specified, however many times it is repeated, in the Master Setup output tab.)

  - The chart shows which outputs are currently set to be generated. The "Formatted epochs" tab shows the outputs in the scale set in the Master Setup output tab whereas the "Voltage output" tab shows the outputs as the raw voltage that will be produced by the DAQ card.

- Setup/Start/Stop/Close/Quit Labview controls
  - SETUP opens the Master Setup window (see Figures A1-2 through A1-4)
  - START begins the input/output operation, which disables accessibility of certain features in the main window until STOP is pressed.
  - STOP ceases the current input/output operation.
  - CLOSE closes the PackIO program
  - If "Quit Labview?" is checked, pressing the CLOSE button will also quit Labview completely.

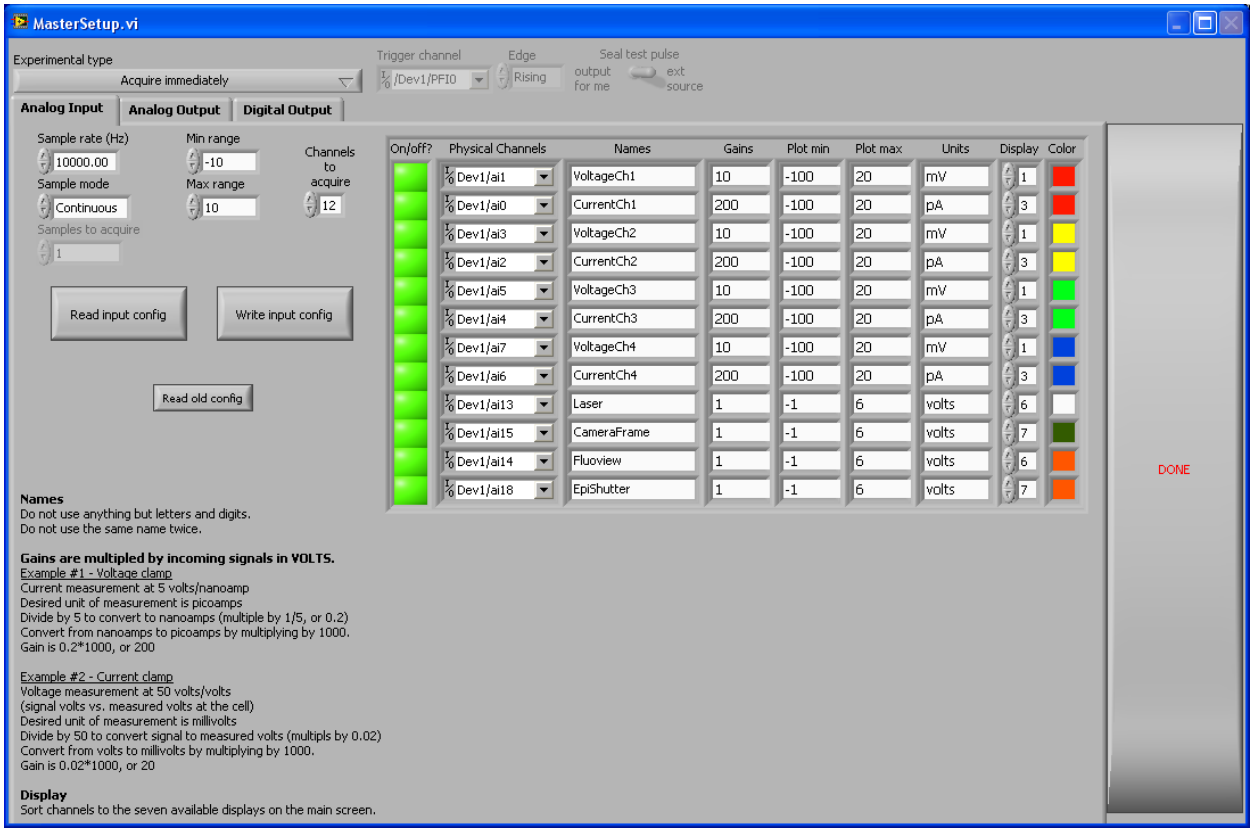

Figure A1-2: The input tab of the Master Setup window.

The Master Setup window controls which channels are recorded, what outputs are generated, and how digital lines are set when the START and STOP buttons are pressed in the main window. It also controls how and which channels are displayed on each of the seven main displays. This window automatically opens when PackIO is initially run. While this window is open, no other Labview windows can be accessed.

At the top of the Master Setup window is a critical dropdown box labeled "Experimental type" which controls what type of acquisition, generation, or both is performed as described below:

- "Acquire immediately": Acquire the specified channels when START is pressed.

- "Acquire upon external trigger": Acquire the specified channels after START is pressed and when the trigger condition, as specified by the controls to the right, are met.
- "Start acquisition and output immediately": Acquire the specified channels and generate the specified outputs when START is pressed.
- "Synchronize acquisition and output on external trigger": Acquire the specified channels and generate the specified outputs after START is pressed and when the trigger condition, as specified by the controls to the right, are met.
- "Acquire immediately and output upon user input": Acquire the specified channels when START is pressed. Generate the specified outputs when the user presses the "Fire output" button in the main window.
- "Acquire immediately and output upon external trigger": Acquire the specified channels when START is pressed. Generate the specified outputs when the trigger condition, as specified by the controls to the right, are met.
- "Acquire immediately and output specified number of times": Acquire the specified channels and generate the specified outputs the number of times specified in the main window output section when START is pressed.

Note that acquisition and generation operations which are synchronized between two separate cards, for example between a PCI-6052E and a PCI-6733, require that a RTSI cable be physically installed to connect the two cards inside the computer and appropriately indicated in the NI-MAX configuration software.

The "Analog Input" tab of the Master Setup window controls how many channels are acquired at which frame rate and within what range (upper left.) Those channels can be turned on and off quickly by clicking the green buttons in the column by each channel. The options in this window are self explanatory, except for the gains, which are multiplied by the raw voltage recorded by the DAQ card. Further details and examples are given in the lower left corner of this window. Configurations (left) can be saved and loaded at a later date so that these options need not be specified each time a similar recording configuration is used.

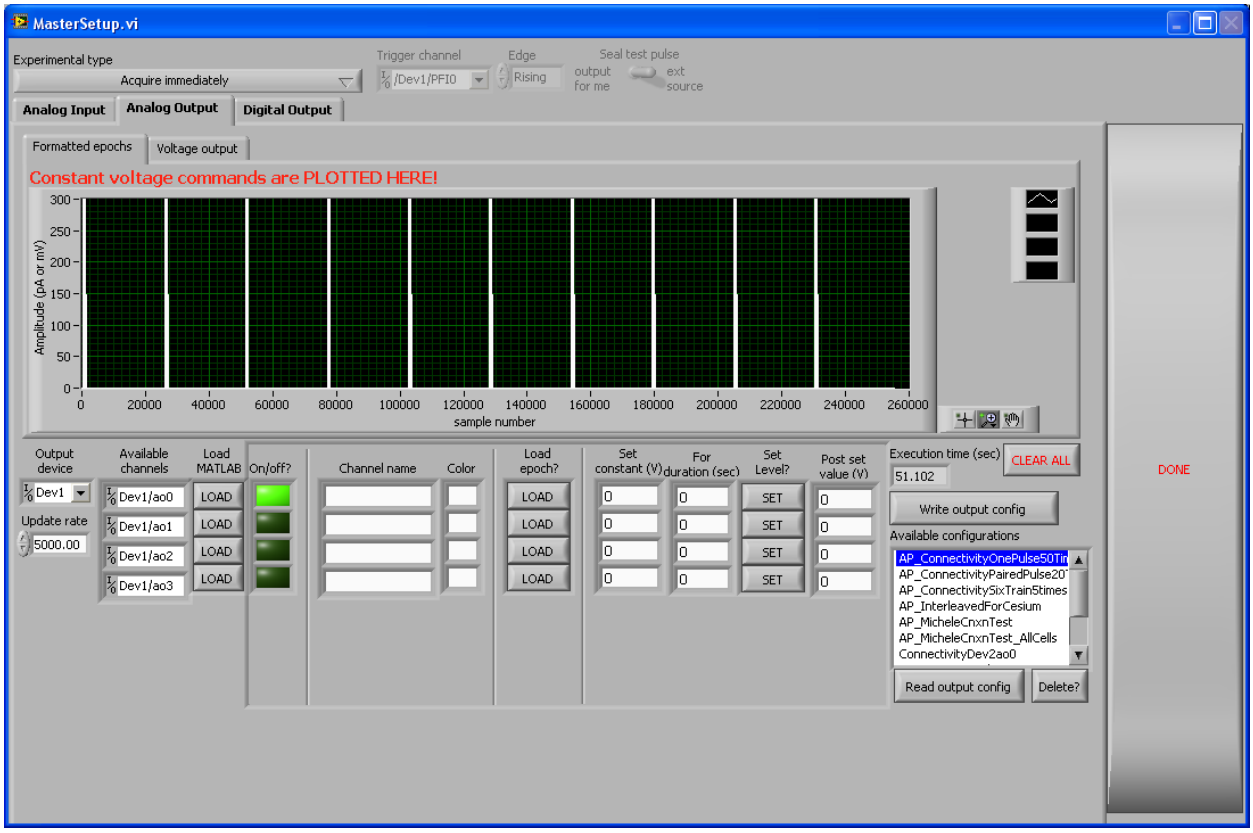

Figure A1-3: The output tab of the Master Setup window.

The display in the top of the output tab of the Master Setup window shows which outputs will be generated. It has two tabs which show the output either in the scale ("Formatted epochs") set by the Epoch Loader (described below), if used, or as raw voltage outputs ("Voltage output").

Selecting an "Output device" on the lower left will automatically update the available channels. Each channel's output can be specified by a MATLAB file, an epoch, or a set voltage:

- MATLAB: Use the following commands in MATLAB to write the vector TestRamp to a file TestRamp.dat which can then be loaded into PackIO on a given channel by clicking the LOAD button in the Load MATLAB column: fid=fopen('TestRamp.dat','w','l');fwrite(fid,TestRamp,'double');fclose(fid);

- Epoch: Clicking LOAD in the "Load epoch?" column will open a program written by Volodymyr Nikolenko, based on pClamp from Axon Instruments, to set various types of stimulation protocols such as steps, ramps, and trains. Gains can be set here so that the user does not have to keep track of voltage to amplitude transformations. Multiple sweeps of a given epoch can also be produced. Individual epochs can be saved for later use.
- Set voltage: A constant voltage can be set for a specified amount of time after which that channel will be set to a given "post set value" after the output is generated.

Configurations can be saved and loaded so that the user does not have to set the appropriate output for each channel every time the same output paradigm is desired.

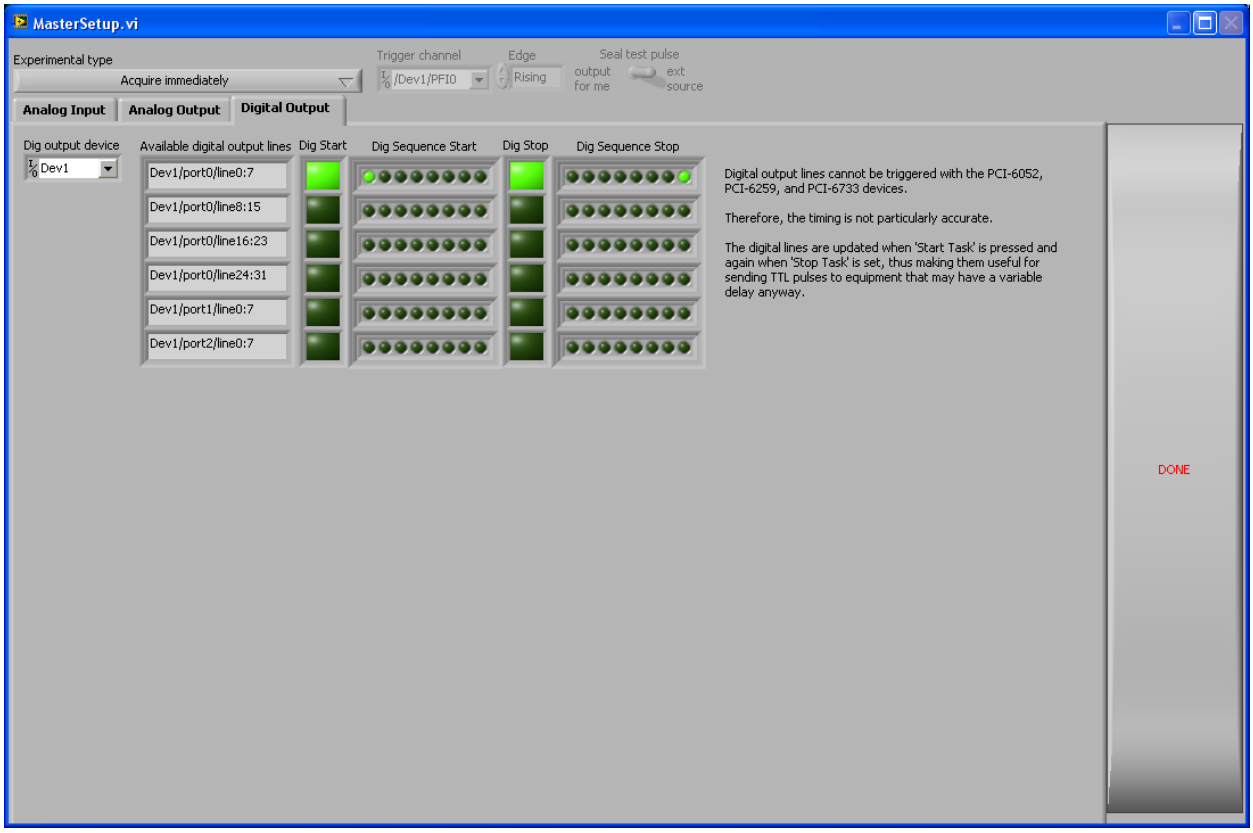

Figure A1-4: The digital output tab of the Master Setup window.

Digital output lines can be updated when the START and STOP buttons are pressed in the main window of PackIO. These are useful for sending TTL pulses to various pieces of equipment which need to be toggled before or after an acquisition or generation operation. Note that these lines are not hardware timed, so the time delays before the lines are updated are variable.

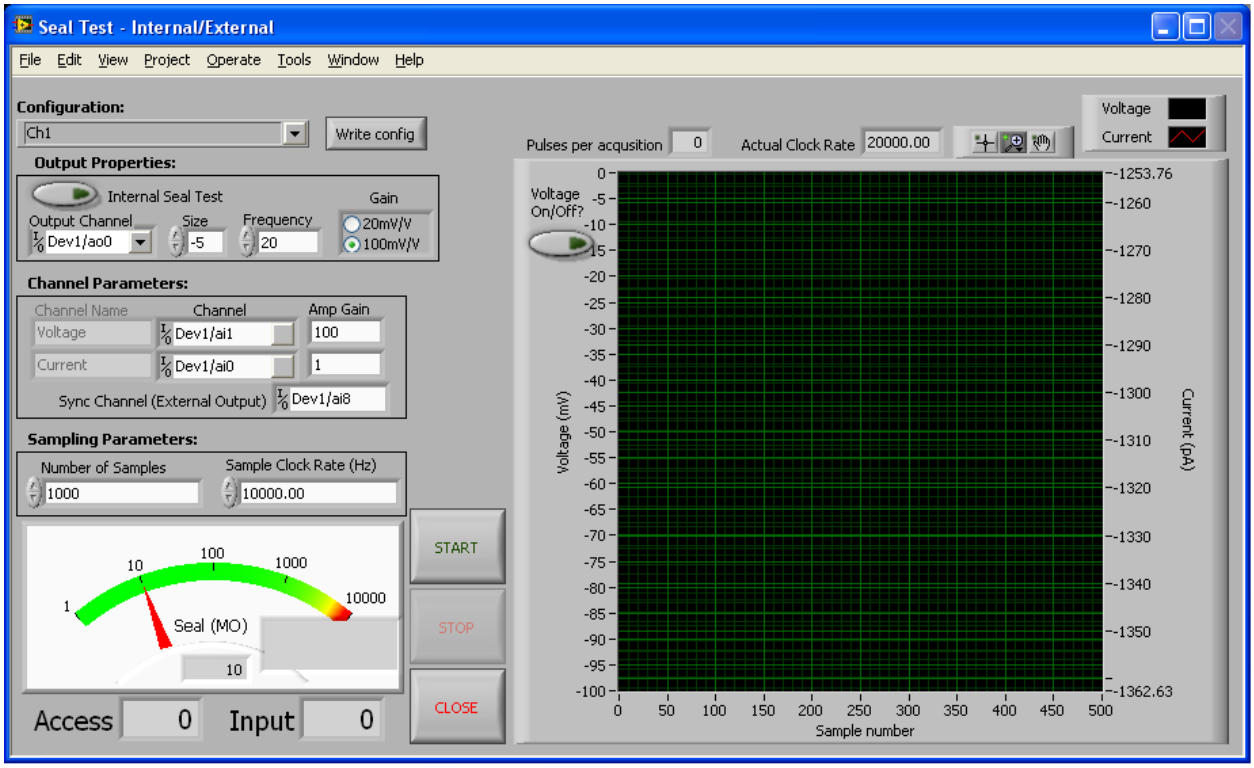

Figure A1-5: The seal test window of PackIO.

The seal test window allows the user to test the access and input resistance of a recorded cell by calculating the response current to square pulse voltage commands. It can be operated in two modes: "Internal" or "External". Internal mode, selected by clicking the "Internal Seal Test" button near the top left of the window, produces pulses of the specified parameters to the selected "Output Channel". This should be connected to the command of the electrophysiological amplifier to drive the cell. If internal mode is not selected, external mode (so called because the seal test pulses are provided by the amplifier) is used. The sync channel from the electrophysiological amplifier should be connected to the selected "Sync Channel (external output)".

The voltage and current channels for the recorded cell need to be appropriately set with the correct gains.

The number of samples acquired per run loop of the seal test subroutine can be set, as well as the sample clock rate at which to acquire samples. The pulses per acquisition and actual clock rate are then displayed above the graph on the right.

The seal resistance is displayed by the needle and the large grey box directly to the left of the START and STOP buttons (which commence and cease acquisition and, if in internal mode, generation of the seal test pulses). The displayed access and input resistances are correct after the appropriate recording configuration is obtained.

The graph on the right shows the response to the seal test pulses. The voltage pulses can also be displayed by toggling the "Voltage On/Off?" button on the upper left of the graph.

Configurations can be saved for later use and loaded by selection from the dropdown box.

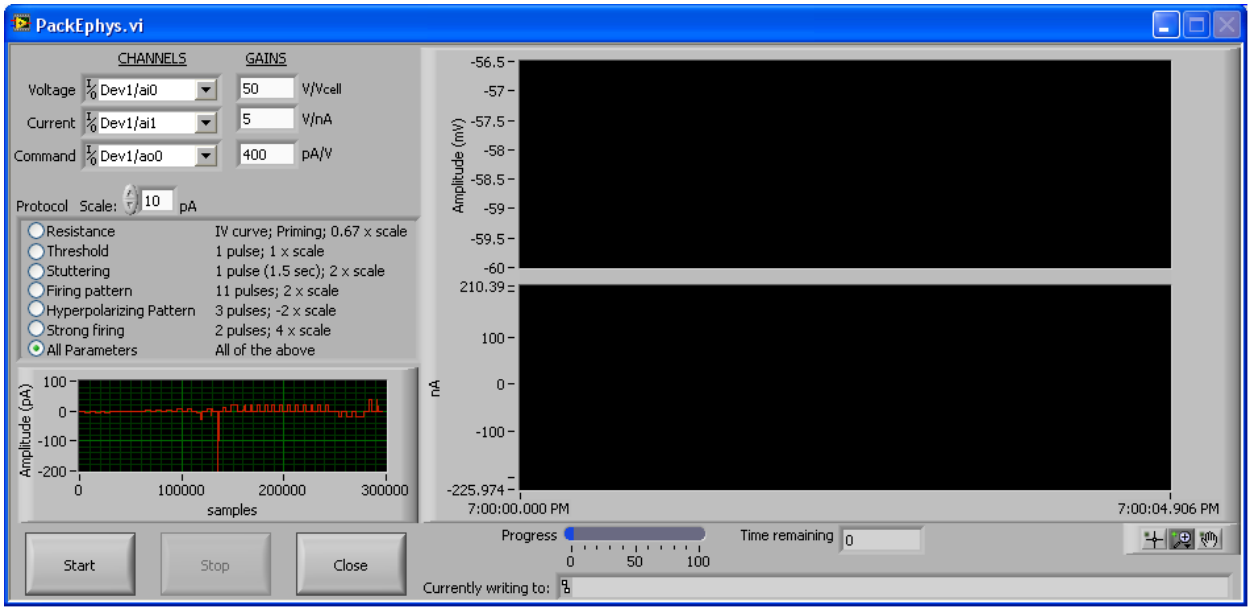

Figure A1-6: The PackEphys window of PackIO.

PackEphys is a program that can automatically probe the electrophysiological characteristics of cells. After setting the correct channels and gains, the user should select the "Threshold" protocol and a low value for the scale factor (which will depend on the cell type) and then click START. After prompting for a file name, the program will output a square pulse of the scaled amplitude. The user should increase the scale until the recorded cell produces at least one action potential. Then the user can select other protocols to apply to the cell at the same scale. A file produced by "All parameters" can be automatically analyzed by PackEphys.m, a MATLAB script available from the Yuste CVS repository in the PaqTools module (http://yustecvs.biology.columbia.edu/cgi-bin/cvsweb/). The electrophysiological characteristics that can be obtained are based on those used in the Markram lab (see Table 3 of Toledo-Rodriguez et. al., Cerebral Cortex, December 2004;14:1310–1327; doi:10.1093/cercor/bhh092):

Markram	PackEphys
e1	Ephys.RestingPotential
e2	Ephys.InitAPDrop
e3	Ephys.AP1ToSteadyDrop
e4	Ephys.AP2ToSteadyDrop
e5	Ephys.MaxRateAPChange
e6	Ephys.AP1Amp
e7	Ephys.AP1Duration
e8	Ephys.AP1HalfWidth
e9	Ephys.AP1RiseTime
e10	Ephys.AP1FallTime
e11	Ephys.AP1RiseRate
e12	Ephys.AP1FallRate
e13	Ephys.AP1fAHP
e14	Ephys.AP2Amp
e15	Ephys.AP2Duration
e16	Ephys.AP2HalfWidth
e17	Ephys.AP2RiseTime
e18	Ephys.AP2FallTime
e19	Ephys.AP2RiseRate
e20	Ephys.AP2FallRate
e21	Ephys.AP2fAHP
e22	Ephys.AP12AmpPercChng
e23	Ephys.AP12DurationPercChng
e24	Ephys.AP12HalfWidthPercChng
e25	Ephys.AP12RiseRatePercChng
e26	Ephys.AP12FallRatePercChng
e27	Ephys.AP12fAHPPercChng
e28	Ephys.InputRPeak
e29	Ephys.InputRSteady
e30	Ephys.RectificationPeak
e31	Ephys.RectificationSteady
e32	Ephys.Sag
e33	Ephys.DeltaTau
e34	Ephys.RampThresh
e35	Ephys.RampfAHP
e36	Ephys.PostBurstMaxAHP
e37	Ephys.PostBurst100msAHP
e38	Ephys.PostBurstTimeToMaxAHP
e39	Ephys.NumSpikesPerPicoAmp
e40	Ephys.AveDelayToFirstSpike
e41	Ephys.StdDelayToFirstSpike
e42	Ephys.AveDelayToSecondSpike
e43	Ephys.StdDelayToSecondSpike
e44	Ephys.AveFirstThreeApISIs
e45	Ephys.SDFirstThreeApISIs
e46	Ephys.AveInitFiringRateAccom
e47	Ephys.AveSteadyFiringRateAccom
e48	NOT INCLUDED
e49	NOT INCLUDED
e50	NOT INCLUDED
e51	Ephys.ISI_CV
e52	Ephys.ISIMedian
e53	Ephys.aveMaxDerivMinusAveDerivISIs
e54	NOT INCLUDED
e55	NOT INCLUDED
e56	NOT INCLUDED
e57	NOT INCLUDED
e58	NOT INCLUDED
e59	NOT INCLUDED
e60	NOT INCLUDED
e61	NOT INCLUDED

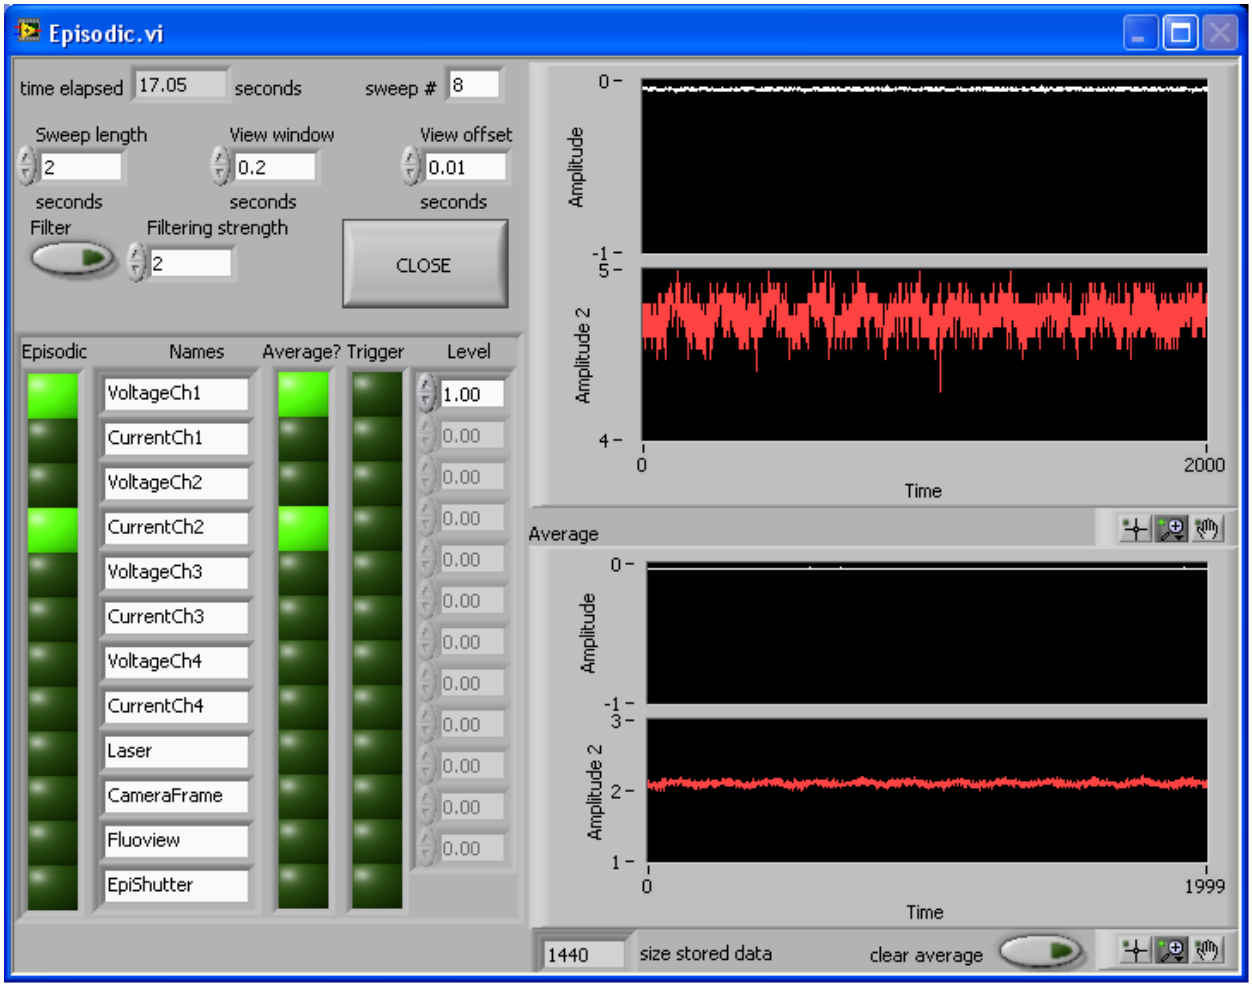

Figure A1-7: The episodic window of PackIO.

Data being acquired can be displayed in an episodic fashion by clicking the "Episodic display" button on the main window of PackIO to bring up the Episodic window.

Available channels will be automatically listed based on the channels that have been set for acquisition in the Master Setup window. Channels the user wishes to view in an episodic fashion should be selected by clicking the episodic button to the left of the name of that channel. An average can be accumulated across sweeps for a given channel by clicking the appropriate button in the "Average?" column. Individual sweeps are

displayed in the two graphs at the top right while accumulated averages are displayed in the two graphs at the lower right.

Episodes can be defined either by a set window or a trigger. If no channels are selected as a trigger by clicking one of the buttons in the "Trigger" column, the episodes will default to occur at the intervals specified by the "sweep length" at the top left of the window. If a trigger is selected, when the trigger level is met on the selected trigger, an episode will start. The length of an episode is set by the "sweep length" (in seconds) and the length of the data displayed is set by the "view window" (in seconds). A "View offset" (in seconds) can be set to control the time the view window begins within a given episodic sweep.

The total time elapsed and episodic sweep number are displayed at the top of the window.