

Scalable Semisupervised Functional Neurocartography Reveals Canonical Neurons in Behavioral Networks

E. Paxon Frady

paxon.frady@gmail.com

*Neurosciences Graduate Program, University of California San Diego,
La Jolla, CA 92093, U.S.A.*

Ashish Kapoor

akapoor@microsoft.com

Eric Horvitz

horvitz@microsoft.com

Microsoft Research, Redmond, WA 98052, U.S.A.

William B. Kristan Jr.

wkristan@ucsd.edu

*Department of Biological Sciences, University of California San Diego,
La Jolla, CA 92093, U.S.A.*

Large-scale data collection efforts to map the brain are underway at multiple spatial and temporal scales, but all face fundamental problems posed by high-dimensional data and intersubject variability. Even seemingly simple problems, such as identifying a neuron/brain region across animals/subjects, become exponentially more difficult in high dimensions, such as recognizing dozens of neurons/brain regions simultaneously. We present a framework and tools for functional neurocartography—the large-scale mapping of neural activity during behavioral states. Using a voltage-sensitive dye (VSD), we imaged the multifunctional responses of hundreds of leech neurons during several behaviors to identify and functionally map homologous neurons. We extracted simple features from each of these behaviors and combined them with anatomical features to create a rich medium-dimensional feature space. This enabled us to use machine learning techniques and visualizations to characterize and account for intersubject variability, piece together a canonical atlas of neural activity, and identify two behavioral networks. We identified 39 neurons (18 pairs, 3 unpaired) as part of a canonical swim network and 17 neurons (8 pairs, 1 unpaired) involved in a partially overlapping preparatory network. All neurons in the preparatory network rapidly depolarized at the onsets of each behavior, suggesting that it is part of a dedicated rapid-response network. This network is likely mediated by the S cell, and we referenced VSD recordings to an activity atlas to identify multiple cells of

interest simultaneously in real time for further experiments. We targeted and electrophysiologically verified several neurons in the swim network and further showed that the S cell is presynaptic to multiple neurons in the preparatory network. This study illustrates the basic framework to map neural activity in high dimensions with large-scale recordings and how to extract the rich information necessary to perform analyses in light of intersubject variability.

1 Introduction

Explaining behavior in terms of the activity of neurons within neuronal circuits is a major goal of systems neuroscience (Alivisatos et al., 2012; Bargmann & Newsome, 2014), but the complexity of the nervous system has limited progress in this endeavor. To tackle this complexity, efforts in large-scale neurocartography (Kasthuri & Lichtman, 2010) have produced detailed maps of neural anatomy and connectivity and have led to numerous insights in neural function (Chalfie et al., 1985; Macosko et al., 2009). However, even the complete wiring diagram of a nervous system (White, Southgate, Thomson, & Brenner, 1986) has not produced a complete understanding of the neural control of behavior (Bargmann & Marder, 2013; Morgan & Lichtman, 2013), because both structure and function are needed (Briggman, Helmstaedter, & Denk, 2011). Further effort has been invested in performing large-scale recordings of neural activity, but developing principled methods for interpreting these high-dimensional data sets has remained an aspirational goal.

The neurons in the leech nervous system are individually identifiable by their homologous anatomy, electrophysiological properties, and relationships to behaviors from ganglion to ganglion and from animal to animal (Muller, Nicholls, & Stent, 1981). However, the majority of the approximately 400 neurons within a single ganglion (Macagno, 1980) remain unidentified, primarily because of intersubject variability (Weeks, 1982; Friesen, 1989a). Many neurons can be identified by eye, through anatomical features alone, but anatomical cues are not enough information to identify the majority of neurons. Information about function is important to identify certain other individual cells, but how does one identify several hundred cells simultaneously? In this study, we used the leech nervous system to develop strategies for identifying and mapping neurons and their functional contributions to behavior based on large-scale recordings of neural activity.

Advances in voltage-sensitive dyes (VSD; Miller et al., 2012) have enabled an unprecedented resolution of imaging leech neural activity—both action potentials and the underlying synaptic potentials—as the neural circuits produce a variety of behavioral motor patterns. The first steps are to use these large-scale recordings to characterize and identify neurons based on their activity patterns and anatomy. This task, however, is not trivial

because of intersubject variability, a major reason that anatomical landmark-based registration schemes fall short of perfect alignment in fMRI (Ng, Abugharbieh, & McKeown, 2009).

Identification becomes even more challenging when dealing with high-dimensional recordings, where many neurons must be identified simultaneously and computer automation is required. Computer automation is necessary to sift through the vast data sets that large-scale recordings produce, and rapid analysis of these data sets provides needed context to guide experiments. Further, it is unclear what information needs to be extracted from the system in order to identify one neuron among hundreds, and statistical feature selection techniques have been developed mainly for low-dimensional classifications (Gareth, Witten, Hastie, & Tibshirani, 2013). This task presents a formidable machine learning challenge in all forms of large-scale spatiotemporal data acquisition, but one that must be overcome if large-scale multicellular recording efforts are to succeed (Bargmann et al., 2014).

We used VSD imaging to record simultaneously from many leech neurons during multiple fictive behaviors, and we developed a machine learning framework for identifying neurons across animals based on their VSD activity patterns. Instead of relying solely on anatomical cues, which do not provide enough information to uniquely identify each neuron, we used both functional and anatomical cues to parse the data set into a “rich medium-dimensional feature space.” We then employed interactive machine learning techniques to piece together a canonical atlas of leech neural activity and handle intersubject variability. Current machine learning techniques still fall short of the capabilities of human pattern recognition, so we turned to semisupervised techniques, which merged algorithms, machine learning, visualization, and human feedback at each stage of the analytical pipeline. We implemented the overall methodological framework within an interactive tool, the “imaging computational microscope” (ICM; Horvitz & Kristan, 2009; Frady & Kristan, 2015), and verified that newly identified neurons could be found across animals and targeted for electrophysiological recording and stimulation.

We used these techniques to identify canonical swim and preparatory networks in the leech by stitching together data collected across eight different animals. We characterized the activity of dozens of previously unidentified neurons during each of these behavioral responses, and we showed that the preparatory network is activated by stimuli that produced several different behavioral outputs, at short latencies, before any behavioral activity pattern is recognizable (Friesen & Kristan, 2007). This analysis predicted a connectivity pattern among the preparatory neurons. We tested some of these predictions electrophysiologically, using the computational microscope to identify and target preparatory neurons, and found the predicted connectivity patterns to be accurate. These analysis techniques can provide information in real time to guide experiments and target neurons

and networks. We propose that this or similar algorithmic frameworks will be necessary for functional neurocartography of more complex systems—of individually identifiable neurons, neuron types, fMRI (Varoquaux et al., 2010) and mesoscale brain regions (Oh et al., 2014).

2 Methods and Results

2.1 Building a Rich Feature Space. Each leech ganglion, to a first approximation, consists of the same homologous set of neurons across animals (Muller et al., 1981; Kristan, Calabrese, & Frieson, 2005), but anatomical variability across ganglia makes identifying most neurons nontrivial. Some of the neurons in the leech ganglion can be recognized by eye based solely on anatomical cues, but the algorithm describing this visual task is unclear and has not been specified. A formalization based on machine learning suggests that human experts identify high-level features from low-level anatomical information. These features are jointly taken into consideration to identify individual neurons. This process can be described as experts crafting a high-dimensional unique (“one-hot”) feature space that is derived from the low-level features of the image, or the anatomical features of each neuron:

$$F_{Neuron} = v(F_{Anatomy}),$$

where F_{Neuron} is a binary high-dimensional feature space in which each dimension corresponds to a neuron category, and v is a function implemented by the visual system of experts that translates $F_{Anatomy}$, the low-level anatomical features of the neurons, into the high-level neuron identity space. For instance, when experts recognize the Retzius neuron from low-level anatomical features (i.e., its position and size in the ganglion), then F_{Neuron} at the index corresponding to the Retzius neuron is 1 ($F_{Retzius}_{Neuron} = 1$) and 0 everywhere else (more generally, F_{Neuron} can represent the probability that the low-level features indicate the identity of the neuron for each neuron, and the vector F_{Neuron} would always sum to 1). The function v implements this transform, but it is not necessarily known. Whatever way experts recognize the neuron, we can say that the neuron is identified if the value of F_{Neuron} is 1 at that neuron’s index (or the probability that it is that neuron).

Most neurons in the leech cannot be distinguished solely on visual cues. For instance the N and T cells are in similar locations and have similar sizes. Although this distinguishes them from other neurons in the ganglion, they are not fully visually distinguishable from each other. In order to fully identify the N or the T cell, an action-potential recording is sufficient, as these cells have distinct action-potential shapes (Muller et al., 1981):

$$F(\{N, T\})_{Neuron} = v(F_{Ephys}).$$

To identify more neurons, we would need a feature that uniquely identifies each neuron. This is the analogous strategy of mapping the unique function of different regions in the human brain: each region corresponds to a particular behavior, and the regions can be mapped by identifying each region's unique behavioral faculty (Gall, 1833). But for many neurons, a uniquely identifying feature is unknown. Even the action-potential shapes of many neurons are too similar to distinguish them from one another (Muller et al. 1981). Further, scalability poses fundamental problems when using unique features to identify hundreds of neurons. To recognize the hundreds of neurons in the leech, hundreds of unique features would be required. The "curse of dimensionality" (Bellman, 1957) exacerbates the problem as each feature must distinguish each neuron among hundreds of distractors, and disambiguating features must be found for every neuron. This makes it a combinatorial impossibility to scale one-hot feature vectors to uniquely identify hundreds of neurons.

An alternative strategy is to build a rich medium-dimensional feature space, where neurons are described based on different combinations of a small number of features rather than each having a unique dimension. For instance, to identify 200 neurons, one would need 200 unique identifying features, whereas only 8 rich binary features would be enough ($2^8 = 256 > 200$). This allows for a much smaller dimensionality of the feature space while still retaining the ability to discriminate a large number of neurons. In order for this to be possible, the medium-dimensional feature space must have certain properties that make it rich. The two main properties required for such richness are the number of categories the feature space can transmit (i.e., the number of neurons that can be identified) and its dimensionality. The amount of information in the unique feature space and the rich feature space is the same, because they transmit the same number of categories. However, the unique space has a dimensionality equivalent to the number of categories, whereas the rich space has a much lower dimensionality. We can make a simple definition of representational richness:

$$richness = \frac{categories}{dimensionality}. \quad (2.1)$$

Three factors affect the number of categories that can be contained in a feature space: the variability, independence, and sparsity of the features (Shannon, 1948; Barlow, 1961; Olshausen & Field, 1996; Penev, 2001). In the continuous case (instead of having only binary features), noise will limit the categorical precision of each feature dimension. Thus, features that have low variability across animals can contain more categories and transmit more information. The feature dimensions for each neuron need to be independent in order to convey information that distinguishes the neurons. Features that are redundant will increase the dimensionality without

increasing the number of categories. And finally, ideal features should be dense; they should convey information about many neurons.

To meet the fundamental needs of the rich feature space and take advantage of its dimensionality, we used large-scale voltage-sensitive dye recordings during multiple behaviors. Theories of population coding suggest that neurons represent information in their joint activities (Pouget & Sejnowski, 1997; Pouget, Dayan, & Zemel, 2000) and that they use multidimensional receptive fields to represent multidimensional information (Mante et al. 2013; Frady & Kristan, 2014). This means that neurons are often multifunctional, and such multifunctional responses are seen in the leech (Briggman & Kristan, 2006). As predicted from population coding theory, the receptive fields of many leech neurons are well spread across input space but slightly overlapping, and this type of receptive field tuning has been previously observed in the leech (Lewis & Kristan, 1998) as well as other organisms (Movshon, Thompson, & Tolhurst, 1978). This overlap property of receptive fields becomes problematic when trying to identify neurons with only one or two dimensions, because the overlap adds to the difficulty of distinguishing neurons. However, in higher-dimensional space, the receptive fields become increasingly spread out, and neurons become more distinct from one another. But, this is true only if each of the higher dimensions has independent information about the identity of the neuron. Luckily, neural coding theory suggests that the neural activity is as independent as possible for different behaviors (Olshausen & Field, 1996). Thus, features describing independent behaviors have high potential for use in crafting a rich feature space.

The isolated leech nervous system can elicit different fictive behaviors depending on where it is stimulated (Briggman & Kristan, 2006). We used VSD imaging (Miller et al., 2012) to optically record from hundreds of neurons in a single midbody ganglion (ganglion 10) during three behaviors: shortening, swimming, and local bending (see Figure 1A). In an intact leech, swimming and shortening are elicited by stimulating the animal near the posterior and anterior, respectively (Palmer, Barnett, Copado, Gardezy, & Kristan, 2014). Fictive shortening was activated by stimulating a nerve (DP) from either ganglion 3 or 7, and swimming was activated by stimulating 11, 14, 17, or the tail brain. Local bending was activated by intracellular stimulation of a single sensory P cell (see Figure 1B; Kristan, 1982). Of the nearly 400 neurons in a leech ganglion, only about one-third have been previously identified (see Figure 1B, colored cells; Muller et al., 1981). About 80% of neurons on the ventral surface can be imaged with VSDs, but due to the variability of the ganglion across animals, as well as variability in dissections and VSD application, a random subset of the neurons was imaged in each experiment. We recorded from eight different animals (A–H) to sample the neurons in the ganglion multiple times. We imaged every ganglion during each behavior one to four times at 50 Hz for 10 seconds, which allowed us to monitor synaptic potentials, action potentials, and

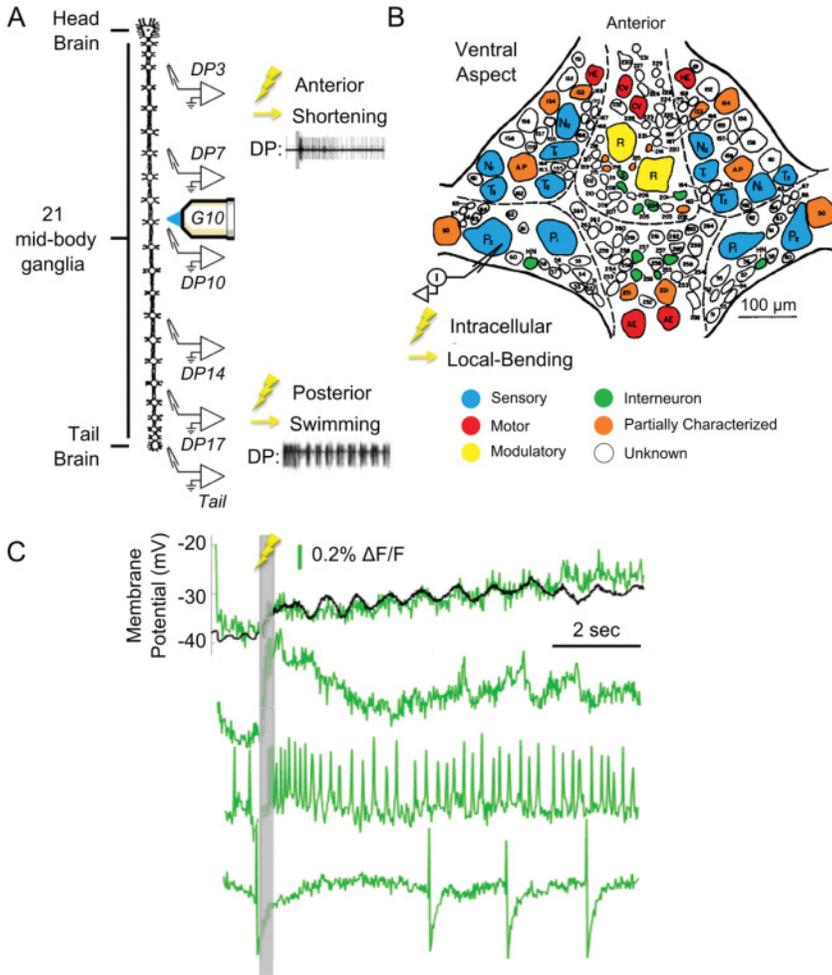


Figure 1: Voltage-sensitive dye imaging during multiple behaviors. (A) The isolated nerve cord of the leech was dissected, and ganglion 10 was prepared for VSD imaging. Anterior or posterior nerves were stimulated with an extracellular suction electrode to elicit shortening or swimming, respectively. (B) Schematic of leech ganglion. Neurons that have been previously identified are shown in color. Local bending is activated by targeted stimulation of a P cell. (C) Voltage-sensitive dye imaging of neurons. A simultaneous intracellular (black) and optical (green) recording is shown (top), which illustrates clear signals of oscillations less than 5 mV. Three other example traces show optically recorded action potentials.

variations in the membrane potential in individual neurons in single trials (see Figure 1C). We refer to an individual recording as a “trial” and the set of recordings from a single animal as an “experiment” or “animal.”

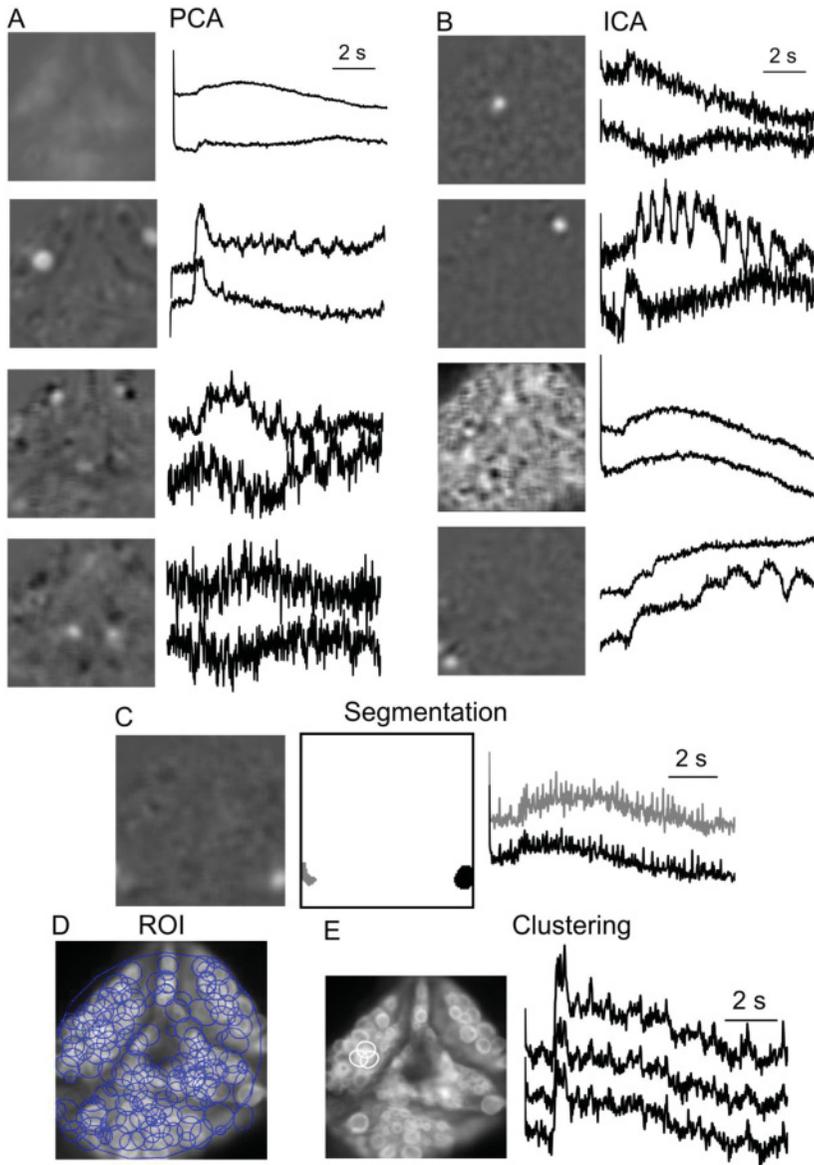
Typically imaging data is decomposed into neuronal responses using manually defined regions of interest (ROIs). New computational techniques using matrix factorizations (PCA-ICA; Mukamel, Nimmerjahn, & Schnitzer, 2009; Hill, Moore-Kochlacs, Vasireddi, Sejnowski, & Frost, 2010) can automatically extract cellular signals directly from imaging data. This technique is superior to ROIs, because it can remove the background and signals that overlap onto the same pixel. To aid in the use of these algorithms and handle their shortcomings, we used the ICM to automatically extract signals from imaging data. The ICM enables investigators to quickly visualize the results of the PCA-ICA extraction (see Figure 2A, B), alter the parameters of the extraction, segment correlated signals (see Figure 2C), automatically generate ROIs (see Figure 2D), cluster decomposed individual sources (see Figure 2E), remove artifact components (e.g., see Figure 2B, bottom two panels), and create activity maps for visualization (see Figure 3).

Both PCA and ICA are linear component decompositions, and when applied to imaging data, they produce several components, each of which has a “source” and a “map” (see Figures 2A and 2B). Sources are the time series of the extracted components, and maps show their spatial distributions. The ICA algorithm extracts both artifacts and neuronal signals as components; the artifacts can be sorted quickly by visual inspection of the sources and maps with ICM. Components that are chosen as signals are called neurons; all others are considered artifacts (two examples of neurons are shown at the top of Figure 2B and two artifacts are shown at the bottom). We analyzed multiple trials of different behaviors by performing “concatenated trial ICA” (ctICA; Frady and Kristan, 2015). In this procedure, we used an image registration algorithm (Evangelidis & Psarakis, 2008) to align multiple trials as if they were obtained from a single imaging acquisition (i.e., they formed a single time series). Because in ctICA the trials are concatenated into a single acquisition, the map of each component is the same for every concatenated trial, and the source is a single time series that extends through all trials. The sources are broken up for each trial to give individual traces, each of which is the activity of a component during a single trial (in Figures 2A and 2B, two example traces are shown for each component from different trials).

We designed computational algorithms based on coherence and factor analysis to extract rich functional features characterizing the main aspects of the shortening, local bending, and swimming behaviors. The shortening response was elicited by stimulating anterior DP nerves (in ganglion 3 or 7) through suction electrodes at 20 Hz for 250 milliseconds with short (1 mS duration) current pulses (Briggman, Abarbanel, & Kristan, 2005). We monitored motor neuron spikes by recording extracellularly from other DP nerves (see Figure 3A, top trace). During shortening, different neurons

showed a variety of responses (see Figure 3A, colored traces) that were well characterized by two factors (see Figure 3B) derived from the neural responses during all shortening trials across animals. The factors were derived by performing PCA on the shortening responses across all trials of shortening (including trials of shortening from other animals not present in this study) and then combining the first four eigenvectors (which explained 74% of the variance) into two factors. Factor 1 corresponds to prolonged depolarizing (positive factor 1) or hyperpolarizing (negative factor 1) responses that last as long as the motor neuronal burst that constitutes fictive shortening. Factor 2 depicts a rapid-onset response that is usually positive, indicating that many neurons rapidly depolarize after each stimulus. Factor coefficients (see Figure 3C; Sh) were computed by fitting each trace with the factors and a fourth-degree polynomial using a linear regression. The polynomial fit was added so some variability in the bleaching and response duration could be ignored while still getting a good fit from the factors. The components are colored based on their two factor coefficients: positive factor 1 increases the green channel, negative factor 1 increases the blue channel, and factor 2 is shown in the red channel. These two coefficients provide a low-dimensional representation of the shortening response. An activity map was generated by coloring each neuronal somata on its ganglionic image using the colors derived from the factor analysis (see Figure 3D). This visualization shows the activity of about 100 neurons and is useful for rapid high-level analysis and visual inspection of large-scale neural activity. All the maps in all figures are labeled by experiment (A–H) and—for shortening and swimming—by stimulus location (ganglion 1–21) on the top right and bottom right, respectively.

To elicit local bending, we activated one of the four mechanosensory P cells in a ganglion with bursts of five to six action potentials in 300 ms, repeated every 2 seconds (Kristan, 1982). This stimulus produced repeated local bend responses, monitored in the DP nerve recordings (see Figure 3E, top trace). The traces from the same neurons shown for shortening (see Figure 3A) indicated that neurons with similar responses in shortening did not necessarily produce similar responses during local bending, illustrating the multifunctional nature of the neural activity. To characterize this behavior compactly, we calculated the coherence of each VSD trace with the intracellular stimulus burst (see Figure 3F; LB). Coherence analysis (Mitra & Bokil, 2007; Bokil, Andrews, Kulkarni, Mehta, & Mitra, 2010) calculates a correlation magnitude and phase in Fourier space for each component, which we plotted in polar coordinates at the stimulus frequency. In these plots, neurons that are functionally excited by the burst of P cell spikes are indicated by values that are nearly in phase (around 0°) with the stimulus (red; see Figure 3F) and inhibition is indicated by coordinates that are more nearly out of phase (around 180°) with the stimulus (cyan; see Figure 3F). We then used these colors to visualize the neurons as an activity map (see Figure 3G; the P cell circled in white was stimulated to elicit the behavior).



We initiated the swimming motor program by stimulating posterior DP nerves (in ganglia 11, 14, 17, or from the tail brain), using the same stimulus parameters as for shortening. Recurring bursts of extracellular spikes at 1–2 Hz by a motor neuron (DE-3) in a DP nerve (see Figure 3H, top) identify the response as swimming (Kristan & Calabrese, 1976). Optical traces from

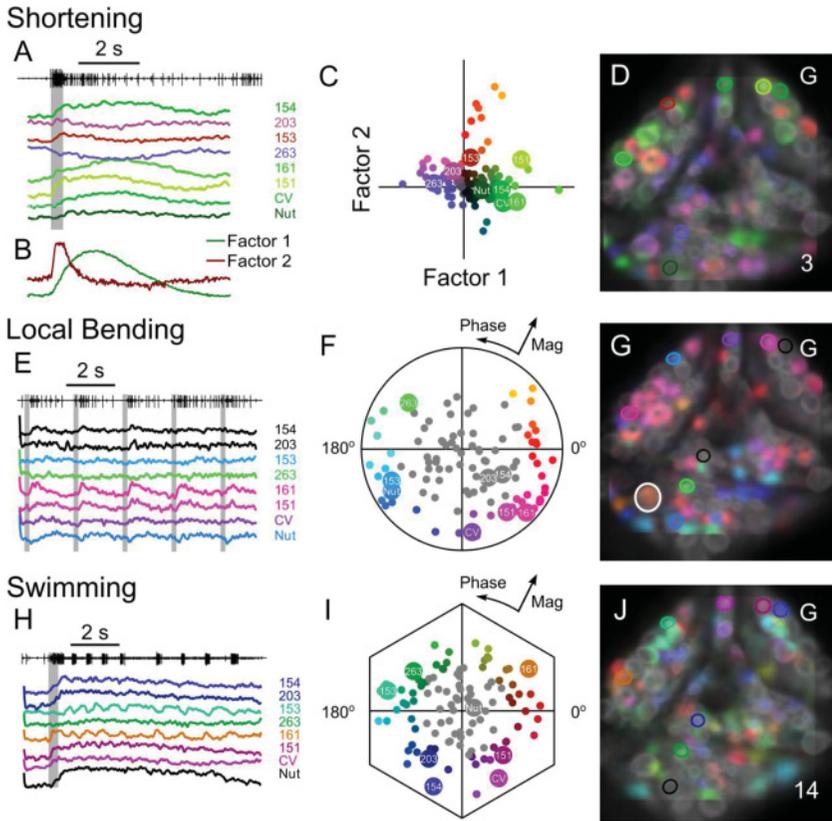
the same neurons as used for local bending and shortening show that the membrane potentials of many neurons oscillated at different phases relative to the spikes in the DP nerve during each cycle of swimming. Coherence analysis correlating the VSD traces with the DE-3 spike bursts (see Figure 3I; Sw) both identified swim related oscillations in neurons and characterized their relative phases. Based on these coordinates, the components were given different colors representing their phases, which were used to color the somata in an activity map (see Figure 3J).

Finally, we derived anatomical features in an automated fashion by segmenting the ICA maps and fitting oval regions of interest (see Figure 2D). We used this to define four dimensions of anatomic features: two dimensions for the neuron's position (Pos) and two dimensions describing the major and minor axis dimensions of the oval (Sz). The features derived from the traces are used to form a rich medium-dimensional feature space by concatenating all of the features for all neurons across all animals into a large matrix:

$$F(n)_{Rich} = [POS(n)_{XY}, Sz(n)_{XY}, Sh(n)_{XY}, LB(n)_{XY}, Sw(n)_{XY}].$$

Each row of this matrix represents the 10 features describing the multifunctional activity and anatomical properties for an individual neuron (the phase and magnitude values determined by coherence analysis of

Figure 2: Semisupervised signal extraction with concatenated-trial ICA (ctICA). (A) PCA was performed on concatenated trials. Four example principal components (PC) are shown. Each PC has a map (left) and multiple traces (right). Two example traces of each PC are shown from different trials. (B) ICA was performed on the top 120 to 150 principal components. The independent component (IC) maps (left) and traces (right) are shown for four examples. The top two components are neurons. The bottom two components are examples of neuropil response and motion artifact, and these components were discarded. (C) The independent component maps (left) were then segmented using a threshold (middle). An example is shown where two neurons are identified within the same independent component, but were separated through spatial segmentation. These example neurons are strongly electrically coupled to each other (Muller et al., 1981), and ICA does not separate them because they have similar activity (right). (D) Regions of interest were automatically generated from the segmentation of the component maps. The traces described in the analysis were direct outputs from ICA, which removes the background and overlapping signals. The ROIs were used only to compute the anatomical features, and they were not used to compute the traces. (E) Three components that are from a single neuron are shown. The algorithm occasionally splits single sources (especially larger neurons), and components with similar locations and correlated traces can be clustered using the GUI.



local-bend and swim behaviors were converted into x, y values). The features were whitened to have a normalized variance.

2.2 Assessing Intrasubject Variability and Feature Selection. For functional features to be useful in identifying homologous neurons across animals, they must have low variability and high independence. We recorded each of the behaviors at least twice in seven of the eight animals and assessed the richness of each feature based on the trial-by-trial variability of the behaviors within animals, in pursuit of independent information about neuron identity.

The ultimate goal is to use the features to identify neurons across animals, and we can assess (the upper bound of) this ability based on the trial-by-trial variability of the neural features within animals. If the variability is too high within animals, there is no hope that the feature will be useful across animals. We assessed the ability of the functional features to discriminate the neurons by analyzing the functional variability of the neurons

across different trials of the same behavior. We built a feature space based on different conjunctions of functional and anatomical features (i.e., some feature subset of F_{Rich}) from one trial and compared the distance of each neuron to itself with the distance to the next closest neuron in a different trial. The discriminability (d^f) is defined as the negative log of the ratio of these distances and is dependent on the feature space used to describe each neuron (f):

$$d_i^f = -\log \left(\frac{\sum_{f1, f2 \in f} \|n_i^{f1} - n_i^{f2}\|_2}{\min_{j \neq i} \|n_i^{f1} - n_j^{f2}\|_2} \right), \quad (2.2)$$

where $f1$ and $f2$ are feature spaces built from different trials of the same behavior. A discriminability greater than 0 means that the neuron was closer to itself than to its next closest neighbor, a promising sign that the

Figure 3: Low-dimensional feature extraction from multiple behaviors. (A) Shortening was activated by stimulating an anterior DP nerve (gray bar). Other DP nerves show motor neuron spikes and were used to read behavioral output (top). Several example VSD traces from identified neurons are shown in color. (B) Neuron traces were decomposed into two factors: a slow factor that follows the motor behavior (factor 1, green) and a fast factor (factor 2, red). (C) The factors were fit to each trace, and the coefficients of the factor fits were plotted. The components were given colors based on these coefficients. The coefficients of the examples from A are shown as large circles with white labels. (D) Activity map of shortening. Each component's map was colored based on the shortening factor coefficients and overlaid onto an image of the ganglion. The example neurons are circled. The top right letter (G) indicates that these data are from animal G, and the bottom left number (3) indicates which ganglion was stimulated. (E) A burst of six action potentials was evoked every 2 seconds to sensory P cells (gray bars) to repetitively elicit the local bend response. The motor neuron spikes of the behavior are seen in the DP nerve (top). The VSD traces of the same example cells are shown in color. (F) The coherence of the traces was calculated against the stimulus at 0.5 Hz (the stimulus frequency), and the coherence phase and magnitude for each trace were plotted on polar axes. The components are colored based on their phase and magnitude. (G) Activity map of local bending. Colors are based on the phase and magnitude of the coherence analysis. The same example cells are circled. A white circle indicates the stimulated P cell. (H) Swimming was activated by stimulating a posterior DP nerve (gray bar). The swim motor pattern was monitored through DP nerve recordings (top). The VSD traces of the same example components are plotted. (I) The coherence values of the traces were calculated against the DE-3 motor neuron spike output at the swim frequency (large spikes in DP nerve recordings). The phase and magnitude were plotted as a polar plot. (J) Activity map of swimming. Same example cells are circled.

feature will be useful for identifying the neuron. Equation 2.2 is an approximate measure related to the sensitivity index (e.g., $\log(d')$), but for high dimensions.

The anatomical features in a single ganglion are exactly the same across different trials, of course, but we know that there is variability across animals. To determine the influence of this on interanimal variability, we added gaussian white noise to the position and size dimensions when calculating the discriminability. The variance of the noise was half the variance of the position and size features across all neurons, which was based on our estimate of the variance across animals, but these results were largely unaffected by the choice in noise variance as long as it was above a minimum threshold (if one considers position or size alone without any noise across trials, then every neuron is discriminable because there is no variability in the feature space, but this does not reflect the variability across animals, which we are trying to approximate).

We then assessed the richness of each feature by building up the feature space with different subsets of features and plotting a histogram of the discriminability metric for every neuron from every animal. As the histogram shifts rightward, more neurons become discriminable. Position (with noise) alone was not very useful in identifying neurons across trials and is reflected by low discriminability (see Figure 4A, left panel). By using a conjunction of anatomic and individual behavioral features, many more neurons can be discriminated (see Figure 4A, middle three panels). In fact, each of the behavioral features (from shortening, local bending, and swimming) added about the same percentage of discriminability. Combining the anatomical features with all three behavioral features nearly doubled the number of discriminable neurons (see Figure 4A, right panel). This was true for every animal, and each animal has approximately the same distribution of discriminability based on the different feature subsets (each animal indicated by a different color in the histograms of Figure 4A).

The intuition from Figure 4A is that as we add more dimensions to the feature space, the distance between the neurons becomes increasingly large, so that they are easier to identify. However, studies of feature selection show that adding redundant features does not help, and that adding noisy features can hinder discriminability (Guyan & Elisseeff, 2003; James, Witten, Hastie, & Tibshirani, 2013; Bermingham et al., 2015). For insights into the true richness of each anatomical and functional feature, it is important to examine how the richness changes when the feature is removed from the full feature space. The independence of the features can thus be assessed, as features carrying redundant information lessen the richness of the representation. The influence of redundancy of a specific feature can be characterized by comparing the richness of the full feature space to the reduced space with that feature removed. If the feature carries independent information, the richness of the reduced feature space will be lower than the full space. However, if the feature has only redundant information, then

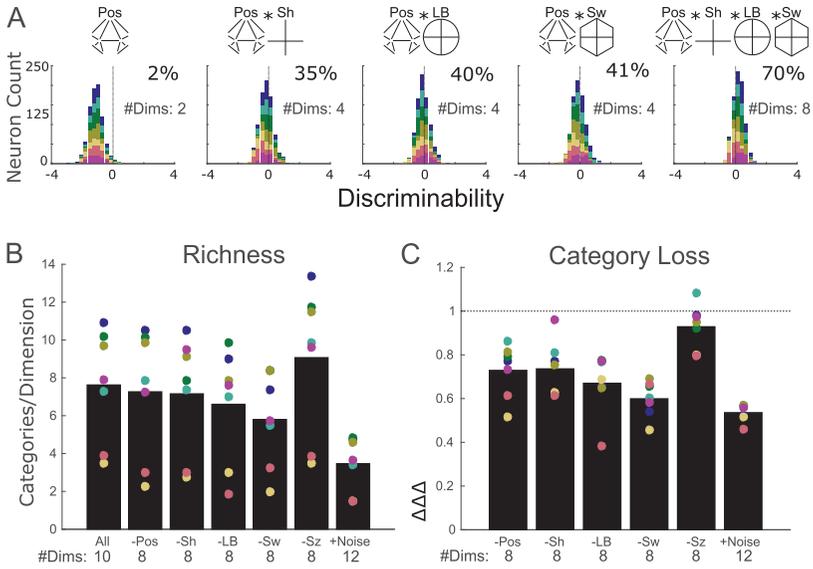


Figure 4: Assessing feature richness based on within-animal variability. (A) The discriminability (see equation 2.2) was measured for each neuron in each animal (animal indicated by color; 722 neurons) based on different subsets of the full feature space and plotted as a histogram. Neurons with discriminability above 0 were closer to themselves in different trials of the same behavior than to the next-nearest distractor neuron. The percentage of neurons with discriminability above 0 is indicated at the top right of each panel. The noisy position feature alone (leftmost panel) discriminates only a few neurons, but in conjunction with functional features, the position feature helps to discriminate many neurons (middle three panels). When multiple functional features and anatomical features are combined, a majority of cells can be discriminated (rightmost panel). (B) The richness of different feature spaces is compared (each animal is individually indicated by the colored circle and the bar is the mean across all animals). Taking out position (-Pos) or the behavioral features (-Sw, -LB, -Sh) reduced the richness of the feature space, whereas removing the Size (-Sz) feature increased the richness. When two dimensions of noise (+Noise) were added to the full feature space, the richness greatly decreased. (C) The change in the number of categories transmitted by each feature space is compared to the full 10-dimensional feature space, which is indicated by the horizontal line with a value of 1. Taking out position or the functional features results in a large reduction in the number of categories transmitted. Taking out the size dimensions of the feature space results in nearly no category loss, indicating that size does not transmit much independent information about neuron identity. Adding two dimensions of noise reduces the number of categories in the feature space by nearly half.

the richness will increase with the removal of the feature because no information is lost and the dimensionality is lowered. We evaluated the richness (see equation 2.1) of each anatomical and functional feature by removing the features, one pair at a time, from the full feature space (see Figure 4B).

To compute richness, we measured the number of categories (in this case, the number of self-similar neurons) that each feature space transmitted. We computed richness for each animal (indicated by the differently colored dots in Figure 4B) by considering every neuron that has a discriminability above 0 as a category in the feature space. We then removed the pairs of dimensions related to the position of the neuron (X and Y locations), swimming, local bending, shortening (the three pairs of behavioral dimensions as defined in Figure 3), and size of each neuron (maximum and minimum diameters), and measured the number of categories that could be distinguished by the remaining eight dimensions. Removing information about the position of the neurons or their activity in any of the three behaviors decreased the richness. This means that each of these features adds independent information that can discriminate the neurons, and fewer categories are transmitted when the features are removed from the full space (see Figure 4C). On the other hand, removing the two size dimensions increased the richness and barely influenced the number of categories being transmitted by the reduced feature space. This observation suggests that the size of the neurons does not help in discriminating among them given the other features. In a sense, if size is totally irrelevant to a neuron's identity, size is simply a noise term.

To measure how noise may influence richness, we added two dimensions of noise to the original 10 dimensions (see Figure 4B, +Noise bar) and measured the resulting richness. In this case, adding noise greatly reduced the richness as well as the number of categories (see Figure 4C, +Noise). We found that adding noise greatly reduced the number of neurons that can be discriminated. This result suggested that noise (i.e., irrelevant dimensions) is not neutral; its presence can significantly reduce the ability to discriminate among neurons. Although all of the information of the rich space is present in the 12-dimensional feature space, the two extra dimensions of noise caused almost half of the neurons to lose their discriminability compared to the 10-dimensional space (represented by the horizontal line at 1 in Figure 4C). The introduction of these noise dimensions dramatically influenced the feature space because there are so many competitors; we used recordings from hundreds of neurons, so the noise dimensions have hundreds of chances to cause confusion in the feature space. These results show how adding extra feature dimensions that do not carry useful information about the identity of the neurons can hurt the potential for identification, and how building a large feature space out of poorly discriminating features makes it harder to identify neurons, especially with many distractors.

This analysis of richness and category discrimination suggested that the size feature carries only redundant information about neural identity and

that the eight-dimensional feature space of [Pos, Sh, LB, Sw] is the richest set of features that provide the best information for identifying the neurons across animals. Therefore, all remaining analyses were performed on this eight-dimensional feature space.

2.3 Assessing Intersubject Variability. We visualized the rich eight-dimensional feature space to assess the consistency of the features across animals (see Figure 5). We computed the factors for each component during two trials of shortening in four different experiments (see Figure 5A; each experiment is indicated by a letter at the top right). For each of the four experiments, we generated two activity maps from the component responses from the two different trials of shortening. The shortening factor coefficients from these two trials are plotted together on a single axis, and the coefficients from each individual neuron are connected by a line segment. A short line indicates small variability in the two trials. Short lines that do not overlap with other lines represent a neuron that has high discriminability in the feature space. Many neurons consistently depolarized (positive factor 1), hyperpolarized (negative factor 1), and/or showed rapid onsets (positive factor 2).

Multiple trials of local bending showed consistent and robust responses from many neurons, as indicated by short, isolated lines in plots of their calculated coherences (see Figure 5B). Stimulating either a ventral or dorsal P cell (indicated by white circles in each activity map) reliably activated a consistent population of neurons. The local bend features (defined in Figures 3E and 3F) were most sensitive to neurons that closely followed the sensory stimulus (see Figure 5B). This feature showed many small isolated lines in two clusters: pink-colored neurons that were functionally excited by the P cell stimulus and cyan/green-colored neurons that were inhibited.

The swimming features (defined in Figures 3H and 3I) identified neurons whose oscillations were phase-locked to the swim motor output (see Figure 5C). These features were so consistent and robust that many neurons could be identified across animals by the swim activity map alone, because the phase of the neuron's oscillation (in conjunction with its somatic position) provided a unique identifying tag. This is reflected in the swim feature plot (see Figure 5C) by the many short, nonoverlapping lines that are spread out through different phases of the swim cycle. Further, swimming was found to be one of the richest features, as it resulted in the greatest reduction in the richness of the feature space when it was removed (see Figure 4B, -Sw).

Based on the visual similarity of the features across animals (see Figure 5 middle) and the discriminability of the neurons (see Figure 4), we found that the eight-dimensional feature space was valuable for identifying neurons across animals. However, the within-animal discriminability of the full feature space (see Figure 4A, right panel) is an upper bound because the

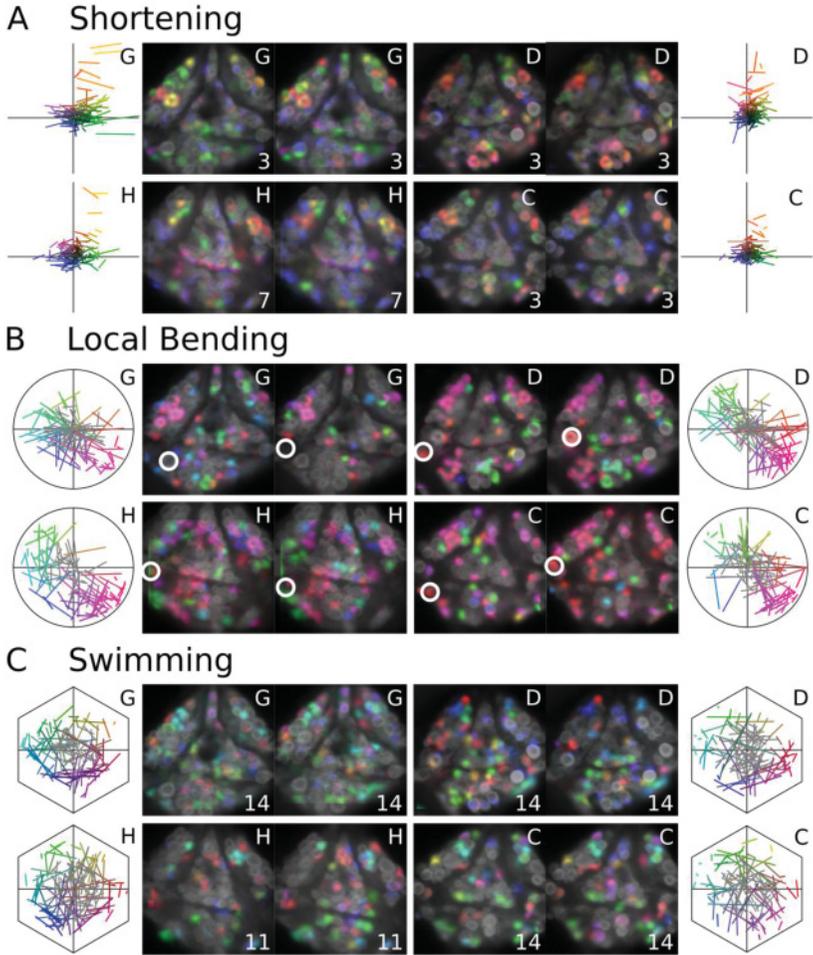


Figure 5: Extracting features within and across animals to visualize functional variability. (A) The shortening feature was extracted from multiple trials of shortening within the same animal. The same neuron's responses during the two trials of shortening are plotted as a line segment, where the ends of the line segment are the shortening factor coefficients from each of the trials. The activity maps of the two trials are shown next to the appropriate coefficient plot. Data from four different animals are shown for comparative purposes. Well-isolated and short line segments indicate that this feature is useful for identifying that neuron. (B) The same analysis was done for the local bending feature. A clear set of neurons could be isolated, as indicated by red/pink and green/cyan somata. In these examples, different P cells were stimulated, showing that many neurons respond consistently to stimulation of any P cell. (C) The same analysis was done for the swim oscillation feature. Many of the neurons show consistent phases of oscillations during swimming, making them easily identifiable.

variability of the features across animals is higher than across trials. We next address the handling of the variability across animals. Although there is clear visual similarity in the activity maps (see Figure 5), characterizing and accounting for intersubject variability is a complex machine-learning problem in need of a solution.

2.4 Managing Intersubject Variability. To help explore the feature space, handle across-animal variability, and identify neurons across animals, we developed semisupervised machine learning algorithms that aid in the registration of homologous neurons across animals to create an atlas. We take an iterative human-computer interactive approach supported by a rich graphical user interface that we implemented within the ICM. In use, users select a neuron at focus of attention in one animal (e.g., colored circles in Figure 6A). The ICM then displays the similarities to neurons from another animal. For instance, of all of the neurons in the ganglion from animal C (see Figure 6B), the somata similarly shaded with a greenish color have similar properties to the green-circled soma in animal H (see Figure 6A). The intensity of the green indicates the degree of similarity. The circled green cell in Figure 6B is the one selected by the automated clustering algorithm as the most likely homologue. The user can then verify, reject, or add matched homologues by giving the neurons labels (see Figures 6C and 6D), and we used the interface to accumulate matches across multiple animals.

In each animal, features were accumulated from several different trials of the three behaviors. To create the full eight-dimensional feature space, multiple trials of the same behavior could be considered. There are several ways to combine the features from multiple trials, and the interface allows users to inspect individual trials or consider parameters and behaviors averaged over multiple trials.

The variability across animals produces artifacts that make identifying neurons difficult. To handle across-animal variability, we worked to allow flexibility in the feature space across animals. Specifically, we developed the weighted correspondence minimization (WCM) algorithm, which optimizes a transform, W , of the feature matrix for each animal to minimize the distance between a set of selected matches.

To find the values for W , we analyzed the relations between cells across animals by comparing their distances in the eight-dimensional feature space. The feature space is transformed into a distance matrix across all animals (matlab: `pdist, squareform`). This translates the eight-dimensional feature matrix into a several-hundred-dimensional distance matrix,

$$D(i, j)_{Rich} = \|W_I F(i)_{Rich} - W_J F(j)_{Rich}\|_2, \quad (2.3)$$

where (i, j) refer to all neurons across all animals and (I, J) refer to the animals that cells (i, j) belong to. Each animal has its own W matrix, which assigns each feature of the feature space a different weight.

can be useful in finding new matches by emphasizing important features or removing common differences across animals.

$$LL = -\log \left(\sum_{(i,j) \in \text{matches}} \frac{\exp(-D(i,j)_{Rich})}{\sum_{k \in J} \exp(-D(i,k)_{Rich})} \right). \quad (2.4)$$

This optimization is analogous to using anatomical landmarks (matches across subjects) and warping a 3D brain scan into a template brain (Ng, Abugharbieh, & McKeown, 2009; Khullar et al., 2011) that is generalized to define landmarks using both anatomical and functional features in a medium-dimensional space rather than a three-dimensional space.

In an iterative procedure, we used the interface to select matches and visualize the relationships of neurons across animals. The set of matches enables the optimization algorithm to find the most important feature

Figure 6: Canonical matching and visualization of medium-dimensional feature space. (A) Three neurons were selected by the user, indicated by the three colored ROIs. (B) For each of the three selected neurons in animal H of panel A, the ICM colored neurons from animal C based on how close each neuron is to the selected neurons in the eight-dimensional feature space. The three selections were visualized through the different R,G,B color channels, and the color intensities indicate the proximity of neurons to the selected ones in the eight-dimensional space. The ICM computed the most likely homologues across animals; the homologues to each of the three circled neurons in panel A are shown as bright ROIs of the corresponding colors. The user can then verify or override the automated matches and assign neurons into categories. (C, D) The matches across animals were given labels, and the neurons that are common matches in multiple animals were accumulated into a label category (only two animals are shown in this figure, but this is done comparing all eight animals). The red labels indicate neurons that were put into an identified neuron category and have been given their category name or number. The black labels indicate neurons that were not identified and have arbitrary numbers. (E) The features of animals H and C are shown in the normalized feature space. Black lines connect chosen matches, and the WCM algorithm was used to emphasize the features that are important to the chosen matches. The warped feature space was projected onto the original feature dimensions (bottom, Warped) illustrating how the warping adjusts the feature space. This emphasizes certain feature dimensions relative to others and allows the user to view different perspectives of the feature space by choosing different sets of matches across animals. (F) The warped space was then visualized and used to find more matches. Notice that the number of lightly colored somata is lower in this panel compared to the prewarped analysis (panel B), indicating that the warping procedure was successful in making the neurons more identifiable. The procedure indicated by the arrows was iterated until many neurons are identified across all eight animals.

dimensions across animals. We employ the Hungarian algorithm (McKay, 1981) at each cycle of the human-computer collaboration to generate the suggested matches, which the user can then accept or reject. Matches are displayed to the user as ROIs (see Figures 6B and 6F). By iterating and comparing each combination of two animals, a set of homologues is accumulated and used as landmarks. The optimization algorithm allows the feature space of each animal to be rescaled and rotated, which removes intersubject variability while emphasizing the important aspects of the feature space indicated by the matches.

We generated a visualization of the warp optimization of the medium-dimensional distance space (see Figure 6E). The position and swim features are stretched out compared to the shortening factors, and the local bend features are compressed along a diagonal axis. These results agree with the richness analysis (see Figure 4) and our visual inspection of the feature variability across animals (see Figure 5): the position of the neurons and their swim coherence were prime indicators of a neuron's identity. The local bend feature separated neurons only by whether they were depolarized (see Figure 5B, pink phase) or hyperpolarized (see Figure 5B, cyan phase) to stimuli, and the algorithm highlighted this feature projection (i.e., the information from local bend is primarily along one dimension). The variability of shortening was the largest across and within animals, making it the least informative feature, so that the influence of shortening parameters was deemphasized (i.e., it is made smaller) by the algorithm. These warping and matching processes were performed iteratively, with the warped feature space visualized by the ICM to aid the detection of new matches (see Figure 6F). The warped feature space highlights the features that were important for the other matches, and visualization of the similarities produced a useful and orthogonal perspective for visualizing the medium-dimensional feature space.

2.5 The Leech Activity Atlas. Using this multiple-iteration human-computer collaboration, we identified 46 neurons (21 bilaterally paired and 4 unpaired) that showed consistent anatomical features and activity patterns during these behaviors across animals (see Figure 7). The results support the involvement of neurons in two canonical networks, one for swimming and a second for preparatory activity—a circuit that has not been previously described. These two networks are partially overlapping: 5 are active only during the preparatory phase, 27 are active only during swimming, and 12 are active during both behaviors. The entries for these identified neurons in Figure 7 were organized by their general anatomical locations and were given names based on an established ganglionic map (Muller et al., 1981). The entries boxed in yellow indicated neurons involved in the preparatory network, and all entries except those indicated by an asterisk (see Figure 7, AP, 212, AE, and S) showed consistent oscillations

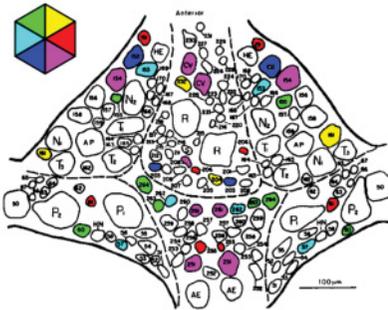
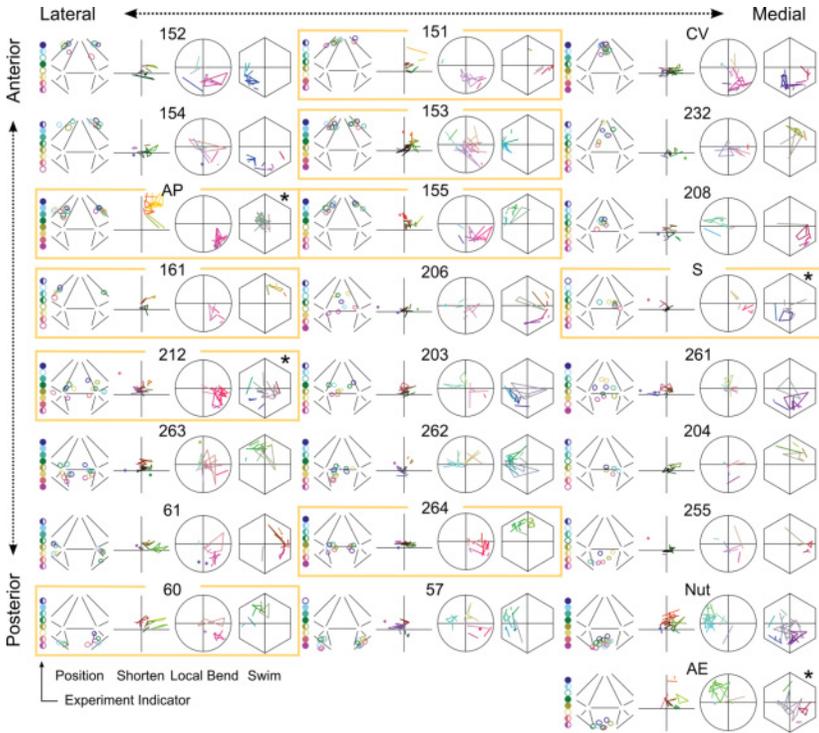
during swimming. Hence, the yellow-boxed neurons without asterisks are neurons that are part of both behavioral networks.

For each identified neuron, we plotted four projections of its unwarped eight-dimensional feature space as an entry in the table of Figure 7. Each plot summarizes the anatomical and functional features of an identified neuron across all trials in all experiments. Each entry in the table is made up of five elements: the experiment indicator (vertical circles), the position indicator (ganglion), the shortening factors (cross), the local bending coherence (circle), and the swimming coherence (hexagon). The experiment indicators (the array of eight vertical circles to the left of each ganglion symbol) represent the experiments in which that neuron was identified. These circles correspond to experiments A to H ordered from top to bottom. Most neurons in the ganglion are bilaterally paired; a solid circle indicates that both members of the pair were identified in a particular experiment, and a half-circle indicates that only one of them was identified or that the neurons are not paired (i.e., cells 232, 208, S, and 204). An open circle means that the neuron was not identified in that experiment. The position indicator shows the location and size of the ROI of the identified component drawn on a schematic ventral view of the ganglion. In addition, the ROI is colored based on its experiment of origin (A–H; the color code for the ROIs are the same as the experiment indicators). For each of the 21 bilateral pairs, both the left and right identified neurons are shown together in the table entry.

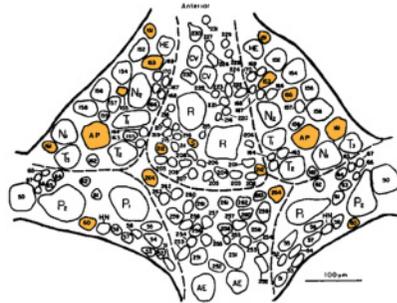
The next three elements in each neuron's description (see Figure 7) show its functional features derived from shortening (cross), local bending (circle), and swimming (hexagon). Data from a single experiment are represented by the vertices of a connected polygon. In one experiment, only a single trial of a behavior was captured; in these cases, a dot represents the coefficients from that trial. When two trials were captured, they are represented as the start and end of a line segment. Likewise, data from three trials are represented by a triangle, four trials by a quadrilateral, and so on, with each corner of the polygon representing feature coefficients from one of the trials. Because each polygon (or dot or line segment) represents an individual experiment, the within-animal variability is represented as the variance of each polygon's vertices, and the variability across animals is the variance between polygons.

The neurons that we identified as part of the canonical swim and preparatory networks were summarized on the standard leech ganglionic map (see Figure 7, bottom; Muller et al., 1981). We attempted to match our identified neurons with those previously characterized. Our data are in good agreement with previous work, but it is not necessarily clear that all neurons identified here truly correspond to neurons with the same labels in the literature (see Table 1).

The extent and nature of the variability in the neuronal responses are best seen in several representative neurons that are most easily identified. The AP and the Nut cell are quite easy to distinguish by eye, because they



Swim Network



Preparatory Network

are typically larger than their neighboring neurons and in characteristic locations. The AP cell is located in the anterior lateral packets and is the largest neuron in the table (ganglion, AP Entry; see Figure 4), which means that it has a high signal-to-noise ratio (SNR) because more pixels are collecting signal. The AP cell showed a consistent rapid depolarization during shortening (cross, AP Entry, see Figure 7), strong depolarizing responses to all P cell input (circle, AP Entry, Figure 7), and no consistent oscillation

during swimming (hexagon, AP Entry, Figure 7). The Nut cell is somewhat smaller than AP but is readily identified by its posterior medial position (ganglion, Nut Entry, Figure 7). The Nut cell variability is much higher than the AP cell, but this amount of variability is more typical of the rest of the neurons identified. The Nut cell usually shows a rapid depolarization during shortening but has some variability (cross, Nut Entry, Figure 7). The Nut cell was usually functionally inhibited by P cell stimulus (circle, Nut Entry, Figure 7) and typically oscillated with peak depolarization in the purple phase during swimming (hexagon, Nut Entry, Figure 7), but occasionally this oscillation was not seen (many line segments do not show significant coherence)—possibly because of true functional variability or because the oscillation was sometimes too subtle for detection in the VSD signal.

Cell 208 is an unpaired neuron that was found in every animal (experiment indicators, 208 entry, Figure 7). One reason that this neuron was found consistently is that it is located in the middle of the ganglion, making it much more likely that it would be in the imaging field of view and on the

Figure 7: Atlas of neurons in canonical swim oscillator and preparatory networks. The identified neurons were given numbers based on their locations and best match to other neurons in the literature. For each entry, the experiment indicator (eight small colored circles, left), the ganglion position (ganglion), the shortening factors (cross), the local-bend coherence (circle), and the swim coherence (hexagon) are shown. The experiment indicator describes the animal in which the neuron was identified; a full circle means that both bilateral pairs of neurons were present in the ganglion. A semicircle means only one bilateral pair was seen or the neuron was not paired, and an empty circle means the neuron was not identified in that particular ganglion. The colors of the experiment indicators are the same colors as the ganglion position ROIs. The position ROIs (ganglion) show the position and size features for each identified cell, and the across-animal variability can be seen by the distribution of ROIs. The shortening factors (cross), local-bend coherence (circle), and swim coherence (hexagon) summarize the activity of each cell during the behaviors (see Figure 3) and show the within- and across-animal variability. For each of these plots, a single connected line or polygon corresponds to a single component across multiple trials. Each corner of the polygon is the feature coefficients from a single trial of a behavior, and the variance within a polygon shows the within-animal variability of the neuron. Each individual polygon is from a different animal, so the variance across polygons shows the across-animal variability. The neurons boxed in yellow are part of the preparatory network, and neurons that did not show a significant swim oscillation have an asterisk above the swim coherence hexagon. The entries in the table are organized based on their anatomical position. The canonical networks are summarized in the maps below the table. The identified neurons that oscillate with swimming are grouped into one of six phases and colored by phase (left). The neurons in the preparatory network are shown on the right.

Table 1: Consolidation with Previous Work.

Neuron ID	Match Reason ^a (Position is always a factor)	Reference	Possible Alternatives
CV	Swim phase aligned with DP -60	Briggman & Kristan, 2006; Brodfuehrer, McCormick, Tapyrik, Albano, & Graybeal, 2008	
255	Swim phase aligned with DP and 208 0	Briggman & Kristan, 2006	
204	Swim phase but DP from different ganglion. +60	Weeks & Kristan, 1978	Phase is unclear
208	Swim phase with DP -60	Briggman & Kristan, 2006; Taylor, Cottrel, Kleinfeld, & Kristan, 2003; Friesen, 1989a, 1989b; Nusbaum, Friesen, Kristan, & Pearce, 1987	
60	Swim phase with 208 -180	Friesen, 1989a, 1989b; Nusbaum et al., 1987	
61	Swim phase with 208 0	Nusbaum, 1986	
161	P cell input	Lockery & Kristan, 1990	162—similar position/input
212	P cell input	Lockery & Kristan, 1990	
153	180 oscillation	www.people.virginia.edu/~wof/leech.html	
151	151 = NS doublet oscillation	Rodriguez, Alvarez, & Szczupak, 2012	

Notes: The neurons previously identified in the literature were compared and consolidated with this work. This is an attempt to keep the nomenclature consistent, and this table summarizes the reasons and references for previously identified neurons.

^aPosition is always a factor.

correct side of the ganglion (ganglion, 208 entry, Figure 7). It did not show consistent patterns during shortening (cross, 208 entry, Figure 7), but it did receive consistent functional inhibition from P cell stimulation (circle, 208 entry, Figure 7) as well as showing consistent oscillations during swimming centered on the red phase (hexagon, 208 entry, Figure 7).

The identified neurons that proved most difficult to detect in this data set were cells 204, S, 206, and 203. These neurons are very small and are consequently captured by fewer pixels, giving them lower signal-to-noise ratios. We were initially surprised that the S cell was not obvious in our data set because it has large overshooting action potentials. However, the S cell action potentials are unusually brief, lasting only a millisecond. By comparison, the Leydig cell action potential lasts 30 milliseconds and can be easily seen in the VSD recordings (see Figure 1C, bottom trace). Because they proved

to be difficult to detect, cells 204, S, 206, and 203 are the least stringent in this data set and not easily distinguished by algorithms. However, on close visual inspection of the features and distance space, enough information was extracted that these neurons were more confidently identified.

Cells 263 and 264 are nearly indistinguishable from each other because their features were similar in all eight dimensions: both were in nearly the same position, both had a small factor 2 during shortening, and they both oscillated in the green phase during swimming (compare entries 263 and 264 in Figure 7). Differences in the local bend response were used as the distinguishing dimension between these two neurons. However, both 264 and 263 appeared to receive small input from some P cell activations, and so it is possible that some of these instances are cross-classified. This feature set distinguished 263 and 264 from others, but more features may be needed to appropriately distinguish them from each other.

The shortening feature showed a large amount of variability within and across animals, but some neurons had consistent patterns during shortening. The AP cell, for instance, showed a clear, rapid depolarization during the shortening behavior, along with several other neurons in the preparatory network (discussed below). A few neurons, such as 57, 261, and 262, showed fairly consistent hyperpolarizing responses during shortening (cross, Figure 7).

The local bend feature was measured by activating a sensory P cell. However, there are four different P cells and local bend trials were all lumped together, even though different P cells are activated within and across animals. Some neurons, such as 208, Nut, and AE, had purely inhibitory responses to all P cell stimulations, while other neurons, such as AP, 212, 264, and 155, had only excitatory responses. Other neurons, such as 57, 203, 152, and 153 had both excitatory (pink phase) and inhibitory (cyan phase) responses.

2.6 Targeted Experiments Using the Activity Atlas. Beyond assisting with the generation of hypotheses about homologous cells, roles, and circuits, an important and promising direction for use of the ICM is to generate guidance on experimentation, including the triaging of efforts to confirm hypotheses. The swim oscillations of neurons were the most robust feature, and many of the swim oscillators can be identified across animals using only position and phase of the oscillation. We performed several electrophysiological verification studies to validate the neurons identified through VSD recordings and show that these neurons can be identified and targeted in novel animals (see Figure 8). We first imaged VSD signals during swimming and created a swim activity map in real time using the ICM (see Figures 8A, 8C, 8E; Frady & Kristan, 2015). We used this activity map to target specific swim oscillators with intracellular electrodes and recorded their membrane potentials during a bout of swimming (see Figures 8B, 8D, and 8F). We performed the same coherence analysis on the intracellular recordings, and

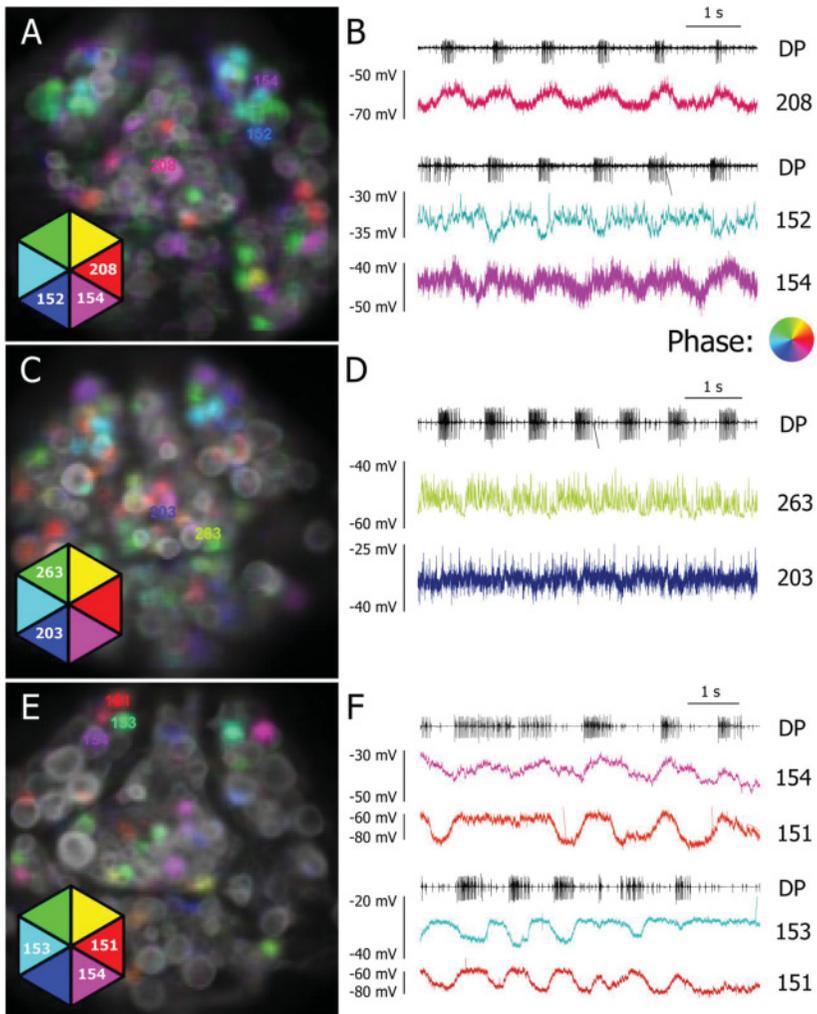


Figure 8: Electrophysiological verification of neurons that have oscillating membrane potentials during swimming. (A) An activity map was generated from a VSD imaging trial of swimming, and three neurons (208, 152, 154) were identified and targeted for electrophysiological recordings. The phase diagram indicates expected phases of targeted neurons. (B) Swimming was activated while intracellular recordings were made to targeted neurons. The coherence of the intracellular recordings was computed, and the traces are colored by the coherence phase (indicated by circle) in the same fashion as the VSD recordings. (C–F) The same verification experiment was carried out in two more animals targeting different swim oscillators. The electrophysiological recordings of these sample neurons are in agreement with the phases predicted by the VSD mapping experiments.

all seven of the unique neurons targeted oscillate at the same phase as predicted by the VSD mapping experiments (phase indicated by color). Cell 153 is the easiest of the swim oscillators to identify, because it has a large 10–20 mV oscillation that is consistently in the antiphase of the DP swim burst (hexagon, 153 entry, Figure 7; Figure 8F). Other neurons, like 152 (see Figure 8B) and 203 (see Figure 8D), have oscillations of only a few millivolts.

We noticed that many of the neurons shared a common set of responses to stimulation: rapid onset excitation during shortening (positive shortening factor 2, cross, Figure 7) and excitation by P cell input (pink phase, circle, Figure 7). The AP cell was the prime exemplar of this feature combination (the cross and circle plots for the AP entry, Figure 7), and many neurons shared this type of activity (indicated by yellow-boxed entries, Figure 7). These feature values meant that these neurons were activated by both local sensory activity (the P cells) and input from distant stimuli that elicited the shortening response. We measured the response latencies of each neuron, visually identified as a sudden change in the VSD traces in response to shortening or swimming stimuli. We saw that many of these neurons showed rapid-onset responses to all stimuli, regardless of which DP nerve was stimulated or which behavior resulted. Thus, many of the neurons showed a rapid sensory-driven response from both local and distant stimuli that is independent of the behavioral response. This result suggested that these neurons are involved in preparing the animal for the rapid execution of a behavior; we hypothesized that these neurons form a canonical “preparatory network” (yellow-boxed entries, Figure 7).

To test this preparatory network hypothesis, we stimulated different DP nerves and plotted the responses to stimuli at different distances from the recording site (see Figures 9A and 9B). The ganglion imaged was always ganglion 10; the stimulus sites were sorted by their distance from ganglion 10, either anteriorly (ganglia 3 and 7) or posteriorly (ganglia 11, 14, 17, and Tail Brain). As the stimulus distance increased away from ganglion 10, the response latency increased (see Figure 9). We measured the responses of the preparatory neurons as well as other neurons not in the preparatory network but which were activated during shortening or swimming behaviors. Three examples of neurons within the preparatory network (see Figure 9A) and three that are not (see Figure 9B) were colored based on their response delays (black circles) according to the color bar above each plot. From these examples, the preparatory network’s neurons appeared to have shorter response latencies than the other neurons.

We visualized the preparatory network by making activity maps of the response latencies during both shortening and swimming behaviors (see Figure 9C). Neurons with response latencies less than 300 ms were colored based on their response latency using the same color code as in Figures 9A and 9B. The activity maps showed several example trials with different behaviors and different stimulated ganglia, in which stimulation of ganglia 3 and 7 elicited shortening and stimulation of 11, 14, 17, and the tail brain elicited swimming. The activity maps within the same animal showed

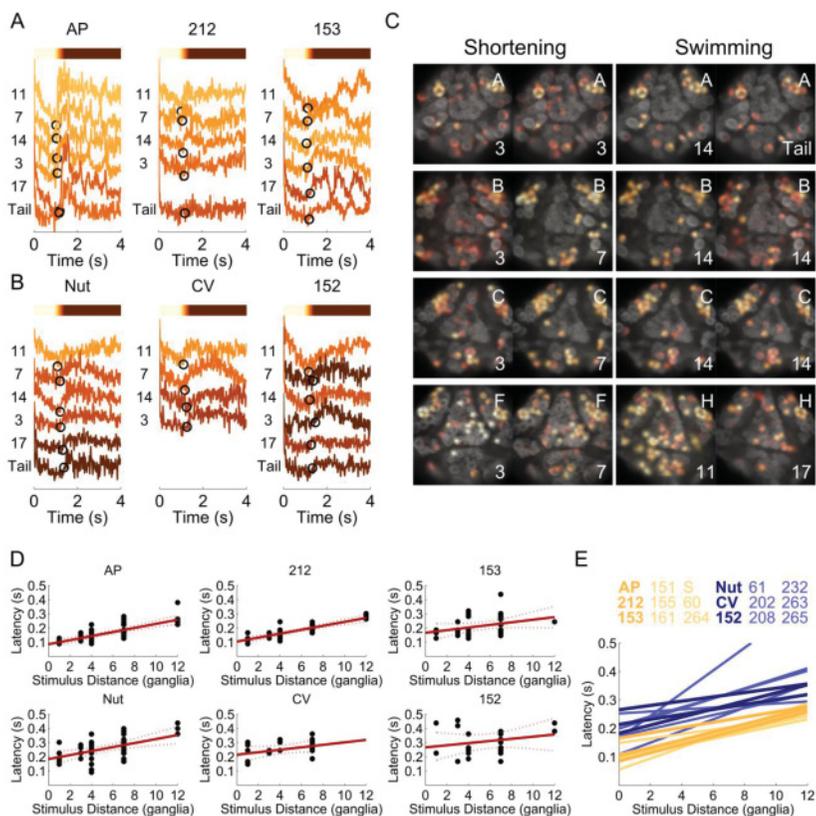


Figure 9: The preparatory network shows rapid responses to the stimulus regardless of behavior. (A, B) Example VSD traces from three identified neurons within the preparatory network (A) and three that are not (B). Each trace is sorted based on the distance the stimulus was from ganglion 10. The number of the ganglion stimulated is shown to the left of each trace. The response latency is shown as a black circle on each trace, which determined the color of the trace (the color bar above the traces indicates the time/color relationship). (C) Activity maps of the response delays for each neuron for several examples of shortening and swimming. Each component was colored by the response delay from one trial. Examples are shown for both shortening and swimming and for different stimulated ganglia. The letters at the top right indicate which ganglion was stimulated, and the numbers to the bottom right indicate which ganglion was stimulated. (D) The response latency of the example neurons from all experiments was plotted against the stimulus distance from the imaged ganglion. The red line is a linear fit to the response latency. (E) The linear fits of the response latency for all nine neurons in the preparatory network were plotted in yellow. The fits of the example neurons that are not part of the preparatory network were plotted in blue.

remarkable consistency regardless of the behavior evoked or the stimulus location, suggesting that the preparatory network is consistently activated by all stimuli that produced any response.

To quantify these impressions of the preparatory network timing, we pooled all data across experiments for each identified neuron and computed linear fits of the response latencies relative to the ganglion stimulated (see Figure 9D). We then plotted the fit line of the response latency versus distance for every neuron in the preparatory network as well as other example neurons (see Figure 9E). This analysis showed that all nine of the neurons in the preparatory network shared a consistent more rapid response than other neurons.

One neuron showed a rapid onset to shortening as well as input from P cells (cell 61, Figure 7), suggesting that this neuron was part of the preparatory network. However, on closer inspection, this neuron did not have responses as rapid as those of the other neurons in the preparatory network, and so we excluded it. Neurons 151 and 153 did show these rapid responses to all sensory stimuli; however, unlike other neurons in the preparatory network, these neurons occasionally showed hyperpolarizing responses to P cell stimulation.

The S cell is a prime candidate for mediating the rapid response of the preparatory network. Each ganglion contains a single unpaired S cell, which forms strong electrical connections to the S cells in the two adjoining ganglia. This connectivity pattern creates a chain of electrically coupled neurons that extends through the entire animal. The S cell network is reminiscent of a giant fiber system or a fast-conducting system (Mistick, 1978), but its causal role in any behavior is unclear (Sahley, Modney, Boulis, & Muller, 1994) and activation of the S cell network does not elicit any particular behavior (Shaw & Kristan, 1999). However, this rapid network may be useful in putting the nervous system and the muscles in a ready state so that one of multiple behaviors can be rapidly executed. From our data, we found the S cell was activated earlier than the other preparatory neurons, so we wanted to test whether the S cell excited other preparatory neurons.

The challenge of this experiment is that the neurons need to be appropriately identified. The S cell can be identified from its unique action-potential characteristics, and an electrophysiological recording was sufficient. Utilizing the activity atlas (see Figure 7), we can find that several preparatory neurons can be well isolated using their positions and the fact that they respond to P cell stimulation, which allowed us to carry out a simple experiment to address the connectivity of the preparatory network. We electrophysiologically activated a P cell (which can be identified by eye) and monitored voltage-sensitive dye activity (see Figures 10A to 10C) to identify preparatory neurons AP, 155, and 153 (see Figures 10A to 10C), aided by the real-time generation of an activity map (see Figure 10C). Injecting pulses of current into the S cell to produce action potentials showed a clear one-for-one EPSP in the AP cell for each S cell spike (see Figure 10D). S cell

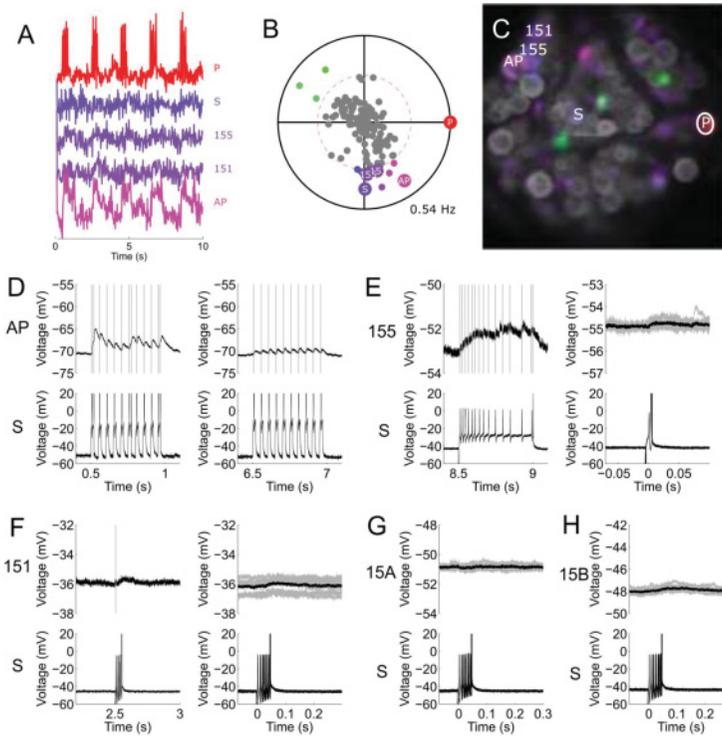


Figure 10: Electrophysiological validation of predicted S cell connections. (A) To identify preparatory network neurons, a P cell was activated with a burst of spikes every 2 seconds and the ganglion was monitored with the VSD. Preparatory neurons S, 155, 153, and AP were identified due to their size, position, and the fact that they received P cell input. (B) The coherence of the components was calculated to quickly identify P cell followers (indicated as violet/purple dots). (C) The activity map highlighted neurons that respond to the P cell stimuli, allowing rapid identification of preparatory neurons of interest. The violet and blue neurons were of special interest; the locations of four of them (AP, 153, 155, S) are indicated. This map was used to target neurons for electrophysiological recordings. (D) The S cell and AP cell were targeted with intracellular electrodes; the AP cell was slightly hyperpolarized to below spike threshold. Two example traces show that every S cell spike produced, one-for-one, large EPSPs in the AP cell. (E) Neuron 155 was then targeted for a paired physiology experiment. Exciting the S cell with current showed a clear excitatory response in the postsynaptic 155, but they were not one-for-one with the S cell spikes and the synaptic potentials were small (about 0.5 mV). (F) Another preparatory network neuron, 153. A small excitatory response can be seen. (G, H) Two other neurons near 153 and 155 were targeted (the locations of their somata is shown as white circles in panel C). These neurons were unidentified but not likely part of the preparatory network. There was very little to no response from S cell stimulation in these neighboring neurons.

spikes produced clear excitatory effects on neurons 155 and 153, but these synaptic potentials were smaller than the EPSPs made by the S cell onto the AP cell, and the one-for-one nature of the S cell inputs onto cells 153 and 155 was less clear (see Figures 10E and 10F). We also targeted two other neighboring neurons, which were not preparatory neurons, and these neurons showed no significant response to S cell activation (see Figures 10G and 10H).

3 Discussion

We used voltage-sensitive dye imaging and electrophysiology (see Figure 1) to identify and map the activity of dozens of neurons in a scalable fashion using information extracted from the multifunctional nature of leech neuronal activity during different behaviors (see Figure 3; Briggman & Kristan, 2006). To do this characterization, we developed computational tools to automatically extract signals from imaging data (see Figure 2; Frady & Kristan, 2015), rapidly create activity maps for visualization (see Figure 3), and build a rich medium-dimensional feature space that summarizes large-scale data and can be used to stitch experiments together (see Figures 4–6). We used this information to identify dozens of neurons in the leech swimming and preparatory networks (see Figure 7), then used electrophysiological recordings to verify the phasing of neurons in the swimming circuit (see Figure 8) and test for connections among neurons that were predicted from the analyses of the preparatory network (see Figures 9 and 10). An interesting feature of mapping leech neurons is that it has a particular one-for-one homology across animals and a stereotypical structure with inter-subject variability, and this type of homology is likely true in other systems (Marder & Bucher, 2007; Bargmann & Avery, 1995; Ng et al., 2009). Below, we discuss the computational microscope—how we used it and how it can be used in other systems—and the significance of the electrophysiological findings.

3.1 Creating an Activity Atlas Using the Computational Microscope.

The framework for activity mapping with the computational microscope consists of three stages of computational algorithms that are refined by user feedback at each stage. In stage 1, we used PCA and ICA (Hill et al., 2010) to extract clusters of pixels as the electrical activity of individual neurons. The PCA-ICA extraction is general for any spatiotemporal data and has previously been demonstrated with calcium imaging (Mukamel et al., 2009), other voltage dyes (Hill et al., 2010), and EEG (Delorme & Makeig, 2004). This PCA-ICA extraction method is easily extended to 3D volume imaging, such as light-sheet (Ahrens, Orger, Robson, Li, & Keller, 2013) or fMRI (Beckmann & Smith, 2004; Varoquaux et al., 2010; Khullar et al., 2011), but larger data sets will require more powerful compute clusters, better memory handling, and parallelized execution (Freeman et al., 2014). Further, these

types of high-dimensional functional imaging data sets will require more sophisticated and standardized formats for extracting the data from raw images.

At stage 2, simple feature detectors were used to characterize the activity of individual neurons with one or two numbers. These low-dimensional representations can be used to create activity maps by false-coloring the neurons and serve as a visualization of large-scale neural activity during individual trials. The feature detectors used here are generalizable to many other types of neural signals, and the ICM includes similar built-in feature detectors (coherence and PCA; Frady & Kristan, 2015). This stage, however, can be reconfigured to be more specific for particular experiments or data sets, and custom feature detectors can be built, incorporated, and visualized using the ICM.

This stage is the primary focus for generalizing this framework to other systems, and creating a rich feature space is the key to creating a useful atlas and registering new data to this atlas. The strategy of finding a unique feature for every neuron is problematic in high-dimensional recordings (Guyan & Elisseeff, 2003). Adding many features together to build a high-dimensional space is useful only if these features contain independent information, and supervised algorithms for feature selection have mainly been developed in the context of a small number of classes (Peng, Long, & Ding, 2005; Rodriguez-Lujan, Huerta, Elkan, & Cruz, 2010; Brown, Pocock, Zhao, & Luján, 2012). If there are many features that do not provide information and appear as noise in identifying most neurons, then these features greatly hurt the richness of the feature space, especially when there are large numbers of potential competitors (see Figure 4).

The statistical analysis for many typical metrics, such as mutual information, become problematic with data sets that have high-dimensional features and classes. The curse of dimensionality requires tremendous amounts of data to overcome the combinatorics of the high-dimensional data sets to accurately compute these values. This data set was too high-dimensional to rely on these classic statistical metrics, and so we turned to semisupervised approaches where we directly visualized and interacted with the machine learning algorithms.

In stage 3, several anatomical and functional features from different behaviors were combined and used for identifying neurons across animals. By iteratively visualizing the medium-dimensional space formed by combining several low-dimensional features and manually refining homologous landmarks across animals, the computational microscope can build a canonical atlas of neural activity for future reference (Frady & Kristan, 2015). The canonical matching algorithms can be extended beyond the one-to-one homologous matching illustrated here. This particular data called for a one-to-one correspondence across animals; a bipartite graph matching algorithm was used to form suggestions, and the warping algorithms were designed to search for one-to-one correspondences. By using an alternative

one-to-many clustering algorithm, such as k-means or hierarchical clustering (Seber, 1984), the framework is generalizable to neuron types (Roux, Stark, Sjulson, & Buzsáki, 2014).

In many studies, the analysis of fMRI data relies on careful anatomical alignment to register the same functional regions to a common set of voxels, and many advanced registration algorithms have been used to warp brains across subjects to a template brain based on anatomy (Dale, Fischl, & Sereno, 1999; Fischl, Sereno, & Dale, 1999; Fischl, Sereno, Tootell, & Dale, 1999; Van Essen, 2005; Yeo et al., 2011). Several studies, however, have pointed out that intersubject variability leads to misregistration of functional regions (Rajkowska & Goldman-Rakic, 1995; Thirion et al., 2007; Ng et al., 2009), and that functional regions are not localized even if anatomical markers are perfectly aligned (Brett, Johnsrude, & Owen, 2002; Crivello et al., 2002). In the leech ganglion, this type of anatomically based spatial warping did not correctly align every individual neuron, and we suggest that the alignment of brain regions suffers a similar type of inter subject variability. This means that anatomically based registration, no matter how sophisticated the warping algorithm, will not be capable of perfectly aligning functional brain regions simply because there is not enough information in the anatomy to identify the brain region.

To characterize and account for this variability, we extracted functional information from a standard set of behaviors and abstracted the anatomical and functional information to a high-dimensional graph-theoretic problem, where anatomy and function are treated mathematically the same. Similar efforts are currently being pursued for creating brain atlases with fMRI (Dosenbach et al., 2007; Varoquaux, Gramfort, Pedregosa, Michel, & Thirion, 2011), and utilizing functional information (i.e. more features), like the default network (Khullar et al., 2011) or functional connectomes (Yeo et al., 2011; Varoquaux & Craddock, 2013; Phlypo, Thirion, & Varoquaux, 2014) are promising steps toward extracting the information needed to identify homologous functional regions.

Our results suggest that combining several functional feature descriptors from several distinct modalities (i.e., forming a rich medium-dimensional feature space) will greatly aid in the identification of homologous functional brain regions and that specific brain regions can be targeted across subjects by first mapping responses to standardized behaviors or stimuli (Poldrack, Halchenko, & Hanson, 2009; Poldrack, 2010). Careful consideration to the features must be given in order to ensure that the medium-dimensional feature space maintains its richness. Key to maintaining richness is to utilize diverse and natural behaviors that engage the entire brain or neural circuit in complex patterns rather than relying on a stimulus that is targeted to a particular neuron or brain region. Combining information from many independent behaviors is an optimal strategy to explore the high-dimensional space spanned by complex systems; machine learning with visualizations can be used to understand these data.

Removal of intersubject differences through warping has been shown to decrease the resolution of specific brain regions (Ng et al., 2009). The weighted correspondence minimization algorithm is a simple version of the many fMRI warping algorithms (Frackowiak et al., 2003) but generalized to higher dimensions. However, here we use warping only to aid in the mapping and registration of individual neurons to the canonical atlas instead of to average data into a canonical response. Combining data across subjects requires new algorithms to appropriately interpret canonical intersubject responses, and algorithms such as probabilistic principal components analysis show promise as new tools to find insights from intersubject data (Kapoor, Frady, Jegelka, Kristan, & Horvitz, 2015).

3.2 Relating Imaging to Behaviors. Because we were searching for the two-thirds of neurons that have not been previously characterized, we largely ignored previously identified neurons (i.e., the colored neurons in Figure 1B). We did, however, include previously identified neurons if they showed behavior-related responses that had not been reported. Among the ignored neurons were the mechanosensory neurons (T, P, and N cells) and the Retzius cells. The mechanosensory neurons were ignored because they were not active during any of the behaviors (except when a P cell was electrically stimulated to produce local bending) and the Retzius cells because their activity in behaviors is well known (Loer & Kristan 1989), and we routinely removed their somata because they often cover the somata of other, smaller neurons that we wanted to study. We paid particular attention to the neurons involved in two behavioral responses, swimming (Kristan et al., 2005), and the preparatory response, a behavior that has been mentioned previously in passing (Esch, Mesce, & Kristan, 2002) but has never been studied directly.

Our approach also proved to be robust to the absence of particular neurons. Based on a great deal of previous work (Kristan et al., 2005), we believe that all leech ganglia—across animals and across segments within each animal—contain the same allotment of neurons. However, because of variability across preparations, for reasons both biological (e.g., variations in cell body locations) and experimental (e.g., variations in the dissection or dye loading of the ganglion), VSD data from any single experiment contained information only from a subset of all neurons in that ganglion. Smaller neurons in the ganglion are also likely to be undersampled because smaller neurons have lower SNR in the optical signals.

Despite this variability, we were able to identify neurons uniquely even when they were sampled only a few times. This is because we relied heavily on visualization and interaction with the machine learning algorithms. For instance, neuron 232 was a challenging neuron to find, but because data from several animals can be compared, there was clear evidence for a match across a few animals. The algorithm suggested other candidates but

could not automatically overcome the full intersubject variability, and this closest match was not always the best. This was because of the low signal-to-noise ratio of this neuron in some animals and some trials. However, on close visual inspection and searching through alternative suggestions, we were able to identify the neuron in seven of the eight animals (see Figure 7).

3.2.1 Swimming Network. Many of the neurons identified as part of the swim network were characterized previously, and we attempted to best match the neuron reported here to those previously found in the literature (see Table 1). This was done by comparing previously reported position and phase information—relative either to motor neuronal impulse burst in a nerve recording or another identified swim oscillator—to the activity atlas, and finding the best match. About one-third of the swim oscillators have been recorded from previously, and the other two-thirds of the network are novel neurons. Previous VSD imaging work (Briggman et al., 2005; Briggman & Kristan, 2006) produced qualitatively similar swim activity maps, but used an older generation of VSD that does not report phase information as precisely. Many of the neurons identified here were likely the same as those in the previous VSD maps.

With experience, visualizations become an essential tool for guiding experiments (Walter et al., 2010). After an activity atlas is created, a few simple visualizations often proved the most efficient way of identifying a particular neuron of interest. For instance, virtually all of the swim oscillators can be identified based on a single swim activity map, because the phases of the swim oscillators, along with soma position, provide enough information to identify these neurons. This characterization proved useful in performing targeted electrophysiological experiments to novel swim neurons. By simply creating an activity map and comparing it with the activity atlas (see Figure 8), neurons of interest can be easily targeted. For more complex systems, the algorithms combined with visualization will prove to be a powerful tool in rapidly identifying and targeting networks of neurons.

3.2.2 Preparatory Network. The preparatory network stood out because it was activated at short latency by every sensory stimulus we presented, including the medial or lateral P cell to trigger local bending, or stimulation of various DP nerves to elicit shortening and swimming. This appeared as correlations in the feature space from two different behaviors. Since shortening factor 2 and local bend excitation are both related to the activity of the preparatory network, they transmit redundant information about the identities of the neurons. The shortening feature is the least rich of the functional features (see Figures 4B), both because of this redundancy affecting factor 2 but also because the variability across animals was much higher than within animals (see Figure 5A). These are reasons that the weighted

correspondence minimization algorithm quashed the shortening features relative to the other features (see Figure 6E). However, these correlations were a clue to the hidden structure of the preparatory network. Because these correlations existed in these independent behaviors, we investigated further and discovered that the correlations were due to an underlying causal phenomenon: preparatory activity.

We investigated the responses of the preparatory neurons in detail during shortening and swimming (see Figure 9). A previous study found that motor neurons were activated about 200 msec before the leech nervous system made the choice to either swim or to crawl (Briggman et al., 2005). The rapid and consistent responses of these motor neurons led to the speculation that leeches have a “do something” network that readies the body musculature for a more rapid induction of a behavioral response (Friesen & Kristan, 2007). (With no muscle contractions, a leech’s body is flaccid; it must co-contract antagonistic muscles to provide a hydrostatic skeleton to be able to make whole-body movements; Kristan et al., 2000.)

The S cell has many properties consistent with a role for it in the preparatory network: it responds to touch anywhere on the body (Mistick, 1974); it is unpaired, but its axon runs the full length of the nervous system (Bagnoli, Brunelli, & Magni, 1975); it makes strong electrical junctions with motor neurons in every segmental ganglion (Gardner-Medwin, Jansen, Taxt, 1973; Magni & Pellegrino, 1978); despite the motor neuron connection, its activity does not cause significant muscle tension, but does make the behavioral responses faster and stronger (Shaw & Kristan, 1999); and it is required for plasticity in the shortening response (Sahley et al., 1994). Based on these properties, we hypothesized that the S cell activated the remaining neurons in the preparatory network. Using imaging, we were able to identify many preparatory neurons and then verified that the S cell is functionally coupled to the AP cell and other preparatory neurons (155, 153) using computationally guided electrophysiology (see Figure 10).

Thus, the preparatory network of the leech functions to ready the animal for movements by bringing both motor neurons and interneurons closer to their firing threshold, getting the system ready for the more precise and complex movements that occur when the pattern generators for other behaviors become active. This function is similar to an interaction between the postural and limb control systems: postural systems are activated well before arm or leg movements begin so that the animal’s balance is not upset by movements of the arms and legs (Massion, Alexandrov, & Frolov, 2004; Cardo & Gurfinkel, 2004). Similar sorts of early, broad activation of sensory systems is seen, too, in response to alerting or warning stimuli, to change the gain of the responses—either up or down—to subsequent stimuli (Weinbach & Henik, 2014). Hence, preparatory activity before the initiation of a directed behavior appears to be a common processing element of many nervous systems.

Appendix: Table of Defined Terms

Definitions	
Rich medium-dimensional feature space	A feature space that carries a lot of information with few dimensions
One-hot	A high-dimensional vector with zero values at every index except one
Curse of dimensionality	To obtain a statistically sound and reliable result, the amount of data needed to support the result grows exponentially with the dimensionality
Trial	A single VSD recording of an individual behavior
Experiment or animal	The set of all VSD recordings in an individual animal
Source	The time series that extends through all trials extracted by PCA or ICA
Trace	The activity of a neuron during a single trial of a behavior
Map	The spatial map indicating the location of a component extracted by PCA or ICA
Neurons	An independent component visually verified as a neuron based on inspecting its map and source
Artifacts	An independent component visually rejected as a neuron
Canonical	Identified features, neurons or networks that are consistent across animals
VSD	Voltage-sensitive dye
GUI	Graphical user interface
ROI	Region of interest
PCA	Principal components analysis
ICA	Independent component analysis
ctICA	Concatenated-trial independent components analysis
ICM	Imaging computational microscope
DP Nerve	Dorsal-posterior nerve that is recorded extracellularly
WCM	Weighted correspondence minimization
EPSP	Excitatory postsynaptic potential

References

- Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M., & Keller, P. J. (2013). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods*, *10*, 413–420.
- Alivisatos, A. P., Chun, M., Church, G. M., Greenspan, R. J., Roukes, M. L., & Yuste, R. (2012). The brain activity map project and the challenge of functional connectomics. *Neuron*, *74*, 970–974.
- Bagnoli, P., Brunelli, M., & Magni, F. (1975). The neuron of the fast conducting system in *Hirudo medicinalis*: Identification and synaptic connections with primary afferent neurons. *Arch. Ital. Biol.*, *113*, 21–43.
- Bargmann, C. I., & Avery, L. (1995). Laser killing of cells in *Caenorhabditis elegans*. *Methods in Cell Biology*, *48*, 225–250.

- Bargmann, C. I., & Marder, E. (2013). From the connectome to brain function. *Nat. Meth.*, *10*(6), 483–490.
- Bargmann, C. I., & Newsome, W. T. (2014). The brain research through advancing innovative neurotechnologies (BRAIN) initiative and neurology. *JAMA Neurol.*, *71*, 675–676.
- Bargmann, C. I., Newsome, W. T., Anderson, A., Brown, E., Deisseroth, K., Donoghue, J., . . . Sanes, J. (2014). BRAIN 2025: A scientific vision. In *Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Working Group Report to the Advisory Committee to the Director, NIH*. Bethesda, MD: National Institutes of Health. <http://www.nih.gov/science/brain/2025/>
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Beckmann, C. F., & Smith, S. M. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans. Med. Imaging*, *23*, 137–152.
- Bellman, R. E. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.
- Bermingham, M. L., Pong-Wong, R., Spiliopoulou, A., Hayward, C., Rudan, I., Campbell, H., . . . Haley, C. S. (2015). Application of high-dimensional feature selection: Evaluation for genomic prediction in man. *Scientific Reports*, *5*, 10312. doi:10.1038/srep10312
- Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S., & Mitra, P. P. (2010). Chronux: A platform for analyzing neural signals. *J. Neurosci. Meth.*, *192*, 146–151.
- Brett, M., Johnsrude, I. S., & Owen, A. M. (2002). The problem of functional localization in the human brain. *Nature Reviews Neuroscience*, *3*(3), 243–249.
- Brodofuehrer, P. D., McCormick, K., Tapyrik, L., Albano, A. M., & Graybeal, C. (2008). Activation of two forms of locomotion by a previously identified trigger interneuron for swimming in the medicinal leech. *Invert. Neurosci.*, *8*, 31–39.
- Briggman, K. L., Abarbanel, H.D.I., & Kristan, W. B. (2005). Optical Imaging of neuronal populations during decision-making. *Science*, *307*, 896–901.
- Briggman, K. L., Helmstaedter, M., & Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, *471*(7337), 183–188. doi:10.1038/nature09818
- Briggman, K. L., & Kristan, W. B. (2006). Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *J. Neurosci.*, *26*, 10925–10933.
- Brown, G., Pocock, A., Zhao, M. J., & Luján, M. (2012). Conditional likelihood maximisation: A unifying framework for information theoretic feature selection. *Journal of Machine Learning Research*, *13*(1), 27–66.
- Cardo, P. J., & Gurfinkel, V. (2004). Motor coordination can be fully understood only by studying complex movements. *Prog. Brain. Res.*, *143*, 29–38.
- Chalfie, M., Sulston, J. E., White, J. G., Southgate, E., Thomson, J. N., & Brenner, S. (1985). The neural circuit for touch sensitivity in *Caenorhabditis elegans*. *Journal of Neuroscience*, *5*(4), 956–964. doi:3981252
- Crivello, F., Schormann, T., Tzourio-Mazoyer, N., Roland, P. E., Zilles, K., & Mazoyer, B. M. (2002). Comparison of spatial normalization procedures and their impact on functional maps. *Human Brain Mapping*, *16*(4), 228–250.

- Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. *NeuroImage*, *9*, 1791-194.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics. *J. Neurosci. Meth.*, *134*, 9–21.
- Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., . . . Petersen, S. E. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proc. Natl. Acad. Sci. USA*, *104*, 11073–11078.
- Esch, T., Mesce, K. A., & Kristan, W. B. (2002). Evidence for sequential decision-making in the medicinal leech. *J. Neurosci.*, *22*, 11045–11054.
- Evangelidis, G. D., & Psarakis, E. Z. (2008). Parametric image alignment using enhanced correlation coefficient maximization. *IEEE Trans. on PAMI*, *30*, 1858–1865.
- Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, *9*, 195–207.
- Fischl, B., Sereno, M. I., Tootell, R.B.H., & Dale, A. M. (1999). High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.*, *8*, 272–284.
- Frackowiak, R.S.J., Friston, K. J., Frith, C., Dolan, R., Price, C. J., Zeki, S., . . . Penny, W. D. (2003). *Human brain function* (Ed.), Orlando, FL: Academic Press.
- Frady, E. P., & Kristan, W. B. (2014). Computation with population codes. In D. Jaeger & R. Jung (Eds.), *Encyclopedia of computational neuroscience* (pp. 1–16). New York: Springer.
- Frady, E. P., & Kristan, W. B. (2015). *The imaging computational microscope*. arXiv: 1502.07009 [q-bio. NC]
- Freeman, J., Vladimirov, N., Kawashima, T., Mu, Y., Sofroniew, N. J., Bennett, D. V., . . . Ahrens, M. B. (2014). Mapping brain activity at scale with cluster computing. *Nature Methods*, *11*(9), 941–950.
- Friesen, W. O. (1989a). Neuronal control of leech swimming movements I. Inhibitory interactions between motor neurons. *J. Comp. Physiol. A*, *166*, 195–203.
- Friesen, W. O. (1989b). Neuronal control of leech swimming movements II. Motor neuron feedback to oscillator cells 115 and 28. *J. Comp. Physiol. A*, *166*, 205–215.
- Friesen, W. O., & Kristan, W. B. (2007). Leech locomotion: Swimming, crawling, and decisions. *Curr. Opin. Neurobiol.*, *17*, 704–711.
- Gall, F. J. (1833). The anatomy and physiology of the nervous system in general, and of the brain in particular, with observations upon the possibility of ascertaining the several intellectual and moral dispositions of man and animal, by the configuration of their heads. *American Journal of the Medical Sciences*, Southern Society for Clinical Investigation.
- Gardner-Medwin, A. R., Jansen, J. K. S., & Taxt, T. (1973). The “giant” axon of the leech. *Acta Physiol. Scand.*, *87*, 30A–31A.
- Gareth, J., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning: With applications in R*. New York: Springer. doi:10.1007/978-1-4614-7138-7
- Guyan, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *JMLR*, *3*, 1157–1182.

- Hill, E. S., Moore-Kochlacs, C., Vasireddi, S. K., Sejnowski, T. J., & Frost, W. N. (2010). Validation of independent component analysis for rapid spike sorting of optical recording data. *J. Neurophysiol.*, *104*, 3721–3731.
- Horvitz, E., & Kristan, W. B. (2009). Toward a computational microscope for neurobiology. In T. Hey, S. Tansley, & K. Tolle (Eds.), *The fourth paradigm: Data-intensive scientific discovery*. Redmond: Microsoft Research.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning with applications in R*. New York: Springer. <http://doi.org/10.1007/978-1-4614-7138-7>
- Kapoor, A., Frady, E. P., Jegelka, S., Kristan, W. B., & Horvitz, E. (2015). Inferring and learning from neuronal correspondences. arXiv:1501.05973 [q-bio.NC]
- Kasthuri, N., & Lichtman, J. W. (2010). Neurocartography. *Neuropsychopharmacology*, *35*(1), 342–343. doi:10.1038/npp.2009.138
- Khullar, S., Michael, A. M., Cahill, N. D., Kiehl, K. A., Pearlson, G., Baum, S. A., & Calhoun, V. D. (2011). ICA-fNORM: Spatial normalization of fMRI data using intrinsic group-ICA networks. *Front. Sys. Neurosci.*, *5*, 1–18.
- Kristan, W. B. (1982). Sensory and motor neurons responsible for the local bending response in leeches. *J. Exp. Biol.*, *96*, 161–180.
- Kristan, W. B., & Calabrese, R. L. (1976). Rhythmic swimming activity in neurones of the isolated nerve cord of the leech. *J. Exp. Biol.*, *65*, 643–668.
- Kristan, W. B., Calabrese, R. L., & Frieson, W. O. (2005). Neuronal control of leech behavior. *Prog. Neurobiol.*, *76*, 279–327.
- Kristan, W. B., Skalak, R., Wilson, R. J. A., Skierczynski, B. A., Murray J. A., Eisenhart, F. J., & Cacciatore, T. W. (2000). Biomechanics of hydroskeletons: Studies of crawling in the medicinal leech. In J. Winters & P. Crago (Eds.), *Biomechanics and neural control of movement* (pp. 206–220). New York: Springer-Verlag.
- Lewis, J. E., & Kristan, W. B. (1998). Representation of touch location by a population of leech sensory neurons. *Journal of Neurophysiology*, *80*(5), 2584–2592.
- Lockery, S. R., & Kristan, W. B. (1990). Distributed processing of sensory information in the leech. II. Identification of interneurons contributing to the local bending reflex. *Journal of Neuroscience*, *10*(6), 1816–1829.
- Loer, C. M., & Kristan, W. B. (1989). Central synaptic inputs to identified leech neurons determined by peripheral targets. *Science*, *244*, 64–66.
- Macagno, E. R. (1980). Number and distribution of neurons in leech segmental ganglia. *J. Comp. Neurol.*, *15*, 283–302.
- Macosko, E. Z., Pokala, N., Feinberg, E. H., Chalasani, S. H., Butcher, R. A., Clardy, J., & Bargmann, C. I. (2009). A hub-and-spoke circuit drives pheromone attraction and social behaviour in *C. elegans*. *Nature*, *458*(7242), 1171–1175. doi:10.1038/nature07886
- Magni, F., & Pellegrino, M. (1978). Neural mechanisms underlying the segmental and generalized cord shortening reflexes in the leech. *J. Comp. Physiol.*, *124*, 339–351.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, *503*, 78–84. <http://doi.org/10.1038/nature12742>
- Marder, E., & Bucher, D. (2007). Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annu. Rev. Physiol.*, *69*, 291–316.

- Massion, J., Alexandrov, A., & Frolov, A. (2004). Why and how are posture and movement correlated? *Prog. Brain Res.*, *143*, 13–27.
- McKay, B. D. (1981). *Practical graph isomorphism*. Congressus Numerantium.
- Miller, E. W., Lin, J. Y., Frady, E. P., Steinbach, P. A., Kristan, W. B., & Tsien, R. Y. (2012). Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires. *PNAS*, *109*, 2114–2119.
- Mistick, D. (1974). Rohde's fiber: A septate axon in the leech. *Brain Res.*, *74*, 342–348.
- Mistick, D. (1978). Neurons in the leech that facilitate an avoidance behavior following nearfield water disturbances. *J. Exp. Biol.*, *75*, 1–23.
- Mitra, P. P., & Bokil, H. (2007). *Observed brain dynamics*. New York: Oxford University Press.
- Morgan, J. L., & Lichtman, J. W. (2013). Why not connectomics? *Nature Methods*, *10*(6) 494–500. <http://doi.org/10.1038/nmeth.2480>
- Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiol.*, *283*, 53–77.
- Mukamel, E. A., Nimmerjahn, A., & Schnitzer, M. J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, *63*, 747–760.
- Muller, K. J., Nicholls, J. G., & Stent, G. S. (1981). *Neurobiology of the leech*. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory.
- Ng, B., Abugharbieh, R., & McKeown, M. J. (2009). Adverse effects of template-based warping on spatial fMRI analysis. In *Proc. SPIE* (p. 7262). Bellingham, WA: SPIE.
- Nusbaum, M. P. (1986) Synaptic basis of swim initiation in the leech III. Synaptic effects of serotonin-containing interneurons (cells 21 and 61) on swim CPG neurones (cells 18 and 208). *J. Exp. Biol.*, *122*, 303–321.
- Nusbaum, M. P., Friesen, W. O., Kristan Jr., W. B., & Pearce, R. A. (1987). Neural mechanisms generating the leech swimming rhythm: Swim-initiator neurons excite the network of swim oscillator neurons. *J. Comp. Physiol. A*, *161*, 355–366.
- Oh, S. W., Harris, J. A., Ng, L., Winslow, B., Cain, N., Mihalas, S., . . . Zeng, H. (2014). A mesoscale connectome of the mouse brain. *Nature*, *508*, 207–214.
- Olshausen, B. A., & Field, D. J. (1996). Natural image statistics and efficient coding. *Network: Computation in Neural Systems*, *7*, 333–339.
- Palmer, C. R., Barnett, M. N., Copado, S., Gardezy, F., & Kristan, W. B. (2014). Multiplexed modulation of behavioral choice. *J. Exp. Bio.*, *217*, 2963–2973.
- Penev, P. S. (2001). Redundancy and dimensionality reduction in sparse-distributed representations of natural objects in terms of their local features. In T. K. Leen, T. G. Dietterich, & V. Tresp., (Eds.), *Advances in neural information processing systems* (pp. 901–907). Cambridge, MA: MIT Press.
- Peng, H., Long, F., & Ding, C. (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *27*(8), 1226–1238.
- Phlypo, R., Thirion, B., & Varoquaux, G. (2014). Deriving a multi-subject functional-connectivity atlas to inform connectome estimation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014* (pp. 185–192). New York: Springer.
- Poldrack, R. A. (2010). Mapping mental function to brain structure: How can cognitive neuroimaging succeed? *Perspectives on Psychological Science*, *5*(6), 753–761. <http://doi.org/10.1177/1745691610388777>

- Poldrack, R. A., Halchenko, Y. O., & Hanson, S. J. (2009). Decoding the large-scale structure of brain function by classifying mental states across individuals. *Psychological Science*, *20*(11), 1364–1372. <http://doi.org/10.1111/j.1467-9280.2009.02460.x>
- Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2) 125–132. <http://doi.org/10.1038/35039062>
- Pouget, A., & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, *9*, 222–237. <http://doi.org/10.1162/jocn.1997.9.2.222>
- Rajkowska, G., & Goldman-Rakic, P. S. (1995). Cytoarchitectonic definition of prefrontal areas in the normal human cortex. II. Variability in locations of areas 9 and 46 and relationship to the Talairach coordinate system. *Cereb. Cortex*, *5*, 323–337.
- Rodriguez, M. J., Alvarez, R. J., & Szczupak, L. (2012). Effect of a nonspiking neuron on motor patterns of the leech. *Journal of Neurophysiology*, *107*, 1917–1924. doi:10.1152/jn.01070.2011
- Rodriguez-Lujan, I., Huerta, R., Elkan, C., & Cruz, C. S. (2010). Quadratic programming feature selection. *Journal of Machine Learning Research*, *11*, 1491–1516.
- Roux, L., Stark, E., Sjulson, L., & Buzsáki, G. (2014). In vivo optogenetic identification and manipulation of GABAergic interneuron subtypes. *Current Opinion in Neurobiology*, *26*, 88–95.
- Sahley, C. L., Modney, B. K., Boulis, N. M., & Muller, K. J. (1994). The S cell: An interneuron essential for sensitization and full dishabituation of leech shortening. *J. Neurosci.*, *14*, 6715–6721.
- Seber, G. A. F. (1984). *Multivariate observations*. Hoboken, NJ: Wiley.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Tech. J.*, *27*, 379–423, 623–656.
- Shaw, B. K., & Kristan, W. B. (1999). Relative roles of the S cell network and parallel interneuronal pathways in the whole-body shortening reflex of the medicinal leech. *J. Neurophysiol.*, *82*, 1114–1123.
- Taylor, A. L., Cottrel, G. W., Kleinfeld, D., & Kristan, W. B. Jr. (2003). Imaging reveals synaptic targets of a swim-terminating neuron in the leech CNS. *J. Neurosci.*, *23*(36), 11402–11410.
- Thirion, B., Pinel, P., Meriaux, S., Roche, A., Dehaene, S., & Poline, J. B. (2007). Analysis of a large fMRI cohort: Statistical and methodological issues for group analyses. *NeuroImage*, *35*, 105–120.
- Van Essen, D. C. (2005). A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex. *NeuroImage*, *28*, 635–662.
- Varoquaux, G., & Craddock, R. C. (2013). Learning and comparing functional connectomes across subjects. *NeuroImage*, *80*, 405–415.
- Varoquaux, G., Gramfort, A., Pedregosa, F., Michel, V., & Thirion, B. (2011). Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. *Inform. Proc. Med. Imag.*, *6801*, 562–573.
- Varoquaux, G., Sadaghiani, S., Pinel, P., Kleinschmidt, A., Poline, J. B., & Thirion, B. (2010). A group model for stable multi-subject ICA on fMRI datasets. *NeuroImage*, *51*, 288–299.

- Walter, T., Shattuck, D. W., Baldock, R., Bastin, M. E., Carpenter, A. E., Duce, S., . . . Heriche, J. K. (2010). Visualization of image data from cells to organisms. *Nature Meth.*, *7*(3).
- Weeks, J. C. (1982). Segmental specialization of a leech swim-initiating interneuron, cell 205. *J. Neurosci.*, *2*, 972–985.
- Weeks, J. C., & Kristan, W. B. (1978). Initiation, maintenance and modulation of swimming in the medicinal leech by the activity of a single neurone. *J. Exp. Biol.*, *77*, 71–88.
- Weinbach, N., & Henik, A. (2014). Alerting enhances attentional bias for salient stimuli: Evidence from a global/local processing task. *Cognition*, *133*, 414–419.
- White, J. G., Southgate, E., Thomson, J. N., & Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos. Trans. R. Soc. Lond. A*, *314*, 1–340.
- Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., . . . Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.*, *106*, 1125–1165.

Received May 4, 2015; accepted March 15, 2016.