State-dependent network interactions differentially gate sensory input at the motor and command neuron level in *Caenorhabditis elegans*

Zachary T. Cecere[1,2,3,6], Kathleen T. Quach[1,2,6], Eviatar Yemini[4,5], Javier J. How[1,2], Tatyana O. Sharpee[1,3], and Sreekanth H. Chalasani[1,2]*

[1]Neurosciences Graduate Program, University of California, San Diego, La Jolla, CA 92093

[2]Molecular Neurobiology Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037

[3]Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037

[4]Department of Biological Sciences, Howard Hughes Medical Institute, Columbia University, New York, NY 10027

[5]University of Massachusetts Chan Medical School, Worcester, MA 01655

[6]These authors contributed equally.

*Corresponding author

Email: schalasani@salk.edu

Author Contributions: Z.T.C. designed experiments. Z.T.C. and J.J.H. performed experiments. E.Y. contributed provided strains. Z.T.C. and K.T.Q. analyzed data. K.T.Q., Z.T.C, and S.H.C. wrote the paper. T.O.S. and S.H.C. supervised research.

Competing Interest Statement: The authors declare no competing interests.

Classification: Biological Sciences, Neuroscience

Keywords: *C. elegans,* whole-brain imaging, network states, sensory gating, soft decision trees

This PDF file includes:

Main Text
Figures 1 to 4

Abstract

Neural responses are influenced by both external stimuli and internal network states. While network states have been linked to behavioral and stimulus states, little is known about how sensory inputs are filtered by whole-brain activity to affect motor and command neurons. Here, we recorded whole-brain activity of *Caenorhabditis elegans* experiencing bacterial food stimuli, and modeled how sensory inputs affect motor and command neurons in a network state-dependent manner. First, we classified active neurons into six functional clusters: two sensory neuron clusters (ON, OFF), and four motor/command neuron clusters (AVA, RME, SMDD, SMDV). Using encoding models, we found that ON and OFF sensory neurons that respond to onset and removal of bacteria, respectively, employ different adaptation strategies. Next, we used decoding models to show that bacterial onset and removal differentially drive AVA and RME cluster activity. To explore state-dependent effects on AVA and RME clusters, we developed a model that identified network states and fitted submodels for each state to predict how each of the six functional clusters drive AVA and RME cluster activity. We also identified network states in which AVA and RME clusters were either largely unperturbed by or receptive to bacterial sensory input. Furthermore, this model allowed us to disentangle the state-dependent contributions of stimulus timescales and bacterial content to neural activity. Collectively, we present an interpretable approach for modeling network dynamics that goes beyond implication of neurons in particular states, and moves toward explicitly dissecting how neural populations work together to produce state dependence.

Significance Statement

A major function of the brain is to transform sensory information into behavior. As the first receiver of sensory input, sensory neuron activity is often most correlated with stimulus features. However, this high-fidelity representation of sensory input becomes diluted as it travels to downstream neurons, where sensory information is integrated with network activity. By the time sensory information reaches motor neurons, it is often difficult to dissociate the influence of sensory input from the influence of network activity. Here, we describe a method that is fully interpretable such that we can show how neural populations on a whole-brain scale interact to produce network states. From there, we can attribute motor neuron activity to network history and sensory input.

Main Text

Introduction

Constant stimuli can have variable influences on neural responses and behavior. On the neural population level, this variability becomes more pronounced as sensory information is transformed by downstream neurons. As the initial receiver of sensory information, sensory neurons exhibit activity that is the most correlated with stimulus features, but are susceptible to sensor errors that can propagate to behavior (1, 2). At the perception level, sensory input can be modulated by attention (3, 4), cognitive load (5), perceptual learning (6, 7), internal noise (8, 9), and internally generated coordinated activity (10). At the motor coordination level, gating of sensory input has been observed occurring in phase with behaviors such as locomotion (11, 12) and active whisking (13). However, little is known about how neural populations interact on a global scale to produce network states that modulate how sensory input is gated at the motor neuron level.

To disentangle the influences of stimuli and internal network state on downstream neural activity, precise stimulus control and whole-brain imaging are both needed to accurately account for experienced stimuli and global network states. Despite advances in modern imaging technology, most studies are limited to imaging small subsets of a brain's total activity (14–16). This problem is alleviated by studying simpler animals like zebrafish, fly larvae, and adult *C. elegans* where neural activity at single-cell resolution can be monitored across the entire brain

(17–20). With only 302 neurons, 189 of which are located in the head, the nematode *C. elegans* is ideally suited for whole-brain functional imaging studies. *C. elegans* whole-brain activity can be monitored both in restrained and freely moving animals (20–22). Moreover, whole-brain imaging of restrained animals in a microfluidic chip (20) allows for precise, fast, and complex stimuli presentations, thereby enabling investigations of stimulus-evoked whole-brain global dynamics underlying sensory coding (23, 24), motor states (25), and physiological states (26, 27). While these studies demonstrate the utility of whole-brain imaging and provide insights into the nematode nervous system, modeling of global network activity has largely focused on identifying which neurons and activity trends underlie particular behavioral or sensory states, but often fall short of explaining how populations work together to influence neural activity.

Here, we imaged the calcium activity of the entire *C. elegans* head while we presented it with rapidly fluctuating sequences of bacterial food stimuli and control buffer. We chose bacterial food instead of single-compound odorants in order to study a complex stimulus that is immediately relevant to *C. elegans*. We correlated activity of individual neurons to show that active neurons can be divided into six functional clusters groups. We then used encoding and decoding models to characterize basic properties of how stimuli drive sensory neurons and motor/command neurons. Finally, we built a hybrid model to identify network states and build submodels to explicitly show how sensory and motor populations drive motor/command neuron activity in each state. This model allowed us to identify network states in which motor/command neurons were either unresponsive or responsive to sensory input. Furthermore, we revealed how stimulus features and sensory context were differentially gated in a state-dependent fashion.

Results

Food-stimulated whole-brain activity reveals six functional cell clusters

We used an automated microfluidic system (28) to simultaneously image calcium activity in *C. elegans* head neurons and present the animal's nose with pulse-based stimulus sequences that rapidly fluctuated between liquid flows of bacterial food stimulus (from channel 1) and control buffer (from channel 2) (Fig. 1A**,** see Materials and Methods: Stimulus delivery). We refer to this as the bacteria↔buffer stimulus sequence. To control for artifacts intrinsic to the microfluidic setup, we also imaged activity while *C. elegans* was presented with a control buffer↔buffer stimulus sequence that fluctuated between two chemically identical buffer flows. This microfluidic system was previously used to show that individual chemosensory neurons detect and respond to bacteria (29, 30). To monitor whole-brain activity (see Materials and Methods: Whole-brain imaging), we used a strain that expressed a genetically-encoded nuclear-localized calcium indicator (GCaMP5K) (25,31). This strain was previously used to obtain whole-brain activity from restrained *C. elegans* during controlled delivery of stimuli (20, 25). In addition to this primary strain, we also recorded calcium activity (GCaMP6s) (32) from a strain that expressed NeuroPAL (24), which labels all *C. elegans* neurons with an invariant multicolor fluorescence map and allows for unambiguous identification of neurons (Fig. 1B-D). We used this as a supplementary strain to confirm cell identity associated with activity patterns found in the primary strain.

We found that active neurons could be divided into six functional clusters based on correlated changes in GCaMP fluorescence: two sensory neuron clusters (ON and OFF) and four motor/command neuron clusters (AVA, RME, SMDD, SMDV) (Fig. 1E–H). We first identified primary sensory neurons by looking for neurons that were either positively correlated (ON cells) or negatively correlated (OFF cells) with bacteria pulse onset during bacteria↔buffer stimulus sequences. Specifically, for each neuron, we first subtracted the changes in fluorescence induced by bacteria removals from the changes in fluorescence induced by bacteria onsets, and then ranked each neuron according to this score (see Materials and Methods: ON and OFF sensory neuron classification). The neurons with the highest and lowest scores were then checked for additional criteria before being classified as ON cells and OFF cells, respectively. ON cells were categorized as those neurons that obviously and immediately increased activity upon all bacteria onsets, and immediately decreased upon bacteria removals (Fig. 1I). Conversely, OFF cells were classified as neurons that decreased activity upon bacteria onsets and increased upon bacteria

removals (Fig. 1I). OFF cells were additionally required to rapidly decrease activity in response to initial bacterial onsets, dropping well below their pre-stimulus baseline, indicating that these sensory neurons are inhibited by bacteria rather than activated by the control buffer flow (Fig. 1I). This OFF cell behavior of being inhibited by a stimulus was absent in the control buffer↔buffer stimulus sequences. Instead, we found only ON cells that were activated by either the onset of the channel 1 buffer (ON-1 cells) or the onset of channel 2 buffer (ON-2 cells) (Fig. S1). Additionally, we identified fewer sensory neurons in animals presented with buffer↔buffer stimulus sequences (Fig. S1A) than with bacteria↔buffer stimulus sequences (Fig. 1F,G). This suggests that additional sensory neurons respond during bacteria ↔ buffer stimulus sequences, compared to buffer↔buffer stimulus sequences. Therefore, bacteria↔buffer stimulus sequences activate ON and OFF sensory neurons that respond to bacteria onset and removal, respectively, rather than to bacteria onset and buffer onset.

The vast majority of neurons did not respond immediately to stimulus changes (Fig. 1F–H). Across these neurons, we observed two pairs of stereotypical activity patterns: a pair of clusters that were anti-correlated and bistable (Fig. 1H, red and blue), and a pair of clusters that were anti-correlated and moderately fast (Fig. 1H, purple and yellow). Using the NeuroPAL–GCaMP6s strain to identify neurons, we found that RME motor neurons and AVE command neurons were anti-correlated and bistable (Fig. 1J), consistent with previous findings (25). We also found that SMDD and SMDV neurons were anti-correlated and exhibited moderately fast dynamics (Fig. 1K). These 4 representative motor and command neurons are associated with forward locomotion (RME), reverse locomotion (AVA), dorsal turning (SMDD), and ventral turning (SMDV) in *C. elegans* (25, 34–37). Using AVA, RME, SMDD, and SMDV as representative neurons, we sorted non-sensory neurons into clusters based on how their activity correlated with these four representative neurons (Fig. 1I–K). Many of the low noise neuronal traces strongly correlated (>85%) with either the identified RME motor neurons or AVA command neurons (Fig. 1F,G). A number of other calcium traces appeared as distorted, noisy versions of AVA or RME (Fig. S2), and the counts shown in the Fig. 1F,G are likely an underestimate of the true number of cells that are highly correlated with AVA and RME. In contrast, SMDV and SMDD were often the sole members of their eponymous clusters, usually with no other neurons that strongly correlated with these neurons (Fig. 1F,G). AVA, and RME were previously shown to have strong positional stereotypy, and thus, after identifying these neurons with NeuroPAL, their activity signature and stereotyped location made them easy to identify in the absence of NeuroPAL. SMDV and SMDD neurons were also readily identified by their previously reported distinctive activity signature (37). Therefore, we were able to subsequently identify all four motor/command neuron clusters in non-NeuroPAL-GCaMP5K animals without a coinciding NeuroPAL map (Fig. 1F,G).

These six cell clusters appeared in every animal, and we could not identify any other cell cluster that appeared consistently across all animals exposed to bacteria↔buffer stimulus sequences. We also observed similar clustering of active neurons in animals exposed to buffer↔buffer stimulus sequences (Fig. S1A). Putative cells not appearing in one of the aforementioned six cell clusters typically resembled noise or noisier versions of the activity patterns exhibited by the six cell clusters. To reduce dimensionality of the dataset with little loss of information (due to high correlation), we averaged across the activity traces of all neurons within a cluster (1I–K). Hence, all subsequent modeling used cluster-averaged representations to simplify model structure.

Bacteria addition and removal differentially drive activity of sensory, AVA, and RME clusters

To analyze how ON and OFF sensory neurons differ in their responses beyond the single-pulse timescale, we built encoding models to predict how sensory neuron activity adapts to repeated stimuli presentations (Fig. 2A, see Materials and Methods: Encoding model). Both ON and OFF cells are driven away from baseline activity (in opposite directions) upon bacteria onset and return to baseline upon bacteria removal. ON and OFF cells habituate to repeated bacteria presentations such that the change (increase for ON, decreases for OFF) in fluorescence from baseline is smaller in subsequent bacteria pulses (Fig. 1I). Here, we explored which timescale of

adaptation best described ON and OFF cells: 1) perfect adaptation, in which neurons cease responding despite persistent stimulation, 2) imperfectly adapting adaptation, in which neurons attenuate but do not fully terminate their responses, and 3) non-adaptation, in which neural responses are unaffected by recent stimulus history.

To examine the timescales of adaptation in ON and OFF cells, we used a cascade model that was previously used to describe adaptation to odors in *C. elegans* sensory neurons (38). This model is built on a cascade of simple ordinary differential equation (ODE) models of the form:

$$dx_i/dt = \tau * ([input] - x_i(t))$$
$$X(t) = x_1(t) - x_2(t)$$

where τ is the time constant that controls how fast a linear temporal filter responds to stimulus . The temporal filter describes how the recent history of the stimulus contributes to the current value of inferred calcium level of the cell, . With τ constrained to be positive, is guaranteed to exhibit perfect adaptation. That is, when encountering a step-change stimulus, will briefly change activity before terminating the response and returning to its baseline level. This model performed well for uncorrelated stimulus patterns but struggled considerably on correlated patterns (38). Kato and colleagues supposed these issues could be mediated by including more than two of the simple ODE filters. In other words, the authors hypothesized that *C. elegans* sensory cells adapt at more than one timescale. To test this hypothesis, as well as to ascertain whether the perfect adaptation assumption is justified, we simplified and generalized this model:

$$\frac{dx_i}{dt} = \tau \cdot ([input] - x_i(t))$$
$$x_i(t) = \sum a_i x_i(t)$$
$$g(t) = (x^p \circledast f_g)(t)$$

where is the inferred calcium level of the cell and is calculated as the sum of the temporal filters (ODE model solutions), is the GCaMP filter, and is a positive value required for the GCaMP transformation. This model learns the coefficients τ and for an arbitrary number of ODE basis function, and thus can learn adaptation on multiple timescales. Moreover, this model can test the effects of perfect and imperfect adaptation on model fit by toggling the coefficient constraint, such that perfect adaptation entails the following constraint:

$$\sum a_i = 0$$

This model formulation produced consistent and robust fits to *C. elegans* sensory neurons and allowed us to test the effects of different model complexities. In this approach, each sensory neuron's fluorescence trace was divided into three contiguous blocks. We used a 3-fold cross validation approach (i.e., two of the blocks are used to fit the model, while the third is used for testing) to assess model performance as a function of the number of basis filters and the type of adaptation. For ON cells, all adapting models significantly outperformed the non-adapting model according to a hierarchical bootstrap ($p < 0.05$ with Bonferroni correction, Fig. 2B, S3A). A perfectly adapting model with one primary filter and one adapting filter performed best, matching previous findings (38). Similarly, OFF cell adapting models outperformed the non-adapting model (Fig. 2C, S3B). However, unlike ON cells, OFF cells were best described by an imperfectly adapting model with one primary filter and one adapting filter, while the perfectly adapting model was the worst-performing model (not significantly better than non-adapting model) (Fig. 2C, S3B). It is possible that OFF sensory cells perfectly adapt over a longer timescale, but the fast

adaptation relevant to this study is imperfect in OFF cells. Using more than two basis functions (one primary filter and one adapting filter) hurt cross-validation performance (data not shown). While previous studies hypothesized that more than two cascade equations were required to model adaptation on multiple timescales (38), here we found that the type of adaptation and not the number of cascade equations to be more important for modeling OFF sensory neurons.

We also compared the timescales of ON and OFF temporal filters. For both ON cells and OFF cells, primary filters are fast and follow the fluctuation of stimuli, while the adapting filters reflect the slower timescale of adaptation. While ON and OFF cells displayed similar timescales in their primary filters, OFF cells have faster adapting filters (higher τ and lower time-to-half-peak) than ON cells (Fig. 2D,E). Thus, the OFF cell rebound was fast relative to ON cell habituation. Moreover, the OFF cell rebound was weak compared to ON cell habituation. In the best-performing OFF cell model, the adapting filter coefficient was typically smaller in magnitude than the primary filter coefficient (Fig. 2E). Thus, the OFF cell adapting filter does little to temper stimulus inhibition of OFF cells. Instead, it seems to be designed to produce fast rebounds to bacteria removal.

Adaptation appears to have different goals in ON and OFF cells. ON cells obey a fairly straightforward perfect adaptation law that can be explained by calcium depletion. OFF cells, on the other hand, exhibit an imperfect, fast rebound strategy. This latter strategy will not efficiently encode stimulus across large concentration scales. It will, however, maintain a higher dynamic range for constrained concentrations scales. We surmise that ON sensory cells may be designed to work across larger concentration scales, engaging specific OFF cells for specific concentration ranges. Overall, we show that ON and OFF sensory neurons have different adaptation kinetics to rapidly fluctuating bacterial stimulus sequences.

To determine how different phases of stimulus presentation affects motor neurons, we next built decoding models to predict stimulus state from motor/command cluster activity (see Materials and Methods: Decoding model). In bacteria↔buffer stimulus sequences, bacteria pulses are effectuated by directing bacteria flow from channel 1 over the *C. elegans* nose (Fig. 2F), while buffer pulses result when buffer from channel 2 flows over the nose and displaces the bacterial flow (Fig. 2G). Since *C. elegans* is positioned asymmetrically in the microfluidic chip relative to the two channels (channel 1 is slightly closer to the nose), we also analyzed animals presented with buffer↔buffer stimulus sequences to control for mechanosensory responses to differences in flow properties between the two channels (Fig. 2H).

Unlike sensory neuron clusters, shifts in motor/command neuron cluster activity do not reliably coincide with stimulus transitions (Fig. 1H–K). Therefore, we used multinomial logistic regression (MLR) to predict the probabilities of a particular stimulus state given motor neuron cluster activity inputs. MLR (39) is a robust classification model that, when combined with class balancing, has a very natural null model: prediction from worm identity only. Here, linear predictions about stimulus state are generated by linearly combining a set of weights with explanatory variables of a given observation:

$$\mathrm{P}(\mathrm{S}_i|X(t) = softmax(\textstyle\sum_{n=1}^{k} \beta_{i,k} X_k(t))$$

where is the stimulus class at time window . The -dimensional captures motor neuron cluster activity at time window along with worm identity information. is made up of the coefficients for stimulus state , which are regularized using an L1 norm and learned using gradient ascents. 16-second time windows of motor/command neuron cluster activity data were divided into thirds (5.33 seconds each). To predict the stimulus state in the middle subwindow (*t*-2.67s to *t*+2.67s, with *t* as the halfway point of the prediction window), motor/command neuron cluster activity data from the first subwindow (*t*-8s to *t-2.67s*) and the last subwindow (*t*+2.67s to *t*+8s) were used for the decoding task (see Materials and Methods: Decoding model).

We used this model to decode neural activity inputs from AVA, RME, SMDD, and SMDV clusters (Fig. 2I). Worm identity was also considered to capture variability across animals. The decoding model predicted four stimulus states for bacteria↔buffer stimulus sequences:

prolonged bacteria, prolonged buffer, bacteria-to-buffer transition, and buffer-to-bacteria transition (Fig. 2I). Importantly, initial onsets and removals of bacteria are included in transition states, but not in prolonged states. For buffer↔buffer stimulus sequences, corresponding states based on channel activation were predicted. In predicting stimulus states associated with bacteria↔buffer stimulus sequences, decoding from the activity of the RME and AVA cluster pair improved performance over decoding from identity alone in both non-NeuroPAL–GCaMP5K (99.3% of bootstraps, Fig. 2J) and NeuroPAL–GCaMP6s (99.4% of bootstraps, Fig. S4A) strains. In contrast, decoding from the activity of the SMDD and SMDV cluster pair did not perform better than the null model (Fig. S4B). Moreover, RME and AVA cluster activity could also be used to predict stimulus state in buffer↔buffer stimulus sequences, outperforming prediction from worm identity alone (Fig. S4C).

The decoding model produced linear temporal filters that described how stimulus states contributed to motor/command neuron cluster activity. Temporal filters for buffer↔buffer stimulus sequences revealed that flow from both buffer channels had similar effects on motor/command cluster activity. Prolonged buffer from either channel was associated with similar transient increases in both AVA and RME cluster activity (Fig. 2K,L). Both types of buffer-to-buffer transitions (channel 1 → channel 2, channel 2 → channel 1) reduced AVA cluster activity while RME cluster activity remained near baseline (Fig. 2M,N). In contrast, temporal filters for bacteria↔buffer stimulus sequences indicated that bacteria and buffer differentially drove AVA and RME cluster activity. Prolonged bacteria was associated with sustained increase in RME cluster activity and strong decrease in AVA cluster activity (Fig. 2O), while prolonged buffer was associated with sustained inhibition of RME cluster activity and low AVA cluster activity (Fig. 2P). The bacteria-to-buffer transition induced a slow increase in RME cluster activity, while AVA cluster activity remained near baseline (Fig. 2Q). The temporal filter predicting bacteria-to-buffer transition from RME cluster activity (Fig. 2Q) resembled a diminished version of the temporal filter predicting prolonged bacteria (Fig. 2O), suggesting that bacteria removal did not immediately alter RME activity. The converse buffer-to-bacteria transition is associated with rapid peak in AVA cluster activity and near-baseline RME cluster activity (Fig. 2R). The temporal filter for predicting buffer-to-bacteria transition from AVA cluster activity (Fig. 2R) is higher in magnitude than the temporal filter predicting prolonged buffer from AVA cluster activity (Fig. 2P), suggesting that bacteria onset has an immediate effect on the AVA cluster. Bacteria↔buffer decoding models were also remarkably similar across non-NeuroPAL–GCaMP5K and NeuroPAL–GCaMP6s strains (Fig. S4D–G). Altogether, these results suggest that bacteria presentation biases the AVA–RME cluster pair towards RME cluster activation (associated with forward locomotion), while bacteria removal biases the cluster pair towards AVA cluster activation (associated with reverse locomotion). Furthermore, motor/command responses to buffer differ depending on whether the overall sensory context also includes bacterial stimuli.

Identification of interpretable network states that vary in sensory gating properties

Based on our previous decoding results that AVA and RME motor/command clusters are influenced by sensory input, we next investigated how global network activity and sensory input drive AVA and RME cluster activity under different network states. We created a hierarchical model that used a soft decision tree (SDT) gating model (40) to identify relevant network states, and then we fitted MLR forecasting submodels for each network state to predict AVA and RME cluster activity (see Materials and Methods: SDT–MLR model). This combination of models, which we refer to as the SDT–MLR model, overcomes the limitation of using a single linear model to describe motor neuron cluster activity. For instance, members of the AVA cluster are bistable: they have upper and lower stable states (boundedness). Also, activity in these cells appears to have momentum: when one of these cells begins a transition between its stable states, it will tend to complete that transition. A single, linear model is unable to describe both momentum and boundedness. An appropriate model should learn the positive correlation between future AVA cluster rise and past AVA cluster increases to capture momentum. At the same time, this positive correlation should weaken and become negative as AVA nears its upper bound as these cells do not rise above that limit. Here, we employed multiple MLR models to capture nonlinear dynamics,

including features like momentum and boundedness. We then used an SDT gating model to route in recent network activity (AVA, RME, SMDD, SMDV clusters) to different network states. Each of these network states is associated with a different MLR model. In so doing, the SDT parceled the space of network trajectories into subspaces in which network evolution can be approximated by linear, probabilistic models. This parcellation and linearization strategy is similar to what has been previously used (41). In our SDT–MLR modeling, there were $M \times N$ MLR submodels, where M is the number of models being compared and N is the number of states parceled by the SDT. If the stimulus effect depends on network history, then we expect that the stimulus filters will differ across MLR submodels. As an end-to-end interpretable distillation of a neural network, the SDT is capable of learning complex, non-linear features. Moreover, since both the SDT and MLR submodels are differentiable, they can be simultaneously fit using gradient-based optimization methods.

We started by comparing forecasting models that predict AVA and RME cluster activity from recent network and sensory neuron activity, with models that predict from network history alone. This allowed us to assess whether sensory input from bacteria↔buffer stimuli were more relevant in some network states compared to others. Using recent network history inputs from AVA, RME, SMDD, and SMDV clusters, the full SDT-MLR model was fitted to a hyperparameter set for predicting RME and AVA calcium change (rise and fall). This model generated a tree with three levels, consisting of top level filters, along with right and left subtree filters (Fig. S5A), and four network states that produced the best results (Fig. 3A, S5B,C, see Materials and Methods: SDT–MLR model). State 1 and 4 were transient states, with state 1 associated with AVA cluster peaks, and state 4 associated with RME cluster peaks (Fig. 3A). In contrast, state 2 and 3 were persistent states (Fig. 3A). State 2 was characterized by high AVA, low RME, low SMDV, and high SMDD cluster activities, while state 3 exhibited the opposite (Fig. 3A). After network states were identified, the SDT was then frozen, and the MLR submodels were fitted to the rest of the data in an out-of-bootstrap cross-validation strategy to assess feature variability (42).

In the out-of-bootstrap cross-validation, inclusion of ON and OFF sensory neuron activity improved overall model performance (summed across states) in >95% of bootstraps. However, when considering model performance for individual network states, inclusion of sensory activity was only useful in states 3 and 4, but provided little predictive value in states 1 and 2 (Fig. 3B), despite the variance in stimulus sequence being similar across all states. Both sensory-responsive states (states 3 and 4) are characterized by low AVA and high RME cluster activities, both of which are associated with forward locomotion (Fig. 3B). Conversely, the sensory-unresponsive states (states 1 and 2) display the opposite activity trends with high AVA and low RME cluster activities, which are associated with reverse locomotion (Fig. 3B). This suggests that sensory input is gated more heavily during reverse locomotion than during forward locomotion, and that forward locomotion is a sensory-responsive behavior.

For each state, the SDT–MLR model generated linear temporal filters that predicted how recent network history from AVA, RME, SMDV, and SMDD clusters affected the probability of AVA and RME activity trends. In general, temporal filters were similar across predicted neural activity associated with forward locomotion (AVA fall, RME rise; Fig. 3C, middle two rows), as well as across activity associated with reverse locomotion (AVA rise, RME fall; Fig. 3C, top and bottom rows). Additionally, temporal filters were similar across the sensory-unresponsive states (state 1 and 2), except that recent SMDV history had a diminished effect on forward-associated activity in state 2. Linear filters across the sensory-responsive states (state 3 and 4) resembled each other, except for how network history drove AVA fall in state 4. Notably, RME cluster history drove AVA and RME cluster activity in general in sensory-unresponsive states, but had a neutral or suppressive effect in sensory-responsive states. Some state-invariant trends included the suppressive effect of SMDD history on AVA rise, as well as the suppressive effect of AVA on RME rise (Fig. 3C). Here, we describe some of the general trends of how network history influenced AVA and RME cluster activity in a state-dependent fashion, but Fig. 3C can be readily interpreted to understand in detail how each motor/command cluster affected AVA and RME cluster activity in each state.

Next, we looked at how sensory input influences particular AVA and RME cluster activity trends under different network states. As expected from Figure 3B, temporal filters predicting AVA and RME cluster activity from ON and OFF sensory activity were generally flat in states 1 and 2, indicating that sensory information was broadly suppressed from motor/command activity in these states (Fig. 3D). In contrast, states 3 and 4 exhibited more temporal filters in which sensory activity either increased or decreased the probability of AVA and RME activity trends (Fig. 3D). The largest effects of sensory input are on reverse-associated motor/command activity (AVA rise, RME fall), which is suppressed by ON activity and elevated by OFF activity (Fig. 3D). Within reverse-associated activity, OFF activity has a greater influence on AVA rise, while ON activity has a greater influence on RME fall (Fig. 3D). Additionally, ON sensory input drives RME rise in state 3 (Fig. 3D), suggesting that forward locomotion that characterizes state 3 can be maintained with ON activation of RME and OFF suppression of AVA. Overall, we show that SDT–MLR models can be used to identify relevant network states, characterize how those states are generated by network history, and delineate the state-dependent effects of stimuli on motor/command neuron activity.

Effect of stimulus timescales and sensory context on AVA and RME clusters

To further deconstruct sensory influences on AVA and RME cluster activity, we divided bacteria↔buffer stimulus sequences into stimulus patterns with either low- or high-variance pulse protocols, and fitted separate SDT–MLR models using corresponding neural activity (Fig. S6, S7). The main difference between low- and high-variance protocols is that alternating stimulus blocks consist of a single long pulse in the former, and composed of multiple short pulses in the latter (see Materials and Methods: Division of stimulus sequences). In general, high-variance pulses generally had more effect on AVA and RME cluster activity than low-variance pulses (Fig. 4A). To compare the overall effect of each network state on sensory gating, we summed the absolute values of magnitudes from all temporal filters within a state (Fig. 4A, last row). While high- and low-variance stimulus pulses had similar overall effects in state 4, high-variance pulses had more influence than low-variance pulses in state 3 (Fig. 4A, last row). Particularly in state 4, low-variance pulses can result in uneven influence from ON and OFF neurons, such as greater ON influence on RME fall and greater OFF influence on AVA rise (Fig. 4A). While state 1 and 2 were initially deemed to be broadly unresponsive to sensory input (Fig. 3B), enriching the model with a subset of pulse lengths revealed some sensory influence in state that was previously masked when all timescales were considered. For example, in state 1, high-variance-pulse ON activity promoted AVA fall, while high-variance-pulse OFF activity suppressed AVA fall (Fig. 4A). These opposite effects likely canceled each other out when all stimulus pulse lengths were considered, thereby resulting in a sensory-unresponsive model prediction (Fig. 3D). Furthermore, segregation of stimuli by low- or high-variance pulses revealed more sensory influence on forward locomotion in states 3 and 4 (Fig. 4A) that was previously undetected in the all-pulse SDT–MLR model (Fig. 3D). Thus, the SDT–MLR model can be used to also detect feature-specific effects of stimuli on motor/command neuron activity.

We next explored the ability of SDT–MLR models to differentiate sensory contexts by comparing buffer↔buffer and bacteria↔buffer stimulus sequences. Since high-variance pulses were shown to have greater overall effect (Fig. 4A), we compared only high-variance pulses from buffer↔buffer and bacteria↔buffer stimulus sequences. In the SDT–MLR model for buffer↔buffer stimulus sequences (Fig. S8), inclusion of buffer↔buffer sensory activity improved overall performance only for state 4 (Fig. S8A). Compared to bacteria↔buffer sensory input, temporal filters predicting AVA and RME cluster activity from buffer↔buffer sensory input typically displayed slower timescales (Fig. S8C). Based on the sum of absolute magnitudes of all temporal filters within a state, buffer↔buffer stimuli generally had more influence than bacteria↔buffer stimuli in states 1 and 2, and less influence in state 3 (Fig. 4B). While the state-agnostic decoding model found no difference between channel 1 and channel 2 buffer flows in how they drove AVA and RME cluster activity (Fig. 2K–L), the SDT–MLR model found that the channel 1 buffer had a greater effect on AVA and RME cluster activity than the channel 2 buffer in states 2, 3, and 4 (Fig. 4B, gray). This bias may be due to channel 1 being closer to the *C. elegans* nose, as previously

surmised. However, this bias disappeared when bacteria was included in the sensory context (Fig. 4B, cyan). Therefore, while channel 2 is chemically the same for both buffer↔buffer and bacteria↔buffer stimulus sequences, the effect that the channel 2 buffer had on AVA and RME cluster activity changed depending on whether the greater sensory context involves switching between chemically identical buffers, or switching between bacterial stimuli and buffer.

Discussion

Here, we demonstrate that the SDT–MLR model can identify interpretable global network states that underlie differential gating of sensory input at the motor and command neuron level. First, we used state-agnostic encoding and decoding models to confirm that bacteria onset and removal differ in how they drive activity of sensory neurons, AVA, and RME clusters. Using the SDT–MLR model, we identified two sensory-unresponsive network states (states 1 and 2) and two sensory-responsive network states (states 3 and 4) in the bacteria↔buffer sensory context. For each network state, we explicitly characterized how the history of each of the four functionally defined motor/command neural populations (AVA, RME, SMDD, SMDV) drive AVA and RME cluster activity to produce each network state. Finally, we used the SDT–MLR model to identify how sensory input in general, as well as stimulus features and sensory context in particular, influence AVA and RME cluster activity.

The encoding model showed that sensory neural responses to the addition of bacterial stimuli (ON) and its removal (OFF) can be modeled by two linear ordinary differential equations using one fast and one slow filter. Linearity in sensory neurons is observed in both vertebrates and invertebrate photoreceptors (43), rat trigeminal neurons (44), and primate vestibular neurons (45), implying that this might be a common phenomenon. We also show that ON sensory neurons perfectly adapt to the bacterial stimulus, consistent with previous studies (38), while OFF neurons adapt imperfectly. Given that our analysis included all sensory neurons responding to bacterial stimuli, we suggest that this might be a general principle of *C. elegans* sensory neurons. Similar differences in adaptation in ON and OFF neurons are also observed in single-cell electrophysiological recordings from vertebrate photoreceptors and olfactory sensory neurons (46). Moreover, these results also provide some hints about the encoding strategies of these two sensory-neuron classes. We suggest that ON neurons encode stimulus over a larger dynamic range compared to OFF neurons. Consistently, we have previously shown that AWA sensory neurons (ON) have a larger dynamic range compared to AWC sensory neurons (OFF) in detecting benzaldehyde (47). Moreover, studies in the vertebrate retina have shown that the dynamic range of the ON pathway is much greater than that of the OFF pathway, likely due to a selective effect of pre-synaptic inhibition on the ON, but not OFF, bipolar cells (48), confirming the validity of our hypothesis.

The decoding model informed which motor/command neural populations were targets of the sensory input, and also served as a state-agnostic model to compare with the state-dependent SDT–MLR model. Both models were used to assess how sensory input affects AVA and RME cluster activity. While the decoding model found no difference in effect between either of the buffer channels in buffer↔buffer stimulus sequences, the SDT–MLR model revealed that buffer↔buffer sensory input does indeed have a channel bias that is more pronounced in some states than others. This suggests that the SDT–MLR model is more sensitive to sensory effect compared to the decoding model. This sensitivity is further amplified by subdividing sensory effects by state, channel, sensory neuron classification (ON/OFF), prediction class (AVA rise, AVA fall, RME rise, RME fall), and stimulus features (high-variance pulses, low-variance pulses). This granular approach allowed us to identify specific sensory effects that were obscured in more general analyses. We found that high-variance and low-variance pulses had similar influences in one state, but sensory input from low-variance pulses were gated more heavily in another state. Additionally, we revealed that state-dependent gating of buffer stimulus is dependent on whether the greater sensory context involves switching between bacteria and buffer, or between buffer and buffer.

The SDT–MLR model differs in both goal and interpretability from recent studies that described network states in *C. elegans* whole-brain activity (25, 27, 37, 49). In these studies, whole-brain activity is analyzed with the purpose of understanding the temporal dynamics of neuronal populations in terms of how the network state evolves over time. Often, probabilities of network state transitions are related to corresponding stimulus or behavioral transitions. A dimensionality reduction technique called principal component analysis (PCA) (50) is used to quantify brain-wide correlations that reflect signals shared by clusters of neurons. These shared signals are referred to as temporal principal components (PCs). A key operation of PCA-based analyses is to transform whole-brain activity to new axes defined by the top PCs that explain the most variance in data. However, the meaningful linear axes that originally described the data are lost in the process. While PCA-based analyses are useful for distinguishing between network states and their transition between each other, these networks states and transitions are described in terms of PCs, which can be difficult to interpret. Nonetheless, PC weights can still be used to identify relevant neurons and activity trends associated with each state (25, 27, 37, 49). In our study, we are less concerned with the probabilities of transitioning between states, and more interested in the within-state conditional probabilities of individual neuronal populations interacting with each other. Instead of using PCA, our model preserves interpretability by using soft decision trees, such that network states and network interactions are always described by the identities of neuronal populations and their corresponding activity patterns. Rather than supplanting PCA-based analyses, our SDT–MLR model serves as a complementary method for focusing on the network interactions within individual network states rather than the temporal dynamics that connect those network states.

Interestingly, the sensory-responsive states identified by the SDT–MLR are characterized by neural activity that has been previously shown to be associated with forward locomotion, while the sensory-unresponsive states are characterized by neural activity associated with reverse locomotion. This is consistent with previous reports showing that inhibition of sensory input occurs at particular phases of the locomotory cycle (11, 12). Inhibition of sensory input during movement may serve to distinguish between external stimuli and self-generated stimuli (51), which is accomplished by integrating sensory inputs with motor inputs (corollary discharge) (52). Since typical *C. elegans* locomotion consists primarily of forward locomotion punctuated by transient reversals (53), one possibility is that sensation of stimuli is suppressed during reversals to temporarily pause processing of stimulus flows until a stable locomotion state is restored. It is important to note that our study was conducted with *C. elegans* trapped in an immobilized position in a microfluidic chip, and therefore behavioral associations were inferred purely from motor and command neuron activity. Moreover, a recent study showed that the set of neurons correlated with AVA differs depending on whether *C. elegans* is immobilized or freely moving (49). Thus, our association of AVA and RME clusters activity with forward and reverse locomotion is tentative and should be confirmed in freely moving animals. However, while imaging freely moving *C. elegans* would provide rich behavioral information that can be added to the SDT–MLR model, complex and precise stimuli presentation can be difficult to achieve when the stimulus target is mobile. There are some efforts to study sensation in freely moving animals (22, 54), and a reasonable balance of behavioral and sensory information richness may be achieved with a microfluidic chip that allows semi-restricted locomotion and somewhat fast waves of liquid stimuli (55).

Overall, we present an approach for understanding how sensory information filters through whole-brain network interactions to affect downstream motor and command neurons in a state-dependent manner. Currently, there is an epistemological bias towards identifying network states that correspond with a particular stimulus or motor state. In contrast, there has been less focus on network states that are defined by altered network interactions. Our computational approaches provide a method for investigating network mechanisms at the level of pairwise interactions between neuronal populations. While our study only looked at the network mechanisms underlying sensorimotor integration, this model can be leveraged to also understand how network inputs are integrated at any network level. More broadly, we suggest that our approach of combining soft decision trees with multinomial logistic regression can be used to

identify relationships, not only in neural networks, but also in cellular signaling pathways, transcription factor networks, and between other complex biological or physical entities.

Materials and Methods

Whole-brain imaging

We used two transgenic strains that expressed GCaMP. The primary strain (ZIM294) expressed GCaMP5K in the nuclei of all neurons (mzmEx199 [*Punc-31::NLSGCaMP5K; Punc-122::GFP*]). To identify neurons associated with activity patterns observed in ZIM294, we used a strain (OH15500) that expressed GCaMP6s and NeuroPAL (*otIs669[NeuroPAL];otIs672[Panneuronal GCaMP6s]*). Cells were identified according to the map described by Yemini and colleagues (24). We monitored changes in GCaMP fluorescence using a Zeiss LSM 880 with Airyscan. Acquisition was done in 2 micron z-steps. In 'Fast' mode, the Airyscan images the entire head of the adult worm at about 1.5 volumes per second. Worms were typically imaged for approximately ten minutes. We then used piecewise rigid registration to remove motion artefacts (56) and non-negative matrix factorization to isolate individual neurons and extract their fluorescence values (33). Out of a total 189 neurons in the head, our approach identified 50-100 neurons per animal.

Stimulus delivery

Day 1 adult animals were washed in M9 and loading into in a microfluidic device that trapped the worm body while exposing only the nose to stimulus flows (28). Animals were also treated with 1.5 mMol of the paralytic tetramisole hydrochloride to suppress most perceivable worm movement. The movement of untreated worms proved too difficult to motion correct. We delivered precise patterns of fluctuating bacteria and M9 buffer liquid flows using a custom designed Arduino device to send pulses to a valve controller. The bacteria solution was prepared as a 1:1 resuspension of a bacterial culture (OD_{600} = 0.4) in M9 buffer as previously described (29). The controller determines whether bacteria or buffer is routed to the nose of the trapped worm or away from the worm. Worms were exposed to binary patterns of bacteria and buffer. A number of different stimulus protocols are used in this study. In the base protocol, the trial is divided up into pulse blocks of ~15 seconds. The pattern is constructed using transition probabilities: p(switch on | off) = 0.2 and p(switch off | on) = 0.4. In the faster protocols, the same switch probabilities are used but the pulse blocks have length ~1.5 seconds. The patterned protocols are effectively the same as the base protocol. The only difference is that their 'stimulus on blocks' are composed of multiple pulses.

GCaMP filters

The GCaMP filter *g(x)* is modeled as a difference of exponentials with parameters matching those of Chen and colleagues (32). This procedure is complicated by the volumetric nature of the imaging data. Consider two sensory neurons with identical calcium dynamics; they respond to stimulus with the same timescale. Neuron A is in imaging slice 0, while neuron B is in slice 7. These two neurons will have the same calcium timings relative to the stimulus. However, Neuron B will appear to have faster response kinetics since it is acquired over half a second later (relative to stimulus onset/removal) compared to neuron A. Thus, the slice in which the neuron appears needs to be considered in the creation of its GCaMP filter:

$$g_f(t) = \int_{t_1}^{t_2} g(x)dx$$
$$t_1 = t + \frac{z}{n_z}$$
$$t_2 = t + \frac{z+1}{n_z}$$

where *g(x)* refers to the difference of exponentials and n_z refers to the number of z slices. The normalized *g(t)* filter is applied via a linear convolution to transition between calcium and GCaMP dynamics.

ON and OFF sensory neuron classification

Sensory neurons were classified as either ON or OFF for neuronal activity collected during bacteria↔buffer stimulus sequences. The change in normalized fluorescence over a series of ≥ 10 stimulus pulses (all trials have at least a few of these). Both the first 10 volumes into a bacteria pulse and the first 10 volumes into a buffer pulse (following bacterial removal) were considered. The following metric was then calculated:

$$rank = \left(\sum_{i}^{P}(f(t_i+10)-f(t_i))\right) - \left(\sum_{j}^{Q}(f(t_j+10)-f(t_j))\right)$$

where f is normalized fluorescence, t is pulse onset, P is the number of bacteria pulses, and Q is the number of buffer pulses. The cells with the highest ranks considered as potential ON cells, and the cells with the lowest ranks were considered as potential OFF cells. Rank cutoffs were selected manually for each trial. ON cells were categorized as those neurons that obviously and immediately increased activity upon bacteria onsets, and immediately decreased upon bacteria removals (Fig. 1I). In contrast, OFF cells were classified as neurons that decreased activity upon bacteria onsets and increased upon bacteria removals. For both ON and OFF cells, baseline (low variance) activity occurred when the stimulus of interest was absent (during buffer pulses). This distinguishes OFF-bacteria sensory neurons from hypothetical ON-buffer sensory neurons, for which baseline activity would occur during bacteria pulses.

Encoding model

The encoding model predicted ON and OFF sensory neuron activity from stimulus features. The core primary sensory model consists of three parts: (1) a set of cascade basis functions, (2) coefficients for the basis functions, and (3) a GCaMP transformation.

$$\frac{dx_i}{dt} = \tau \cdot ([input] - x_i(t))$$
$$x_i(t) = \sum a_i x_i(t)$$
$$g(t) = (x^p \circledast f_g)(t)$$

where is the inferred calcium level of the cell and is calculated as the sum of the temporal filters (ODE model solutions), is the GCaMP filter, and is a positive value required for the GCaMP transformation. The fitting procedure learns values for the time constant τ and the basis coefficients . This model uses the solutions to these ordinary differential equations as basis functions. For a given model instantiation, each equation is assigned a value for τ. After the ordinary differential equations are solved analytically, the model can be linearized, allowing for robust estimates of the a_i coefficients (see Methods). We then used a random search strategy to obtain estimates of τ. Moreover, by toggling the coefficient constraint, we can test whether perfect adaptation is necessary to predict sensory neuron activity. This constraint yields perfect adaptation as all the ordinary differential equation basis functions saturate at the same value ( ).

We used a random search strategy to find the correct set of basis functions. The strategy chooses *N* basis time constants for an *N* cascade model and an initial estimate of *p* power. It solves the basis equations analytically for a given stimulus pattern and produces an initial estimate for the coefficients by fitting to a linearized approximation of the neuron's calcium trace. In this approximation, the raw sensory neuron fluorescence trace is deconvolved using the Richardson-Lucy method (57) and taken to the *(1/p)* power. The initial estimate for the coefficients is the solution to the resulting linear regression equation. Finally, this system fits the full model (free variables consist of basis coefficients and *p* power) is fit to the raw fluorescence trace using gradient descent. We then repeat this process for a large number of random searches to define the basis functions.

Decoding model

The decoding model predicted stimulus states activity from stimulus features. For both the decoding and prediction analyses, we split the data into contiguous blocks of ~10 second (16 volume) duration. Within each block, subwindows were created in a rolling fashion. For instance, for an *N* volume prediction window, there are 16-*N* legal, overlapping subwindows within each block. It should be noted that there is no overlap between prediction windows of adjacent blocks. Test/train sampling is done at the block level, guaranteeing no train/test prediction overlap. This system captures sharp transitions while measuring targets over multiple time bins, thereby limiting noise in the targets. We chose a length of 16 volumes because this captures the entirety of motor neuron event initiation (for all motor neuron classes). This causes bootstrap sampling to be performed at the event level, stopping a small, handful of events from dominating model outcomes. We did not test other blockstrap sizes.

We used Gaussian basis functions for both decoding and prediction tasks. This involves filtering network history and stimulus data through these basis functions before being fed into decoding and prediction models. We chose specific gaussians (parametrized by mean and variance) by hand-tuning model performance on the hyperparameter set.

Given the choice of 16 volume blocks and 8 volume time windows, there are 16-8=8 prediction windows within each time block. Unlike the prediction task, the decoding task centers the prediction windows relative to the input data. For instance, in one decoding task, RME/AVA cell cluster data from *t*+4-8 to *t*+4+8 volumes is used to predict whether the stimulus is on/off/altered from *t* to *t*+8 volumes. The unit of 8 volumes was chosen because many of the pulse protocols use 10 volume pulses as a base unit. Thus, using 8 volumes guarantees a fair number of samples of the on stimulus class. We did not test other volume lengths.

In the decoding analyses, we predict five classes of stimulus patterns from network activity (58). Prediction from worm identity alone serves as the null model for all decoding analyses. In order to ensure best performance for the null model, we balanced the classes within each worm. For each worm, if there are N occurrences of class A in the training set, there are N occurrences of class A in the test set. We ensured this by randomly removing prediction windows.

SDT–MLR model

We used different L1 norms for each of the multinomial logistic regression input classes: AVA/RME/SMDV/SMDD terms, ON/OFF cell terms, and worm identity terms. The worm identity terms essentially make the forecasting models into a random-intercept model (59). Initial experiments found no benefit in random-intercept style models. The number of network states and number/shape of gaussian filters were additional hyperparameters. A combination of hand-turning and grid search were used on the hyperparameter set to find good regularizers. This hyperparameter set was also used for the soft decision tree, which could also be thought of as an additional hyperparameter in this study. These hyperparameters were frozen on the out-of-bootstrap analysis (42).

In the forecasting task, we predicted changes in GCaMP fluorescence from time t to t+T1 using network history from t-T2 to t and sensory neuron activity from t-T2+T1 to t+T1. We used a length of twenty-four volumes for T2 for all prediction analyses. We treated the length of T1 as an additional hyperparameter. As in the decoding analysis, these prediction windows are contained inside larger data blocks. We performed train/test sampling at the block level. We chose sixteen seconds (24 volumes) for T1 in all models, as 16 seconds is sufficient to capture the majority of command neuron events.

RME and AVA activation and inactivation is clearly probabilistic. We used a multinomial logistic regression (MLR) as the base model for motor-neuron activity prediction. In order to use MLR, we discretized GCaMP fluorescence activity in every prediction window. This is done by subtracting the average GCaMP fluorescence over the prediction window by the GCaMP fluorescence level at the beginning of that window.

RME and AVA neurons have non-linear calcium dynamics. Two of these features are boundedness and momentum. Members of the AVA cluster exhibit boundedness: their activity is

limited to a range between their upper and lower stable states. Also, activity in these neurons appears to have momentum: when one of these neurons begins a transition between the stable states, it will tend to complete that transition. A single, linear model is unable to describe both momentum and boundedness. It must learn the positive correlation between a future AVA rise and past AVA activity to capture momentum. However, this positive correlation should weaken and become negative as AVA nears its upper bound, since AVA activity does not rise above this bounded limit. Thus, a gating model is required to capture the change in this positive correlation.

Here, we employ multinomial logistic regression (MLR) models to capture nonlinear dynamics, allowing us to capture features like momentum and boundedness. A gating model is used to divide the space of network histories into subspaces. Each of these different subspaces is associated with a different MLR model. This approach with two submodels can simultaneously model both momentum and boundedness in AVA. In one theoretical solution, one MLR submodel is only active when the AVA cell cluster is near its lower bound, while another is active when AVA near its upper bound. The first submodel learns a positive correlation between a past AVA increase in activity and future AVA rise (momentum), while the second learns a weak positive or negative correlation (boundedness).

We use a hybrid model combing soft decision trees and multinomial logistic regression (SDT–MLR) method to divide the network trajectory space into different subspaces. We assumed that GCaMP fluorescence changes in each of these subspaces can be well-described by an MLR. A soft decision tree is a form of oblique decision tree that is end-to-end differentiable (40). Each branch of the soft decision tree is a different logistic regression model on the same input vector, which outputs a left vs right probability. These left and right probabilities are multiplied by the predicted class probabilities of the corresponding left and right subtrees. In this manner, soft decision trees are essentially hierarchical filters that can be learned through gradient descent. The outputs of the soft decision tree weigh the different MLR models. For instance, a soft decision tree with depth 2 and width 2 will have 4 output states. Each of these output states are associated with a different MLR. All of these MLRs are trained against the entire dataset. However, the data points are weighted by the particular soft decision tree output leaf. Hence, different MLR models will focus on different subsets of the data.

We found that averaging predictions across SDT–MLR models improves cross-validation performance. This averaging is done at two levels. First, within a SDT–MLR model, the predictions of each MLR are weighted by the soft decision tree and averaged. Second, these averaged predictions are further averaged across several SDT–MLR models. Here, twenty-five SDT–MLR models were fit separately to the hyperparameter set. The best hybrid model was chosen for analysis on the cross-validation set. For training on the cross-validation set, the Soft Trees were frozen; only the MLRs were trained. Hence, training is convex (weighted multinomial logistic regression) on the cross-validation set. Freezing the soft decision during cross-validation allows us to easily align model data across bootstraps. This, in turn, gives us information on the variance of different features of the MLR models.

In our exploration of the hyperparameter set, we found that SDT–MLR models tend to converge on bad solutions if not regularized. These bad solutions are characterized by poor training and test set performance as well as state imbalance. The Soft Tree assigns most data points to one of its submodels, while its other MLR submodels are trained on very small subsets of the data resulting in state imbalance. We solved this issue by maximizing entropy regularizer H(X):

$$V_k(X) = \frac{1}{N}\sum_{i}^{N} S_k(X_i)$$

$$H(X) = \sum_{k} V_k(X) log\left(\frac{1}{V_k(X)}\right)$$

where X_i is the ith input data point in the minibatch. S_k is the probability assigned to the kth state/submodel by the soft decision tree. Thus, V_k is the average probability of state k across the

minibatch. Therefore, *H(X)* is high when all states are equally represented in the minibatch. It should be noted that this regularizer does not directly penalize high state probabilities.

Division of stimulus sequences

To compare the effect of short versus long stimulus pulses in the SDT–MLR model, complete stimulus sequences were divided into low-variance (long pulses) and high-variance (short pulses) stimulus patterns. Low-variance stimulus patterns were constructed by dividing the trial into pulse blocks of ~15 seconds, with the transition probabilities p(switch on | off) = 0.2 and p(switch off | on) = 0.4. High-variance stimulus patterns were also constructed by dividing the trial into bacteria or buffer blocks with the same transition probabilities as for low-variance stimulus patterns. The difference between low- and high- variance stimulus patterns is that the bacteria block consists of a series of sub-pulses rather than a single constant pulse. These sub-pulses are ~3 seconds bacteria and ~3 seconds buffer.

Acknowledgments

We thank Manuel Zimmer, Oliver Hobert and CGC (Caenorhabditis Genetics Center) for worm strains and Uri Manor, Tong Zhang and the Waitt Advanced Biophotonics Center for advice and assistance with our imaging experiments. We also thank Saket Navlakha, Ryan Rowekamp, Javier How, Molly Matty, Jessica Haley, Michael Rieger, and Eric Edsinger for discussions and comments on the manuscript. This work was supposed by grants from the National Institutes of Health NIH 5R01MH096881 (S.H.C); 5T32DK007328-35, 5T32DK007328-37, 5T32MH015174-37 and 5T32MH015174-38 (E.Y). U19NS112959, P30AG068635, AHA-Allen Initiative in Brain Health and Cognitive Impairment award 19PABH134610000 (T.O.S.). Z.T.C. was supported by a fellowship from the Dan and Marina Lewis Foundation. K.T.Q. is supported by a Postdoctoral fellowship from the Paul F. Glenn Foundation. CGC is funded by the NIH Office of Research Infrastructure Programs (P40 OD010440).

References

1. J. M. Beck, W. J. Ma, X. Pitkow, P. E. Latham, A. Pouget, Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* **74**, 30-39 (2012).
2. L. C. Osborne, S. G. Lisberger, W. Bialek, A sensory source for moto variation. *Nature* **437**, 412-416 (2005).
3. J. H. Reynolds, D. J. Heeger, The normalization model of attention. *Neuron* **61**, 168-185 (2009).
4. J. H. Reynolds, T. Pasternak, R. Desimone, Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703-714 (2000).
5. P. Sörqvist, S. Stenfelt, J. Rönnberg, Working memory capacity and visual–verbal cognitive load modulate auditory–sensory gating in the brainstem: Toward a unified view of attention. *J. Cogn. Neurosci.* **24**, 2147-2154 (2012).
6. A. Schoups, R. Vogels, N. Qian, G. Orban, Practising orientation identification improves orientation coding in V1 neurons. *Nature* **412**, 549-553 (2001).
7. A. F. Teich, N. Qian, Learning and adaptation in a recurrent model of V1 orientation selectivity. *J. Neurophysiol.* **89**, 2086-2100 (2003).
8. A. A. Faisal, L. P. Selen, D. M. Wolpert, Noise in the nervous system. *Nat. Rev. Neurosci.* **9**, 292-303 (2008).
9. M. N. Shadlen, K. H. Britten, W. T. Newsome, J. A. Movshon, A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci.* **16**, 1486-1510 (1996).
10. Z. W. Davis, L. Muller, J. Martinez-Trujillo, T. Sejnowski, J. H. Reynolds, Spontaneous travelling cortical waves gate perception in behaving primates. *Nature* **587**, 432-436 (2020).

11. K. T. Sillar, A. Roberts, A neuronal mechanism for sensory gating during locomotion in a vertebrate. *Nature* **331**, 262-265 (1988).
12. M. H. Ouellette, M. J. Desrochers, I. Gheta, R. Ramos, M. Hendricks, A Gate-and-Switch Model for Head Orientation Behaviors in Caenorhabditis elegans. *eNeuro* **5** (2018).
13. E. Eggermann, Y. Kremer, S. Crochet, C. C. Petersen, Cholinergic signals in mouse barrel cortex during active whisker sensing. *Cell reports* **9**, 1654-1660 (2014).
14. J. M.Blackwell, M. N. Geffen, Progress and challenges for understanding the function of cortical microcircuits in auditory processing. *Nat. Commun.* **8**, 1-9 (2017).
15. M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, *et al.*, Neural population dynamics during reaching. *Nature*, **487**, 51-56 (2012).
16. R. C. Williamson, B. Doiron, M. A. Smith, M. Y. Byron, Bridging large-scale neuronal recordings and large-scale network models using dimensionality reduction. *Curr. Opin. Neurobiol.* **55**, 40-47 (2019).
17. M. B. Ahrens, J. M. Li, M. B. Orger, D. N. Robson, A. F. Schier, *et al.*, Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471-477 (2012).
18. M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, P. J. Keller, Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat Methods* **10**, 413-420 (2013).
19. W. C. Lemon, S. R. Pulver, B. Höckendorf, K. McDole, K. Branson, *et al.*, Whole-central nervous system functional imaging in larval Drosophila. *Nat. Commun.* **6**, 1-16 (2015).
20. T. Schrödel, R. Prevedel, K. Aumayr, M. Zimmer, A. Vaziri, Brain-wide 3D imaging of neuronal activity in Caenorhabditis elegans with sculpted light. *Nat. Methods* **10**, 1013-1020 (2013).
21. J. P. Nguyen, F. B. Shipley, A. N. Linder, G. S. Plummer, M. Liu, *et al.* (2016). Whole-brain calcium imaging with cellular resolution in freely behaving Caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, *113*(8), E1074-E1081.
22. V. Venkatachalam, N. Ji, X. Wang, C. Clark, J. K. Mitchell, *et al.,* Pan-neuronal imaging in roaming Caenorhabditis elegans. *Proc. Natl. Acad. Sci. U. S. A.* **113**, E1082-E1088 (2016).
23. J. J. How, S. Navlakha, S. H. Chalasani, Neural network features distinguish chemosensory stimuli in Caenorhabditis elegans. *bioRxiv* [Preprint] (2020). https://www.biorxiv.org/content/10.1101/2020.02.18.955245v1 (accessed 31 October 2021).
24. E. Yemini, A. Lin, A. Nejatbakhsh, E. Varol, R. Sun, *et al.*, NeuroPAL: a multicolor atlas for Whole-Brain neuronal identification in C. elegans. *Cell* **184**, 272-288 (2021).
25. S. Kato, H. S. Kaplan, T. Schrödel, S. Skora, T. H. Lindsay, *et al.*, Global brain dynamics embed the motor command sequence of Caenorhabditis elegans. *Cell*, **163**, 656-669 (2015).
26. A. L. Nichols, T. Eichler, R. Latham, M. Zimmer, A global brain state underlies C. elegans sleep behavior. *Science* **356** (2017).
27. S. Skora, F. Mende, M. Zimmer, Energy scarcity promotes a brain-wide sleep state modulated by insulin signaling in C. elegans. *Cell Rep.* **22**, 953-966 (2018).
28. N. Chronis, M. Zimmer, C. I. Bargmann, Microfluidics for in vivo imaging of neuronal and behavioral activity in Caenorhabditis elegans. *Nat. Methods* **4**, 727-731 (2007).
29. A. J. Calhoun, A. Tong, N. Pokala, J. A. Fitzpatrick, T. O. Sharpee, *et al.* Neural mechanisms for evaluating environmental variability in Caenorhabditis elegans. *Neuron* **86**, 428-441 (2015).
30. S. H. Chalasani, N. Chronis, M. Tsunozaki, J. M. Gray, D. Ramot, *et al.*, Dissecting a circuit for olfactory behaviour in Caenorhabditis elegans. *Nature* **450**, 63-70 (2007).
31. J. Akerboom, N. Carreras Calderón, L. Tian, S. Wabnig, M. Prigge, *et al.* Genetically encoded calcium indicators for multi-color neural activity imaging and combination with optogenetics. *Front Mol Neurosci.* ***6*** (2013).
32. T. W. Chen, T. J. Wardill, Y. Sun, S. R. Pulver, S. L. Renninger, *et al.*, Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295-300 (2013).

33. E. A. Pnevmatikakis, D. Soudry, Y. Gao, T. A. Machado, J. Merel, *et al.* Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron*, *89*(2), 285-299.
34. Hendricks, M., Ha, H., Maffey, N., Zhang, Y. (2012). Compartmentalized calcium dynamics in a C. elegans interneuron encode head movement. *Nature* **487**, 99-103 (2016).
35. Y. Shen, Q. Wen, H. Liu, C. Zhong, Y. Qin, *et al*., An extrasynaptic GABAergic signal modulates a pattern of forward movement in Caenorhabditis elegans. *Elife* **5**, e14197 (2016)..
36. A. Kocabas, C. H. Shen, Z. V. Guo, S. Ramanathan, Controlling interneuron activity in Caenorhabditis elegans to evoke chemotactic behaviour. *Nature* **490**, 273-277 (2012).
37. H. S. Kaplan, O. S. Thula, N. Khoss, M. Zimmer, Nested neuronal dynamics orchestrate a behavioral hierarchy across timescales. *Neuron* **105**, 562-576 (2020).
38. S. Kato, Y. Xu, C. E. Cho, L. F. Abbott, C. I. Bargmann, Temporal responses of C. elegans chemosensory neurons are preserved in behavioral dynamics. *Neuron* **81**, 616-628 (2014).
39. D. W. Hosmer, S. Lemeshow, R. X. Sturdivant. *Applied logistic regression* (Wiley, 2000)
40. N. Frosst, G. Hinton, Distilling a neural network into a soft decision tree. *arXiv* [Preprint] (2017). https://arxiv.org/abs/1711.09784 (accessed 31 October 2021).
41. S. Linderman, A. Nichols, D. Blei, M. Zimmer, L. Paninski, Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in C. elegans. *BioRxiv* [Preprint] (2019). https://www.biorxiv.org/content/10.1101/621540v1 (accessed 31 October 2021).
42. I. Tsamardinos, E. Greasidou, G. Borboudakis, Bootstrapping the out-of-sample predictions for efficient and accurate cross-validation. *Machine Learning* **107**, 1895-1922 (2018).
43. J. Pahlberg, A. P. Sampath, Visual threshold is set by linear and nonlinear mechanisms in the retina that mitigate noise: How neural circuits in the retina improve the signal-to-noise ratio of the single-photon response. *Bioessays* **33**, 438-447 (2011).
44. L. M. Jones, D. A. Depireux, D. J. Simons, A. Keller, Robust temporal coding in the trigeminal system. *Science* **304**, 1986-1989 (2004).
45. C. Massot, A. D. Schneider, M. J. Chacron, K. E. Cullen, The vestibular system implements a linear–nonlinear transformation in order to encode self-motion. *PLoS Biol.* **10**, e1001365 (2012).
46. G. De Palo, G. Facchetti, M. Mazzolini, A. Menini, V. Torre, C. Altafini, Common dynamical features of sensory adaptation in photoreceptors and olfactory sensory neurons. *Sci. Rep.*, **3**, 1-8 (2013).
47. S. G. Leinwand, C. J. Yang, D. Bazopoulou, N. Chronis, J. Srinivasan, *et al*., Circuit mechanisms encoding odors and driving aging-associated behavioral declines in Caenorhabditis elegans. *Elife* **4**, e10181 (2015).
48. B. T. Sagdullaev, M. A. McCall, P. D. Lukasiewicz, Presynaptic inhibition modulates spillover, creating distinct dynamic response ranges of sensory output. *Neuron* **50**, 923-935 (2006).
49. K. M. Hallinen, R. Dempsey, M. Scholz, X. Yu , A. Linder, *et al.*, Decoding locomotion from population neural activity in moving C. elegans. *Elife* **10**, e66135 (2021).
50. I.T. Jolliffe, Principal Component Analysis, Second Edition (Springer, 2002).
51. H. Straka, J. Simmers, B. P. Chagnaud, A new perspective on predictive motor signaling. *Curr. Biol.* **28**, R232-R243 (2018).
52. T. B. Crapse, M. A. Sommer, Corollary discharge across the animal kingdom. *Nat. Rev. Neurosci.* **9**, 587-600 (2008).
53. N. A. Croll, Components and patterns in the behaviour of the nematode Caenorhabditis elegans. *J. Zoology* **176**, 159-176 (1975).
54. V. Susoy, W. Hung, D. Witvliet, J. E. Whitener, M. Wu, *et al.*, Natural sensory context drives diverse brain-wide activity during C. elegans mating. *Cell* **184**, 5122-5137 (2021).

55. D. R. Albrecht, C. I. Bargmann, High-content behavioral analysis of Caenorhabditis elegans in precise spatiotemporal chemical environments. *Nat. Methods* **8**, 599-605 (2011).
56. E. A. Pnevmatikakis, A. Giovannucci, NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *J Neurosci Methods.* **291**, 83-94 (2017).
57. D. A. Fish, A. M. Brinicombe, E. R. Pike, J. G. Walker, Blind deconvolution by means of the Richardson–Lucy algorithm. *JOSA A* **12**, 58-65 (1995).
58. L. Paninski, J. Pillow, J. Lewi, Statistical models for neural encoding, decoding, and optimal stimulus design. *Prog Brain Res.* **165**, 493-507 (2007).
59. D. Bates, M. Mächler, B. Bolker, S. Walker, Fitting linear mixed-effects models using lme4. *arXiv* [Preprint] (2014). https://arxiv.org/abs/1406.5823 (accessed 31 October 2021).

Figures and Tables

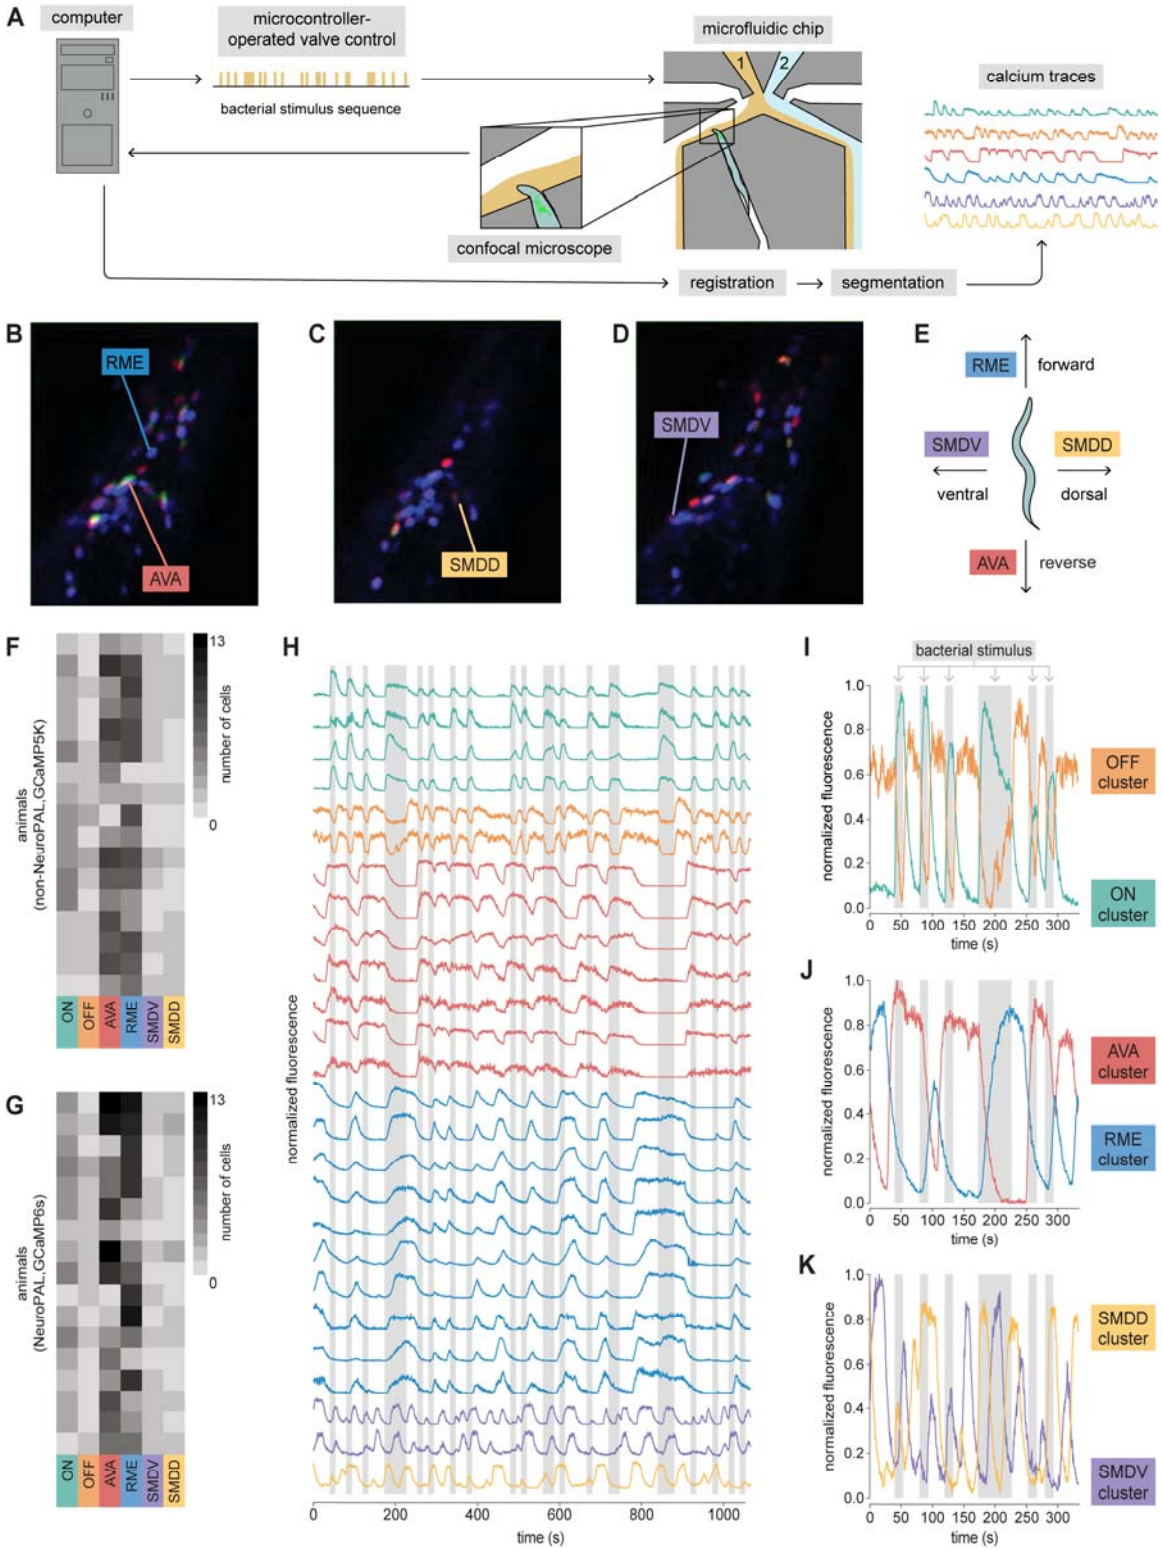

Figure 1. Food-stimulated whole-brain activity reveals six functional clusters. (A) Stimulus presentation and imaging setup. A computer precisely controls delivery of a stimulus sequence that alternates between variable-length pulses of two liquid flows: bacterial food stimulus (gold) and control buffer (light blue). This stimulus sequence is presented to the nose of a *C. elegans* animal that is restrained in a microfluidic chip. Volumes of the *C. elegans* head are acquired and subsequently processed to acquire calcium traces. (B–D) Identification of neurons in the NeuroPAL–GCaMP6s strain. Some photobleaching occurred due to acquisition after calcium imaging (see Methods: Cell Identification). (B) RME motor neuron and AVA command neuron. (C) SMDD motor neuron. (D) SMDV motor neuron. (E) Neurons and their associated direction of locomotion. (F–G) Number of neurons in each functional cluster for all animals, based on > 85% correlation with representative neurons. (F) non-NeuroPAL–GCaMP5K animals. (G) NeuroPAL–GCaMP6s animals. (H) Calcium activity of low-noise active neurons for a single worm. Gray shading represents bacterial stimulus pulse duration. (I–K) Calcium activity averaged across all traces within a functional cluster. Colors of average calcium traces corresponds with colors of individual traces in (H). (I) OFF and ON clusters. (J) AVA and RME clusters. (K) SMDD and SMDV clusters.

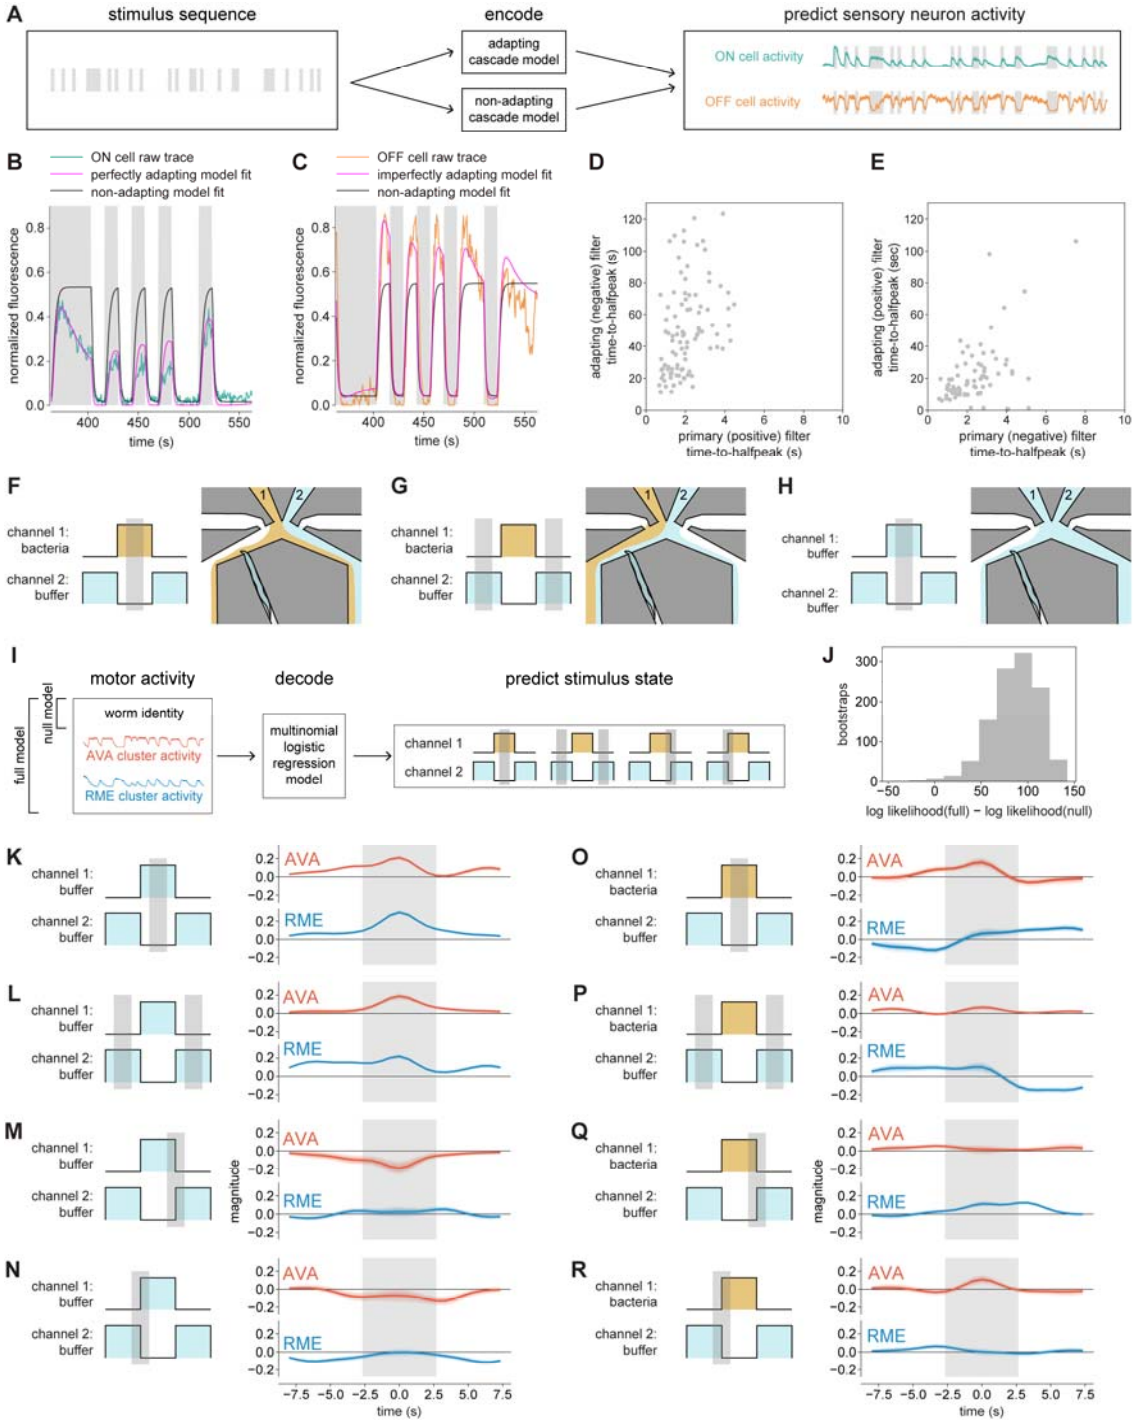
A
stimulus sequence
encode
adapting cascade model
non-adapting cascade model
predict sensory neuron activity
ON cell activity
OFF cell activity
B
ON cell raw trace
perfectly adapting model fit
non-adapting model fit
normalized fluorescence
time (s)
C
OFF cell raw trace
imperfectly adapting model fit
non-adapting model fit
normalized fluorescence
time (s)
D
adapting (negative) filter time-to-halfpeak (s)
primary (positive) filter time-to-halfpeak (s)
E
adapting (positive) filter time-to-halfpeak (sec)
primary (negative) filter time-to-halfpeak (s)
F
channel 1: bacteria
channel 2: buffer
G
channel 1: bacteria
channel 2: buffer
H
channel 1: buffer
channel 2: buffer
I
motor activity
decode
predict stimulus state
null model
full model
worm identity
AVA cluster activity
RME cluster activity
multinomial logistic regression model
channel 1
channel 2
J
bootstraps
log likelihood(full) − log likelihood(null)
K
L
M
N
O
P
Q
R
channel 1: buffer
channel 2: buffer
channel 1: bacteria
channel 2: buffer
AVA
RME
magnitude
time (s)

Figure 2. Bacteria onset and removal differentially drive activity of sensory, AVA, and RME clusters. (A) Stimulus sequence features are fed into an encoding model (see Materials and Methods: Encoding model) to predict sensory cluster activity. (B–C) Representative examples of sensory neurons raw traces compared with model predictions. (B) ON cell. (C) OFF cell. (D–E) Time-to-half-peak distributions for the best performing (D) ON and (E) OFF cell adapting models. (F–H) Stimulus states (left, highlighted in gray) and their corresponding flow configurations relative to the *C. elegans* nose. (F) Bacteria from channel 1 flows over nose. (G) Buffer from channel 2 flows over nose. (H) As a control, buffer emanates from both channel 1 and channel 2. (I) AVA and RME cluster activity, as well as worm identity, is fed into a decoding model (see Materials and Methods: Decoding model) to predict different stimulus states (highlighted in gray). In the null model, only worm identity is used. (J) Out-of-bootstrap cross-validation performance of the full decoding model that includes AVA and RME cluster activity. (K–R) Temporal filters predicting stimulus states from AVA and RME cluster activity for buffer↔buffer stimulus sequences (K-N) and bacteria↔buffer sequences (O–R). Gray shading represents the prediction time window (with 0 s as the halfway point of the prediction window), such that preceding time represents baseline activity and subsequent time represents delayed effects. Median bootstrap temporal filters are plotted, with graded shading indicating 50%, 75%, and 90% of bootstraps.

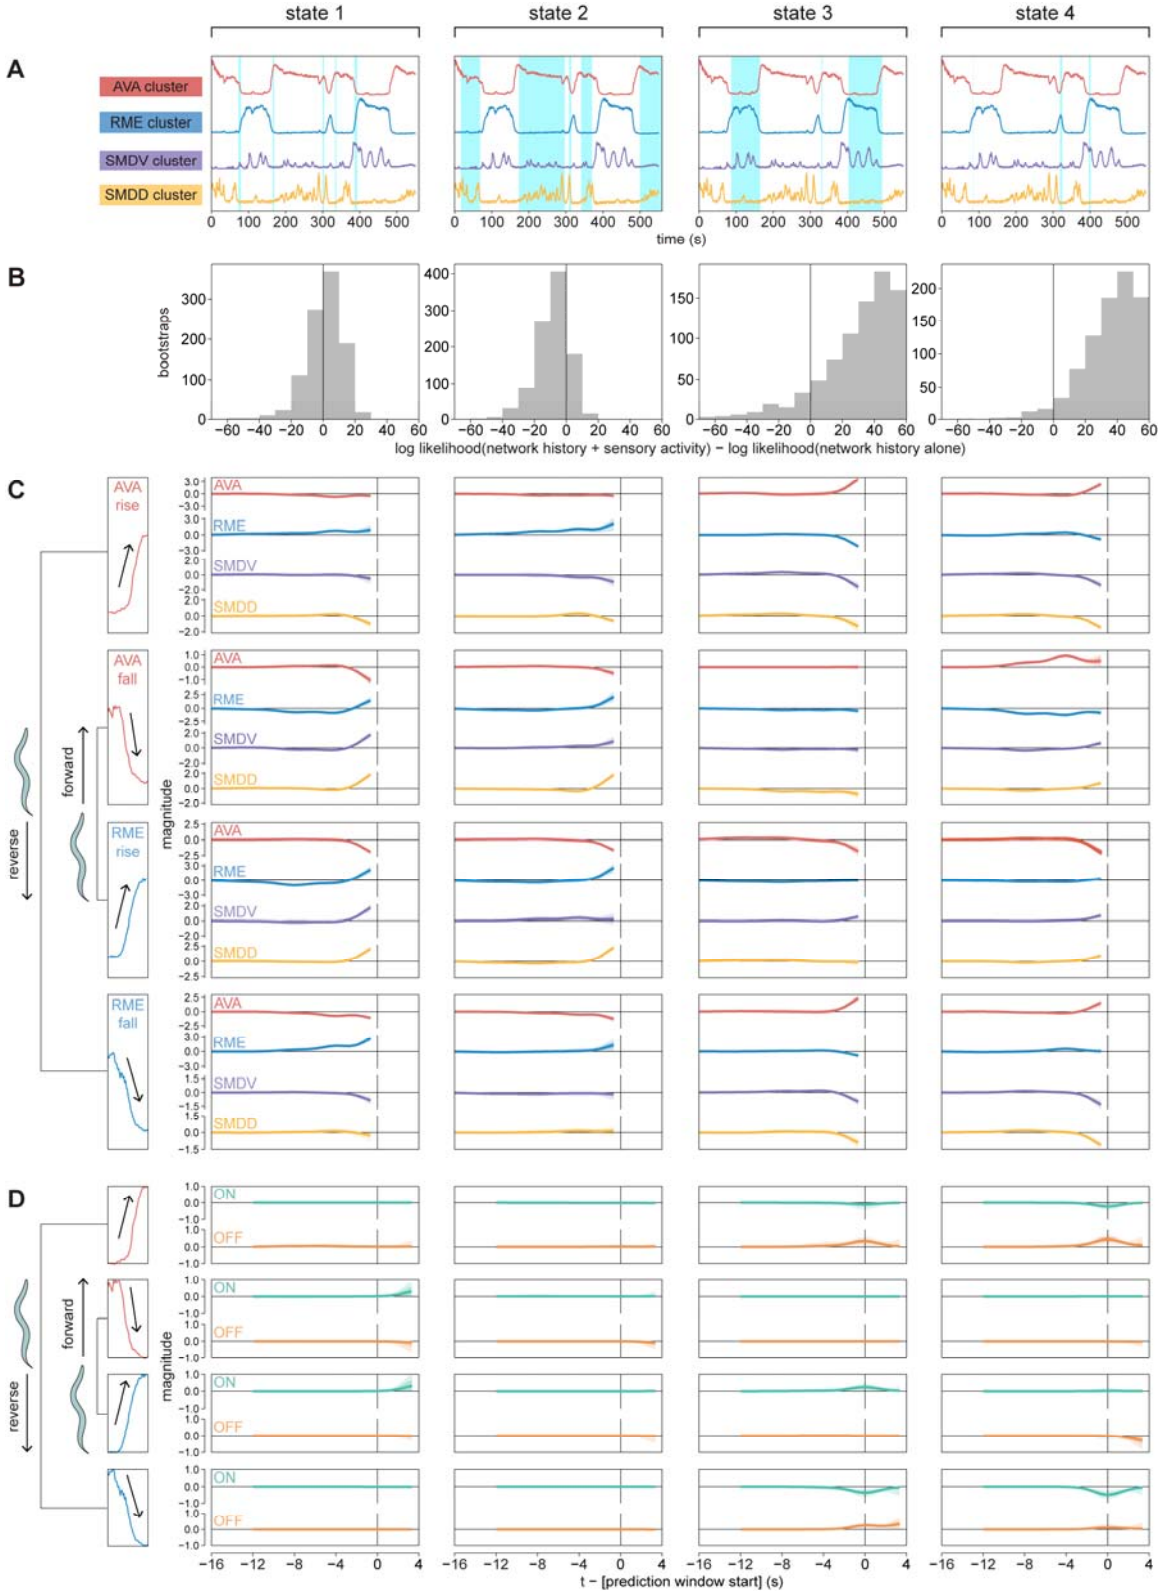
state 1
state 2
state 3
state 4
A
AVA cluster
RME cluster
SMDV cluster
SMDD cluster
time (s)
B
bootstraps
log likelihood(network history + sensory activity) − log likelihood(network history alone)
C
AVA rise
AVA fall
RME rise
RME fall
forward
reverse
magnitude
AVA
RME
SMDV
SMDD
D
ON
OFF
t − [prediction window start] (s)

Figure 3. Identification of interpretable network states that vary in sensory gating properties. (A) The top four network states identified by the soft decision tree gating model portion of the SDT-MLR forecasting model (see Fig. S5, Materials and Methods: SDT-MLR model). Cyan shading represents the time windows for which a particular state's probability exceeds 0.75. (B–C) Temporal filters predicting rise and fall of AVA and RME cluster activity from (B) recent history of AVA, RME, SMDV, and SMDD cluster motor/command activity; and (C) ON and OFF cell cluster sensory activity. (B–C) Median bootstrap temporal filters are plotted, with graded shading indicating 50%, 75%, and 90% of bootstraps. (D) Difference in out-of-bootstrap cross-validation performance between models that included both network history and sensory activity and models that that only included network history.

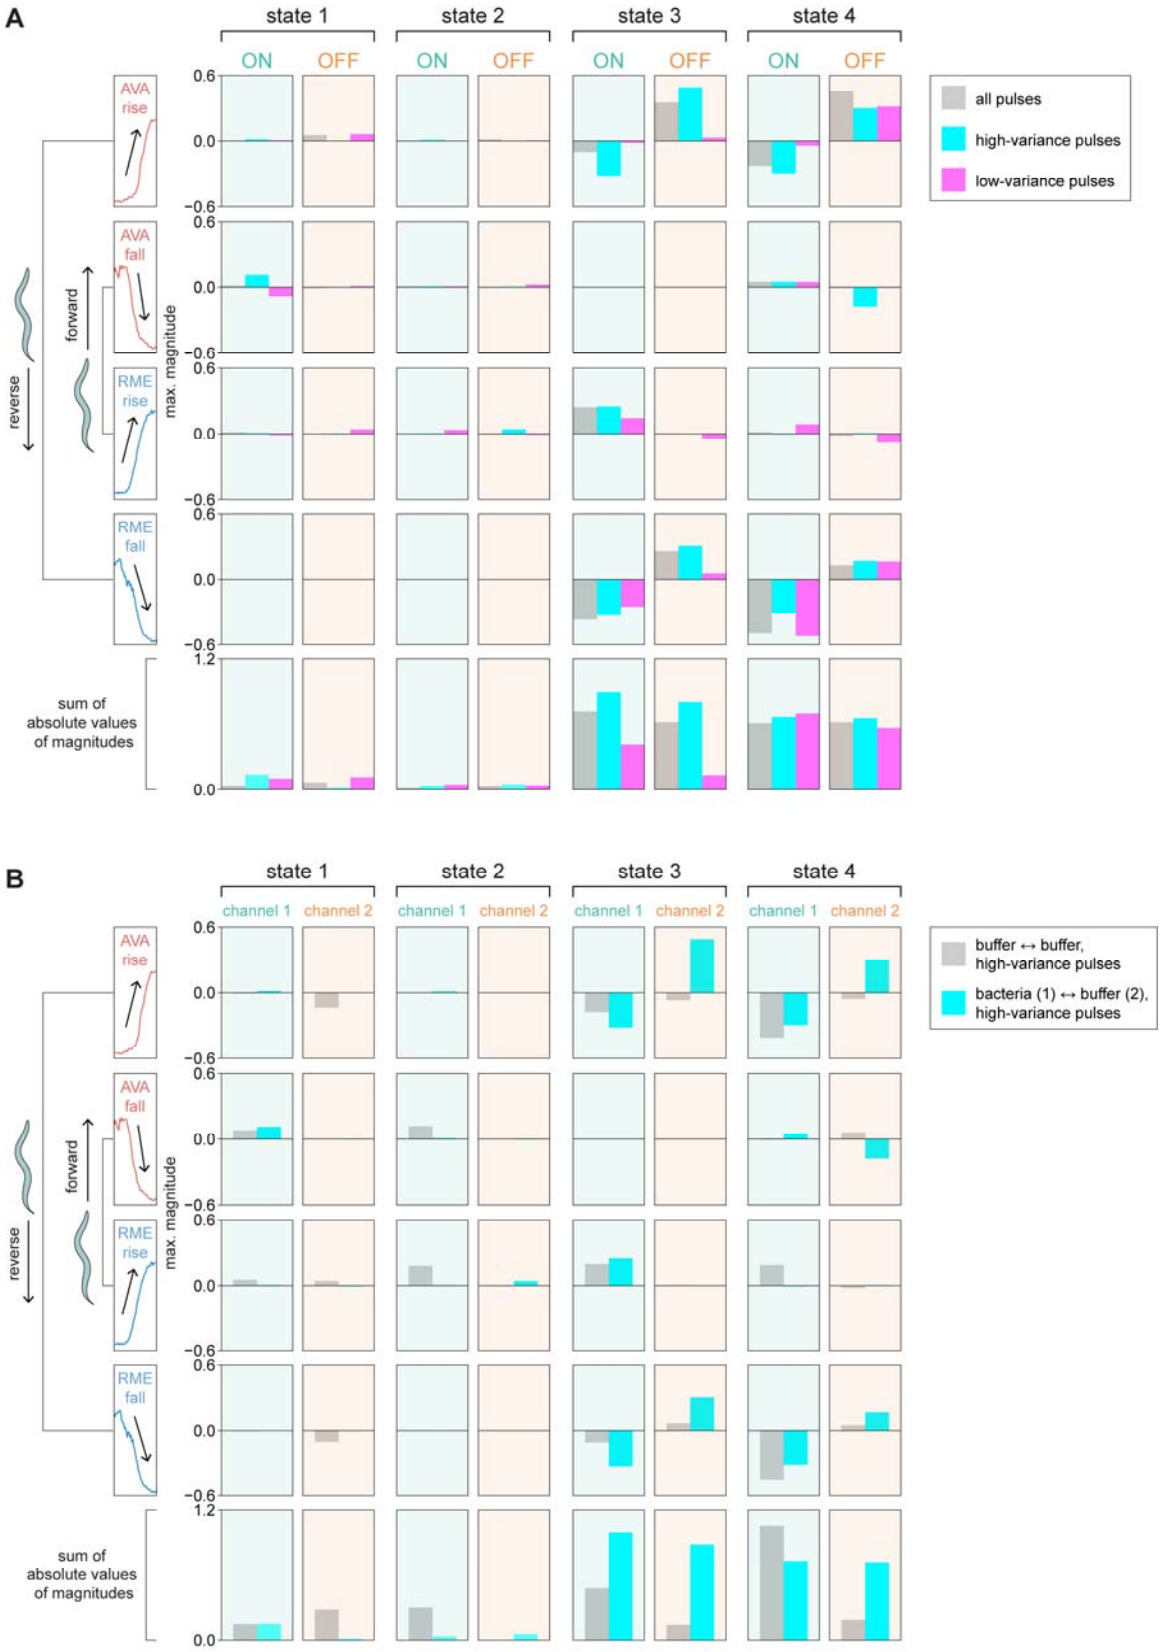
A
state 1
state 2
state 3
state 4
ON
OFF
AVA rise
AVA fall
RME rise
RME fall
forward
reverse
max. magnitude
sum of absolute values of magnitudes
all pulses
high-variance pulses
low-variance pulses
B
channel 1
channel 2
buffer ↔ buffer, high-variance pulses
bacteria (1) ↔ buffer (2), high-variance pulses

973

Figure 4. Effect of stimulus timescales and sensory context on AVA and RME clusters. (A–B) Maximum magnitudes of temporal filters predicting rise and fall of AVA and RME activity from sensory neuron clusters. Maximum magnitudes are calculated from *t*-12s to *t*-0s relative to the start of the prediction window (*t*). Sums of absolute values of magnitudes measure the overall sensory influence within a state. (A) Comparison of all stimulus pulse lengths, only high-variance pulses, and only low-variance pulses. (B) Comparison of buffer↔buffer and bacteria↔buffer stimulus sequences. Only high-variance pulses were compared. Channel 1 and 2 are the same as ON and OFF, respectively, for bacteria↔buffer stimulus sequences, as seen in (A).