Unsupervised learning of control signals and their encodings in *C. elegans* whole-brain recordings

Charles Fieseler[a], **Manuel Zimmer**[b], **and J. Nathan Kutz**[c]

[a, 1]Department of Physics, University of Washington, Seattle, WA 98195; [b]Research Institute of Molecular Pathology, Vienna, Austria; [c]Department of Applied Mathematics, University of Washington, Seattle, WA 98195

This manuscript was compiled on August 26, 2019

Recent whole brain imaging experiments in *C. elegans* have revealed that the neural connectomic dynamics live on a low dimensional manifold with stereotyped transitions between behaviors. Typical theoretical efforts use data to produce a set of local linear models characterizing the data, but it is unknown how a single, global neural network model can generate the observed dynamics. We propose instead a control framework to achieve a global model of the dynamics, whereby underlying linear dynamics is actuated by sparse control signals. The method learns the control signals in an unsupervised way from data, then uses *Dynamic Mode Decomposition with control* (DMDc) to create the first global, linear dynamical system that can reconstruct whole-brain imaging data. These internally generated control signals are shown to be implicated in transitions between behaviors. In addition, we analyze the time-delay encoding of these control signals, both reproducing known neural encodings and showing that these transitions can be predicted from previously unknown neurons. Moreover, our decomposition method allows one to understand the observed nonlinear global dynamics instead as linear dynamics with control. The proposed mathematical framework is generic and can be generalized to other neurosensory systems, potentially revealing transitions and their encodings in a completely unsupervised way.

C elegans | Control | DMD | ...

The nematode *Caenorhabditis elegans* (*C. elegans*) is an ideal model organism as it is comprised of only 302 sensory, motor and inter-neurons whose stereotyped electrophysical connections (i.e. its connectome) are known from serial section electron microscopy (1). Indeed, *C. elegans* is perhaps the simplest biophysical organism to display many of the hallmark features of high-dimensional networked biological systems, including the manifestation of low-dimensional patterns of activity associated with functional behavioral responses. Thus the nervous system encodes behavior by reducing the high-dimensional representation of the environmental stimulus into a much lower representations of motor command (2–6). Low dimensional representations have been separately considered in posture (behavioral) analysis (7, 8) as well as in the static analysis of calcium imaging data (9, 10). Understanding the computational processing that transforms sensory input into motor representations requires the ability to record the activity of sensory neurons, decision-making circuits, and motor circuits in a behaving animal, something that is now largely available with modern imaging of *C. elegans*. Real-time, whole-brain imaging of these non-spiking neurons allows for data-driven discovery of the governing dynamics of the system and the low-dimensional manifold (coordinates) on which neural activity exists (11). In this work, we exploit this new, whole-brain imaging technology to posit a data-driven model of neurosensory integration in *C. elegans*, showing that a global, linear control framework alone explains and reproduces much of the activity of the network.

It has long been observed that *C. elegans* produces a small number of stable discrete behaviors (e.g. forward and backward motion, and turns), and that these behaviors change both spontaneously and very quickly in response to external stimuli and/or stimulation of even a single neuron (12–14). A potential dynamical systems explanation for this observation is that of discrete behaviors as fixed points on an underlying manifold with some transition signals that move the system between them. A purely linear model cannot produce multiple fixed points, but switching (hybrid) linear dynamical systems methods (15–18) circumvent this by segmenting the dynamics into patches with different dynamics (and thus different fixed points) in each patch. An alternate method uses different phase loops and the phase along them to predict behavior, producing conserved dynamics in a special phase space (19). Recent efforts have also attempted to explicitly model the nonlinear connectomic dynamics (6, 20–25), but this has currently been limited to subsets of neurons and has moreover had difficulty capturing multiple behaviors. This work instead focuses on how a single, global, neuron-level model with simple and interpretable additions can capture the nonlinear dynamics

Significance Statement

Biological organisms are very well adapted to the environments they live in, and can perform tasks that even our most advanced engineered cannot. For example, the soil-dwelling nematode *C. elegans* lives in a noisy environment made of vastly differing materials and viscosities. Sensory stimuli come from many different channels, and yet this "worm" is able to consistently integrate them and react appropriately, all with only 302 neurons. We use data-driven techniques to analyze real-time neuronal, whole-brain recordings in order to understand the separation of the system into a simple intrinsic linear dynamics and a more complex set of internally generated control signals. Our 3-step modeling framework can reproduce whole-brain imaging datasets from initial conditions and an interpretable control signal using Dynamic Mode Decomposition with Control (DMDc). Biologists can use each step of our framework to learn, possibly in new organisms: 1) where behavioral transitions occur; 2) about the complexity of different behaviors; and 3) which neurons produce the transition signals. Theorists can use this framework to: 1) build low-dimensional models characterizing connectomic dynamics; 2) identify distinct system states; and 3) build closed-loop feedback models.

The authors declare no conflict of interest.

[1]To whom correspondence should be addressed. E-mail: konda uw.edu

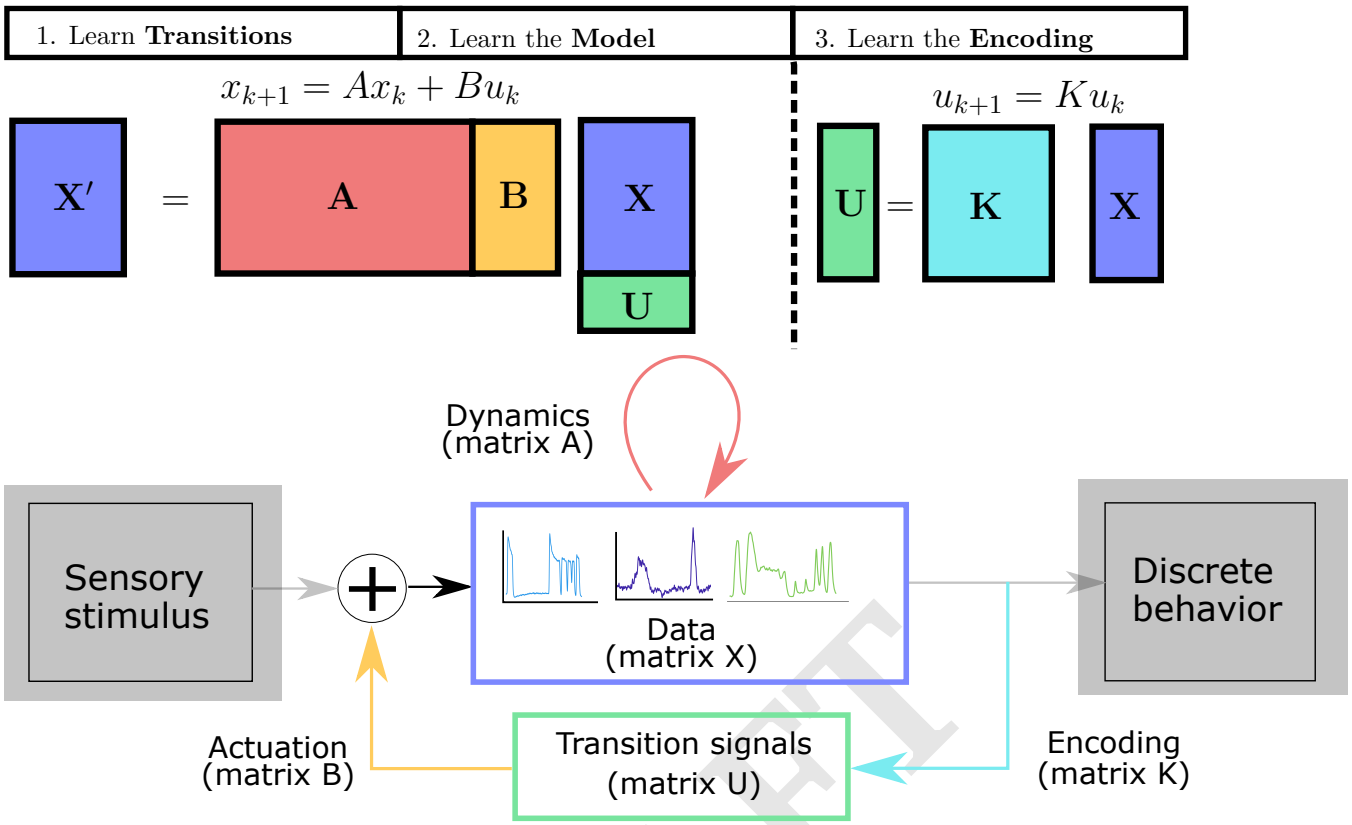

Fig. 1. A 3-step framework for modeling neurosensory integration. 1) Transition signals are learned from data with an assumption of linear dynamics. 2) A DMDc model is learned which uses dynamics, transition signals, and actuation. These are global models, and are capable of reconstructing much of the data dynamically from an initial state. 3) Where and at what timescales control signals are encoded in the neural activity is studied using sparse linear models.

simply by appropriate framing as a control problem.

The recent availability of real-time calcium imaging data allows for a neuron-level data-driven approach. A full model of *C. elegans* neural activity should describe how multiple states are produced in a single network, and how dynamics operating at multiple scales are integrated to produce the states and transitions between them. This can be naturally expressed mathematically using control theory, with the data-driven method of *Dynamic Mode Decomposition with control* (DMDc) (4, 26) providing a regression framework for approximating linear control laws. We propose a mathematical framework for building such a model via a step-by-step analysis of the required components necessary for DMDc, generating interpretable and testable hypotheses at each step. This framework is able to 1) learn known and novel transition signals; 2) reconstruct entire datasets, including with multiple states, demonstrating that additional nonlinearities are not needed to describe many of the interactions in the system; and 3) analyze the timescales and locations of where these transition signals are encoded. We provide code written in MATLAB (27) for a full analysis pipeline that uses raw data and, if available, external behavioral labels to discover both the intrinsic dynamics and the effects of control on the state of the system.

Data-Driven Methods

Our analysis relies on two established mathematical methods: DMDc and sparse optimization. A brief summary of each is given below.

Dynamic Mode Decomposition with control. Our data-driven strategy is based upon the *dynamic mode decomposition* (DMD). DMD provides a linear model for data matrices constructed using temporal snapshots of the state space, $\mathbf{X} = [\mathbf{x}_1\ \mathbf{x}_2\ \dots\ \mathbf{x}_{m-1}]$ and $\mathbf{X}' = [\mathbf{x}_2\ \mathbf{x}_3\ \dots\ \mathbf{x}_m]$ where $\mathbf{x}_j = \mathbf{x}(t_j)$. Specifically, it finds the best fit linear dynamical system

$$\mathbf{X}' = \mathbf{A}\mathbf{X} \qquad [1]$$

There are a number of variants for computing $\mathbf{A}$ (4), with the *exact DMD* simply positing $\mathbf{A} = \mathbf{X}'\mathbf{X}^\dagger$ where † denotes the Moore-Penrose pseudo-inverse.

DMDc (26) capitalizes on all of the advantages of DMD and provides the additional innovation of being able to disambiguate between the underlying dynamics and actuation. For a control input matrix $\mathbf{U} = [\mathbf{u}_1\ \mathbf{u}_2\ \dots\ \mathbf{u}_{m-1}]$ where $\mathbf{u}_j = \mathbf{u}(t_j)$, DMDc regresses instead to the linear control system

$$\mathbf{X}' = \mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{U}. \qquad [2]$$

Note that DMDc uses only snapshots in time of the state space and control input, making it compelling for systems whose governing equations are unknown. The DMDc equation is graphically represented in Fig. 1. The governing matrices ($\mathbf{A}$ and $\mathbf{B}$) along with the control signal ($\mathbf{U}$) produce a predictive model, such that the state of the system far in the future can be predicted. For instance, the third time step can be estimated from the first via:

$$\mathbf{x}_3 = \mathbf{A}(\mathbf{A}\mathbf{x}_1 + \mathbf{B}\mathbf{u}_1) + \mathbf{B}\mathbf{u}_2 \qquad [3]$$

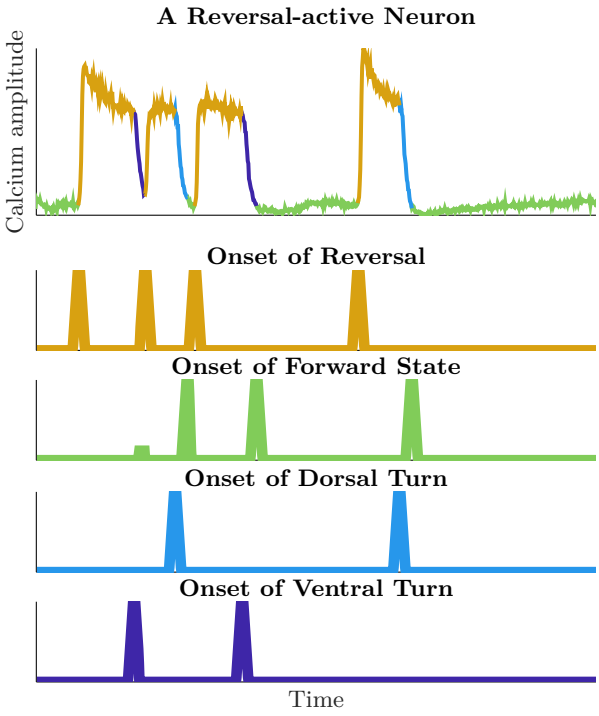

Fig. 2. Transition signals in *C. elegans*: Top: A calcium imaging trace of a neuron connected with the discrete reversal behavior. Behavioral labels are determined by experimentalists, as described in (9). Green=Forward; Yellow=Reversal; Dark Blue=Ventral Turn; Light Blue=Dorsal Turn. Below: These labels can be reframed as "onset" signals, and are characteristically sparse in time.

Figure 4 refers to "full reconstructions" of the datasets, i.e. the prediction of snapshots up to 3000 time steps in the future given the initial data snapshot ($\mathbf{x}_1$) and the full time series of the control signals ($\mathbf{U}$).

Learning control signals via sparse optimization. The DMDc algorithm requires knowledge of the linear control signals $\mathbf{U}$. Expert-identified state labels and an example neuron that displays strong state-dependent behavior are shown in figure 2. However, these are only available because of the decades of *C. elegans* experimental work identifying 1) discrete behavioral states and 2) the command neurons for each activity. For new organisms, and in order to generate hypotheses about potential new states in *C. elegans*, the unsupervised problem, i.e. learning the signal directly from data, is of critical interest.

DMDc (2) can be thought of as an error minimization problem over the dynamics matrices, $\mathbf{A}$ and $\mathbf{B}$. If the control signal is unknown, the minimization must be extended to the control signal $\mathbf{U}$ itself. However, there is now a trivial solution where the control signal dominates the model: $\mathbf{X}_2 = \mathbf{B}\mathbf{U}$ with $\mathbf{A} = 0$. For this reason, an assumption must be made about the control signals. In this case, the statement that these signals are sparse is directly biologically interpretable, and means that the transitions between states should be rare as a percentage of frames. This "sparsity constraint" can be directly expressed in the language of optimization, i.e. the ℓ_0 norm:

$$\min_{\mathbf{A},\mathbf{B},\mathbf{U}} \left[||\mathbf{A}\mathbf{X}_1 + \mathbf{B}\mathbf{U} - \mathbf{X}_2||_2 + \lambda ||\mathbf{U}||_0 \right] \quad [4]$$

Directly solving this optimization problem is extremely difficult, although there are efficient algorithms in certain cases

Algorithm 1 Unsupervised Learning of Control Signals

```
procedure LearnControllers(r)
    U_0 := InitializeU(r)
    S := InitializeSparsityPattern(U_0)
    for i ← 1, MaxIter do
        A, B = SolveAB(X, U_{i-1})          ▷ Solves eq. 2
        U_i = SolveU(X, A, B)
        S = UpdateSparsityPattern(S, U_i)
        U_i(S) = 0
```

(28). More recently, a convex relaxation of the ℓ_0 to an ℓ_1 norm is often solved (29), though this has been recently shown to make mistakes (30). We use a different approximation, the sequential least squares thresholding algorithm (31), which has been shown to converge to the minima of the original ℓ_0 problem (32). The code is outlined in algorithm 1 and more detail is given in the supplement. The matrix $\mathbf{U}$ in this algorithm is additionally constrained to be positive, for better interpretability as "on" transition signals.

Variable selection via sparse linear models. If internally generated control signals are present, then there are two options: they are random and fundamentally unpredictable, or they are encoded in the network. Although there is almost certainly some amount of stochasticity in the true biological system, any encoding at all can be used to study the initiations and precursors of the behavior. Mathematically, this is a variable selection problem: given the data, which few neurons predict the transitions? In this paper we additionally use *time-delay embedding* where data from further in the past is utilized:

$$\mathbf{U} = \mathbf{K}_1\mathbf{X}_1 + \mathbf{K}_2\mathbf{X}_2 + ... \quad [5]$$

There are multiple methods that are often used to perform this variable selection task (33). However, these methods may make mistakes in their selections (30), and in general it is unclear how unique the selection is. The behaviors of *C. elegans* have been well studied, and each onset is associated with well-known neurons. Variable selection methods will almost certainly discover these well-known neurons, but by exploring further in the "elimination path", less obvious encodings can be discovered. Algorithmically, this is the sequential removal of the most important neuron for all time delays, and then a re-fitting of the sparse model. If the quality of the reconstruction does not degrade along the elimination path, the signal ($\mathbf{U}$) must be distributed throughout the data ($\mathbf{X}$). The quality of signal reconstruction is defined here as the number of false positives and false negatives in the reconstructed signal. Event detection is defined as a number of frames above a hard threshold, as shown in Fig. 5 and discussed in the supplement.

Results

Known transitions are discovered and characterized. Experimentalists have long separated behavior into discrete behaviors through careful study of individual neurons. However, open questions remain about the number of behaviors that exist and how discrete they are. Some works have posited up to six forward motion states and three reversal states, multiple turning subtypes, and even a continuum of behaviors (34). As Fig. 3 shows, using unsupervised optimization three behavioral onsets can be discovered: Reversal, and Dorsal and

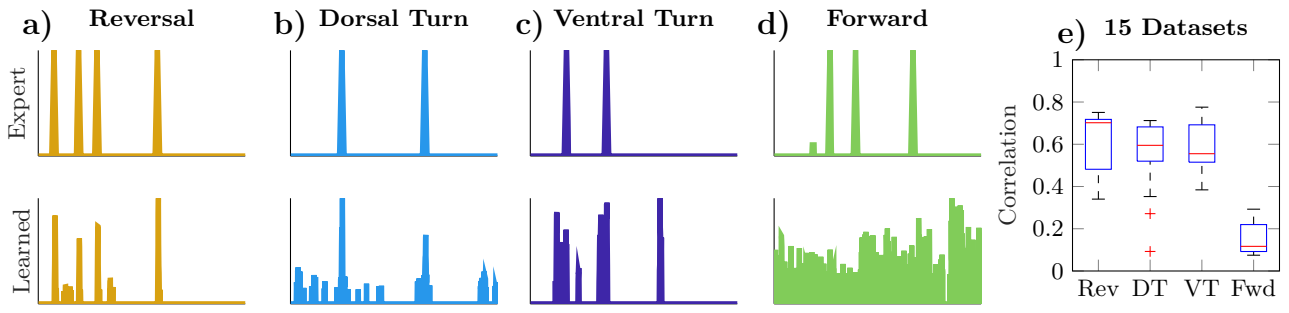

Fig. 3. Control signals can be learned from data via algorithm 1. a-d) The onset of well-known states as determined by experts (above) and as learned (below). e) Correlation between expert and learned signals across 15 individual datasets. Reversals (Rev), Dorsal (DT) and Ventral Turns (VT) are consistently learned, but Forward state (Fwd) onsets are never significant.

Ventral turns. In particular, the single Reversal onset signal for each individual suggests that this transition is fundamentally the same within individuals, with variability produced by activation amplitude not a different direction in neuron-space.

However, in no individuals could a signal correlated to the onset of Forward motion be discovered. This lack of a discovery can be interpreted as the underlying neural mechanism producing forward motion being fundamentally different. It is well known that different, dedicated sub-networks of neurons are active in forward and backwards motion (1), but they are often modeled as mirror images of each other (25). This result implies that for the onset of these behaviors, forward motion is significantly more complex than reversals, not simply that the underlying networks are physically distinct.

A global, linear system with control reconstructs entire time series. The manifold observed in *C. elegans* neural dynamics cannot be described by a purely linear model due to the presence of multiple stable global behaviors, as shown in Fig. 4.b. Specifically, linear models can only admit a single fixed state. However, the majority of neurons can be reconstructed using our *controlled, global, linear* dynamical system due to the sparse transition signals as shown in Fig. 4.c for expert hand-labeled signals and Fig. 4.d for signals learned from data. Each time snapshot of this data is reconstructed analogously to equation 3, and then projected onto the two dominant PCA modes of the original data so that each panel in Fig. 4.a-d is in the same coordinate space. Because this is a global linear model that uses a single framework for the entire state space, the need for additional nonlinear modeling can be constrained to particular groups of neurons and well-defined time windows.

In particular, across individuals the reversal class of neurons is captured very well by the supervised control signal as shown in Fig. 4.j and thus, up to encoding the transition signal itself, the relevant subnetwork does not appear to require nonlinearities. This means that future efforts related to nonlinear modeling should concentrate on the small window of time during the onset of the behavior, instead of the entire neural trace where linear models hold. In addition, the type of nonlinearity required to more fully model this class of neurons is characterized: fast and short-lived spike-like activations.

Turns are also largely captured, as shown by the high correlation for the light and dark blue boxplots. The neurons involved in turning have a large number of smaller events, as shown in the SMDDL reconstruction Fig. 4.f; these do not lead to one of the four state transitions identified by experimentalists in this dataset (CITE timescale nesting paper?), but may correspond to an additional state as discussed in the supplement. However, the unsupervised method does pick up on these smaller events and reconstructs them well Fig. 4.h, but over all datasets there is much more variability as shown in Fig. 4.j.

The last group of neurons, those related to forward motion, has a very large variability of correlation between the data and reconstructions, implying that this state requires nontrivial nonlinearities throughout the time series to capture. It may be continuously parametrized instead of a simple "on" transition signal at the onset, for example by speed, steering (35), or tracking (36). Some recent experimental work (37) characterizes this asymmetry between Forward and Reversal states as due to intrinsic bias towards the Forward state, and this result is consistent with that interpretation but adds that the Forward state is significantly more complex.

To further characterize the effects of the control signals on the ability of this framework to capture the neural dynamics, partial models were created with a subset of control signals. Expert-labeled partial models are shown in Fig. 4.k. Adding Reversal-onset signals alone does not produce a model that captures the data better than a straight-line fit to the data, but the combination of Reversal and Turning signals is significantly better. The subsequent additional of Forward control signals is, remarkably, useless and is another line of evidence showing that this behavior is truly different from the simpler Reversal state.

Transitions are encoded in previously unknown neurons. Having shown the control signals to contribute significantly to the reconstruction of the data, we reconstruct the control signals themselves using time-delay data matrices and sparse linear models as shown in step 3 of Fig. 1 according to equation 5. As described in (10), each of the four interpretable transition signals shown in figure 3 are hand-labeled using the activity of certain well-known neurons. Thus, it is not surprising that these signals can be reconstructed from data when those well-known neurons are included. In particular, as they were used to define the Dorsal Turn behavioral states, an excellent validation is that the SMDDL/R pair of left/right neurons consistently encodes this control signal, as Fig. 5.a shows.

However, as the elimination path is explored further, it is revealed that these well-known neurons can be eliminated from the sparse models and the transition signals can still be

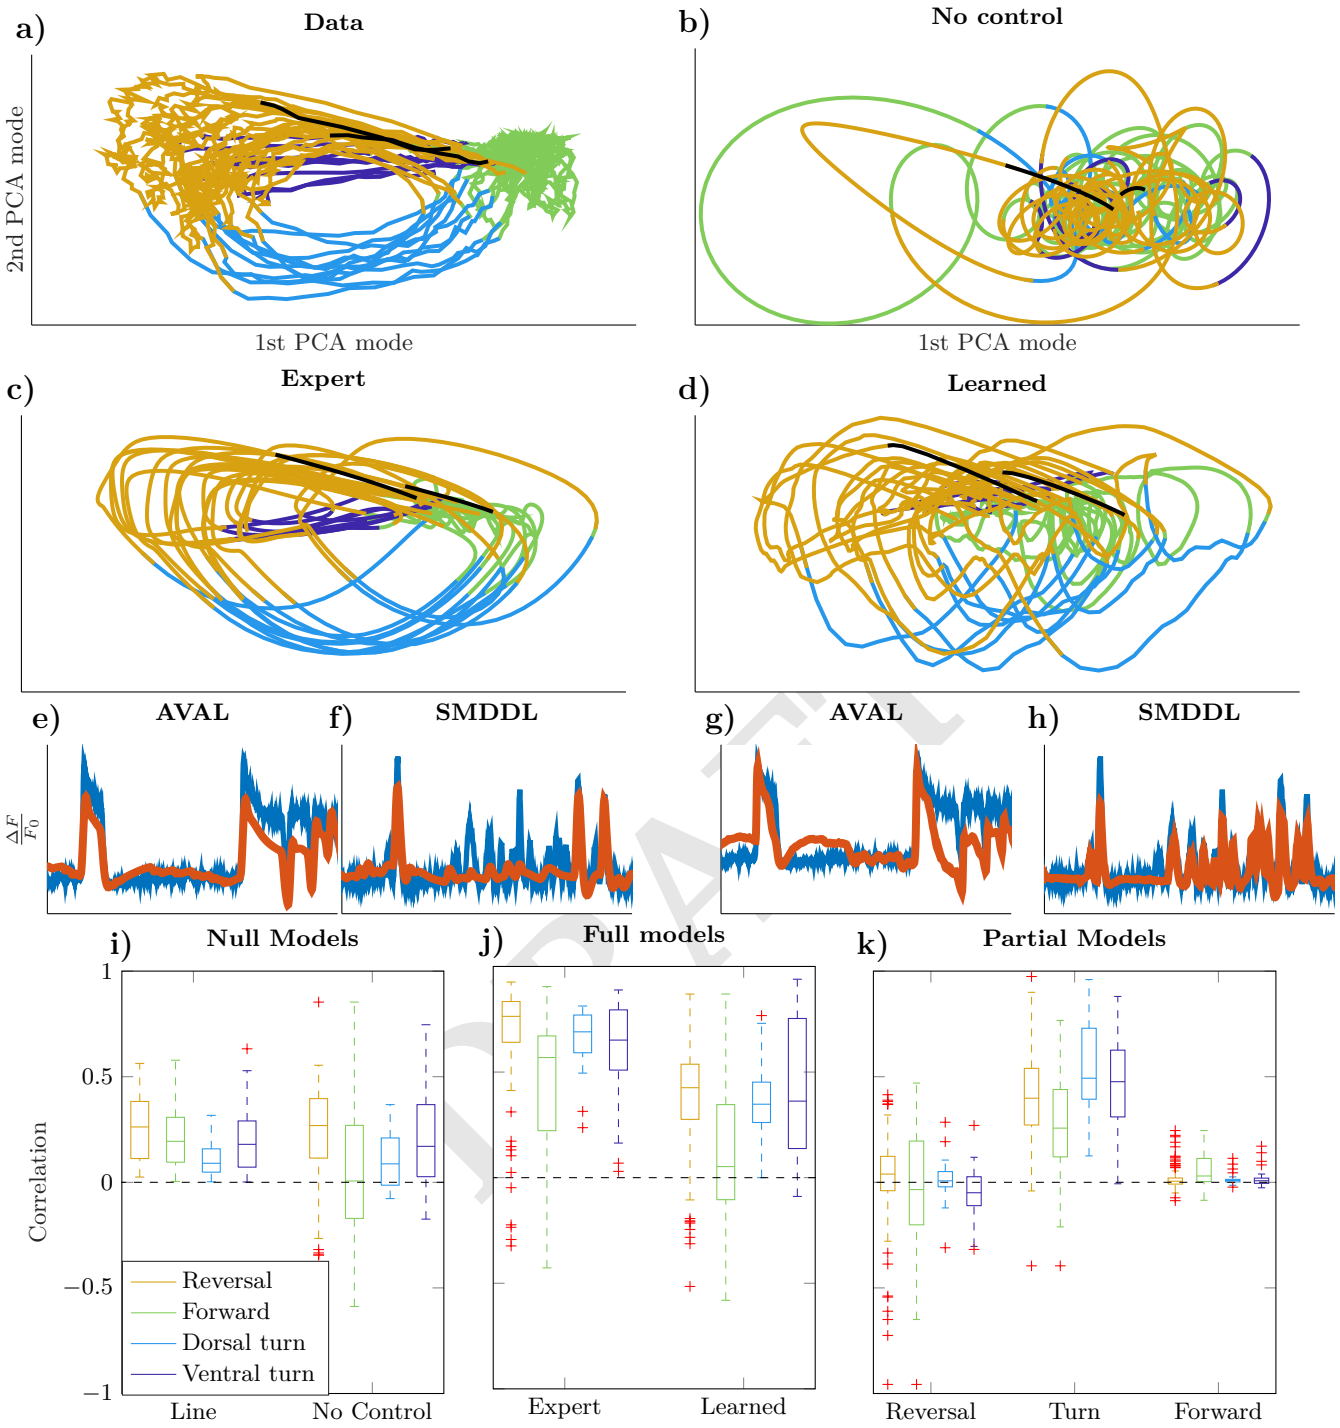

Fig. 4. 2d PCA projections of a) data, b) an uncontrolled "null" model, c) a "supervised" model using expert-determined control signals, and d) an "unsupervised" model that uses control signals learned via algorithm 1. The governing equations matrices in are all learned from data, either uncontrolled (b, Equation 1) or controlled (c-d, Equation 2). These data are color-coded by state: Black for Unknown, Yellow for Reversal neurons, Green for Forward, and Light (Dark) Blue for Dorsal (Ventral) turns. e-f) Example neuron datasets with reconstructions from the supervised model. A reversal-active (AVAL) and a Dorsal-Turn-active neuron (SMDDL) are shown. g-h) The same neurons shown with reconstructions from the unsupervised model. i-k) Correlations across datasets between data and reconstructions, split up into 4 different neuron groupings for interpretability. i) Baseline null models. The left-hand side is simply fitting a straight line to a neural trace. The right-hand side corresponds to the uncontrolled model in panel (b), and is generally worse than a straight-line fit. j) Full models with either expert/supervised control signals, (c) above, or learned/unsupervised control signals, (d) above. For each neuron grouping the expert signals produce significantly better fits. k) Partial supervised models, as more signals are added. Shown are additive improvements, i.e. how much better each partial model is than the one immediately to the left. Specifically, "baseline" of a straight-line fit is subtracted from the Reversal (left-hand side) set, and the Reversal + Turn model correlations are subtracted from the cumulative Forward (right-hand side) set, which has Reversal + Turn + Forward control signals.

reconstructed as shown in 5.b. Indeed, Fig. 5.a and and 5.b look nearly identical, and Fig. 5.c quantifies this using the percentage of false positives and negatives. Fig. 5.c also shows more of the elimination path and when the reconstructions

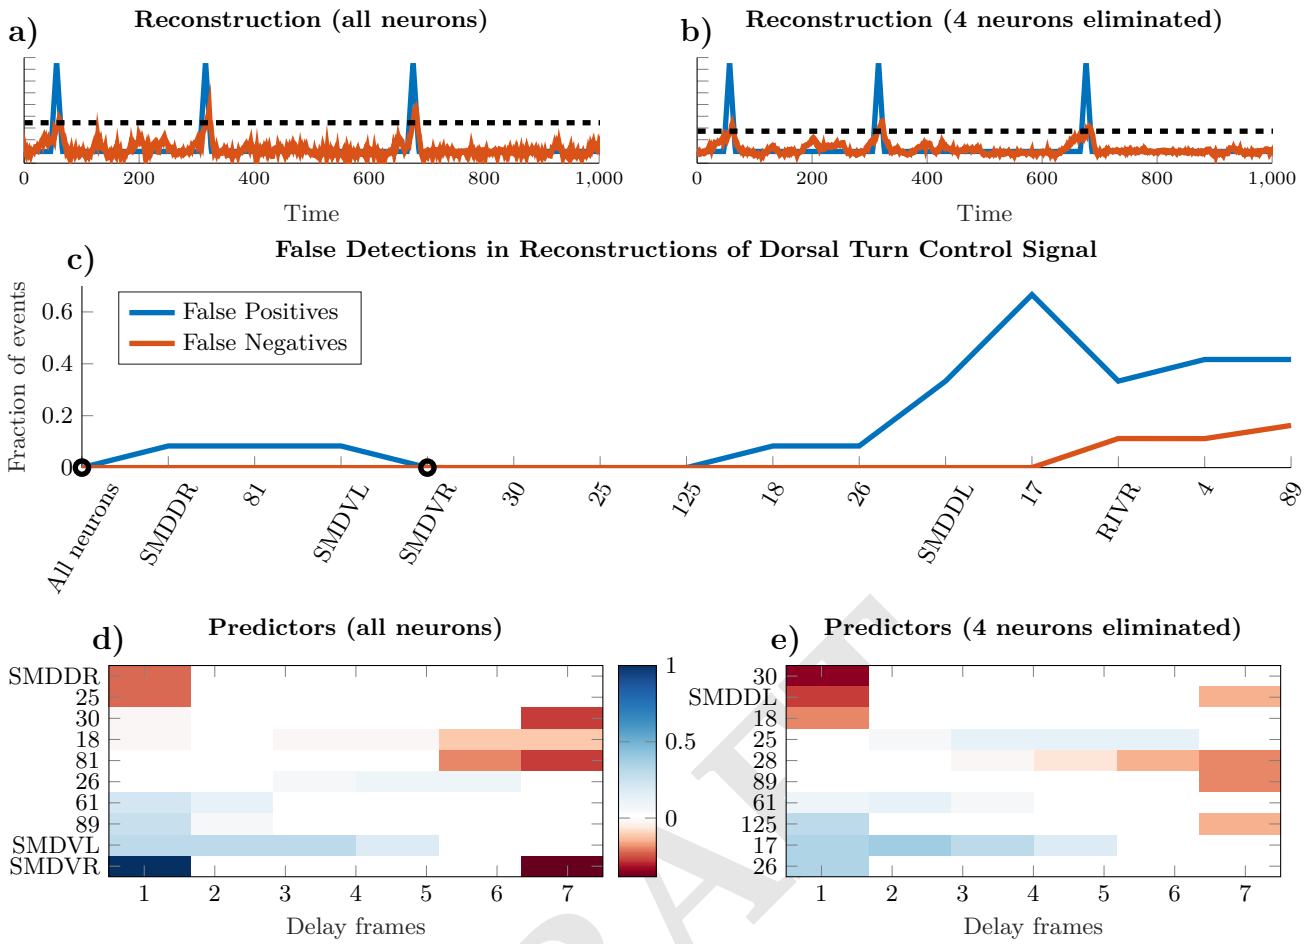

Fig. 5. a-b) Signal reconstructions via linear encoding on the data including time delays, with all neurons (a) or 4 neurons removed (b). Event detection is determined via a simple hard threshold for each signal, as shown by the dotted line. c) Neurons are eliminated in order of the largest magnitude given to them by the linear model. The number of false detections increases significantly only after 8 neurons have been removed. d-e) The weights given to the top 10 most important neurons for different iterations. Even though the prediction is very good as measured by the false positive and negative rates, many of the neurons determined to be important are different.

finally break down. Fig. 5.d and 5.e show the how **K** matrices in equation 5 change as neurons are removed. Taken together, these results reveal previously unknown neurons that can successfully predict control signals shown to be important to reconstructing the full neural manifold. However, only rows with names are neurons that have been connected to the stereotyped *C. elegans* connectome and can thus be identified across individuals; rows with numbers cannot be so compared. This work identifies sets of unknown neurons that could be investigated in further experiments, and the timescale at which they are relevant.

Discussion

We have presented the first data-driven model that uses a single set of intrinsic dynamics that can reconstruct the multiple behavioral regimes present in a real animal and transitions between them. The fact that this controlled linear model accurately reproduces both short and long time-scale dynamics places clear restrictions on the need, specifically the lack thereof, for nonlinearities in this system, and provides hypotheses about the neurons that may contain those nonlinearities and their role in the global dynamics of the system. In addition, we have embedded this model in a mathematical framework of feedback and control, which can be generalized to other organisms or to include hypothesized nonlinearities.

Much excitement has been generated by the availability of the *C. elegans* physiological connectome, and one hope of data-driven modeling efforts is to produce a functional connectome that can complement the physiological data. The DMDc in this paper is similar to several algorithms in the engineering literature that attempt similar network reconstruction tasks, namely System Identification (38). One strategy to fully disambiguate the effects of the intrinsic dynamics and the external control signals uses known external perturbations should be applied and the system response measured. Such perturbations are not generally available in biological systems and thus the data collected are "uninformative" (39) in the sense that the underlying structure cannot be determined.

A limitation of this model is that it is not generative; it cannot be used to predict a system response that includes transitions to novel stimuli. To accomplish this, the transition signals must be written as a function of the data. Step three of our method does this with a linear encoding and demonstrates that the signals can be successfully reconstructed with all neurons to a certain level of accuracy. If this level of accuracy were sufficient, then the system would be fully linear and an

uncontrolled model would produce a good reconstruction, as is clearly not the case. Recent methods for incorporating nonlinearities into controlled systems (e.g. (31, 40)) have the potential to create a fully closed-loop feedback system and this is an active area of further research.

A potential criticism of this method is that we have used discrete labeled states in our model, despite ongoing debate regarding how uniform "states" in *C. elegans* are across instances, and if they should be subdivided or are simply continuous (34). We have contributed to this debate by providing evidence that the reversal and turn states in fact appear to be simple and have well-defined initiation signals, but that the forward "state" is much more complex. We argue that this is an example of a strength of this methodology: the fact that a state cannot be reconstructed gives additional information about that state, and about its complexity in relation to other states.

An alternate approach to modeling complex systems in order to understand structure is to use locally linear models (15–18). In this methodology, the initial network as described by the matrix $\mathbf{A_i}$ is replaced by a new matrix, $\mathbf{A_{i+1}}$, at certain change points. These have achieved great success in reconstructing nonlinear datasets and is an active field in machine learning research. However, it is difficult to interpret what such a replacement of the underlying dynamics would mean biologically, particularly if many separate matrices $\mathbf{A_i}$ are required. On the other hand, the language of control theory from engineering meshes directly with the biological intuition that certain states are initiated by relatively unique signals produced by a small number of neurons. We believe that our framework for constructing a single, global model of the dynamics of this neural system is promising not only in its ready generalizability to include nonlinearities, but also in its biological interpretability.

We have produced the first, to our knowledge, global data-driven model of both the intrinsic and control dynamics of *C. elegans*. We hope this work can contribute to the realization of fully in-silico ablation and actuation experiments, the holy grail of *C. elegans* simulation.

References.

1. White JG, Southgate E, Thomson JN, Brenner S (1986) The structure of the nervous system of the nematode caenorhabditis elegans. *Philos Trans R Soc Lond B Biol Sci* 314(1165):1–340.
2. Roberts WM, et al. (2016) A stochastic neuronal model predicts random search behaviors at multiple spatial scales in c. elegans. *Elife* 5.
3. Liu H, Kim J, Shlizerman E (2017) Functional connectomics from data: Probabilistic graphical models for neuronal network of c. elegans. *arXiv preprint arXiv:1711.00193*.
4. Kutz JN, Brunton SL, Brunton BW, Proctor JL (2016) *Dynamic mode decomposition: data-driven modeling of complex systems*. (SIAM) Vol. 149.
5. Kunert-Graf JM, Shlizerman E, Walker A, Kutz JN (2017) Multistability and long-timescale transients encoded by network structure in a model of c. elegans connectome dynamics. *Frontiers in computational neuroscience* 11:53.
6. Fieseler C, Kunert-Graf J, Kutz JN (2018) The control structure of the nematode caenorhabditis elegans: Neuro-sensory integration and proprioceptive feedback. *Journal of Biomechanics* 74:1–8.
7. Stephens GJ, Johnson-Kerner B, Bialek W, Ryu WS (2008) Dimensionality and dynamics in the behavior of c. elegans. *PLoS computational biology* 4(4):e1000028.
8. Stephens GJ, de Mesquita MB, Ryu WS, Bialek W (2011) Emergence of long timescales and stereotyped behaviors in caenorhabditis elegans. *Proceedings of the National Academy of Sciences* 108(18):7286–7289.
9. Kato S, et al. (2015) Global brain dynamics embed the motor command sequence of caenorhabditis elegans. *Cell* 163(3):656–669.
10. Nichols AL, Eichler T, Latham R, Zimmer M (2017) A global brain state underlies c. elegans sleep behavior. *Science* 356(6344):eaam6851.
11. Venkatachalam V, et al. (2016) Pan-neuronal imaging in roaming caenorhabditis elegans. *Proceedings of the National Academy of Sciences* 113(8):E1082–E1088.
12. Kocabas A, Shen CH, Guo ZV, Ramanathan S (2012) Controlling interneuron activity in caenorhabditis elegans to evoke chemotactic behaviour. *Nature* 490(7419):273.
13. Leifer AM, Fang-Yen C, Gershow M, Alkema MJ, Samuel AD (2011) Optogenetic manipulation of neural activity in freely moving caenorhabditis elegans. *Nature methods* 8(2):147.
14. Nagel G, et al. (2005) Light activation of channelrhodopsin-2 in excitable cells of caenorhabditis elegans triggers rapid behavioral responses. *Current Biology* 15(24):2279–2284.
15. Linderman S, Adams R (2014) Discovering latent network structure in point process data in *International Conference on Machine Learning*. pp. 1413–1421.
16. Linderman SW, et al. (2016) Recurrent switching linear dynamical systems. *arXiv preprint arXiv:1610.08466*.
17. Costa AC, Ahamed T, Stephens GJ (2019) Adaptive, locally linear models of complex dynamics. *Proceedings of the National Academy of Sciences* 116(5):1501–1510.
18. Linderman SW, Nichols AL, Blei DM, Zimmer M, Paninski L (2019) Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans. *bioRxiv* p. 621540.
19. Brennan C, Proekt A (2019) A quantitative model of conserved macroscopic dynamics predicts future motor commands. *eLife* 8:e46814.
20. Kunert J, Shlizerman E, Kutz JN (2014) Low-dimensional functionality of complex network dynamics: Neurosensory integration in the caenorhabditis elegans connectome. *Physical Review E* 89(5):052805.
21. Mujika A, Leškovskỳ P, Álvarez R, Otaduy MA, Epelde G (2017) Modeling behavioral experiment interaction and environmental stimuli for a synthetic c. elegans. *Frontiers in neuroinformatics* 11:71.
22. Costalago-Meruelo A, et al. (2018) Emulation of chemical stimulus triggered head movement in the c. elegans nematode. *Neurocomputing* 290:60–73.
23. Izquierdo EJ (2018) Role of simulation models in understanding the generation of behavior in c. elegans. *Current Opinion in Systems Biology*.
24. Gleeson P, Lung D, Grosu R, Hasani R, Larson SD (2018) c302: a multiscale framework for modelling the nervous system of caenorhabditis elegans. *Philosophical Transactions of the Royal Society B: Biological Sciences* 373(1758):20170379.
25. Boyle JH, Berri S, Cohen N (2012) Gait modulation in c. elegans: an integrated neuromechanical model. *Frontiers in computational neuroscience* 6:10.
26. Proctor JL, Brunton SL, Kutz JN (2016) Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems* 15(1):142–161.
27. Fieseler C (2019) Celegansmodel.
28. Jewell S, Witten D (2018) Exact spike train inference via 0 optimization. *The annals of applied statistics* 12(4):2457.
29. Donoho DL (2006) For most large underdetermined systems of linear equations the minimal 1-norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* 59(6):797–829.
30. Su W, Bogdan M, Candes E, , et al. (2017) False discoveries occur early on the lasso path. *The Annals of Statistics* 45(5):2133–2150.
31. Brunton SL, Proctor JL, Kutz JN (2016) Sparse identification of nonlinear dynamics with control (sindyc). *IFAC-PapersOnLine* 49(18):710–715.
32. Zhang L, Schaeffer H (2018) On the convergence of the sindy algorithm. *arXiv preprint arXiv:1805.06445*.
33. Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1):267–288.
34. Szigeti B, Deogade A, Webb B (2015) Searching for motifs in the behaviour of larval drosophila melanogaster and caenorhabditis elegans reveals continuity between behavioural states. *Journal of the Royal Society Interface* 12(113):20150899.
35. Iino Y, Yoshida K (2009) Parallel use of two behavioral mechanisms for chemotaxis in caenorhabditis elegans. *Journal of Neuroscience* 29(17):5370–5380.
36. Luo L, Clark DA, Biron D, Mahadevan L, Samuel AD (2006) Sensorimotor control during isothermal tracking in caenorhabditis elegans. *Journal of experimental biology* 209(23):4652–4662.
37. Kawano T, et al. (2011) An imbalancing act: gap junctions reduce the backward motor circuit activity to bias c. elegans for forward locomotion. *Neuron* 72(4):572–586.
38. Ljung L (2001) System identification. *Wiley Encyclopedia of Electrical and Electronics Engineering*.
39. Angulo MT, Moreno JA, Lippner G, Barabási AL, Liu YY (2017) Fundamental limitations of network reconstruction from temporal data. *Journal of the Royal Society Interface* 14(127):20160966.
40. Williams MO, Kevrekidis IG, Rowley CW (2015) A data–driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science* 25(6):1307–1346.

ACKNOWLEDGMENTS. Please include your acknowledgments here, set in a single paragraph. Please do not include any acknowledgments in the Supporting Information, or anywhere else in the manuscript.