

Mapping Sub-Second Structure in Mouse Behavior

Highlights

- Computational modeling reveals structure in mouse behavior without observer bias
- Mouse behavior appears to be composed of stereotyped, sub-second modules
- From this perspective, new behaviors result from altering both modules and transitions
- Unsupervised analysis reveals how genes and neural activity impact behavior

Authors

Alexander B. Wiltschko,
Matthew J. Johnson, Giuliano Iurilli, ...,
Victoria E. Abraira, Ryan P. Adams,
Sandeep Robert Datta

Correspondence

srdatta@hms.harvard.edu

In Brief

Mouse behavior appears inherently divided into brief modules of 3D motion. This sub-second structure reveals the influence of the environment, genes and neural activity on action.



Mapping Sub-Second Structure in Mouse Behavior

Alexander B. Wiltschko,^{1,2} Matthew J. Johnson,^{1,2} Giuliano Iurilli,¹ Ralph E. Peterson,¹ Jesse M. Katon,¹ Stan L. Pashkovski,¹ Victoria E. Abraira,¹ Ryan P. Adams,² and Sandeep Robert Datta^{1,*}

¹Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA

²School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

*Correspondence: srdatta@hms.harvard.edu

<http://dx.doi.org/10.1016/j.neuron.2015.11.031>

SUMMARY

Complex animal behaviors are likely built from simpler modules, but their systematic identification in mammals remains a significant challenge. Here we use depth imaging to show that 3D mouse pose dynamics are structured at the sub-second time-scale. Computational modeling of these fast dynamics effectively describes mouse behavior as a series of reused and stereotyped modules with defined transition probabilities. We demonstrate this combined 3D imaging and machine learning method can be used to unmask potential strategies employed by the brain to adapt to the environment, to capture both predicted and previously hidden phenotypes caused by genetic or neural manipulations, and to systematically expose the global structure of behavior within an experiment. This work reveals that mouse body language is built from identifiable components and is organized in a predictable fashion; deciphering this language establishes an objective framework for characterizing the influence of environmental cues, genes and neural activity on behavior.

INTRODUCTION

Innate behaviors are sculpted by evolution into stereotyped forms that enable animals to accomplish particular goals (such as exploring or avoiding a predator). Ultimately understanding how neural circuits create these patterned behaviors requires a clear framework for characterizing how behavior is organized and evolves over time. One conceptual approach to addressing this challenge arises from ethology, which proposes that the brain builds coherent behaviors by expressing stereotyped modules of simpler action in specific sequences (Tinbergen, 1951).

Although behavioral modules have traditionally been identified, one at a time, through careful human observation, recent technical advances have enabled more comprehensive characterization of the components of behavior. For example, in invertebrates, behavioral modules and their associated transition probabilities can now be discovered systematically through automated machine vision, clustering, and classification algorithms (Berman et al., 2014; Croll, 1975; Garrity et al., 2010; Ste-

phens et al., 2008, 2010; Vogelstein et al., 2014). Furthermore, identifying behavioral modules and transition probabilities has uncovered context-specific strategies used by invertebrate brains to adapt behavior to changes in the environment, which include both the emission of new behavioral modules (such as when the animal switches from “exploring” to “mating”) and the generation of new behaviors through re-sequencing existing modules. In *C. elegans*, for example, neural circuits that respond to appetitive cues alter transition probabilities between a core set of locomotor-related behavioral modules, thereby creating new behavioral sequences that enable taxis toward attractive odorants (Gray et al., 2005; Pierce-Shimomura et al., 1999). Similar observations have been made for sensory-driven behaviors in fly larvae (Garrity et al., 2010).

Comparable systematic approaches to discovering behavioral modules have not yet been implemented in mice. Instead, traditional behavioral classification approaches have been instantiated in silico, enabling machine vision algorithms to replace tedious and unreliable human scoring of videotaped behavior (de Chaumont et al., 2012; Jhuang et al., 2010; Kabra et al., 2013; Weissbrod et al., 2013). These approaches are powerful and have been successfully used to quantify components of innate exploratory, grooming, approach, aggressive, and reproductive behaviors. However, because they depend upon the prior specification, by human observers, of what constitutes a meaningful behavioral module, the insight from these methods is bounded by human perception and intuition. Currently available approaches therefore focus on identifying a small number of pre-specified modules within a given experiment, rather than on discovering new behavioral modules (which potentially encapsulate novel patterns of action), describing the global structure of behavior, or predicting future actions based upon those in the past.

Systematically describing the structure of behavior in mice—and understanding how the brain alters that structure to enable adaptation—requires overcoming three challenges. First, it is not clear which features are important to measure when identifying candidate behavioral modules. Mice interact with the world by expressing complex 3D pose dynamics, but because these are difficult to capture, most current methods track 2D parameters such as the position, velocity, or 2D contour of the mouse (de Chaumont et al., 2012; Gomez-Marin et al., 2012; Jhuang et al., 2010; Kabra et al., 2013; Spink et al., 2001; but see Ouyang et al., 2011). Second, given that behavior evolves on many timescales in parallel, it is not clear how to objectively identify the relevant spatiotemporal scales at which to modularize behavior. Finally, even stereotyped modules of behavior exhibit

moment-to-moment and animal-to-animal variability (Colgan, 1978). This variability raises significant challenges for identifying the number and content of behavioral modules, and with associating observed actions with specific behavioral modules.

Here we describe a novel method, based upon recent advances in machine vision and learning, that identifies behavioral modules and their associated transition probabilities without human supervision. This approach uses 3D imaging to capture the pose dynamics of mice as they freely behave in a variety of experimental contexts; these data reveal a surprising regularity that appears to divide mouse behavior into recognizable behavioral motifs that are organized at the sub-second timescale. A computational model then takes advantage of the observed fast temporal structure to describe mouse behavior as a series of reused modules, each a brief and stereotyped 3D trajectory through pose space that is connected in time to other modules through predictable transitions. We use this combined 3D imaging/modeling approach to explore how the underlying structure of behavior is altered after distinct environmental, genetic, or neural manipulations and show that this method can detect both predicted changes in action and new features of behavior that had not been previously described. This work reveals that defining behavioral modules based upon structure in the 3D behavioral data itself—rather than using a priori definitions for what should constitute a measurable unit of action—can yield key information about the elements of behavior, offer insight into adaptive behavioral strategies used by mice, and enable discovery of subtle alterations in patterned action.

RESULTS

3D Imaging Captures Inherent Structure in Mouse Pose Dynamics

We wished to develop a method that would allow unsupervised phenotyping of mice based upon patterns of 3D movement. However, it is not clear whether spontaneous behaviors exhibited by mice have a definable underlying structure that can be used to characterize action as it evolves over time. To ask whether such a structure might exist, we measured how the shape of a mouse's body changes as it freely explores a circular open field (Experimental Procedures). We used a single depth camera placed above the arena to capture these 3D pose dynamics and then extracted the image of the mouse from the arena; corrected imaging artifacts due to parallax effects; centered, and aligned the mouse along the inferred axis of its spine; and then quantified how the mouse's pose changed over time (Figures 1A and S1; Movie S1).

Plotting these 3D data over time revealed that mouse behavior is characterized by periods during which pose dynamics evolve slowly, punctuated by fast transitions that separate these periods; this pattern appears to divide the behavioral imaging data into blocks typically lasting 200–900 ms (Figures 1B and S2A). This temporal structure is apparent in the raw imaging pixels (Figure 1B, top), in the inferred shape of the mouse's spine (Figure 1B, middle), and in dimensionally reduced data that randomly samples from the depth images on the sensor (Figure 1B, bottom; see Supplemental Experimental Procedures). This structure is absent in data acquired from anesthetized or

dead mice but is present for the entire duration of each experiment in the pose dynamics data of mice exploring behavioral arenas of different shapes and after exposure to a wide variety of sensory cues, suggesting that it is a fundamental and ubiquitous feature of mouse behavior (data not shown; Figure S2).

To characterize this fast temporal structure, we performed three separate quantitative analyses. First, approximate boundaries between blocks in the behavioral imaging data were identified by a changepoints algorithm, which revealed that the mean block duration was about 350 ms, roughly matching the timescale of the blocks apparent upon visual inspection (Figures 1B and 1C; Supplemental Experimental Procedures). Second, we performed temporal autocorrelation analysis on the pose dynamics data, which demonstrated that autocorrelation in the mouse's pose largely dissipated after 400 ms ($\tau = 340 \pm 58$ ms, Figures 1C and S2B). This pattern of autocorrelation reflects specific behavioral dynamics organized at sub-second timescales, as it was destroyed by shuffling the behavioral data at timescales of 500 ms or less, and was not observed in synthetic mouse behavioral data designed to evolve with either random walk or Levy flight characteristics (Figures 1C and S2C). Third, we used a Wiener filter analysis to compare power spectral densities taken from live and dead mice; this approach identifies frequencies that must be changed in imaging data taken from a dead mouse to match the power spectrum of a live mouse. Nearly all of the frequency content differentiating behaving from dead mice was concentrated between 1 and 6 Hz (measured by spectrum ratio, mean 3.75 ± 0.56 Hz, Figures 1C and S2B). Taken together, the qualitative appearance of block structure in the pose dynamics data, along with the convergent results obtained with these three quantitative analyses, demonstrate that mouse pose dynamics exhibit structure at the sub-second timescale.

The observed temporal structure within the pose dynamics data suggests a timescale at which continuous behavior may be naturally segmented into meaningful components, as visual inspection of 3D movies revealed that each of the sub-second blocks of behavior appears to encode a recognizable action (e.g., a dart, a pause, the first half of a rear; see Movie S2). To explore the possibility that these sub-second actions are stereotyped (and therefore reproducibly performed at different times during an experiment), the 3D mouse imaging data was subjected to wavelet decomposition followed by principal-component analysis (PCA), which transformed each block in the pose dynamics data into a continuous trajectory through principal component (PC) space (Figures 1D). By scanning the behavioral data using a template matching method (using Euclidean distance among the first 10 PCs, which explain 88% of the data variance) (Figure S2D), additional instances were identified in which each template action was reused (Figure 1D; Supplemental Experimental Procedures; see Figure S3 for additional examples). These anecdotal observations suggest that mice create complex behaviors through the serial expression of stereotyped and reused behavioral modules (Tinbergen, 1951).

Mouse Behavior Can Be Described and Predicted with Modules and Transitions

Although our analysis suggests a timescale at which behavioral modules might exist, and qualitative inspection of 3D video

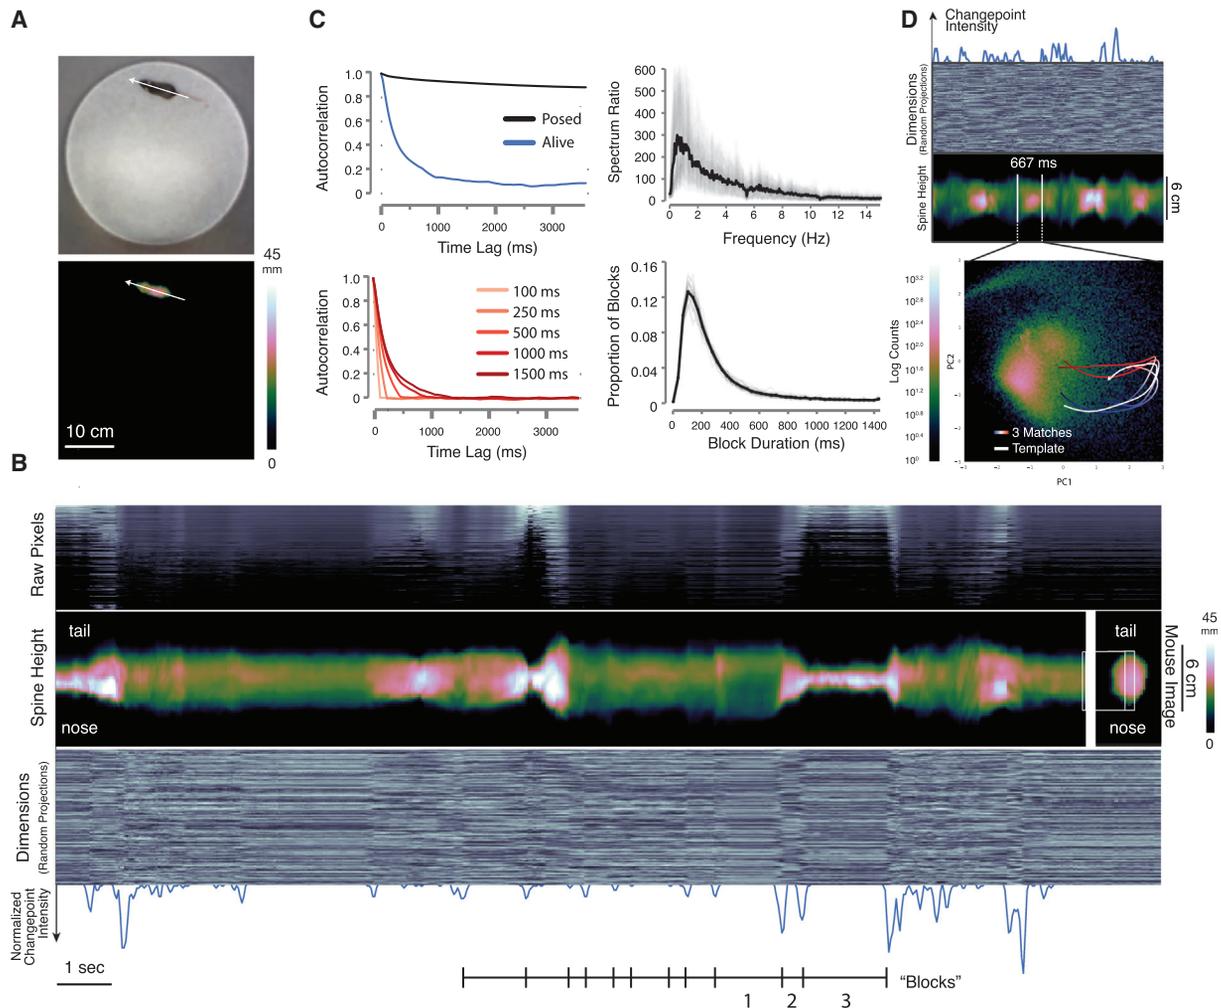


Figure 1. Depth Imaging Reveals Block Structure in 3D Mouse Pose Dynamics Data

(A) Mouse imaged in the circular open field with a standard RGB (top) and 3D depth camera (bottom, mm = mm above floor). Arrow indicates the inferred axis of the animal's spine.

(B) Raw pixels of the extracted and aligned 3D mouse image (top panel, sorted by mean height), compressed data (bottom panel, 300 dimensions compressed via random projections arrayed on the y axis, pixel brightness proportional to its value), and height at each inferred position of the mouse's spine (middle panel, "spine" data extracted from the mouse as indicated on the right, mm = mm above floor) each reveal sporadic, sharp transitions in the pose data over time. Note that the cross-sectional profile of the spine with respect to the camera varies depending upon the morphology of the mouse; when reared, this profile becomes smaller, and when on all fours, it becomes larger. Changepoints analysis (bottom panel, blue trace = normalized changepoint probability) identifies approximate boundaries between blocks. Blocks encode a variety of behaviors including locomotion (1), a pointing rear with the mouse's body elongated with respect to the sensor (2), and a true rear (3).

(C) Upper left: Autocorrelation analysis performed on the top ten principal components (PCs) of the pose dynamics data reveals that temporal correlation in the mouse's pose stabilizes after about 400 ms ($\tau = 340 \pm 58$ ms).

(C) Lower left: Shuffling behavioral data in blocks of 500 ms or shorter destroys autocorrelation structure (shuffle block size indicated).

(C) Upper right: Spectral power ratio between behaving and dead mice (mean plotted in black, individual mice plotted in gray) reveals most frequency content is represented between 1 and 6 Hz (mean = 3.75 ± 0.56 Hz).

(C) Lower right: Changepoints-identified block duration distribution (mean = 358 ms, SD 495 ms, mean plotted in black, individual mice in gray, $n = 25$ mice, 500 total min imaging).

(D) Projecting mouse pose data (top panels, random projections and spine data depicted as in [B]) into PC space [bottom]) reveals that blocks of pose data encode reused trajectories (density of all recorded poses colormapped behind trajectories). Tracing out the path associated with a block highlighted by changepoint detection (top) identifies a trajectory through PC space (white). Similar trajectories identified through template matching (time indicated as progression from blue to red) are superimposed. Note that this procedure uses the first ten PCs to identify matched trajectories, although only the first two PCs are depicted here.

and PCA trajectories is consistent with the possibility that sub-second blocks of behavior correspond to reused modules, current methods do not allow for the systematic identification of

candidate behavioral modules in mice. Indeed, available approaches neither reveal whether dividing behavior into modules actually helps to explain the overall structure of behavior nor

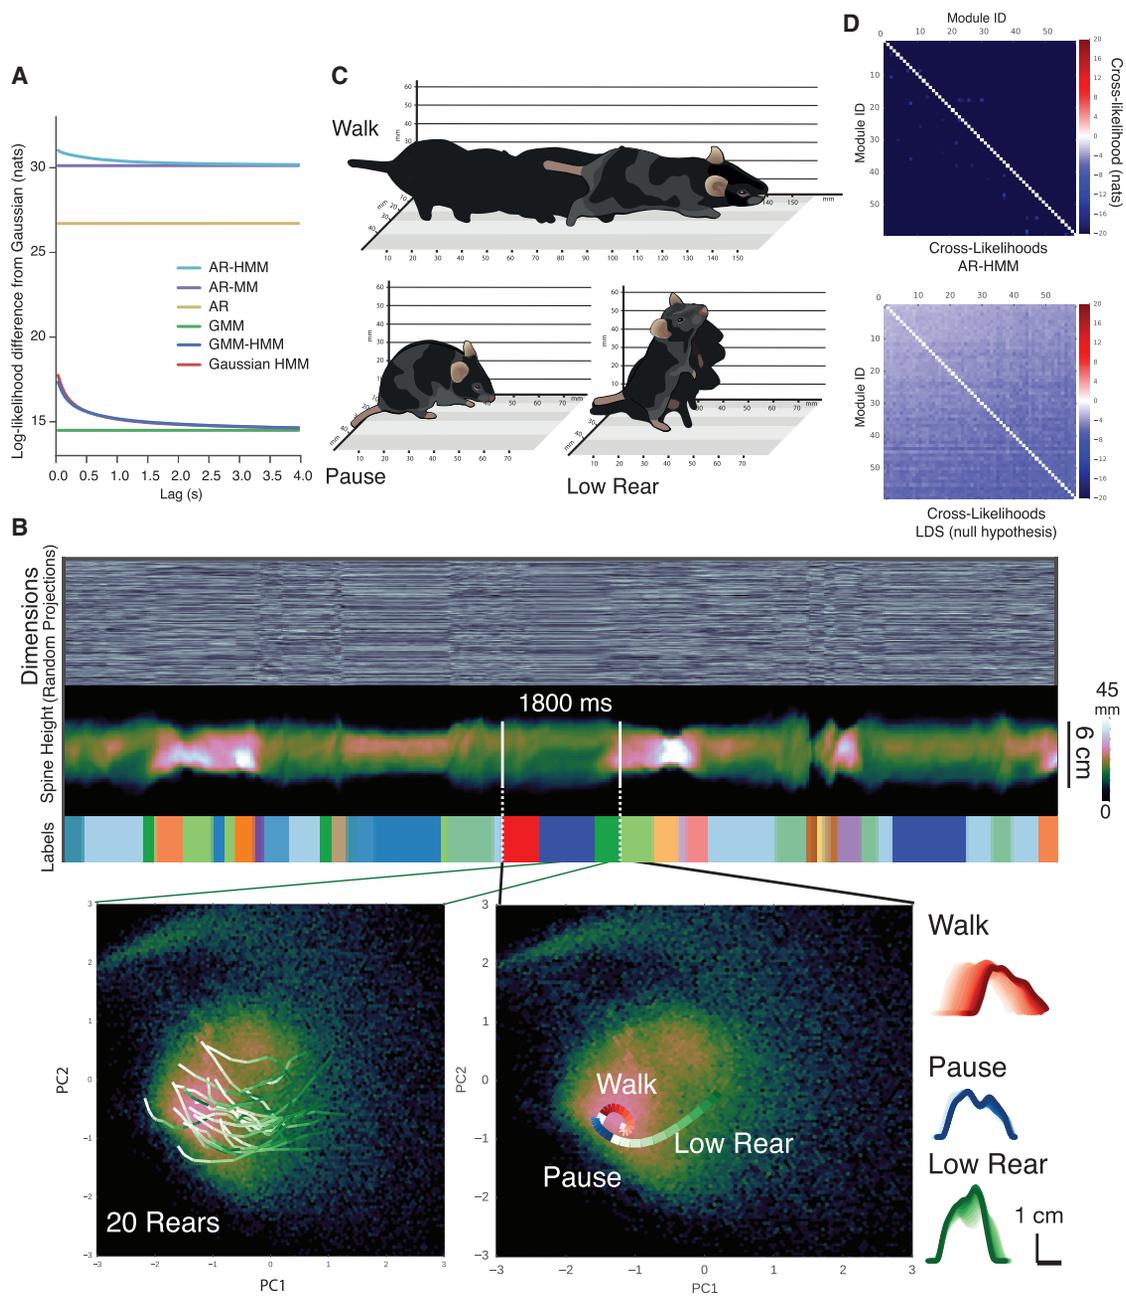


Figure 2. Reused Behavioral Modules within Mouse Pose Dynamics Data

(A) Predictive performance comparison of computational models describing possible structures for behavior (details of each model and the comparison metric provided in [Supplemental Experimental Procedures](#)). Models range from a Gaussian model (which proposes that mouse behavior is built from modules, each a single Gaussian in pose space) to an AR-HMM (which proposes that mouse behavior is built from modules, each of which encodes an autoregressive trajectory through pose space, and which transition from one to another with definable transition statistics; AR-MM, autoregressive mixture model; AR, autoregressive model; GMM, Gaussian mixture model; GMM-HMM, GMM hidden Markov model; Gaussian HMM, Gaussian hidden Markov model). Model performance plotted as the log likelihood (y axis) ascribed to held-out test data at some time lag (x axis) into the future (after subtracting Gaussian model performance).

(B) The AR-HMM parses behavioral data into identifiable modules (top panels—marked “labels,” each module is uniquely color coded). Multiple data instances associated with a single behavioral module (encoding a rear) each take a stereotyped trajectory through PC space (bottom left, trajectories progress from white to green over time, see [Movie S4](#)); multiple trajectories define behavioral sequences (bottom center, see [Movie S6](#)). Each trajectory within a sequence encodes a different elemental action (side-on view of the mouse calculated from depth data, bottom right, time indicated as increasingly darker lines, from module start to end).

(legend continued on next page)

identify the most likely number of modules expressed within any given dataset, or the content and durations of those modules.

To address these issues, we built a family of computational models, each of which proposes a unique underlying structure for behavior, and asked which of these models best predicts the pose dynamics of freely behaving mice (Figures 2A and S4); we reasoned that the model that most closely fit behavioral data (to which the model had not been exposed) would reveal key features of the underlying organization of behavior and could be used to characterize its components. After pre-processing the imaging data, the top ten PCs of the data (Figure S4E) were used to fit each model; this use of PCs (which directly reflect the pixel data) as a basis for modeling minimized potential biases from feature engineering. Models were fit using Bayesian nonparametric and Markov Chain Monte Carlo techniques that can automatically identify structure within large datasets, including the optimal state number for a given dataset and model (see Supplemental Experimental Procedures; Figure S4). Each model was trained on one set of pose dynamics data and then tested for its ability to explain a separate set of held-out data, a metric canonically used to compare unsupervised learning models (Hastie et al., 2009).

Our simplest model proposed that each behavioral module is a single fixed pose with no defined transition probabilities between the modules. We iteratively added structure to build increasingly complex models, which incorporated modules with more elaborate internal structures (ranging from mixtures of poses to smooth pose trajectories), allowed predictable transitions between specific modules (by embedding the modules within a Markov model), or both (Figure S4). Where relevant, the fitting procedures were explicitly focused to search for behavioral modules at the sub-second timescale matching the temporal structure identified using our model-free methods; this approach provided an important—and previously unavailable—constraint, given the multiple possible timescales upon which behavior evolves simultaneously.

The model that best fit previously unseen behavioral data describes mouse behavior as a sequence of modules (each capturing a brief motif of 3D body motion) that switch from one to another at the sub-second timescale identified by our model-free analysis of pose dynamics (Figures 2A and S4D; Supplemental Experimental Procedures). We refer to this model as an AR-HMM, as each behavioral module was modeled as a vector autoregressive (AR) process capturing a stereotyped trajectory through pose space, and the switching dynamics between different modules were modeled using a hidden Markov model (HMM). In other words, the model is a hierarchical description of behavior, with the “internals” of each module reflecting the mouse’s pose dynamics over short timescales, and the longer-timescale relationships between behavioral modules (i.e., the possible module sequences) governed by the transition probabilities specified by an HMM (Figures S4B and S4C). The observation that the AR-HMM outperforms alternative models

(Figure 2A) demonstrates that modularity and transition structure at fast timescales are critical for describing mouse behavior, a key prediction from ethology.

Model-Identified Behavioral Modules Are Stereotyped and Distinct

The AR-HMM systematically identifies modules and their transition probabilities from behavioral data without human supervision; this suggests that the AR-HMM can be used to identify behavioral modules and their associated transition probabilities—and thereby expose the underlying structure of behavior—in a wide variety of experimental contexts. We therefore performed a series of control analyses to establish whether the AR-HMM can indeed reliably identify behavioral modules encoding repeatedly used and stereotyped motifs of distinguishable behavior that are organized at sub-second timescales.

Although the AR-HMM is tuned to identify modules at a particular timescale, it is possible that after training the model could fail to effectively capture temporal structure in behavior. However, the AR-HMM successfully identified modules at the fast behavioral timescale defined by the model-free methods, as the distribution of module durations was similar to the duration distribution for changepoints-identified blocks (Figure S5A). Importantly, the ability of the AR-HMM to identify behavioral modules depended upon the inherent sub-second organization of mouse pose data, as shuffling the behavioral data in small chunks (i.e., <500 ms) substantially degraded model performance, whereas shuffling the data in bigger chunks had little effect (Figure S5B).

We then asked whether model-identified modules encode repeatedly used and stereotyped motifs of behavior. The pose trajectories associated with a specific model-identified behavioral module took similar paths through PC space, and visual inspection of the 3D movies associated with multiple instances of this module revealed they all encoded a stereotyped rearing behavior (Figure 2B and 2C; Movie S3). In contrast, data instances drawn from different behavioral modules traced distinct (and stereotyped) paths through PC space (Figure 2B; see Figure S5C for additional examples). Furthermore, visual inspection of the 3D movies assigned to different modules reveals that each encoded a coherent pattern of 3D motion that post hoc can be distinguished and labeled with descriptors (see Movie S4 for “walk,” “pause,” and “low rear” modules depicted in Figure 2C, as well as additional examples).

The modules identified by the AR-HMM are distinct from each other, as a cross-likelihood analysis demonstrated that the imaging data associated with a given module are best assigned to that module, and not to any of the other behavioral modules (Figure 2D; Supplemental Experimental Procedures). In contrast, the AR-HMM failed to identify any well-separated modules in a synthetic mouse behavioral dataset that lacks modularity but otherwise matches all multidimensional and intertemporal correlations of the real data, demonstrating that the AR-HMM does not discover modularity where none exists (Figure 2D).

(C) Isometric-view illustrations of the 3D imaging data associated with walk, pause, and low rear modules (also see Movie S4).

(D) Cross-likelihood analysis depicting the likelihood that a data instance assigned to a particular module is well-modeled by another module. Cross-likelihoods were computed for the open field dataset (above, see Supplemental Experimental Procedures, units are nats, where e^{nats} is the likelihood ratio) and for module-free synthetic data whose autocorrelation structure matches actual mouse data (below).

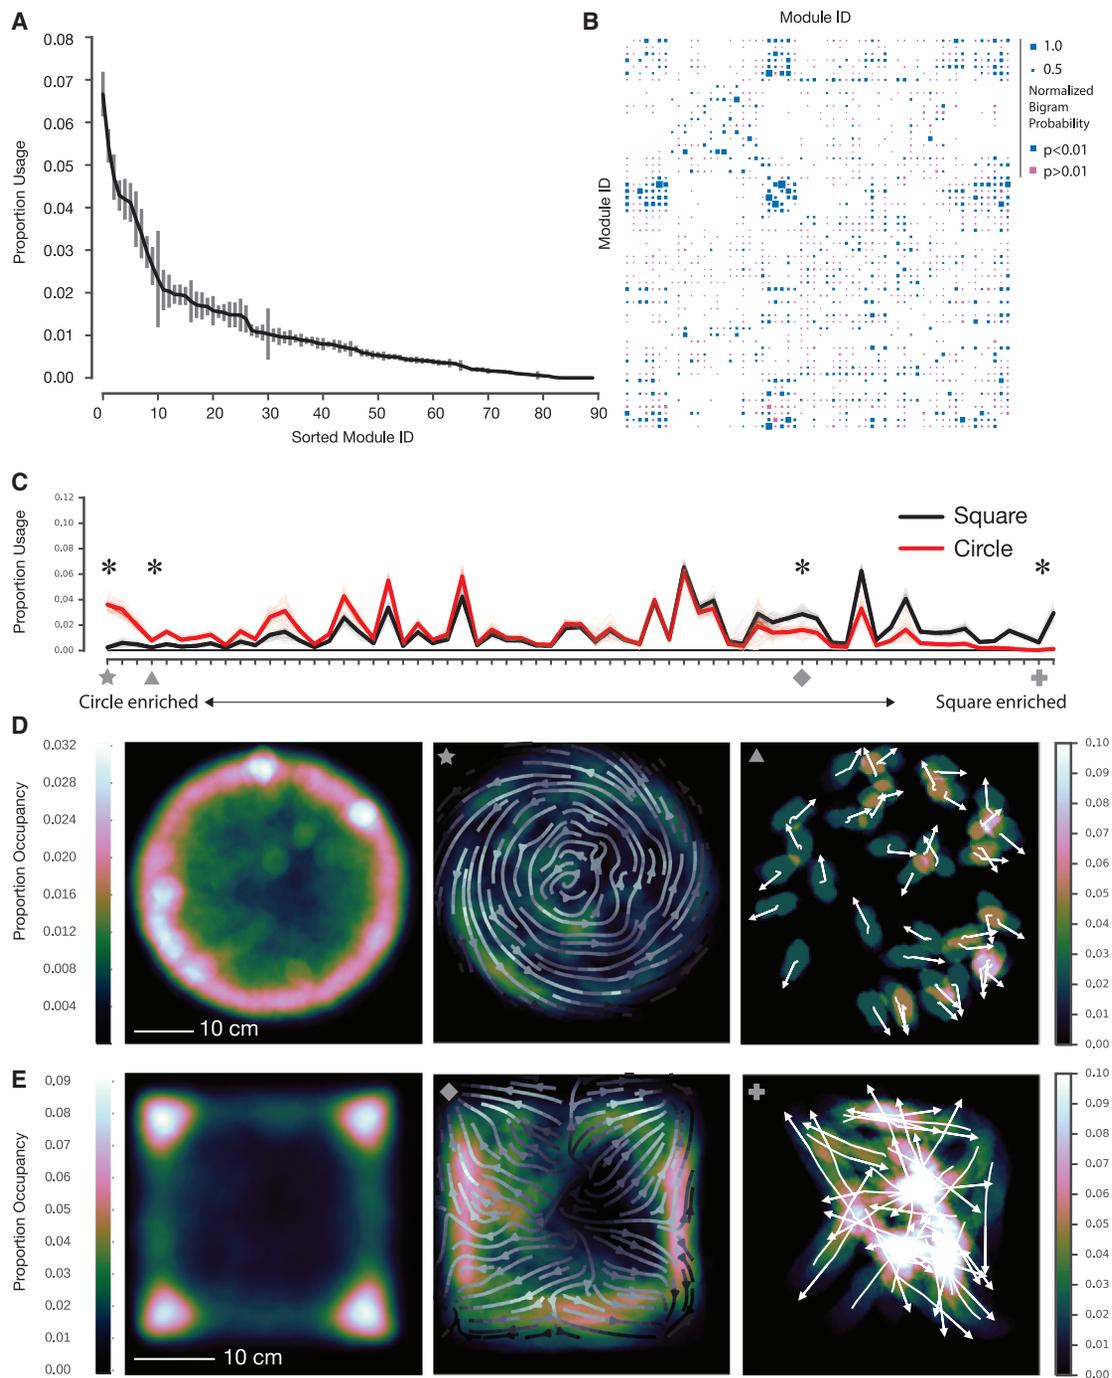


Figure 3. The Physical Environment Influences Module Usage and Spatial Pattern of Expression

(A) Modules identified by the AR-HMM, sorted by usage ($n = 25$ mice, 20 min per mouse, data from circular open field, error bars are SEs calculated using bootstrap estimation, $n = 100$ bootstrap estimates, see [Figure S5E](#) for Bayesian credible intervals).

(B) Hinton diagram depicting the probability that any pair of modules is observed as an ordered pair (p values calculated via bootstrap estimation and color coded); modules were sorted by spectral clustering to emphasize neighborhood structure.

(C) Module usage, sorted by context (with “circularness” on left, overall usages differ significantly, $p < 10^{-15}$, Hotelling two-sample t-squared test, see [Supplemental Experimental Procedures](#) for sorting details). Mean usages across animals depicted with dark lines, with bootstrap estimates depicted in fainter lines ($n = 100$). Marked modules discussed in main text and shown in (D): star, circular wall-hugging locomotion (“thigmotaxis”); triangle, outward-facing rears; diamond, square thigmotaxis; cross, square dart; see [Movie S7](#). Usage for all marked modules significantly modulated by context (indicated by asterisk, Wald Test, Holm-Bonferroni adjusted $p < 0.006$ for square dart, otherwise $p < 10^{-5}$).

(legend continued on next page)

Furthermore, rerunning the AR-HMM training process from random starting points generated highly similar behavioral modules ($R^2 = 0.94 \pm 0.03$, $n = 15$ restarts, [Supplemental Experimental Procedures](#)); in comparison, models with lower held-out likelihood scores had lower consistency ($R^2 < 0.4$), suggesting that these alternatives fail to reliably identify underlying structure in behavior. The AR-HMM output was also robust to the specific training data used, as models created from subsets of a larger dataset representing a single experiment were highly similar to each other ($R^2 > 0.9$). These findings demonstrate that the AR-HMM converges on a consistent set of behavioral modules regardless of the specific training data (within a given experiment) or how the model is initialized.

Finally, the modules of behavior and the associated transition probabilities identified by the AR-HMM appear to fully capture the richness of mouse behavior, as a 3D movie of a behaving mouse generated by a trained AR-HMM was qualitatively difficult to distinguish from a 3D movie of behavior exhibited by a real animal ([Movie S5](#); [Figure S4E](#)); in contrast, movies generated by more poorly performing models appeared discontinuous and were easily distinguished from actual mice (data not shown). Taken together with our model-free and model-based analyses and controls, the observation that after training our model can synthesize a convincing replica of 3D mouse behavior from learned modules and transition probabilities is consistent with the hypothesis that mouse behavior is organized into distinct sub-second modules that are combined to create coherent patterns of action.

Using the AR-HMM to Characterize Baseline Patterns of Behavior

The AR-HMM identifies two key features of mouse behavior (from the perspective of 3D pose dynamics): which behavioral modules are expressed during behavior and how those modules transition into each other over time to create action. The AR-HMM identified ~ 60 reliably used behavioral modules from a circular open field dataset, which is representative of normal mouse exploratory behavior in the laboratory (51 modules explained 95 percent of imaging frames, and 65 modules explained 99 percent of imaging frames) ([Figures 3A](#), [S5D](#), and [S5E](#); subjective categorization of the 51 most-used modules is shown in [Figure S6A](#)). These modules were connected to each other over time in a highly non-uniform manner, with each module preferentially linked to some modules and not others ([Figures 3B](#) and [S5F](#); average node degree after thresholding transitions that occur with $<5\%$ probability, 4.08 ± 0.10). This specific transition pattern among modules constrained the module sequences observed in the dataset ($\sim 17,000 \sim 75,000$ possible sequences of three modules [“trigrams”] given the total data size), demonstrating that certain module sequences were favored over others and that mouse behavior is therefore predictable (per frame en-

tropy rate without self-transitions 3.78 ± 0.03 bits, with self-transitions 0.72 ± 0.01 bits, entropy rate in a uniform matrix 6.02 bits; average mutual information without self-transitions 1.92 ± 0.02 bits, with self-transitions 4.84 bits ± 0.03 bits; see [Movie S6](#) for multiple examples of the sequence of three modules depicted in [Figure 2B](#)). Note that while estimating coarse changes in specific trigram frequencies is possible, accurately estimating higher-order k -gram transition statistics is difficult, as the amount of data required grows exponentially with k . We therefore focus our analysis on lower-order statistics such as module usage frequencies and temporal interactions between pairs of modules.

Using the AR-HMM to Characterize the Nature of Behavioral Change

We tested whether the AR-HMM could effectively capture changes in behavior (both predicted and unpredicted) elicited by a range of simple experimental interventions designed to probe the influence of the environment, genes, or neural circuit activity on behavior. We first asked how mice behavior adapts to changes in apparatus shape. We imaged mice within a small square box and then co-trained our model with both the circular open field and square box data, thereby enabling direct comparisons of modules and transitions under both conditions; the modules identified by this co-training procedure did not erroneously lump together data that would otherwise be distinguishable, as each behavioral module’s mean pose trajectory was stereotyped across the different experimental arenas (data not shown). Although mice tended to explore the corners of the square box and the walls of the circular open field, the overall usage of most modules was similar between these apparatuses, consistent with exploratory behavior sharing many common features across arenas ([Figure 3C](#)).

However, the AR-HMM also identified a small number of behavioral modules that were deployed selectively in just one context, suggesting that different physical environments drive expression of new behavioral modules ([Figure 3C](#)). For example, one circular arena-specific module encoded a behavior in which mice walk near the arena wall with a body posture that matches the curvature of the wall, while within the square box mice expressed a context-specific module that encodes a dart out of the center of the square ([Figures 3D](#) and [3E](#)). Several behavioral modules were also differentially enriched (but not exclusively expressed) in one context or the other. In the circular arena, for example, mice preferentially executed a rear characterized by pointing outward while pausing near the center of the open field, whereas in the smaller square box mice preferentially expressed a high rear in the corners of the box ([Figure 3D](#); data not shown).

These results demonstrate that the AR-HMM can effectively capture predictable changes in behavior resulting from altering the physical environment (like walking along a curved wall or

(D) Occupancy plot of mice in circular open field (left, $n = 25$, 500 min total) indicating cumulative spatial positions across all experiments. Occupancy plot depicting deployment of circular thigmotaxis module (middle, average orientation across the experiment indicated as arrow field) and circle-enriched outward-facing rear module (right, orientation of individual animals indicated with arrows).

(E) Occupancy plot of mice in square box (left, $n = 15$, 300 min total) indicating cumulative spatial positions across all experiments. Occupancy plot depicting a square-enriched thigmotaxis module (middle, average orientation across the experiment indicated as arrow field) and square-specific darting module (right, orientation of individual animals indicated with arrows).

rearing in a corner). Importantly, these results also demonstrate that the AR-HMM can unmask arena-specific patterns of behavior that are expressed in the center of both arenas, away from the physically constraining walls (like the darting across the center of the square box and the outward pointing behavior expressed in the circle); this surprisingly suggests that arena shape influences mouse behavior in a manner that extends significantly beyond the predictable changes in action at the walls themselves. Taken together, these experiments reveal that the AR-HMM can suggest strategies used by the mouse brain to adapt to new physical environments: in the case of a change in environmental geometry, this strategy includes the recruitment of a limited set of context-specific behavioral modules into baseline patterns of action, and a broad rewriting of where in space modules are expressed with respect to the arena boundaries.

The small number of behavioral modules distinguishing the circular and square arenas suggests only modest differences in the global pattern of behavior in these two experiments. To ask how the AR-HMM captures changes in the underlying structure of behavior after an overt change in behavioral state, we exposed mice to an ethologically relevant olfactory cue, the aversive fox odor trimethylthiazoline (TMT), which was delivered to one quadrant of the square box via an olfactometer. This odorant profoundly changes mouse behavior, inducing odor investigation, escape, and freezing behaviors that are accompanied by increases in corticosteroid and endogenous opioid levels (Fendt et al., 2005; Wallace and Rosen, 2000). Consistent with these known effects, mice sniffed the odor-containing quadrant and then avoided the quadrant containing the predator cue, displaying prolonged periods of immobility traditionally described as freezing behavior (Figure 4A and S6B).

Given that TMT-induced behaviors are dramatically different than those observed at baseline, one might predict that TMT should induce new behavioral modules that underlie the generation of these new actions. Surprisingly, the AR-HMM revealed that the TMT-induced suite of new behaviors was best explained by the same set of behavioral modules that were expressed during normal exploration; several modules were upregulated or downregulated after TMT exposure, but new modules were not introduced or eliminated relative to control (Figure 4B).

We therefore asked whether the changes in observed behavior were the consequence of altered connections between behavioral modules. Plotting the module transitions altered after exposure to TMT defined two neighborhoods within the behavioral statemap; the first included an expansive set of transitions that was modestly downregulated by TMT, and the second included a focused set of transitions that was upregulated by TMT (Figure 4C). Many of these newly interconnected modules encoded different forms of freezing behavior (average velocity 0.14 ± 0.54 mm/s, for other modules 34.7 ± 53 mm/s) (Movie S8; Figure S6B). In addition, the TMT-initiated modulation of transition probabilities altered the expression of specific behavioral sequences; for example, the most commonly observed sequence of three freezing modules was expressed 716 times after TMT exposure (in 300 min of imaging), as opposed to just 17 times under control conditions (in 480 min of imaging). The stimulus-evoked rewriting of transition probabilities was accom-

panied by an increase in the overall predictability of mouse behavior (per frame entropy rate fell from 3.92 ± 0.02 bits to 3.66 ± 0.08 bits without self-transitions, and from 0.82 ± 0.01 bits to 0.64 ± 0.02 bits with self-transitions) consistent with the mouse enacting an avoidance strategy that was more deterministic in nature than locomotor exploration.

Proximity to the odor source also governed the usage of specific behavioral modules (Figures 4D and 4E). For example, a set of freezing-related modules tended to be expressed in the quadrant most distant from the odor source, while the expression of an investigatory rearing module (whose overall usage was not altered by TMT) was specifically enriched within the odor quadrant (Figures 4D and 4E; Movie S9).

Although TMT is known to induce dramatic changes in behavior, it has not been possible to systematically identify those specific behavioral features altered in response to this odorant or to place those altered features in context with normal patterns of exploration. Analysis by the AR-HMM suggests that the strategy used by the mouse brain to adapt to the presence of TMT overlaps with—and yet is distinct from—that used to accommodate physical changes in the environment. As was true for a changed physical environment, exposure to TMT alters the spatial deployment of modules and sequences to support particular patterns of action; in contrast, the complete cohort of behaviors elicited by TMT, including seemingly “new” behaviors such as freezing, are the consequence of altered transition structure between individual modules. Behavioral modules are not, therefore, simply reused over time but instead act as flexibly interlinked components whose specific sequencing and deployment in space has profound consequences for the generation of adaptive behavior.

Sub-Second Architecture of Behavior Reflects the Influence of Genes and Neural Activity

As described above, the AR-HMM shows the fine-timescale structure of behavior to be sensitive to persistent changes in the physical or sensory environment. However, manipulation of individual genes or neural circuits influences behavior across a range of spatiotemporal scales and with variable penetrance and reliability; these changes in behavior may or may not be effectively captured by a classification method designed to characterize the sub-second structure of action. We therefore directly asked whether the AR-HMM could systematically reveal the behavioral consequences of manipulating the mouse genome or activity within the nervous system.

To explore these possibilities, we first used the AR-HMM to characterize the phenotype of mice mutant for the retinoid-related orphan receptor 1β (*Ror1\beta*) gene, which is expressed in neurons in the brain and spinal cord. This mouse was selected for analysis because adult homozygous mutant animals permanently exhibit abnormal gait, which would be expected to be observed during a brief open-field experiment (André et al., 1998; Eppig et al., 2015; Liu et al., 2013; Masana et al., 2007). Analysis with the AR-HMM revealed that littermate control mice are nearly indistinguishable from fully inbred C57/Bl6 mice, whereas homozygous mutant mice express a unique behavioral module encoding a waddling gait (Figure 5A and 5C; Movie S10; see Supplemental Experimental Procedures for

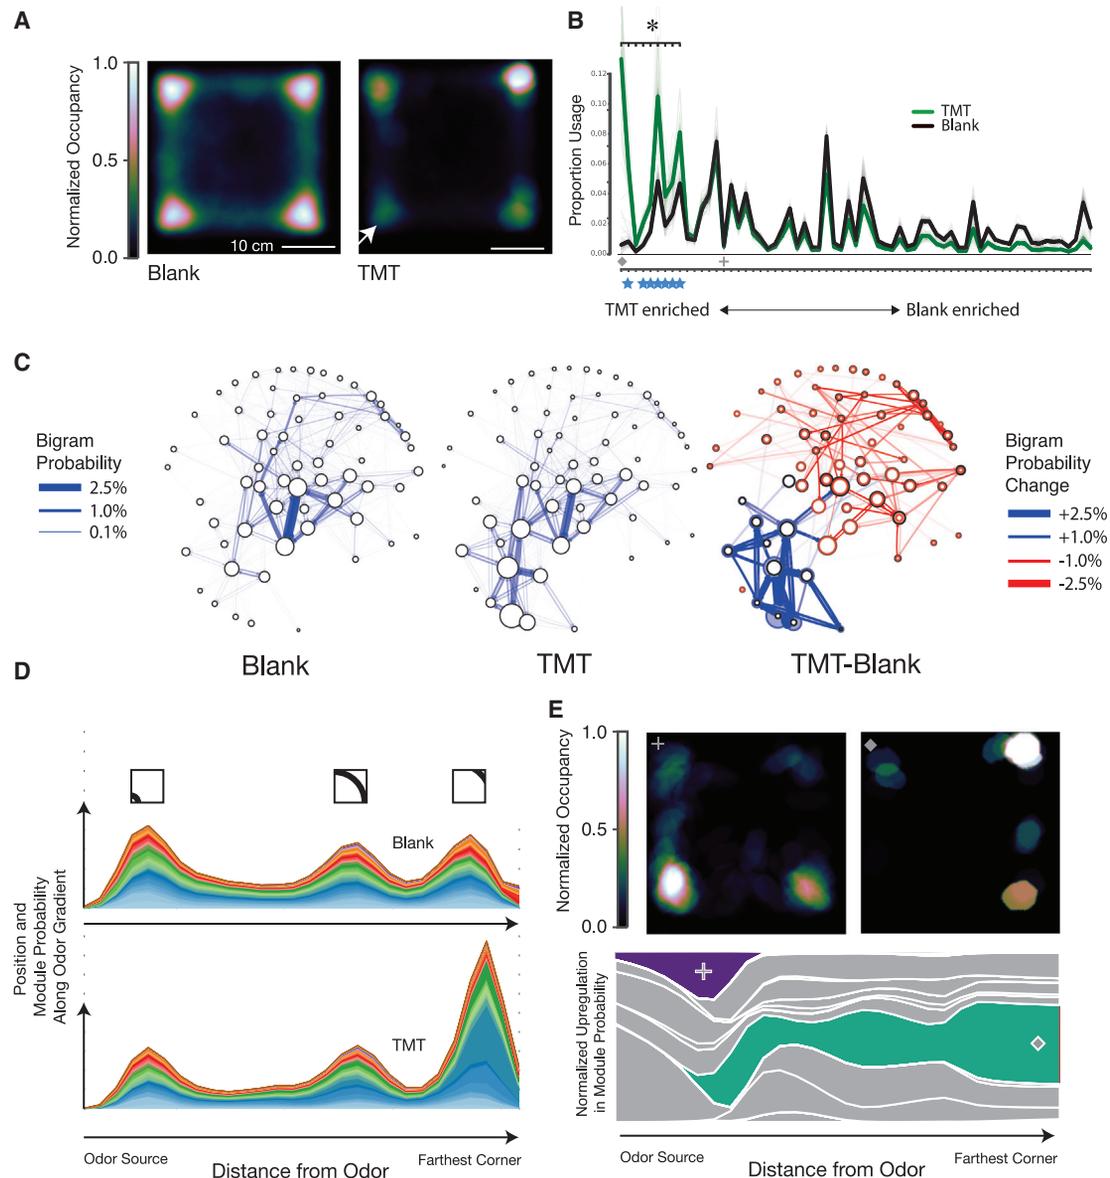


Figure 4. Odor-Driven Innate Avoidance Alters Transition Probabilities

(A) Occupancy plot of mice under control conditions ($n = 24$, 480 min total) and after exposure to the fox-derived odorant TMT (5% dilution, $n = 15$, 300 min total, model co-trained on both conditions) in the lower left quadrant (arrow). Plots normalized such that maximum occupancy = 1.

(B) Module usage plot sorted by “TMT-ness” (dark lines depict mean usages, bootstrap estimates depicted in fainter lines, sorting as in Figure 3). Marked modules discussed in main text and (E): cross, sniff in TMT quadrant; diamond, freeze away from TMT. See Movies S8 and S9. Blue stars indicate freezing modules. Asterisk indicates statistically significant regulation (Wald test, Holm-Bonferroni corrected, $p < 10^{-4}$).

(C) Left and middle: Behavioral state maps for mice exploring a square box before and after TMT exposure, with modules depicted as nodes (usage proportional to the diameter of each node), and bigram transition probabilities depicted as directional edges. Graph layout minimizes the length of edges and is seeded by spectral clustering to emphasize local structure.

(C) Right: Statemap depiction of the difference between blank and TMT. Usage differences are indicated by the newly sized colored circles (upregulation indicated in blue, downregulation indicated in red, previous usages in control conditions indicated in black). Altered bigram transition probabilities are indicated in the same color code; only those significant transition probabilities ($p < 0.01$) are depicted.

(D) Mountain plot depicting the joint probability of module expression and spatial position, plotted with respect to the TMT corner (x axis); note that the “bump” two-thirds of the way across the graph occurs due to the two corners equidistant from the odor source (see inset for approximate position in square box, modules are color coded).

(E) Occupancy plot (upper) indicating spatial position at which mice after TMT exposure emit an investigatory sniffing module (left) or a pausing module (right, see Movie S8). Mountain plot (lower) indicating the differential deployment of these two modules (purple, green; other modules in gray) with respect to distance from the odor source.

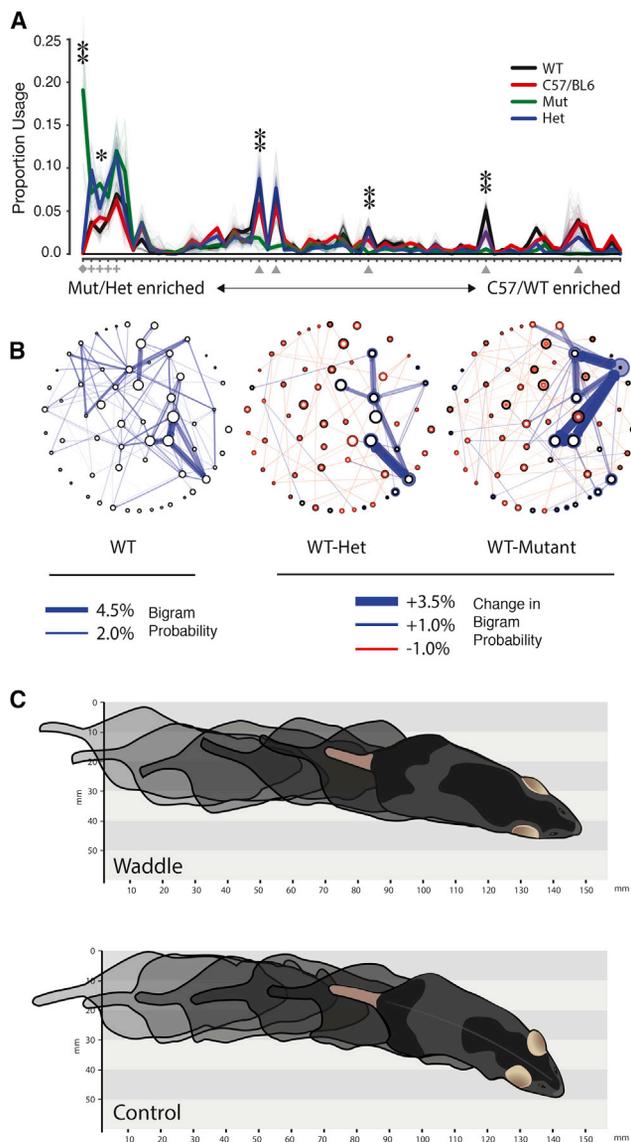


Figure 5. The AR-HMM Disambiguates Wild-Type, Heterozygous, and Homozygous *Ror1β* Mice

(A) Usage plot of modules exhibited by littermate *Ror1β* mice ($n = 25$ C57/BL6, $n = 3$ +/+, $n = 6$ ±, $n = 4$ −/−, open field assay, 20 min trials), sorted by “mutantness” (sorting and depiction as in Figure 3). Wild-type module usage is not statistically different from C57 but differs significantly from homozygote and heterozygote (Hotelling two-sample t-squared test, $p < 10^{-15}$). Marked modules described in main text: diamond, waddle; triangle, normal locomotion; cross, pause. Single asterisk indicates significant usage difference between mutant and wild-type, $p < 0.05$; double asterisk indicates $p < 0.01$ under Wald test, Holm-Bonferroni corrected.

(B) State map depiction of baseline OFA behavior for +/+ animals as in Figure 4C (left); difference state maps as in Figure 4C between the +/+ and +/- genotype (middle) and +/+ and -/- genotype (right); all depicted transitions that distinguish genotypes are statistically significant, $p < .01$.

(C) Illustration of the “waddle” module in which the hind limbs of the animal are elevated during walking (see Movie S10).

statistical details, data were co-trained with both circular open field and TMT datasets to facilitate comparisons with C57 animals). Conversely, the expression of five behavioral modules encoding normal forward locomotion in wild-type and C57 mice was downregulated in *Ror1β* mutants (Figure 5A, average during-module velocity = 114.6 ± 76.3 mm/s).

Previously unobserved phenotypes in the *Ror1β* mutant mice were also uncovered by the AR-HMM, as the expression of a set of four modules that encoded brief pauses and headbobs was upregulated in mutant mice (Figure 5A, average during-module velocity = 8.8 ± 25.3 mm/s); this pausing phenotype had not been previously reported in the literature. Furthermore, heterozygous mice—which have no reported phenotype (André et al., 1998; Eppig et al., 2015; Liu et al., 2013; Masana et al., 2007), exhibit wild-type running wheel behavior (Masana et al., 2007), and appear normal by eye—were also found to express a fully penetrant mutant phenotype: they overexpressed the same set of pausing modules that were upregulated in the full *Ror1β* mutants while failing to express the more dramatic waddling phenotype (Figure 5A). Differences between wild-type animals and both heterozygotes and mutants were also observed in transition probabilities associated with these pausing modules (Figure 5B).

The AR-HMM therefore describes the pathological behavior of *Ror1β* mice as the combination of a single neomorphic waddling module, decreased expression of normal locomotion modules, and increased expression of a small group of physiological modules encoding pausing behaviors; heterozygous mice express a defined subset of these behavioral abnormalities, whose penetrance is not intermediate but equals that observed in the mutant. These results suggest that the sensitivity of the AR-HMM allows fractionation of severe and subtle behavioral abnormalities within the same litter of animals, facilitates comparisons among genotypes and enables discovery of new phenotypes.

In the case of the *Ror1β* animals, a permanent alteration in DNA sequence is translated into an ongoing change in the overall sub-second statistical structure of behavior, one that is expressed continuously over the lifetime of the animal. We also wished to characterize how transient changes in activity in specific neural circuits influence the moment-to-moment structure of behavior; furthermore, given that the relationship between neural circuit activity and behavior can vary on a trial-to-trial basis, we wanted to probe the ability of the AR-HMM to afford insight into the probabilistic relationships between neural circuit activity and behavior.

To address these questions, we unilaterally expressed the light-gated ion channel Channelrhodopsin-2 in a subset of layer 5 corticostriatal neurons in the right hemisphere and assessed behavioral responses before, during, and after 2 s of light-mediated activation of motor cortex (Glickfeld et al., 2013). At negligible power levels, no light-induced changes in behavior were observed, whereas at the highest power levels, the AR-HMM identified two behavioral modules whose expression was reliably induced by the light, as on nearly every trial either one or the other module was expressed (Figure 6A). As would be expected, both of these modules encode forms of spinning-to-the-left behavior, and neither of these modules was expressed during normal mouse locomotion (Figure 6B; Movie S11). In addition, we noted that approximately 40 percent of the time the overall

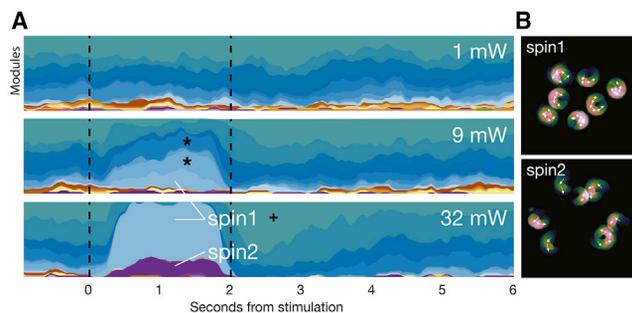


Figure 6. Optogenetic Perturbation of Motor Cortex Yields Both Neomorphic and Physiological Modules

(A) Mountain plot depicting the probability of expression of each behavioral module (assigned a unique color on the y axis) over time (x axis), with light stimulation indicated by dashed vertical lines (each plot is the average of 50 trials). Note that modest variations in the baseline pattern of behavior, due to trial structure, are captured before light onset. Stars indicate two modules expressed during baseline conditions that are also upregulated at intermediate powers (11 mW) but not high powers (32 mW, Wald test, Holm-Bonferroni adjusted $p < 10^{-5}$); cross indicates pausing module upregulated at light offset (Wald test, Holm-Bonferroni adjusted, $p < 10^{-5}$).

(B) Average position of example mice (with arrows indicating orientation over time) of the two modules induced under the highest stimulation conditions (see [Movie S11](#)). Note that (A) and (B) are generated from one animal and that the observed modulations are representative of the complete dataset ($n = 4$ mice, model was trained separately from previous experiments).

pattern of behavior did not return to baseline for several seconds after the end of optogenetic stimulation. This deviation from baseline was not due to continued expression of the same spinning modules that were triggered at light onset; instead, mice often paradoxically expressed a pausing module at light offset (average during-module velocity = 0.8 ± 7 mm/s, [Figure 6A](#), see cross).

The behavioral changes induced by high-intensity optogenetic stimulation were reliable, as on nearly every trial the animal emitted one of the two spinning modules. To ask whether the AR-HMM could characterize the pattern of behavior observed when behavioral modules are generated more probabilistically, the levels of light stimulation were reduced; we identified conditions under which one of the two spinning modules was no longer detected, and the other was expressed in only 25 percent of trials. Under these conditions, the AR-HMM detected the upregulation of a distinct set of behavioral modules, each of which was expressed in a fraction of trials ([Figure 6A](#), see stars). Unlike the spinning modules triggered by high-intensity stimulation, these modules were not neomorphic; rather, these modules were normally expressed during physiological exploration and encoded distinct forms of forward locomotion behavior (data not shown). Interestingly, although each of these individual light-regulated modules was emitted probabilistically on any given trial, low-intensity neural activation reliably influenced behavior across all trials when the behavioral modules were considered in aggregate ([Figure 6A](#)).

Thus, both low- and high-intensity optogenetic stimulation dramatically alters behavior, but the identity of the induced behaviors and their relative probabilities of expression vary across light levels. This effect would not have been apparent without the

ability of the AR-HMM to distinguish new from previously expressed behavioral modules and to identify the specific behavioral module induced at light onset on each trial; indeed, under low-light stimulation, the behavioral phenotype apparent by eye on any given trial often appears to be an extension of normal mouse exploratory behavior. Taken together, these data demonstrate that the AR-HMM can identify and characterize both obvious and subtle optogenetically induced phenotypes, distinguish “new” optogenetically induced behaviors from upregulated expression of “old” behaviors, and reveal the trial-by-trial relationships between neural activity and action.

DISCUSSION

It has been long hypothesized that innate behaviors are composed of stereotyped modules and that specific sequences of these modules encode coherent and adaptive patterns of action ([Bizzi et al., 2000](#); [Brown, 1911](#); [Drai et al., 2000](#); [Lashley, 1967](#); [Sherrington, 1907](#); [Tinbergen, 1951](#)). However, most efforts to explore the underlying structure of mouse behavior have relied on ad hoc definitions of what constitutes a behavioral module and have focused on specific behaviors rather than systematically considering behavior as a whole. As a consequence, we lack insight into the global organization of mouse behavior, the relationships between currently expressed actions and past or future behaviors, and the strategies used by the brain to generate behavioral change. Furthermore, we lack a comprehensive framework for characterizing the influence of individual genes or neural circuits on behavior.

Here we use 3D imaging to identify a sub-second spatiotemporal scale at which mouse behavior may be organized. Using this finding as a constraint, we then built a family of computational models, each of which represents a different hypothesis for the potential structure of behavior, and compared the ability of these models to predict and explain mouse behavior. The best performing model (the AR-HMM) specifically searches for modularity at sub-second timescales similar to those observed in the pose dynamics data and quantitatively describes behavior as a series of sub-second modules with defined transition probabilities. This combined 3D imaging/modeling approach can be used to automatically identify the behavioral modules expressed during a variety of experiments and to systematically discover how the architecture of behavior is altered as the mouse adapts to a changing world.

The AR-HMM Automatically and Systematically Captures Known and New Phenotypes

Our experiments reveal that the AR-HMM identifies both predicted changes in action and new features of behavior that had not been previously described. The high sensitivity of the AR-HMM—illustrated by the identification of a previously undescribed phenotype in heterozygous *Ror1 β* mice—raises the possibility that the AR-HMM could be useful for extracting subtle phenotypes from mouse models, including those in which single gene alleles are mutated; such patterns of mutation are common in human disease but are rarely explored in mouse models. Furthermore, the ability of the AR-HMM to place changes in behavior into context—illustrated by the ability of the AR-HMM

to identify the modules expressed after modest optogenetic stimulation on a trial-by-trial basis—suggests that the AR-HMM may be useful for relating unreliable or noisy patterns of neural activity to the probabilistic expression of specific actions. The AR-HMM (or similar approaches based upon unsupervised machine learning) may therefore be useful for discovering the behavioral consequences of a wide variety of experimental manipulations, particularly those enabled by the ever-growing toolbox of gene editing and optogenetic techniques.

The AR-HMM Identifies Possible Mechanisms for Behavioral Change

The AR-HMM suggests three regulatory strategies that may be used by the brain to implement behavioral adaptation. First, behavioral modules and their transitions appear to be selectively—and to some extent independently—vulnerable to alteration. None of the physiological and pathological deviations from baseline described here, from freezing to a waddling gait, caused global changes in the underlying structure of action; instead, new behaviors were well-described as changes in a small number of specific modules or transitions. This suggests that the brain can focally manipulate individual modules or transitions to generate new behaviors, and furthermore that mice can accommodate pathological actions (such as waddling) without catastrophic alterations in behavioral patterning.

Second, dramatically new behaviors can be created by altering the transitions statistics between modules alone—without invoking new behavioral modules—as was observed in animals exposed to TMT. This strategy has been shown to underlie sensory-driven negative and positive taxis behaviors in bacteria, worms, and flies (Berg and Brown, 1972; Garrity et al., 2010; Gray et al., 2005; Pierce-Shimomura et al., 1999); here we show that this strategy is conserved in mice and used by vertebrate nervous systems to create complex patterns of action in response to an external cue. In the specific case of the TMT response, the induced behaviors extend beyond taxis to adaptations like freezing, suggesting that the restructuring of transition probabilities may be a general mechanism for creating new patterns of action.

Third, we find that modulation of where in space behavioral modules are expressed supports the generation of specific adaptive behaviors. For example, the rearing module used by mice to investigate TMT is not significantly upregulated or resequenced relative to control, and yet its spatial pattern of expression in the quadrant containing TMT facilitates detection (and therefore avoidance) of the aversive odorant. Characterizing where in space behavioral modules are expressed also reveals that changes in physical context—such as the difference between a square and circular arena—elicit “state”-like changes in mouse behavior that extend beyond predictable changes in action at apparatus boundaries. Because the training data for the AR-HMM do not include any explicit allocentric parameters (such as the mouse’s spatial position within the apparatus), the ability of the model to uncover meaningful relationships between allocentric space and egocentric pose dynamics is an important validation that its segmentation of behavior is informative.

Potential Neural Underpinnings for Modules and Transitions

The observations generated by the AR-HMM lead to several predictions about the neural control of behavior. One such prediction is that specific behavioral modules and their associated transition probabilities will have explicit neural correlates, whose pattern of activity should reflect the ~2–5 Hz timescale at which modules are expressed; it is tempting to speculate that neural correlates representing transition probabilities between behavioral modules will be encoded in higher-order neural circuits tasked with behavioral sequencing, while neural correlates for the behavioral modules themselves might be encoded in central pattern generators or related circuit motifs in the brainstem or spinal cord. The relevant neural circuits may include both evolutionarily ancient regions of the brain involved in releasing innate behaviors, such as the amygdala, hypothalamus, and brainstem, as well as other areas, such as the striatum, that regulate fine-timescale behavioral sequencing (Aldridge et al., 2004; Swanson, 2000). Testing these predictions will require simultaneous characterization of neural activity and assessment of behavior; future embellishments of the modeling approach described here may allow for inference of joint structure between dense neural and behavioral data, and therefore be useful for revealing mechanistic relationships between the dynamics of neural activity and action.

Unsupervised Behavioral Characterization via Modeling of 3D Data: Strengths, Caveats, and Future Directions

Although there are multiple possible approaches to acquiring 3D pose dynamics data, we chose to implement our imaging system using a single standard depth camera, because such cameras are widely available, adaptable to a variety of different experimental circumstances, and can characterize behavior under most lighting conditions and in animals with any coat color (due to the use of infrared light as an illumination source). The “unsupervised” modeling approach taken here is also transparent, insofar as the assumptions of the model are explicitly stated (with no constraints supplied by the researcher other than the structure of the model and the specification of a single parameter that acts as a tunable “lens” to focus the model on behavior at a particular timescale). This makes clear the precise bounds of human influence on the output of the AR-HMM and insulates key aspects of that output from the vagaries of human perception and intuition. Given that inter-observer reliability in scoring even single mouse behaviors can be low (from 50 to 70 percent, with reliability falling as the number of scored behaviors increases), developing methods free from observer bias is essential for informatively characterizing mouse behavior (Garcia et al., 2010).

In addition, the generative modeling and inferential fitting methods described herein offer several practical advantages over the approaches typically used to analyze mouse behavior (Crawley, 2003), including the explicit time-series modeling of behavioral data (as opposed to simple clustering of dimensionally reduced data); the ability to directly inspect and explore each behavioral module; the flexible discovery of previously unobserved behavioral modules (rather than characterization of behavior from the perspective of “known” phenotypes); and

the ability to generate synthetic behavioral data, thereby allowing quantitation of how well a given model predicts the structure of behavior. This quantitative framework for comparing the performance of different methods for dividing up and measuring behavior is critical for advancing behavioral neuroscience, as it enables objective evaluation of alternative models for behavior, and assessment of future extensions that incorporate inevitable improvements in camera resolution, model structure, and fitting procedures.

On the other hand, the conclusions drawn here regarding the underlying structure of behavior (including its timescale) are limited by the simplicity of our experimental manipulations, which were designed to expose differences in motor outputs. Furthermore “mouse behavior” as described by the model is restricted to the imaged pose dynamics of the mouse at a particular spatiotemporal resolution and within a controlled laboratory experiment. Because the AR-HMM directly models the pixel data (after pre-processing), comparisons can only be made between mice of roughly similar size and shape. In addition, there are clearly important physical features of mouse behavior (operating at a variety of spatiotemporal scales) not captured in the pose data and therefore not modeled—these range from individual joint dynamics and paw position to sniffing, whisking, and breathing. In the future, complementary datastreams that capture different facets of behavior could be integrated with 3D pose data to generate more comprehensive behavioral models.

The modeling approach itself also has several important limitations. The AR-HMM uses Bayesian nonparametric approaches to identify the most likely number of modules that describe behavior at a particular temporal scale. However, this insight also comes at a cost: as the amount of data fed to the algorithm increases, the number of discovered modules necessarily rises. This challenge parallels the well-described phenomenon in ethology in which the number of discovered behaviors increases in subjects that have been observed either more frequently or for longer (Colgan, 1978). This monotonic (although sublinear) relationship between data size and the number of discovered modules limits comparisons of behavior across experiments without co-training models, as was done here. One potential approach to address this challenge could be the incorporation of a “canonical” behavioral dataset against which additional data could be compared; such a framework (in which the “canonical” modules are either frozen or flexible) may enable analysis of new behavioral experiments within a fixed frame of reference.

Finally, the AR-HMM cannot explicitly disambiguate those features of behavior that are the consequence of the biomechanics of the mouse—for example, transitions between specific modules that are impossible due to physical constraints—from those that are the consequence of the action of the nervous system. Given that the nervous system and the body in which it is embedded co-evolved to facilitate action, ultimately disentangling the relative contributions of each to the organization of behavior may be difficult (for discussion of this issue, see, for example, Tresch and Jarc, 2009). However, changes in the structure of behavior that are induced by experimental intervention arise principally from the action of the nervous system; the

observed context-dependent flexibility of the transition statistics between modules, taken with the ability of the mouse to emit new behavioral modules in response to internal or external cues, together suggest that—at the spatiotemporal scale captured by our methods—the nervous system plays a key role in regulating the overall structure of behavior.

Mouse Body Language: Syllables and Grammar

Candidate behavioral modules have been recognized in a variety of different contexts and on a wide range of spatiotemporal scales, and accordingly, researchers have given them a diverse set of names, including motor primitives, behavioral motifs, motor synergies, prototypes, and movemes (Anderson and Perona, 2014; Flash and Hochner, 2005; Tresch et al., 1999). The behavioral modules identified by the AR-HMM here find their origin in switching dynamics that are expressed on timescales of hundreds to milliseconds. Mouse behavior is clearly also organized at the varied and interdigitated timescales at which internal state (e.g., neural activity, endocrine function and development) and external state (e.g., daily, monthly, seasonal, and annual variation in behavior) unfold. The behavioral modules we have characterized therefore likely exist at an intermediate hierarchical level within the overall structure of behavior, albeit one that captures many of the behavioral changes induced by experimental manipulations carried out at both short and long timescales.

The modules identified by the AR-HMM do not exist in isolation; instead, they are given behavioral meaning through a transition structure that governs their sequencing. The observation of both modularity and transition structure within the pose dynamics of mice suggests strong analogies to birdsong, which is also hierarchically organized and composed of identifiable modules whose sequence is governed by definable transition statistics; importantly, birdsong has also been well described using generative modeling and low-order Markov processes (Berwick et al., 2011; Markowitz et al., 2013; Wohlgemuth et al., 2010). By analogy to birdsong, we therefore propose to refer to the modules we have identified as behavioral “syllables” and the statistical interconnections between these syllables as behavioral “grammar.” Such a grammatical structure has been previously proposed for restricted subsets of mouse behavior (such as grooming) in which behavioral modules were defined on an ad hoc basis (Berridge et al., 1987; Fentress and Stilwell, 1973); through unsupervised identification of behavioral syllables, here we show that the notion of a regulatory grammar is general and can be used to explain a wide variety of behavioral phenotypes. As is true for birdsong, the grammar we describe is highly restricted in nature (as only low-order interactions are modeled) and lacks the richness and flexibility of context-dependent grammars that have been explored in human language (Berwick et al., 2011). Despite this limitation, the experiments described herein expose an underlying structure for mouse body language organized at the sub-second timescale; this structure encapsulates mouse behavior (as detected by a depth sensor) within a given experiment and reveals a balance between stochasticity and determinism that is dynamically modulated as the mouse varies its pattern of action to adapt to challenges in the environment.

EXPERIMENTAL PROCEDURES

All error bars indicated in the paper are SEM as determined by bootstrap analysis unless noted otherwise. For complete details on methods used, please consult [Supplemental Experimental Procedures](#).

SUPPLEMENTAL INFORMATION

Supplemental Information includes seven figures, eleven movies, and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2015.11.031>.

A video abstract is available at <http://dx.doi.org/10.1016/j.neuron.2015.11.031#mmc14>.

ACKNOWLEDGMENTS

We thank Michael Greenberg, Vanessa Ruta, David Ginty, Jesse Gray, Mike Springer, Aravi Samuel, Chris Harvey, Bernardo Sabatini, Rachel Wilson, Andreas Schaefer, and members of the Datta Lab for helpful comments on the manuscript; Joseph Bell for useful conversations; and Alexandra Nowlan and Christine Ashton for laboratory assistance. We thank David Roberson and Clifford Woolf for assistance with behavioral experiments and Ofer Mazor and Pavel Gorelik from the Research Instrumentation Core Facility for engineering support. We thank Sigrid Knemeyer for illustration assistance. Core facility support is provided by NIH grant P30 HD18655. A.B.W. is supported by an NSF Graduate Research Fellowship and is a Stuart H.Q. & Victoria Quan Fellow. G.I. is supported by the Human Frontiers Science Program. M.J.J. is supported by a fellowship from the Harvard/MIT Joint Grants program. S.L.P. is supported by an NSF Graduate Research Fellowship and is a Stuart H.Q. & Victoria Quan Fellow. R.P.A. is supported by NSF IIS-1421780. S.R.D. is supported by fellowships from the Burroughs Wellcome Fund, Searle Scholars Program, the Vallee Foundation, the McKnight Foundation, and the Khodadad Program; by the Fidelity Biosciences Research Initiative; by grants DP2OD007109 and RO11DC011558 from the NIH; and by the SFARI program and the Global Brain Initiative from the Simons Foundation. Code is available upon request (datta.hms.harvard.edu/research.html).

Received: June 12, 2015

Revised: September 13, 2015

Accepted: November 18, 2015

Published: December 16, 2015

REFERENCES

- Aldridge, J.W., Berridge, K.C., and Rosen, A.R. (2004). Basal ganglia neural mechanisms of natural movement sequences. *Can. J. Physiol. Pharmacol.* **82**, 732–739.
- Anderson, D.J., and Perona, P. (2014). Toward a science of computational ethology. *Neuron* **84**, 18–31.
- André, E., Conquet, F., Steinmayr, M., Stratton, S.C., Porciatti, V., and Becker-André, M. (1998). Disruption of retinoid-related orphan receptor beta changes circadian behavior, causes retinal degeneration and leads to vacillans phenotype in mice. *EMBO J.* **17**, 3867–3877.
- Berg, H.C., and Brown, D.A. (1972). Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking. *Nature* **239**, 500–504.
- Berman, G.J., Choi, D.M., Bialek, W., and Shaevitz, J.W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *J. R. Soc. Interface.* **11**, <http://dx.doi.org/10.1098/rsif.2014.0672>.
- Berridge, K.C., Fentress, J.C., and Parr, H. (1987). Natural syntax rules control action sequence of rats. *Behav. Brain Res.* **23**, 59–68.
- Berwick, R.C., Okanoya, K., Beckers, G.J.L., and Bolhuis, J.J. (2011). Songs to syntax: the linguistics of birdsong. *Trends Cogn. Sci.* **15**, 113–121.
- Bizzi, E., Tresch, M.C., Saltiel, P., and d'Avella, A. (2000). New perspectives on spinal motor systems. *Nat. Rev. Neurosci.* **1**, 101–108.
- Brown, T.G. (1911). The intrinsic factors in the act of progression in the mammal. *Proc. R. Soc. Lond. B* **84**, 308–319, <http://dx.doi.org/10.1098/rspb.1911.0077>.
- Colgan, P.W. (1978). *Quantitative ethology* (John Wiley & Sons).
- Crawley, J.N. (2003). Behavioral phenotyping of rodents. *Comp. Med.* **53**, 140–146.
- Croll, N.A. (1975). Components and patterns in the behaviour of the nematode *Caenorhabditis elegans*. *J. Zool.* **176**, 159–176.
- de Chaumont, F., Coura, R.D.-S., Serreau, P., Cressant, A., Chabout, J., Granon, S., and Olivo-Marin, J.-C. (2012). Computerized video analysis of social interactions in mice. *Nat. Methods* **9**, 410–417.
- Drai, D., Benjamini, Y., and Golani, I. (2000). Statistical discrimination of natural modes of motion in rat exploratory behavior. *J. Neurosci. Methods* **96**, 119–131.
- Eppig, J.T., Blake, J.A., Bult, C.J., Kadin, J.A., Richardson, J.E., and Group, M.G.D.; Mouse Genome Database Group (2015). The Mouse Genome Database (MGD): facilitating mouse as a model for human biology and disease. *Nucleic Acids Res.* **43**, D726–D736.
- Fendt, M., Endres, T., Lowry, C.A., Apfelbach, R., and McGregor, I.S. (2005). TMT-induced autonomic and behavioral changes and the neural basis of its processing. *Neurosci. Biobehav. Rev.* **29**, 1145–1156.
- Fentress, J.C., and Stilwell, F.P. (1973). Letter: Grammar of a movement sequence in inbred mice. *Nature* **244**, 52–53.
- Flash, T., and Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Curr. Opin. Neurobiol.* **15**, 660–666.
- Garcia, V.A., Crispim Junior, C.F., and Marino-Neto, J. (2010). Assessment of observers' stability and reliability - a tool for evaluation of intra- and inter-concordance in animal behavioral recordings. *Conf. Proc. IEEE Eng. Med. Biol. Soc. 2010*, 6603–6606.
- Garrity, P.A., Goodman, M.B., Samuel, A.D., and Sengupta, P. (2010). Running hot and cold: behavioral strategies, neural circuits, and the molecular machinery for thermotaxis in *C. elegans* and *Drosophila*. *Genes Dev.* **24**, 2365–2382.
- Glickfeld, L.L., Andermann, M.L., Bonin, V., and Reid, R.C. (2013). Corticocortical projections in mouse visual cortex are functionally target specific. *Nat. Neurosci.* **16**, 219–226.
- Gomez-Marin, A., Partoune, N., Stephens, G.J., Louis, M., and Brembs, B. (2012). Automated tracking of animal posture and movement during exploration and sensory orientation behaviors. *PLoS ONE* **7**, e41642.
- Gray, J.M., Hill, J.J., and Bargmann, C.I. (2005). A circuit for navigation in *Caenorhabditis elegans*. *Proc. Natl. Acad. Sci. USA* **102**, 3184–3191.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition (Springer).
- Jhuang, H., Garrote, E., Mutch, J., Yu, X., Khilnani, V., Poggio, T., Steele, A.D., and Serre, T. (2010). Automated home-cage behavioural phenotyping of mice. *Nat. Commun.* **1**, 68.
- Kabra, M., Robie, A.A., Rivera-Alba, M., Branson, S., and Branson, K. (2013). JAABA: interactive machine learning for automatic annotation of animal behavior. *Nat. Methods* **10**, 64–67.
- Lashley, K.S. (1967). The problem of serial order in behavior. In *Psycholinguistics: A book of readings*, L.A. Jeffress, ed.
- Liu, H., Kim, S.-Y., Fu, Y., Wu, X., Ng, L., Swaroop, A., and Forrest, D. (2013). An isoform of retinoid-related orphan receptor β directs differentiation of retinal amacrine and horizontal interneurons. *Nat. Commun.* **4**, 1813.
- Markowitz, J.E., Ivie, E., Kligler, L., and Gardner, T.J. (2013). Long-range order in canary song. *PLoS Comput. Biol.* **9**, e1003052.
- Masana, M.I., Sumaya, I.C., Becker-Andre, M., and Dubocovich, M.L. (2007). Behavioral characterization and modulation of circadian rhythms by light and melatonin in C3H/HeN mice homozygous for the RORbeta knockout. *Am. J. Physiol. Regul. Integr. Comp. Physiol.* **292**, R2357–R2367.
- Ou-Yang, T.-H., Tsai, M.-L., Yen, C.-T., and Lin, T.-T. (2011). An infrared range camera-based approach for three-dimensional locomotion tracking and pose reconstruction in a rodent. *J. Neurosci. Methods* **201**, 116–123.

- Pierce-Shimomura, J.T., Morse, T.M., and Lockery, S.R. (1999). The fundamental role of pirouettes in *Caenorhabditis elegans* chemotaxis. *J. Neurosci.* *19*, 9557–9569.
- Sherrington, C. (1907). *The Integrative Action of the Nervous System*. J. Nerv. Ment. Dis.
- Spink, A.J., Tegelenbosch, R.A., Buma, M.O., and Noldus, L.P. (2001). The EthoVision video tracking system—a tool for behavioral phenotyping of transgenic mice. *Physiol. Behav.* *73*, 731–744.
- Stephens, G.J., Johnson-Kerner, B., Bialek, W., and Ryu, W.S. (2008). Dimensionality and dynamics in the behavior of *C. elegans*. *PLoS Comput. Biol.* *4*, e1000028.
- Stephens, G.J., Johnson-Kerner, B., Bialek, W., and Ryu, W.S. (2010). From modes to movement in the behavior of *Caenorhabditis elegans*. *PLoS ONE* *5*, e13914.
- Swanson, L.W. (2000). Cerebral hemisphere regulation of motivated behavior. *Brain Res.* *886*, 113–164.
- Tinbergen, N. (1951). *The study of instinct* (Clarendon Press).
- Tresch, M.C., and Jarc, A. (2009). The case for and against muscle synergies. *Curr. Opin. Neurobiol.* *19*, 601–607.
- Tresch, M.C., Saltiel, P., and Bizzi, E. (1999). The construction of movement by the spinal cord. *Nat. Neurosci.* *2*, 162–167.
- Vogelstein, J.T., Park, Y., Ohyama, T., Kerr, R.A., Truman, J.W., Priebe, C.E., and Zlatic, M. (2014). Discovery of brainwide neural-behavioral maps via multi-scale unsupervised structure learning. *Science* *344*, 386–392.
- Wallace, K.J., and Rosen, J.B. (2000). Predator odor as an unconditioned fear stimulus in rats: elicitation of freezing by trimethylthiazoline, a component of fox feces. *Behav. Neurosci.* *114*, 912–922.
- Weissbrod, A., Shapiro, A., Vasserman, G., Edry, L., Dayan, M., Yitzhaky, A., Hertzberg, L., Feinerman, O., and Kimchi, T. (2013). Automated long-term tracking and social behavioural phenotyping of animal colonies within a semi-natural environment. *Nat. Commun.* *4*, 2018.
- Wohlgemuth, M.J., Sober, S.J., and Brainard, M.S. (2010). Linked control of syllable sequence and phonology in birdsong. *J. Neurosci.* *30*, 12936–12949.

Neuron, Volume 88

Supplemental Information

Mapping Sub-Second Structure in Mouse Behavior

Alexander B. Wiltschko, Matthew J. Johnson, Giuliano Iurilli, Ralph E. Peterson,
Jesse M. Katon¹, Stan L. Pashkovski, Victoria E. Abraira, Ryan P. Adams,
and Sandeep Robert Datta

Supplemental Data

Figure S1. Image extraction workflow. Related to Figure 1. Workflow (shown here with text descriptions of each step **above** and graphics corresponding to each step **below**) used to convert depth images of freely-behaving mice into mouse pose dynamics data in which each mouse is aligned along the axis of its spine (see Movie S1 for both the raw input data and the product of this process). This process is described in detail in the Supplemental Experimental Procedure. In brief, after recording depth data within a behavioral apparatus (**A**), background values are computed (**B**) and missing pixel values are imputed (**C**). The mouse is then resampled into real-world coordinates to correct for parallax artifacts (**D-E**), and the background contributed by the apparatus is subtracted away (**F**). These images are then filtered, (**G**) and the image of the mouse is identified and extracted as a region of interest (**H-I**). The unrotated image is then aligned longitudinally based on fitting an ellipse to the boundary of the image (the major axis of which we refer to as “the axis of its spine) (**J**), enabling quantitative analysis of egocentric pose dynamics while the animal freely translates in space. This pre-processed imaging data is then subject to wavelet decomposition and PCA for dimensionality reduction before submission to the modeling algorithms.

Figure S2. Block structure, autocorrelation and compressibility in mouse depth imaging data. Related to Figure 1. **A.** Block structure is present in random projections data, spine data and raw pixel data derived from aligned mouse pose dynamics data (random projections and spines are depicted as in Fig. 1A, raw data in bottom panels are heatmapped based upon imaging pixel value; values are sorted by mean height, and rarely-used pixels are omitted). Each dimension of random projections data uses all pixels from the original mouse image. Live mice

exhibit significant block structure in imaging data (left panels), while dead mice do not (right panels). Smoothing refers to the modest median filtration applied during pre-processing, which does not visibly affect the structure of the data (Supplemental Experimental Methods). **B.** Dimensionality reduction does not significantly affect autocorrelation structure or spectral content of mouse pose dynamics data, nor does the particular experimental setup. The result of the Wiener filtering analysis (**left**) is similar for both a 10-dimensional PCA and a random projection representation of depth data, as well as across experimental conditions (OFA = open field assay, TMT Odor = square box with the fox odor TMT in one quadrant, Blank Odor = square box without introduced odorants, RorB = OFA data with mice mutant, heterozygous or wild-type at the Ror1 β locus, Optogenetics = optogenetic experiment in the OFA with unilateral motor cortex implantation of an optical fiber). The autocorrelation of random projections and PCA data (**right**), both of which represent the same depth data, exhibit approximately the same falloff, demonstrating that data compression does not influence fine-timescale autocorrelation structure in the imaging data. The particular arena in which the mouse was recorded, or the experimental manipulation to which the mouse was subjected also does not affect the autocorrelation structure, indicating it is a core feature of our data. Note that the slightly different autocorrelation structure observed in the optogenetics data can be explained by the presence of an optical fiber implanted in the head of the animal, which alters the imaged pose dynamics of the animal. **C.** This correlation structure is not observed if mouse poses evolve as if taking a Levy flight (**left**) or random walk (**right**), suggesting that live mice express a specific sub-second autocorrelation structure. The alpha parameter and sigma parameter control the rate of dispersion for data generated from the Levy flight and Brownian motion processes, respectively. **D.** The

first 10 principal components capture 88 percent of the variance in raw mouse imaging data. This number of dimensions was used for data analysis in the AR-HMM.

Figure S3. Additional template-matched trajectories in pose space. Related to Figure 2.

Scanning the behavioral data using a template matching method identified additional instances in which given stereotyped motifs were reused (with time proceeding from blue to red). Six additional templates and their nine closest matched trajectories are plotted in the first two dimensions of principal components. The complete pose data in principal component space is represented as a background density plot. Note that 10 PCs were used to identify examples matching the seed templates, although only 2 PCs are depicted here for clarity.

Figure S4. Modeling mouse behavior. Related to Figure 2 and Supplemental Experimental Procedures. **A.** Sketch of the part of the data processing pipeline that includes model fitting. The mouse is first identified, extracted and aligned from depth camera video frames along the major axis of the mouse, which approximates its spine. Parallax artifacts are removed. The tracked mouse images are then subjected to a wavelet decomposition, and then reduced to a principal component representation, yielding a multidimensional time series. This principal component time series is then modeled with one of a family of models, which is then fit using Gibbs sampling iterations; this procedure segments the principal component time series into reused behavioral modules. **B.** The workflow for fitting behavioral data with the AR-HMM, the best-performing of the family of models composed to represent different hypotheses about the underlying structure of behavior. The input to the workflow extracted and aligned mouse images, and the output is a label sequence equal in length to the original data timeseries that assigns a

behavioral module identity for every recorded frame of behavioral data. This process is described in detail in the Supplemental Experimental Procedures. (i) First, after pre-processing image data are compressed from 3600 to 10 dimensions via wavelet decomposition and PCA. (ii) Data are randomly split into “train” and “test” segments (in a 70:30 ratio), in order to evaluate the performance of the model on unseen data, and to guard against overfitting. (iii) The parameters of the model are randomly initialized (parameters describing each module’s characteristic trajectory through pose space, duration distribution, and transition statistics relative to other modules), and fit iteratively with Gibbs sampling. Gibbs sampling alternates between several updates: the algorithm first attempts to segment the imaging data into modules given a fixed set of transition statistics and a fixed description of the AR parameters that describe any given module, and then the algorithm switches to fixing the segmentation and updating the transition matrix and the AR parameters. Filled circles represent fixed parameters, while open circles represent parameters that are varied in a particular iteration. (iv) The product of this process is the parameters described above: “AR Matrices,” the AR parameters that describe the specific trajectory through pose space that defines any given behavioral module, “Duration Distributions,” the parameters describing the duration distribution for each behavioral module (in model classes in which this distribution is explicitly modeled), “Transition Matrix,” the probability that any given module will precede or follow any other module, and “State Labels,” the behavioral module inferred for each frame of 3D video data. (v) The quality of the fit in terms of predicting behavior is evaluated using the likelihood of data that was held out from training. C. A schematic representation of the aspects of behavior captured by the AR-HMM. The AR-HMM takes in a sequence of mouse behavioral data (fed into the model as 10 PCs, which are derived from pixel data), and generates a sequence of re-used labels over time

(“Labels”), such that every moment of the mouse’s behavior is associated with a single behavioral module. Each behavioral module is internally modeled as an autoregressive process that can be visualized as a trajectory through pose space (“AR Process 1,” AR Process 2,” etc.). The transition probabilities between modules are represented in a transition matrix, which may be represented as a state map, where each behavioral module is a node, and the probability of transition between two nodes is represented by the thickness of a directed arrow between those nodes (“Transition Probabilities”). **D.** The AR-HMM outperforms alternative models as measured by next-frame average held-out log likelihood. **E.** Data generated from a fit AR-HMM model are qualitatively similar to PCs derived directly from imaging data.

Figure S5. Characterizing behavioral modules identified by the AR-HMM. Related to

Figure 2 and Supplemental Experimental Procedures. A. Duration distributions for

change-point-identified blocks and AR-HMM-identified modules are similar. **B.** Shuffling

behavioral data at fast timescales lowers AR-HMM performance, as measured by the ratio of

held-out likelihood before and after shuffling; destroying temporal information at a timescale

longer than 500 ms has essentially no effect on model performance, while shuffling at faster

timescales catastrophically prevents effective behavioral prediction. **C.** Data instances drawn

from different behavioral modules trace distinct paths through PCA space. Six separate

representative behavioral modules are plotted in the first two principal components, overlaid on a

density indicating the distribution of all recorded poses. Note that the data represented here are

plotted in 2D; pose dynamics are significantly more separated in 10D (see Fig 2D). **D.** 95 percent

of frames were explained by only 51 behavioral modules; 99 percent of frames were explained

by 65 behavioral modules in the open field dataset (error bars indicate \pm S.E. calculated with

bootstrap samples). **E.** Modules (X axis) sorted by usage (Y axis) with Bayesian credible intervals indicated as error bars, which are smaller than SEs computed with the nonparametric bootstrap (see Fig. 2). **F.** Module interconnectivity is sparse. With no thresholding, each module is interconnected with on average 16.85 ± 0.95 other modules, with the rate of interconnectivity falling sharply with even modest thresholding (Y axis indicates interconnectivity, \pm SE calculated via the nonparametric bootstrap, X axis indicates thresholding applied to transition probabilities). This is consistent with sparse temporal interconnectivity between behavioral modules. These transition probabilities are distinct from bigram probabilities (e.g., Fig. 3B & 4C); bigram probabilities represent the frequency with which one module occurs after another as a proportion of all pairs of temporally adjacent modules. Transition probabilities, in contrast, represent the frequency with which one module occurs after the other as a proportion of all instances of only the preceding module.

Figure S6. Subjective and objective analysis of AR-HMM identified behavioral modules.

Related to Fig. 4. A. Histogram depicting the number of modules that fall into one of five coarse labels of behavior, as scored by a human observer; this coarse subjective description is intended to offer a sense of the distribution of behaviors parsed by the AR-HMM. Note that these module descriptors are completely subjective, and applied post-hoc after modeling. Furthermore, inspection revealed that the modules within each broad category could be further distinguished from the others within that category; the “rear” category, for example, included a wide variety of behaviors in which the mouse’s head was transiently or perdurantly elevated, including true rears, elevated pointing behaviors, head-bobs, and sniffs. **B.** Histogram depicting the average

velocity of the modules that were interconnected after TMT exposure (red), which individually and collectively encode “freezing,” compared to all other modules in the dataset.

Figure S7. Sensitivity analysis of model-free parameters. Related to Supplemental

Experimental Procedures. A. To denoise data from the Kinect a filtering approach was used in which a median filter was applied iteratively both in space and in time; this approach has been shown to effectively maintain data structure while smoothing away noise (see Supplemental Experimental Procedures). To determine filter settings, we imaged dead mice that were differentially posed in rigor mortis; ideal filter settings would distinguish mice that were posed differently, but be unable to distinguish data from the same mouse. Filter settings are indicated as ((pixels), (frames)) with the numbers within each parenthesis referring to the iterative settings for each round of filtering. To assess filter performance, a within/between pose correlation ratio (Y axis) was computed, in which the mean spatial correlation for all frames of the same pose was divided by the mean spatial correlation for all frames of different poses. This revealed that light filtering (with settings ((3), (3,5))) optimized discriminability in the data. **B.** Optimal parameters for changepoints analysis were identified by grid scanning. By maximizing the number of changepoints identified in a live mouse, while fixing the number in a dead mouse to be zero, clear optimal values were identified via grid scanning for sigma and H (left two panels). This changepoint ratio was not highly sensitive to K; a setting of 48 (at the observed maximum) was therefore chosen.

Movie S1. Related to Fig. 1A. Imaging the three-dimensional pose dynamics of freely-behaving mice in the open field with depth cameras. A depth camera placed above the arena

captures the three-dimensional shape of the mouse at 30 frames-per-second as it freely behaves (**right**, height is colormapped as indicated). These raw imaging data are pre-processed (filtered, background subtracted and parallax corrected) and submitted to a machine learning algorithm that aligns the animal along the axis of its spine, thereby enabling quantitative measurements of how its pose dynamics evolve over time (**left**, see Experimental Procedures for details). Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S2. Related to Fig. 1D. Changepoints analysis identifies approximate boundaries between behavioral motifs. Three-dimensional movie of a freely-behaving mouse, with changepoints indicated on upper left. Periods of behavior between the changepoints, which identifies approximate boundaries between blocks, appear to capture behavioral motifs. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S3. Related to Fig. 2B, 2C. The AR-HMM identifies a stereotyped and reused behavioral module that encodes low rearing. Three separate 3D behavior movies (**top left, top right, bottom left**) taken from a single mouse in a single experiment, each depicting a separate instance in which the AR-HMM identified a particular sequence of frames as encoding the same behavioral module. In these movies, the ongoing expression of the module is indicated by a dot being placed over the mouse; imaging frames before the dot onset and after the dot offset are associated with different behavioral modules that are distinct in all three movies, as is the allocentric position of the mouse within the arena. Nevertheless, during expression of the module (i.e., when the dot is over the animal), the mouse in all three examples emits a stereotyped motif of three-dimensional motion that can be identified in these movies as a low rear. By collating and

overlaying 20 of these individual movies (high-likelihood examples taken from 20 different experiments using 20 individual mice) into a single movie (**lower right**), the stereotyped nature of this behavioral module is clear. Note that this “crowd” movie is timelocked such that all the modules are expressed at approximately the same time during the movie; this timelocking was achieved by aligning the middle frame of the module across all examples. Despite this timelocking, the duration of each instance of the behavioral module is slightly different; this is because the AR-HMM defines a unique duration distribution for every behavioral module, rather than a single duration. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S4. Related to Fig. 2C. Behavioral modules are distinct, stereotyped, and reused motifs of behavior. “Crowd” movies (see legend for Movie S3) depicting 10 distinct behavioral modules automatically identified by the AR-HMM, each labeled with a descriptor; three of these are from the trigram depicted in Fig. 2. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S5. Related to Fig. S4E, 2D. The synthetic output of the AR-HMM appears similar to the pose dynamics data exhibited by real mice. A “dream” (left) and real (right) mouse in aligned pose space during normal exploration. The images on the left were generated by training an AR-HMM on open field behavioral data and then, after training, having the model emit synthetic data representing its best estimate of mouse behavior (Supplemental Experimental Procedures). Note that there is some sensor noise apparent in the real mouse (right) that is not modeled by the AR-HMM, but that otherwise these movies are visually very similar.

Movie S6. Related to Fig. 2B. Module sequences encode exploratory behaviors. Three-dimensional imaging of three separate examples of the sequence of three modules (“walk,” “pause” and “low rear”) from Fig. 2 being expressed by mice during free behavior. Note that, in this case, when the dot comes over the animal (see legend for Movie S3 for details) the sequence is being expressed, rather than a single module. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S7. Related to Fig. 3D. Modules encoding wall-hugging behaviors. “Crowd” movie (see legend to Movie S3 for details) of an open-field specific behavioral module (**left**) encoding a behavior in which mice locomote while hugging the walls of the apparatus, referred to as thigmotaxis. In this case, the module can be seen to encode locomotion with a body habitus that is slightly curved to match the curve in the circular walls of the arena. A similar thigmotaxis module can be seen in the square (**right**); in this case, however, the syllable is also used during straight patterns of locomotion in both the square box and the circular open field arena. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S8. Related to Fig. 4E. TMT-regulated behavioral sequences encode freezing behavior. “Crowd” movies of three behavioral modules that are placed into series after exposure to the aversive fox odor TMT (see legend to Movie S3 for details); these trigrams reveal that this sequence of behavioral modules encodes freezing behavior. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S9. Related to Fig. 4D,4E. TMT regulates where behavioral modules are expressed within an experimental apparatus. “Crowd” movie (see legend to Movie S3 for details) of two separate behavioral modules whose usage is not altered by the aversive fox odor TMT, but whose pattern of allocentric deployment is changed. An investigatory sniffing module is expressed near the TMT quadrant (**left**), and a pausing module is expressed away from the TMT quadrant (**right**). Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S10. Related to Fig. 5A, 5C. *Ror1β* mice express a waddling module. “Crowd” movie (see legend to Movie S3 for details) of a *Ror1β*–specific behavioral module that encodes a waddling gait; note that upon the expression of the module the animal’s hindquarters waggle during locomotion. Color scale bar indicates the height of any given pixel above the floor of the arena (mm).

Movie S11. Related to Fig. 6B. Optogenetically-induced spinning behavior. “Crowd” movie (see legend to Movie S3 for details) of the most upregulated behavioral module in response to optogenetic stimulation of the left motor cortex. This module is not expressed during normal behavior, but encodes pathological spinning behaviors that are associated with strong unilateral stimulation of corticostriatal neurons.

Supplemental Experimental Procedures

Methods Summary

Experiments, unless otherwise noted, were performed on C57/BL6 male mice, aged 6-8 weeks, on the dark cycle of a reverse 12 hour light / 12 hour dark cycle. All mice were videotaped with a Microsoft Kinect depth camera and custom recording software. Recorded depth video was subjected to custom algorithms to extract, parallax correct and align the 3D image of the mouse, and then the time-series data were dimensionally compressed using a wavelet transformation and PCA before being subjected to both model-free and model-based analysis. Autocorrelation analysis was performed using correlation distance between time-lagged data points to correct for sensor-specific noise. Spectral analysis was carried out using the Welch periodogram method, and approximate changepoints were identified using a filtered derivative algorithm. Template matching was achieved using a minimal Euclidean distance criterion. Established Bayesian time-series modeling methods were modified to identify behavioral modules. The core AR-HMM was built from three components: 1) an observation distribution, describing the fast-timescale pose dynamics of each behavioral module; 2) a duration distribution, to allow variability in the length of emission of each behavioral module; and 3) a transition distribution (when relevant), which summarizes the interconnectivity over time of identified behavioral modules. Each behavioral module has its own distinct pose dynamics, its own characteristic duration distribution, and its own transition structure in relation to all other behavioral modules. Models were fit using Markov Chain Monte Carlo methods such as Gibbs sampling. The quality of fit was evaluated by testing the model on previously held-out data, which allows direct comparison between alternative model constructions.

Statistical significance of differences in behavioral module usage and transition usage were evaluated by computing both bootstrap samples over the data (on a per-mouse basis) and Bayesian credible intervals. Group comparisons were performed using the Hotelling t-squared statistic, single hypothesis tests using the Wald test. Both tests were adjusted with the Holm-Bonferroni step-down procedure to correct for multiple comparisons.

Mouse Strains, Housing and Habituation

Unless otherwise noted, all experiments were performed on 6-8 week old C57/BL6 males (Jackson Laboratories). Mice from the *Ror1 β* and *Rbp4-Cre* strains (Jackson Laboratories) were habituated and tested identically to the reference C57/BL6 mice. Mice were brought into our colony at 5 weeks of age, where they were group-housed for one week in a reverse 12 hours light/12 hours dark cycle. On the day of testing, mice were brought into the laboratory in a light-tight container, where they were habituated in darkness for 20 minutes before testing.

Behavioral Assays: Innate Exploration

For the open field assay (OFA), mice were habituated as noted above and then placed in the middle of a circular 17" diameter enclosure with 15"-high walls (US Plastics), immediately after which 3D video recording was begun. Mice were allowed to freely explore the enclosure for the 20 minute experimental period. Mice whose behavior was assessed in a square box were handled identically to the OFA.

Behavioral Assays: Stimulus-Driven Innate Behaviors

To assess innate behavioral responses to volatile odorants, an odor delivery system was developed that spatially isolates odors in specific quadrants of a square box. Each 12" x 12" box was constructed of 1/4" black matte acrylic (Altech Plastics), with 3/4" holes patterning the bottom of the box in a cross formation, and a 1/16" thick glass cover (Tru Vue). These holes were tapped and connected via PTFE tubing to a vacuum manifold (Sigma Aldrich) that provides negative pressure to isolate odors within quadrants. Odor was injected into the box through 1/2" NPT-3/8" pipe-fittings (Cole-Parmer). Filtered air (1.0 L/min) was blown over odorant-soaked blotting paper (VWR) placed at the bottom of Vacutainer syringe vials (Covidien). The odorized airstream was then passed through corrugated PTFE tubing (Zeus) into one of the four pipe-fittings in a corner of the odor box. The ability of the odor box to isolate odors within specified quadrants was verified by visualizing vaporized odor or smoke through sheet illumination of the box with a low-power handheld HeNe laser. This approach allowed us to tune the vacuum flow and odor flow rates to achieve odor isolation, which was verified using a photoionization device (Aurora Scientific). To eliminate the possibility of cross contamination between experiments, the odor boxes were soaked in a 1% Alconox solution overnight, then thoroughly cleaned with a 70% ethanol solution. Mice were habituated to the experimental room for 30 minutes before the initiation of the experiment. Under control conditions, dipropylene glycol with air (1.0 L/min) was delivered to each of the four corners of the apparatus before a single mouse was placed in the center of the box and allowed to freely explore while 3D video records were acquired for 20 minutes. The same cohort of animals was tested for odor responses by subsequently repeating the experiment with odorized air delivered to one of the four quadrants. All 3D video recordings were performed in total darkness. TMT (Pherotech) was used at 5% concentration.

Behavioral Assays: Optogenetics

Four adult male *Rbp4*-Cre (The Jackson Laboratory) mice were anesthetized with 1.5% isoflurane and placed in a stereotaxic frame (Leica). Microinjection pipettes (O.D. 10-15 μm) were inserted into the left motor cortex (coordinates from Bregma: 0.5 AP, -1 ML, 0.60 DV). 0.5 μl of AAV5.EF1a.DIO.hChR2(H134R)-eYFP.WPRE.hGH ($\sim 10^{12}$ infectious units/mL, Penn Vector Core) was injected in each mouse over 10 minutes followed by an additional 10 minutes to allow diffusion of viral particles away from the injection site. After the injection, a bare optical fiber with a zirconia ferrule (O.D. 200 μm , 0.37 numerical aperture) was inserted 100 μm above the injection site and secured to the skull with acrylic cement (Lang). Twenty-eight days following the viral injection, mice were placed in a circular arena and the optical implant was coupled to a laser pump (488 nm, CrystaLaser) via a patch-chord and a rotary joint (Doric Lenses). The laser was directly controlled from a PC. After allowing the mouse 20 minutes of familiarization to the arena, optogenetic stimuli were delivered. The laser power, the pulse width, the inter-pulse interval and the inter-train interval were controlled by custom-made software (NI Labview). Each train of laser pulses consisted of 30 pulses (pulse width: 50 ms) at 15 Hz. The interval between successive trains was set to 18 seconds. 50 trains were delivered for each laser intensity. The animal was progressively exposed to higher laser intensities over the course of the experiment.

Camera Setup and Initialization

Mice were tracked in 3D using a Kinect for Windows (Microsoft). A boom tripod (Manfrotto) was used to suspend the camera above the recording arena, affording a stable top-down view of the mouse. The Kinect has a minimum working distance (in Near Mode) of 0.5

meters; by quantitating the number of missing depth pixels within an imaged field, the optimal sensor position was found to be between 0.6 and 0.75 meters depending on ambient light conditions and assay material. Thermal effects can cause degradation in the signal coming off the sensor; after each recording session, the camera was therefore unplugged from its power source to allow the sensor to deactivate and cool. The Kinect sensor has a small indicator LED that was masked using electrical tape so as to not emit spurious illumination.

Data Acquisition

Data from the Kinect was sent to an acquisition computer (hand-assembled, 16GB RAM, Intel i7, 512GB SSD) via USB. Custom Matlab software was used to interface the Kinect via the official Microsoft .NET API that retrieves depth frames at a rate of 30 frames per second and saves them in raw binary format (16-bit signed integers), along with the timestamp of each frame, in milliseconds, to an external hard-drive. After acquisition, we indicate a region-of-interest (ROI) delimiting the boundaries of the experimental arena using Matlab GUI tools, and save the polygon specifying the ROI containing the arena.

Data Preprocessing

All data analysis performed with Python scientific computing tools and OpenCV for image analysis (Pedregosa et al., 2011; Perez and Granger, 2007; van der Walt et al., 2011) on r3.8xlarge EC2 virtual machines running an AMI customized from a base Ubuntu 14.04 Linux AMI. Tracking the evolution of an imaged mouse's pose over time requires identifying the mouse within a given video sequence, segmenting the mouse from the background (in this case the apparatus the mouse is exploring), orienting the isolated image of the mouse along the axis of

its spine, correcting the image for perspective distortions, and then compressing the image for processing by the model.

To isolate our analysis to the experimental arena in which the mouse is behaving, a region-of-interest (ROI) was identified for further analysis that was manually traced along the outside edge of any imaged arena; pixels outside this ROI were by default set to zero to prevent spurious object detection. The raw imaging data were filtered with an iterative median filter (Arias-Castro and Donoho, 2007), which is well suited to removing correlated noise from the sensor within the Kinect. Filter settings were identified that maximized between-pose discriminability while minimizing within-pose differences by using dead mice in rigor mortis that were placed into distinct poses (Fig. S7A); based upon this analysis, a three pixel median filter over space, and then a three-then-five frame iterative median filter over time was applied to the imaging data.

To subtract the background image of the arena, the median value of the first 30 seconds of data from any imaging stream were taken, and subtracted from all video frames; any spurious values less than zero were by default reset to zero. The image was then binarized and any objects that did not survive three iterations of morphological opening were eliminated. The mouse was defined as the largest object within the arena that survived the subtraction and masking procedures. The centroid of the mouse was identified as the center-of-mass of the preprocessed image; an ellipse was fit to its contour to detect its overall orientation. In order to orient the mouse along the axis of the spine, a random forest classifier was trained on a set of manually oriented mouse images to facilitate automated identification of the head and tail. Although this algorithm was nearly always correct, its output was manually supervised to ensure accuracy. Note that this procedure does not directly identify the spine, but that the ellipse-based alignment

procedure has the effect of aligning the animal along this anatomic axis. During episodes when the mouse was reared directly upwards (and therefore appears as a circle against the sensor, which is orthogonal to the axis of the mouse's spine under these conditions) neither a human nor our algorithm can identify the spinal axis; because we model pose dynamics based upon dimensionally-reduced representations of the pixels (rather than extracted scalar features, like orientation or length), the presence of this symmetry does not adversely affect model performance, and synthetic data generated from the model appropriately captures these rearing episodes without “flipping” the mouse along its spinal axis.

The position of the mouse within the arena enabled extraction of allocentric data about the mouse, including the centroid, head and tail positions of the mouse, orientation, length, width, height, and each of their first derivatives with respect to time. Critically, none of these scalar parameters are fed to the model; we obtain these data for the sake of convenience and to understand how the egocentric behaviors captured by the model might relate to allocentric parameters like position or velocity.

To correct perspective distortion in the X and Y axes, a tuple of (x,y,z) coordinates was generated for each pixel in real-world coordinates, and then those coordinates were resampled to fall on an even grid in the (x,y) plane using Delaunay triangulation. This procedure corrected for the slight differences in the appearance of the mouse when at the left or right edges of the arena; these differences derive from parallax effects caused by imaging the animal from a single fixed vantage point from above, but are normalized by the interpolation procedure. All extracted images and scalars were saved in HDF5 format. The extraction pre-processing which transforms raw data into the extracted HDF5 format runs at approximately real-time (e.g. 20 minutes of

preprocessing is required for a 20 minute recording session) on a 32-core r3.8xlarge EC2 instance.

The full extracted image size of a mouse covered an area of 120 mm², where each pixel is 2 mm on a side; each frame to be submitted for analysis contained 3,600 16-bit integers (which were cast to floating point before any image analysis) in a 60x60 array. The information captured in this image was often either highly correlated (neighboring pixels) or uninformative (pixels on the border of the image that never contain the mouse's body). To both reduce redundant information and make modeling computationally tractable, each image was dimensionally reduced. A five-level wavelet decomposition was first applied, thereby transforming the image into a representation in which each dimension captured and pooled information at a single spatial scale; in this transformation, some dimensions coded explicitly for fine edges on the scale of a few millimeters, while others encoded broad changes over spatial scales of centimeters (Mallat, 1989). This wavelet decomposition expanded the dimensionality of the image; to reduce this dimensionality principal components analysis was applied to these vectors, and projected the wavelet coefficients into ten dimensions, which captures 88% of total variance (Fig. S2D) and which were submitted to our modeling algorithms. Principal components were built using a canonical dataset of 25 C57 BL/6 mice, aged 6 weeks, recorded for 20 minutes each in the OFA, and all datasets were projected into this common pose space.

For some model-free depictions of behavioral data, the random projections technique was used to reduce the data dimensionality. This approach also has the benefit of representationally normalizing the data to changes in the size of the mouse on the sensor. For example, when a mouse rears up the number of pixels that describe the mouse changes drastically, leading to obvious striations in the raw pixel data over time that may visually overstate the overall change

in the behavior of the mouse. Because, as described below, each random projection dimension takes into account all of the pixels within the 60x60 array, representations of the random projections data are not dominated by these effects, enabling more holistic inspection of changes in the data over time. Random projections is an approach that produces new dimensions derived from an original signal, with dimensionality D_{orig} , by randomly weighting each original dimension, and then summing each dimension according to that weighting, producing a single number per data point. This procedure is repeated several times, with new random weightings, to produce a set of "randomly projected" dimensions. The Johnson–Lindenstrauss lemma shows that distances between points in the original dataset with dimensionality D_{orig} are preserved in the randomly projected dimensions, D_{proj} , where $D_{\text{proj}} < D_{\text{orig}}$. As random projections is primarily employed as a visualization aid, we set $D_{\text{proj}} \ll D_{\text{orig}}$, where $D_{\text{proj}}=300$, $D_{\text{orig}}=3600$. For visualization purposes we do not sort the random projections dimensions in any way.

Model-Free Data Analysis Methods

Standard methods for computing autocorrelation were modified, as sensor-specific noise results in a monotonically decreasing autocorrelogram, even for a mouse that is posed in rigor mortis. Average correlation of all 10 dimensions of the mouse's pose data was therefore used as the comparator between time-lagged versions of the time-series signal in question, resulting in a nearly flat autocorrelation function of value ~ 1.0 for a dead mouse, and a declining autocorrelation function for a behaving mouse. The recording of the dead mouse was performed after the animal had entered rigor mortis; the posed dead mouse was placed in the open field arena and recorded for 5 minutes. The rate at which this autocorrelogram declines in a behaving

mouse may be characterized as a time-constant, τ , of an exponentially decaying curve. τ was fit using the Levenberg-Marquardt algorithm (non-linear least squares) using the SciPy optimization package.

Power-spectral density (PSD) analysis was performed on behavioral data to further analyze its time domain structure (Welch, 1967) using a Wiener filter (Wiener, 2013), and implemented by taking the ratio of the PSD of a behaving mouse over the PSD of a dead mouse. The PSD was computed using the Welch periodogram method, which takes the average PSD over a sliding window across the entire signal; a 33-second window was used, sliding in 0.8 second increments.

Plotting the mouse depth data over time (as either raw data or as random projections) yields obvious striations, each a potential boundary between blocks or modules. To automate the identification of these approximate boundaries between blocks, a simple changepoint identification technique was used, called the filtered derivative algorithm (Basseville and Nikiforov, 1993). Briefly, the algorithm calculated the derivative of the per-frame unit-normalized random projections with a lag of $k=4$ frames; this signal was binarized using a threshold $h=0.15$ mm, summed across each of $D=300$ random projection dimensions, and then the resulting one-dimensional (1D) signal was smoothed with a Gaussian filter with a standard deviation of $\sigma=0.43$ frames. Changepoints were identified as the local maxima of this smoothed 1D time-series. As specified above, this procedure depends in part upon the values of the parameters k , h and σ ; these were identified as those values that maximize the number of changepoints in the behaving mouse while yielding no change points in a dead mouse. Optimizing the difference between live and dead mice was largely insensitive to k , but sensitive to h and σ ; however, grid searching revealed clear optimal values for h and σ , which are

used in the analysis above (Fig. S7B). Note that changepoints methods provide access to the overall timescale of behavior, but the specific boundaries between blocks identified by these approaches are highly approximate, as they only consider local data structure and ignore higher-order interactions between modules (including transition structure).

To ask whether any reasonably long snippet of behavior (greater than just a few frames) was ever "repeated" (without reliance on a underlying model for behavior), a template matching procedure was devised to identify similar trajectories through PCA space; here similarity was defined both quantitatively (minimal Euclidean distance in pose PC space) and qualitatively (via inspection of spines and PC trajectories). To identify similar trajectories, the Euclidean distance between some target snippet, the "template", and every possible snippet of equal length was computed. The most similar snippets were selected, ignoring snippets discovered that were shifted less than 1 second from each other (to ensure selection of behavioral snippets that occur distanced in time from each other).

Data Modeling and Fitting Methods

The modeling approach used herein takes preprocessed depth-camera video data (e.g., the 10 top PCs) as input and yields a fit model as output (Fig. S4A), where the fit model includes estimates of several key parameters, including the number of behavioral modules observed within a given set of experimental data, the autoregressive parameters that describe the motion expressed by the mouse within a given module, and the transition parameters that describe how often any particular module precedes or follows any other module; furthermore, for each video frame an assignment is made to the most likely associated behavioral module. To model pre-processed data, we modified established and well-documented modeling and fitting approaches

that have provided insight into complex and dense datasets with time-series structure (Fox et al., 2008; Fox et al., 2009). These approaches include Bayesian nonparametric methods, which allow identification of the most likely number of behavioral modules and lags in the AR process without human supervision. For convenience and clarity, we provide a summary of these methods and their specific application to the 3D behavioral data below, together with formal mathematical descriptions of our modeling and fitting procedures.

In brief, three-dimensional pose data were modeled using generative probabilistic modeling, which is often used to model complex dynamical processes (Bishop, 2006; Koller and Friedman, 2009; Murphy, 2012). By choosing probability models that reflect stylized versions of the physical dynamics that might give rise to the 3D pose dynamics data within a given experiment, this framework allowed the instantiation and testing of hypotheses about different organizations of behavior. These hypotheses ranged from simple progressions of poses described as Gaussians in pose space to the hierarchical AR-HMM that includes both autoregressive states and switching dynamics. Furthermore, using a generative framework for modeling behavior addresses a key limitation in alternative methods used for behavioral classification: distinct models for behavior can now be quantitatively and objectively compared based upon each model's ability to describe the statistics of held-out test data.

To test the hypothesis that behavior is composed of a series of behavioral modules that are stereotyped in form and which are interconnected in time with defined transition statistics, a family of discrete-time hidden Markov models (HMMs) were instantiated (Bishop, 2006; Koller and Friedman, 2009; Murphy, 2012). Here, at each point in time (e.g., for every frame of imaging data), an HMM posits that the mouse is within a discrete state that can be given a label. Each hidden Markov state corresponds to a prototypical brief three-dimensional motif of motion

the animal undertakes while within that state. The Markovian state of the dynamical process is then composed of both the latent discrete component, which identifies the behavioral mode of the animal, and a number of lagged values of the observation sequence, which are used by the autoregressive model to predict the short-timescale behavior of the animal based on the behavioral mode. This model structure is often called a switching vector-autoregressive (SVAR) model or autoregressive HMM (AR-HMM) (Fox et al., 2009; Murphy, 2012). Different experimental conditions were allowed to share the same library of state-specific VAR dynamics but learned their own transition patterns as well as any unique VAR dynamical modes, allowing the model to reveal changes in the parameters due to changes in the experiment.

Compositionally altering model structure allowed objective comparisons between different underlying models for how mouse behavior is organized over time. Removing the discrete switching dynamics captured in the transition matrix and replacing them with a mixture model generated an alternative model in which the distribution over each discrete state does not depend on its previous state. This would be the case if animals had a set of behavioral modules from which to choose, and the likelihoods of expressing any given one of them did not depend on the order in which they appear. This simplification resulted in the autoregressive mixture model (AR-MM). Alternatively, replacing the conditionally autoregressive dynamics with simple state-specific Gaussian emissions resulted in a Gaussian-emission HMM (G-HMM); this model explored the hypothesis that each behavioral module is best described by a simple pose, rather than being a dynamical trajectory. Applying both simplifications yielded a Gaussian mixture model (G-MM), in which behavior is simply a sequence of poses over time in which the probability of expressing any given pose does not depend on the prior pose. Removing the switching dynamics yielded a pure autoregressive (AR) or linear dynamical system (LDS)

model, in which behavior was described as a trajectory through pose space without any reused discrete behavioral modules at all. By fitting each of these alternative models and comparing model performance, the model structure that best explains the behavioral dynamics captured by the 3D imaging was identified (in this case the parent AR-HMM); this model (of the alternatives) thus best reflected the underlying organization of action (Fig. S4).

To fit models we utilized algorithms that have been well validated in a variety of contexts, including specifically for fitting models of the class used herein (Bishop, 2006; Koller and Friedman, 2009; Murphy, 2012; Robert, 2013, r14413). In brief, to estimate the parameters of any given model, approximate Bayesian inference was performed using Gibbs sampling, a Markov Chain Monte Carlo (MCMC) inference algorithm (Robert and Casella, 2013). The Gibbs sampling algorithm employed here has a natural alternating structure (Fig. S4C), directly analogous to the alternating structure of the popular expectation-maximization (EM) algorithm. Applied to the AR-HMM, the algorithm first segments the imaging data into modules given a fixed set of transition statistics and a fixed description of the AR parameters that describe any given module, and then the algorithm switches to fixing the segmentation and updating the transition matrix and the AR parameters. These two steps are alternated as an iterative process.

To employ Bayesian inference methods to fit the model, unknown quantities, including the transition matrix and the autoregressive parameters that describe each state, were treated with a uniform representation as latent random variables. In particular, weak prior distributions (except for the sticky parameter) were placed on these quantities and their posterior distributions after conditioning on observed 3D imaging data were investigated. For the autoregressive parameters, we used Automatic Relevance Detection (ARD) algorithms, which incorporate a prior with a Lasso-like penalty to encourage uninformative lag indices to have their

corresponding regression matrix coefficients tend to zero (Gelman et al., 2003). For the transition matrix, a Hierarchical Dirichlet Process (HDP) prior was used, which served to regularize the number of discrete latent states. In addition, the transition matrix prior also included a sticky bias, which is a single nonnegative number that controlled the tendency of the discrete states to self-transition (Fox et al., 2008; Fox et al., 2009). Because this parameter controlled the timescale of the inferred switching dynamics, this parameter was set such that the output of the model inference algorithms matches (as closely as possible) the model-free duration distribution determined by our changepoint analysis. In this sense the sticky bias acts as a kind of “lens” which focuses the model on the organization of behavior at a specific temporal scale. It is important to note that despite the fact that this parameter was tuned — indeed this parameter defined our prior over the timescale of behavior — it was not clear *a priori* that the model could identify a duration distribution that was similar to the changepoints distribution (or that a duration distribution mode would match the specific sub-second timescale found in the model-free analysis) for any sticky parameter value.

Model Training and Parameter Tests

Datasets from the open-field, odor, and genetic manipulation experiments were modeled jointly in a single model to increase statistical power and facilitate comparisons between results. Because the neural implants associated with the optogenetics experiment modestly altered the physical profile of the animal, these data were modeled separately. In all experiments, the first 10 principal components for the wavelet decomposition coefficients of each frame of each imaged mouse were gathered. Data were then randomly subdivided and assigned either a “train” or a “test” label, with 70% of the data assigned to the training subset, and 30% assigned to the test

subset. The mice labeled “test” were held-out from the training process, and used to test generalization performance via measurement held-out likelihood. This approach allowed us to directly compare algorithms whose composition reflected different underlying structures for behavior.

Models were trained on data using the procedures described above; modeling was robust to both initialization settings and to parameter and hyperparameter settings (with the exception of the sticky bias parameter κ , see below). Specifically, the number of lags used in our AR observation distributions and the number of used states in our transition matrix with an HDP prior was found to be robust to the particular hyperparameter settings on both. The hyperparameters of our sparsifying ARD prior were varied by several orders of magnitude, and held-out likelihood, the number of used lags, and the number of used states varied negligibly. The hyperparameters of our HDP prior were also varied by several orders of magnitude and again no change to the number of used states or held-out likelihood was observed. All jointly-trained data shared observation distributions, but each treatment class was allowed its own transition matrix. Each model was updated through 1000 iterations of Gibbs sampling; upon the last iteration of Gibbs sampling the model output was saved; all further analysis was performed on this final update.

The “stickiness” of the duration distribution of our behavioral modules — defined by the κ setting of the model — influenced the average duration of behavioral modules discovered by the AR-HMM; this allowed us to control the temporal scale at which behavior was modeled. As discussed in the main text, autocorrelation, power spectral density, and the changepoint algorithm identified switching dynamics at a specific sub-second temporal scale (as encapsulated by the changepoints duration distribution and reflected by the spectrogram and autocorrelogram).

The kappa stickiness parameter of the time-series model was therefore empirically set to best match the duration distribution discovered by changepoint detection. To find the kappa setting at which these distributions were best matched, the Kolmogorov-Smirnov distance was minimized between the inter-changepoint interval distribution and the posterior behavioral module duration distribution through a dense grid search. Using alternative distance metrics did not affect the results.

Evaluating the Log-likelihood of Held-out Test Data

A crucial feature of a useful model is its ability to “generalize”, meaning its ability to usefully describe phenomena in similar data that it has not previously seen. This is intimately related to the problem of over-fitting; if a model performs well on training data, but poorly on some new data, then the model has been over-fit. The higher the generalization performance, which was measured as the log-likelihood of held-out test data, the less the model has been over-fit. Model comparisons were performed by evaluating the log-likelihood of held-out test data under the AR-HMM and each member of the family of alternative models. For each Markov Chain Monte Carlo (MCMC) sample of the AR-HMM transition matrix and autoregressive parameters, the log likelihood of a held-out data sequence was computed by summing over every possible hidden state sequence the joint probability of the data sequence and that hidden state sequence (Bishop, 2006).

Predictive comparison metric to stratify model capacities

We use two metrics to compare model performance: the aggregate held-out likelihood score used to compare models in Figure S4D, and a temporal horizon-based predictive

comparison metric in Figure 2A. While the held-out likelihood score used in Fig. S4D compares the models’ overall prediction performance, the metric shown in Figure 2A shows separately the benefit of each component of the model’s structure.

There are two possible versions of the horizon-based prediction metric, both based on predictive likelihood but differing in what data is used to make the prediction:

1. Given the data up to time t , score how well the model predicts what happens at time $t+k$ (for various $k>0$ and averaging over t). Symbolically, this metric is

$$\log p(y_{t+k} | y_1, \dots, y_t)$$

2. Given the data up to time t *and an instant of data before time $t+k$ for the short timescale model*, score how well the model predicts what happens at time $t+k$. More precisely, an “instant of data” is taken to be the amount of data used in the short-timescale AR model. Symbolically, this metric for the AR-HMM, AR-MM, and AR models is

$$\log p(y_{t+k} | y_1, \dots, y_t, y_{t+k-3}, y_{t+k-2}, y_{t+k-3})$$

where the short-timescale data is not included for the G-HMM, G-MM, or Gaussian models (i.e. the “pose” models).

Metric #1 shows overall prediction capability, much like the comparison in Figure S4D, and hence lumps together the benefit from the AR model component and that from the discrete hidden state model component. However, Metric #2 actually separates these effects out: the amount by which the AR-HMM beats the AR-MM is precisely due to the discrete Markov state, since both models use the same short-timescale information. We therefore use Metric #2 for Fig. 2. Metric #2 demonstrates that the AR-MM persistently beats the AR model, corresponding to the fact that it is a more expressive dynamics model for short timescales. Metric #2 also shows

the G-HMM and G-MM (i.e. the “pose models”) as persistently lower because they do not model the short timescale information. Finally, we note that Metric #2 can be computed exactly, while #1 requires approximation with Monte Carlo or importance sampling methods.

Analysis of the Model Consistency

To evaluate model consistency the following procedure was adopted, which enabled comparisons between model fits. For any two fits, each containing a set of behavioral modules, the modules between the two fits were matched by the Frobenius norm between the AR matrices for each module, and then the Hungarian algorithm (also called Munkres matching) was used to generate a mapping that associated a module in the first fit with a single module in the second fit. The number of frames each fit assigned to the module was then plotted as a single point, where the x-axis represents frames assigned to that module in the first fit, and the y-axis represents frames assigned to that module in the second fit. We evaluate how well two fits matched as the goodness-of-fit R^2 value between these points and the $y=x$ line. Both algorithmic consistency, in which the data do not change, but the random seed that initializes the model is different, and data consistency, in which different bootstrapped datasets are used to train the model, were evaluated with this procedure. Results for algorithmic consistency are described in the main text ($R^2 = 0.93 \pm 0.03$). To test data consistency, we randomly sampled a subset of $n=24, 20, 10$ and 5 mice from a dataset with 25 mice (20 minutes each in the OFA), and refit the AR-HMM model 15 times. For $n=24$: $R^2 = 0.95 \pm 0.02$; $n=20$: $R^2 = 0.93 \pm 0.04$; $n=10$: $R^2 = 0.92 \pm 0.04$; $n=5$: $R^2 = 0.88 \pm 0.09$.

Analysis on Model Output

Fit models included a number of states that were infrequently used (see Fig. S5E). Because we wished to focus on those changes in behavior that are most significant, we concentrated our analysis on the set of modules that explain 95 percent of the total frames in our dataset (e.g., the top 51 modules in the dataset) and in the main text only describe modules within this group. Visual inspection reveals that many of the frames assigned to states above this threshold are corrupted by sensor-specific noise that is not eliminated by our pre-processing. Because the model attempts to cluster together data with similar structure, this thresholding procedure has a practical de-noising effect and mitigates ambiguity injected into the imaging by the depth camera itself. Future use of our approach with different depth cameras may require distinct pre-processing and thresholding methods, as alternative cameras will have unique noise characteristics.

To represent the grammatical relationship between behavioral syllables, two analyses were performed. First, the probabilities of transition from any given module to any other module (in a pairwise fashion) were computed; second, the bigram probability was computed, which is the probability that two syllables were found occurring one after the other (a “bigram” of modules), as a fraction of all observed bigrams. To calculate this value for each pair (i,j) of modules, first a square $n \times n$ matrix, A , was created where n is the number of total modules in the label sequence. The label sequences that were saved at the last iteration of Gibbs sampling were then scanned, incrementing the entry $A[i,j]$ for every time syllable i directly preceded syllable j . At the end of the label sequence, each entry was divided by the number of total bigrams observed.

In order to visually organize those modules that were specifically up-regulated or selectively expressed as a result of a manipulation, a selectivity index was assigned to each

module. Specifically, where $p(\text{condition})$ indicates the percent usage of a module in a condition, modules in the circular open field versus square box comparison were sorted by $(p(\text{circle}) - p(\text{square}) / (p(\text{circle}) + p(\text{square})))$. In the comparison between blank odor and fox odor (TMT), modules were sorted by $(p(\text{tmt}) - p(\text{blank})) / (p(\text{tmt}) + p(\text{blank}))$. In the comparison of *Ror1 β* mutant, heterozygous mutant, wildtype and C57/BL6, modules were sorted by $(p(\text{het}) + p(\text{mut})) / (p(\text{wt}) + p(\text{c57}))$.

Statemap Visualizations

To lay out each graph in a reproducible manner (up to global rotation of the figure), the position of the nodes was initialized using the spectral layout algorithm and fine-tuned node positions using the Fruchterman-Reingold iterative force-directed layout algorithm (Koren, 2005); both algorithms were used as implemented in the NetworkX and GraphViz software package. To simplify the visual representation, only edges for which the bigram value was more than some multiple of bootstrap-estimated standard error units (indicated in figure legends) above zero were drawn.

Hinton Diagram

To succinctly display all bigram probabilities between all syllables, the full bigram probability matrix was displayed using a Hinton diagram. In this representation, the value at each index of a rectangular matrix is shown as a 2D geometric object (such as a square or circle) at the spatial index of the matrix's element, where the area of the geometric element is proportional to the value of the matrix at that index. This has the property of representing larger values as more visually salient than smaller values. Additionally, we encoded additional properties of the matrix,

such as the statistical significance of a particular bigram, using color of each geometric element. Because the indices of syllables in our AR-HMM model are chosen randomly by our model, we resorted the rows and columns of the bigram probability matrix using spectral clustering to better represent syllables which tend to transition amongst themselves. This re-sorting did not affect the data itself, only its presentation.

Computing Entropy Rate and Mutual Information

Let A denote the estimated transition matrix, so that $A_{ij} = P[x_{t+1} = j \mid x_t = i]$ for $i, j = 1, 2, \dots, n$ and all time indices t , and let π denote the steady-state distribution of the Markov chain, so that $\pi_i = \Pr[x_t = i]$. The entropy rate on the transition matrix was calculated as $-\sum_{i,j} \pi_i A_{ij} \ln A_{ij}$. The mutual information between states was calculated as $-\sum_i \pi_i \ln \pi_i + \sum_{i,j} \pi_i A_{ij} \ln A_{ij}$.

Cross-likelihood Analysis of AR-HMM States

To evaluate whether discovered AR-HMM states represent meaningful clusters of data in the underlying dataset, we measured how well, on average, each state's autoregressive parameters explain its assigned data segments compared to how well, on average, they explain data segments assigned to other states. This cross-likelihood measure can also be viewed in terms of likelihood ratio testing: given a data segment to assign to one of two possible states, if that data segment was assigned according to a likelihood ratio test, the corresponding cross-likelihood values measure how far on average such data segments are from the likelihood ratio test decision boundary. Cross-likelihoods near or above 1 (i.e., when their logarithms are near or above 0) indicate that the corresponding states' autoregressive models can explain each others'

data segments; the states' assigned data segments cannot be readily or reliably discriminated, at least using the states' autoregressive parameters and a likelihood ratio test. Conversely, likelihoods well below 1 indicate that the states' assigned data segments are easily discriminated.

For any two states $i, j \in \{1, 2, \dots, n\}$ with assigned data segments

$\{y_k^{(i)}\}_{k=1}^{N_i} = \{y_{t_a:t_b}: x_{t_a:t_b} = i, x_{t_a-1} \neq i \neq x_{t_b+1}\}$ and $\{y_k^{(j)}\}_{k=1}^{N_j} = \{y_{t_a:t_b}: x_{t_a:t_b} = j, x_{t_a-1} \neq j \neq x_{t_b+1}\}$ and autoregressive parameters $\theta^{(i)}$ and $\theta^{(j)}$, respectively, we computed the (i, j) entry of the log cross-likelihood matrix as

$$CL_{ij} \stackrel{\text{def}}{=} \frac{1}{N_j} \sum_{k=1}^{N_j} \log \frac{p(y_k^{(j)} | \theta^{(i)})}{p(y_k^{(j)} | \theta^{(j)})}$$

While these cross-likelihood ratios can be interpreted on an absolute scale in terms of classification of data segments using likelihood ratio testing, it is also useful to compare this measure to a null model (i.e., a reasonable baseline hypothesis that should not have meaningful discrete state structure). This baseline was constructed by first fitting with Gibbs sampling a linear dynamical system (LDS) with 60-dimensional latent state to the same open field dataset on which AR-HMM was fit; while an LDS can capture the correlation structure in the data, both in time and across components, its only state is continuous and thus it by definition lacks modularity. Synthetic data were generated from the fit LDS; an AR-HMM was then fit to the synthetic data, thereby producing the corresponding cross-likelihood plot from the inferred states. In this specific case the AR-HMM hyperparameters were adjusted to enable more than one state to be inferred; otherwise, the Bayesian inference procedure only instantiates one autoregressive state to explain the LDS synthetic data, a result that is consistent with our non-parametric methods appropriately recognizing state number in the underlying data.

Visualizing Synthetic Mice

To generate synthetic mice from the AR-HMM and each model in the family of alternative models, the sticky kappa biasing term in each model was replaced with an explicit duration model so as to more accurately represent the timescales in each model's inferred label sequence. For each model fit, a corresponding autoregressive hidden semi-Markov model (AR-HSMM) was instantiated with the same autoregressive and transition parameters but with a negative binomial duration model for each state, which was also fit using Gibbs sampling. Each AR-HSMM generates a 10D principal coordinate sequence; using the principal components analysis fit to the original data, these are mapped into images to produce the "dream mouse" videos.

Statistical Methods

To evaluate the model's performance from a frequentist perspective, the inference algorithm is treated as a (stochastic) point estimation procedure, and how that point estimate might change if the estimation procedure were applied to a new random dataset was explored. In particular, a bootstrap procedure was used to estimate the distribution of these point estimates under new random datasets (see below). In addition, tools from hypothesis testing were used to summarize the degree to which the measured module usage differences could be explained by random fluctuation or finite-sample effects. To evaluate the model's performance from a Bayesian perspective, the posterior uncertainty estimated by our MCMC inference algorithm was reported in terms of both raw samples and coordinate-wise posterior standard deviation error bars.

For the purposes of evaluating the procedure as an estimator that might be applied to new random datasets, the last MCMC sample was treated as a (stochastic) point estimate. To estimate the distribution of that point estimate under new random datasets, a bootstrap was performed: the underlying population distributions over sequences for each group (e.g., blank, TMT) were approximated by their empirical distributions, and sampling with replacement from those distributions generated new random datasets. Under the assumption that the data are generated from a stationary random process, this bootstrap procedure provides a consistent estimate of the distribution of the quantities estimated by the model inference, one that has been well validated in a wide range of contexts (Wasserman, 2013).

To solve the label matching problem in the bootstrap samples, the new fits were initialized from the first fit on the full (un-bootstrapped) dataset; this way the variations in the bootstrap fits reflected the variations due to new random datasets and not any label matching effects (see the model consistency section above for methods that enable independent evaluation of label matching in fits with distinct random initializations). Figures 3-5 show the bootstrap samples and resulting standard error estimates for the module usages in each experiment. In the main text, all additive statistical errors and error bars are bootstrap-estimated standard errors (\pm SE) unless otherwise indicated. We also computed Bayesian uncertainty estimates from MCMC samples. In Fig. S5E we show the last 1000 sampled module usage vectors in a 1250 iteration run of our MCMC algorithm. The small variation indicates that the posterior mode being explored by the MCMC algorithm is highly peaked.

To summarize the degree to which the differences in module usage vector estimates might be explained by random fluctuation and finite sample sizes, null hypotheses were formulated in terms of asymptotic normality assumptions on the distributions of the estimates,

which was validated using Probability-Probability plots (data not shown). To define a single overall null hypothesis that the usage vectors of two populations are the same up to statistical variation, let u and v denote vectors of estimated usage proportions for two populations (e.g., blank and TMT). We test the null hypothesis that $u - v$ is distributed according to a multivariate normal distribution

$$H_0 : u - v \sim \text{Normal}(0, \Sigma)$$

where Σ is estimated from the bootstrap samples. We used a standard test statistic, Hotelling's T-squared statistics, for H_0 .

In addition to testing overall hypotheses about whether module usage vectors are the same in two populations, we also formulated p-values for hypotheses on the individual module usages to bound the (estimated) Type I family-wise error rate. To define these families of module-wise null distributions and p-values, asymptotic normality assumptions were again employed: letting u and v denote the estimated module usage vectors for two populations and letting u_i and v_i denote their i th coordinates, respectively, for each module i we define the hypothesis

$$H_0^{(i)} : u_i - v_i \sim \text{Normal}(0, \sigma_i^2)$$

where the σ_i^2 are estimated from the bootstrap samples. For each such hypothesis an un-adjusted p-value was calculated using a simple Wald test. To adjust the p-values, the Holm-Bonferroni step-down procedure was used, which is a slightly less conservative variant of the Bonferroni correction with the same family-wise error rate guarantees.

The p-values in the optogenetic stimulation experiment were computed differently from the statistics described above. Let $u_i^{(t)}$ denote the frequency (across the 50 trials for a single fit) with which module i is used at frame index t relative to stimulus onset for $t \in [-1, 6]$ seconds

(with 30 frame indices per second). We modeled each $u_i^{(t)}$ as a random variable with a distribution that may depend on the stimulus; that is, we partitioned the time relative to stimulus onset into three subsets corresponding to times when the laser is on ($0 \leq t < 2$), times just after the laser turns off ($2 \leq t < 3$), and all other times. The family of hypotheses that

$$H_0^{(i, \text{on vs off})} : u_i^{(t_1)} - u_i^{(t_2)} \sim \text{Normal}(0, \sigma_{i, \text{on vs off}}^2) \quad 0 \leq t_1 < 2, t_2 \in [-1, 0) \cup (3, 6]$$

$$H_0^{(i, \text{offset vs off})} : u_i^{(t_1)} - u_i^{(t_2)} \sim \text{Normal}(0, \sigma_{i, \text{offset vs off}}^2) \quad 2 \leq t_1 < 3, t_2 \in [-1, 0) \cup (3, 6]$$

was then tested, where each null distribution's variance was estimated from the sample variance across time indices. P-values were then computed using a t-test and adjusted with the Holm-Bonferroni step-down procedure.

Technical definitions of the generative models

In this section, the priors and generative models implemented in this paper are formally defined following standard notation. These generative models are adaptations of models successfully used previously to characterize structure in complex datasets (Fox et al., 2009; Murphy, 2012).

Prior on the transition matrix

A sticky HDP prior was placed on the transition matrix π with concentration parameters $\alpha, \gamma > 0$ and sticky parameter $\kappa > 0$

$$\begin{aligned} \rho_j &\stackrel{\text{iid}}{\sim} \text{Beta}(1, \gamma) & \beta_i &= (1 - \rho_i) \prod_{j < i} \rho_j \\ \pi_i &\stackrel{\text{iid}}{\sim} \text{DP}(\alpha\beta + \kappa\delta_i) & i &= 1, 2, \dots \end{aligned}$$

where δ_{ij} is 1 when $i = j$ and is 0 otherwise and π_i denotes the i th row of π . Gamma priors are placed on α and γ , setting $\alpha \sim \text{Gamma}(1,1/100)$ and $\gamma \sim \text{Gamma}(1,1/100)$.

Generation of the discrete state sequence

Given the transition matrix, the prior on a discrete state sequence x was

$$x_t \sim \pi_{x_{t-1}} \quad t = 2, 3, \dots, T$$

where x_1 is generated by the stationary distribution under π .

Prior on the autoregressive parameters

The autoregressive parameters $\theta = \{\theta^{(i)}\}_{i=1}^{\infty} = \{A^{(i)}, b^{(i)}, \Sigma^{(i)}\}$ for each state $i = 1, 2, \dots$ were sampled from a Matrix Normal Inverse-Wishart prior :

$$(A, b), \Sigma \sim \text{MNIW}(v_0, S_0, M_0, K_0)$$

or equivalently

$$\begin{aligned} \Sigma &\sim \text{InvWishart}(v_0, S_0) \\ \text{vec}((A, b)) &\sim \text{Normal}(\text{vec}(M_0), \Sigma \otimes K_0) \end{aligned}$$

where \otimes denotes a Kronecker product and (A, b) denotes the matrix formed by appending b to A as a column. Note that the noise covariance matrices $\Sigma^{(i)}$ are dense full-rank matrices, not restricted to diagonal matrices (i.e. uncorrelated noise models), and there is more than enough data to learn these fully general noise models. In addition, a block ARD prior (Fox et al., 2009; Pole et al., 1994) on K_0 is used to encourage uninformative lags to be shrunk to zero:

$$K_0 = \text{diag}(k_1, \dots, k_{KD}) \quad k_i \stackrel{\text{iid}}{\sim} \text{InvGamma}(1/25, 1/25)$$

Generation of the 3D pose sequence principal components

Given the autoregressive parameters and discrete state sequence, the data sequence y was generated according to an affine autoregression :

$$y_t \sim \text{Normal}(A^{(x_t)}\tilde{y}_{t-1} + b^{(x_t)}, \Sigma^{(x_t)}) \quad t = K + 1, K + 2, \dots, T$$

where the covariance matrices Σ are dense and \tilde{y} denotes a vector of K lags:

$$\tilde{y}_t \stackrel{\text{def}}{=} [y_{t-K}^T \quad y_{t-K+1}^T \quad \dots \quad y_{t-1}^T]^T$$

The alternative models are special cases of the AR-HMM and were constructed by adding constraints. In particular, the Gaussian-emission HMM (G-HMM) corresponds to constraining $A^{(i)} = 0$ for each state index i . Similarly, the autoregressive mixture (AR-MM) and Gaussian mixture (GMM) correspond to constraining the transition matrix to be constant across rows, $\pi_{ij} = \pi_{i'j} = \pi_j$ for each i and i' , in the AR-HMM and G-HMM, respectively.

Technical description of the inference algorithms

As discussed above, the Gibbs sampling inference algorithm alternated between two principal stages: updating the segmentation of the data into modules given a fixed transition matrix and autoregressive parameters, and updating the transition matrix and autoregressive parameters given a fixed segmentation. These fitting procedures are adaptations of those successfully used previously to fit generative models including the AR-HMM (Fox et al., 2008; Fox et al., 2009). Mathematically, updating the segmentation sampled the label sequence x conditioned on the values of the data y , the autoregressive parameters θ , and the transition matrix π ; that is, sampling the conditional random variable $x \mid \theta, \pi, y$. Similarly, updating the transition matrix and autoregressive parameters given the segmentation sampled $\pi \mid x$ and $\theta \mid x, y$, respectively.

For inference in the AR-HMM the weak limit approximation to the Dirichlet process was used, in which the infinite model was approximated by a finite one. That is, choosing some finite approximation parameter L , β and π were modeled using finite Dirichlet distributions of size L

$$\beta \sim \text{Dir}(\gamma/L, \dots, \gamma/L)$$

$$\pi_k \sim \text{Dir}(\alpha\beta_1, \dots, \alpha\beta_j + \kappa\delta_{kj}, \dots, \alpha\beta_L)$$

where π_k denotes the i th row of the transition matrix. This finite representation of the transition matrix allowed the state sequence x to be resampled as a block and for large L provides an arbitrarily good approximation to the infinite Dirichlet process.

Using a weak limit approximation, the Gibbs sampler for the AR-HMM iterated resampling the conditional random variables

$$x \mid \pi, \theta, y \quad \theta \mid x, y \quad \text{and} \quad \beta, \pi \mid x$$

For simplicity, throughout this section notation for conditioning on hyperparameters and the superscript notation for multiple observation sequences is suppressed.

Sampling $x \mid \pi, \theta, y$

Sampling the state labels x given the dynamical parameters, π and θ , and the data y corresponds to segmenting the 3D video sequence and assigning each segment to a behavioral mode that describes its statistics.

Given the observation parameters θ and the transition parameters π , the hidden state sequence x is Markov with respect to a chain graph. The standard HMM backward message passing recursions are

$$B_t(k) = p(y_{t+1:T} \mid \theta, \pi, x_t = k)$$

$$= \sum_{j=1}^K p(x_{t+1} = j | x_t = k, \pi) p(y_{t+1} | x_{t+1} = j, \theta) B_{t+1}(j)$$

for $t = 1, 2, \dots, T - 1$ and $k = 1, 2, \dots, K$, where $B_T(k) = 1$ and where

$y_{t+1:T} = (y_{t+1}, y_{t+2}, \dots, y_T)$. Using these messages, the conditional distribution of the first state x_1 , marginalizing over all the future states $x_{2:T}$ is

$$p(x_1 = k | \pi, \theta, y) \propto p(x_1 = k | \pi) p(y_1 | x_1 = k, \theta) B_1(k)$$

which can be sampled efficiently. Given a sampled value \bar{z}_1 , the conditional distribution of the second state x_2 is

$$p(x_2 = k | \pi, \theta, y, x_1 = \bar{z}) \propto p(x_2 = k | x_1 = \bar{z}_1, \pi) p(y_2 | x_2 = k, \theta) B_2(k).$$

Therefore after passing HMM messages backward the state sequence can be recursively sampled forwards.

Sampling $\theta | x, y$

Sampling the autoregressive parameters θ given the state sequence x and the data sequence y corresponds to updating each mode's dynamical parameters to describe the 3D video data segments assigned to it.

To resample the observation parameters θ conditioned on a fixed sample of the state sequence x and the observations y one can exploit conjugacy between the autoregressive likelihood and the MNIW prior. That is, the conditional also follows the MNIW distribution:

$$p(A^{(k)}, \Sigma^{(k)} | x, y, S_0, \nu_0, M_0, K_0) = p(A^{(k)}, \Sigma^{(k)} | S_n, \nu_n, M_n, K_n)$$

where (S_n, ν_n, M_n, K_n) are posterior hyperparameters that are functions of the elements of y assigned to state k as well as the preceding lagged observations:

$$S_n = S_0 + S_{yy^\top} + (M_0 K_0^{-1} M_0^\top - M_n K_n^{-1} M_n^\top)$$

$$M_n = (M_0 K_0^{-1} + S_{\tilde{y}\tilde{y}^\top}) K_n$$

$$K_n = (K_0^{-1} + S_{\tilde{y}\tilde{y}^\top})^{-1}$$

$$\nu_n = \nu_0 + n$$

where

$$\begin{aligned} S_{yy^\top} &= \sum_{t:x_t=k} y_t y_t^\top & S_{\tilde{y}\tilde{y}^\top} &= \sum_{t:x_t=k} \tilde{y}_t \tilde{y}_t^\top \\ S_{y\tilde{y}^\top} &= \sum_{t:x_t=k} y_t \tilde{y}_t^\top & n &= \#\{t: x_t = k\}. \end{aligned}$$

Therefore resampling $\theta | x, y$ includes three steps: collecting statistics from the data assigned to each state, forming each state's posterior hyperparameters, and updating each state's observation parameter by simulating a draw from the appropriate MNIW. Simulating $(A, \Sigma) \sim \text{MNIW}(S_n, \nu_n, M_n, K_n)$ proceeds as

$$\begin{aligned} \Sigma &\sim \text{InvWishart}(S_n, \nu_n) \\ A &= M_n + \Sigma^{\frac{1}{2}} G K_n^{-\frac{1}{2}} \quad \text{where } G_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0,1). \end{aligned}$$

Sampling $\beta, \pi | x$

Sampling the transition parameters π and β given the state sequence x corresponds to updating the probabilities of transitions among behavioral modules to reflect the transition patterns observed in the state sequence. Updating β encouraged redundant behavioral modes to be pruned from the model, while updating each π_{ij} fit the transitions observed from state i to state j .

Resampling the transition parameters β and π , which are draws from the weak limit approximation to the (sticky) HDP, was performed using an auxiliary variable sampling scheme. That is, $\beta, \pi | x$ was generated by first sampling auxiliary variables $m | \beta, x$. Then $\beta, \pi | x, m$ was generated by first sampling from the marginal $\beta | m$ and then the conditional $\pi | \beta, x$.

The matrix of transition counts in the sampled state sequence x is

$$n_{kj} = \#\{t: x_t = k, x_{t+1} = j, t = 1, 2, \dots, T - 1\}.$$

Suppressing conditioning notation for simplicity, the auxiliary variables

$m = \{m_{kj}: k, j = 1, 2, \dots, K\}$ are sampled via

$$m_{kj} = \sum_{l=1}^{n_{kj}} b_{kjl} \quad \text{where} \quad b_{kjl} \stackrel{\text{iid}}{\sim} \text{Bernoulli} \left(\frac{\alpha \beta_j}{\alpha \beta_j + \kappa} \frac{\alpha \beta_j + \kappa \delta_{kj}}{\alpha \beta_j + l + \kappa \delta_{kj}} \right)$$

where $\text{Bernoulli}(p)$ denotes a Bernoulli random variable that takes value 1 with probability p and takes value 0 otherwise. Note that the update for the HDP-HMM without a sticky bias corresponds to setting $\kappa = 0$ in these updates.

Given the auxiliary variables, the update to β is a Dirichlet-multinomial conjugate one, where

$$\beta | m \sim \text{Dir} \left(\frac{\gamma}{K} + m_{.1}, \frac{\gamma}{K} + m_{.2}, \dots, \frac{\gamma}{K} + m_{.K} \right)$$

where $m_{.j} = \sum_{k=1}^K m_{kj}$ for $j = 1, 2, \dots, K$. The update to $\pi | \beta, x$ is similar, with

$$\pi_k | \beta, x \sim \text{Dir}(\alpha \beta_1 + n_{k1}, \dots, \alpha \beta_j + n_{kj} + \kappa \delta_{kj}, \dots, \alpha \beta_K + n_{kK}).$$

Supplemental References

Arias-Castro, E., and Donoho, D.L. (2007). Does Median Filtering Truly Preserve Edges Better Than Linear Filtering?

Basseville, M., and Nikiforov, I.V. (1993). Detection of abrupt changes (Prentice-Hall Publishing).

Bishop, C.M. (2006). Pattern Recognition and Machine Learning (Springer).

Fox, E.B., Sudderth, E.B., Jordan, M.I., and Willsky, A.S. (2008). An HDP-HMM for Systems with State Persistence. In Proc International Conference on Machine Learning.

Fox, E.B., Willsky, A.S., Sudderth, E.B., and Jordan, M.I. (2009). Sharing Features among Dynamical Systems with Beta Processes. Advances in Neural Information Processing Systems 22 (NIPS 2009), 549-557.

Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. (2003). Bayesian Data Analysis, Second Edition (Chapman & Hall/CRC Texts in Statistical Science). (29 July 2003).

Koller, D., and Friedman, N. (2009). Probabilistic Graphical Models (MIT Press).

Koren, Y. (2005). Drawing graphs by eigenvectors: theory and practice. Computers & Mathematics with Applications 49, 1867-1888.

Mallat, S.G. (1989). A theory for multiresolution signal decomposition: the wavelet representation. Pattern Analysis and Machine Intelligence, IEEE Transactions on 11, 674-693.

Murphy, K.P. (2012). Machine Learning (MIT Press).

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., *et al.* (2011). Scikit-learn: Machine Learning in Python. The Journal of Machine Learning Research 12, 2825-2830.

Perez, F., and Granger, B.E. (2007). IPython: a system for interactive scientific computing. *Computing in Science & Engineering*.

Pole, A., West, M., and Harrison, J. (1994). *Applied Bayesian Forecasting and Time Series Analysis* (CRC Press).

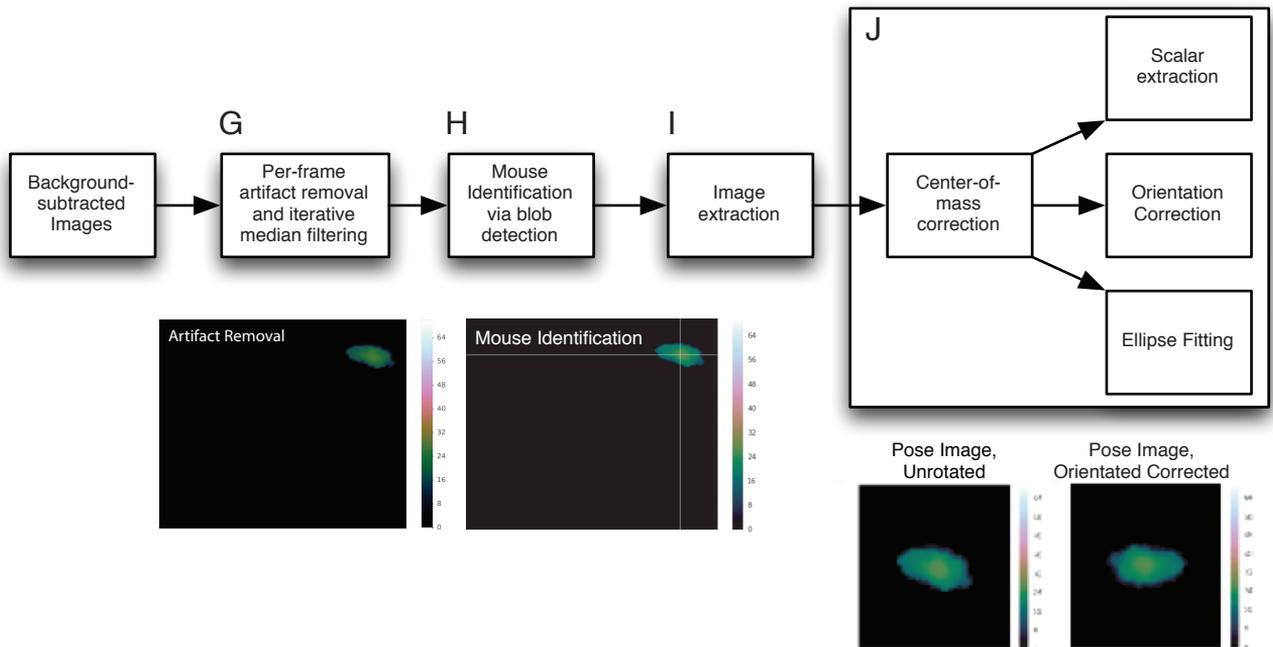
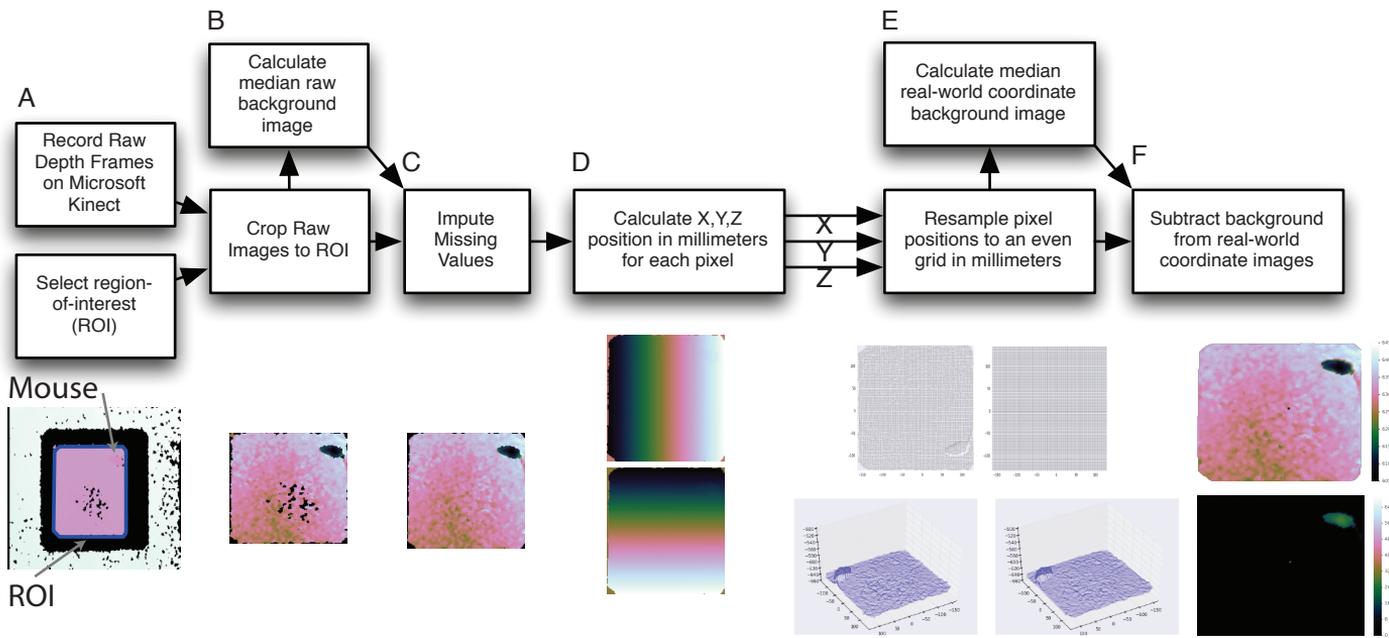
Robert, C., and Casella, G. (2013). *Monte Carlo Statistical Methods* (Springer Science & Business Media).

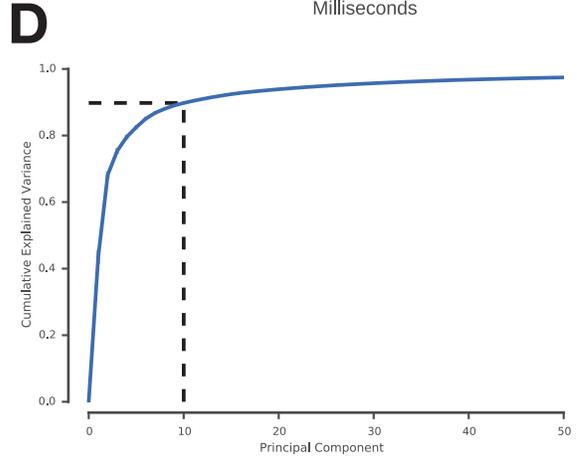
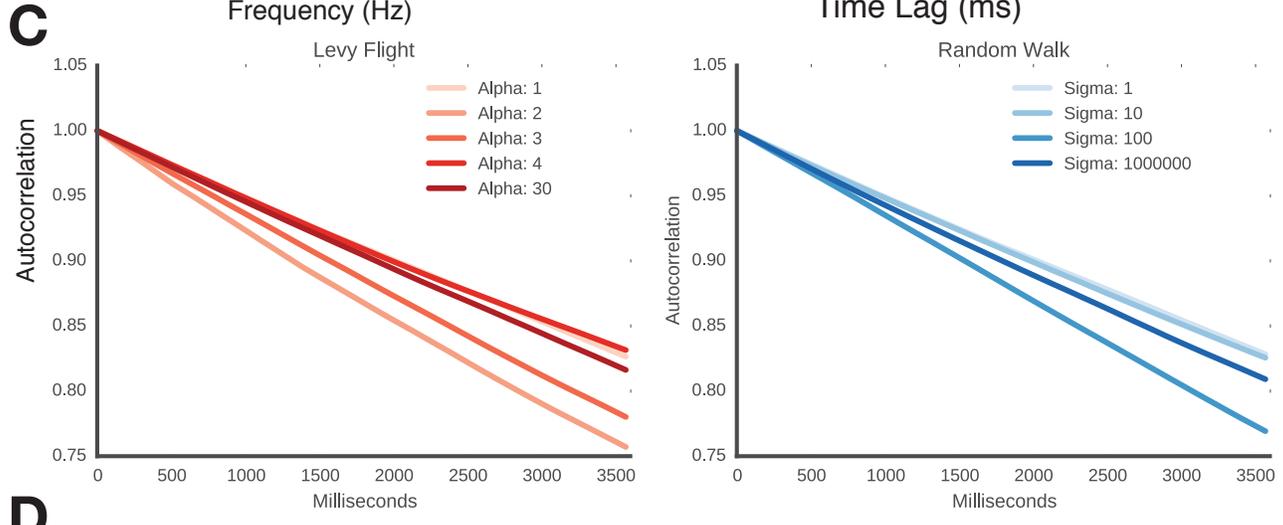
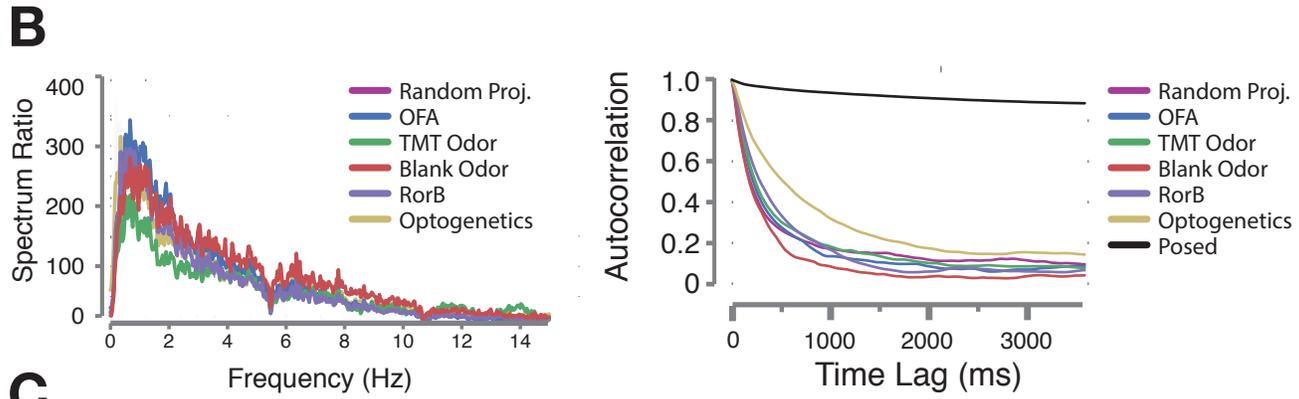
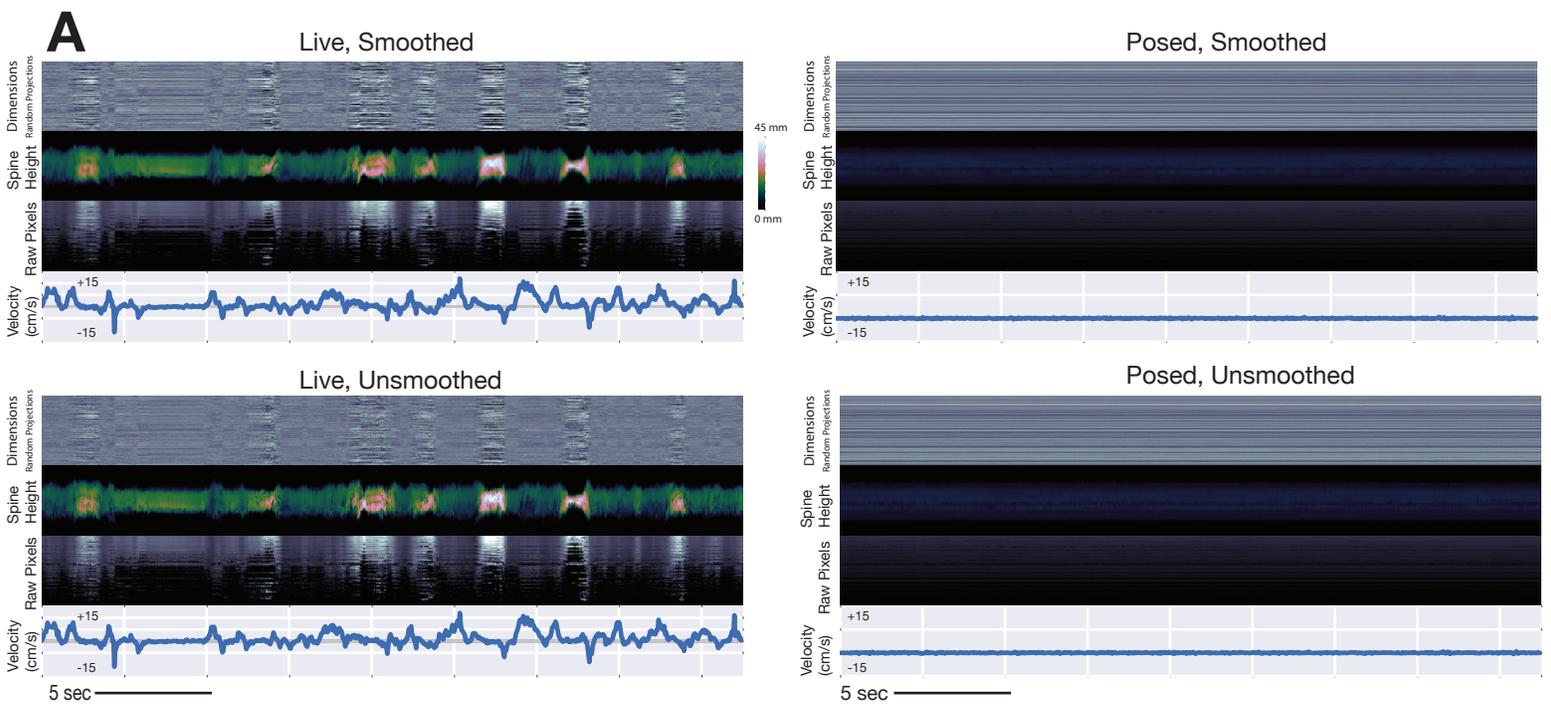
van der Walt, S., Colbert, S.C., and Varoquaux, G. (2011). The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering* *13*, 22-30.

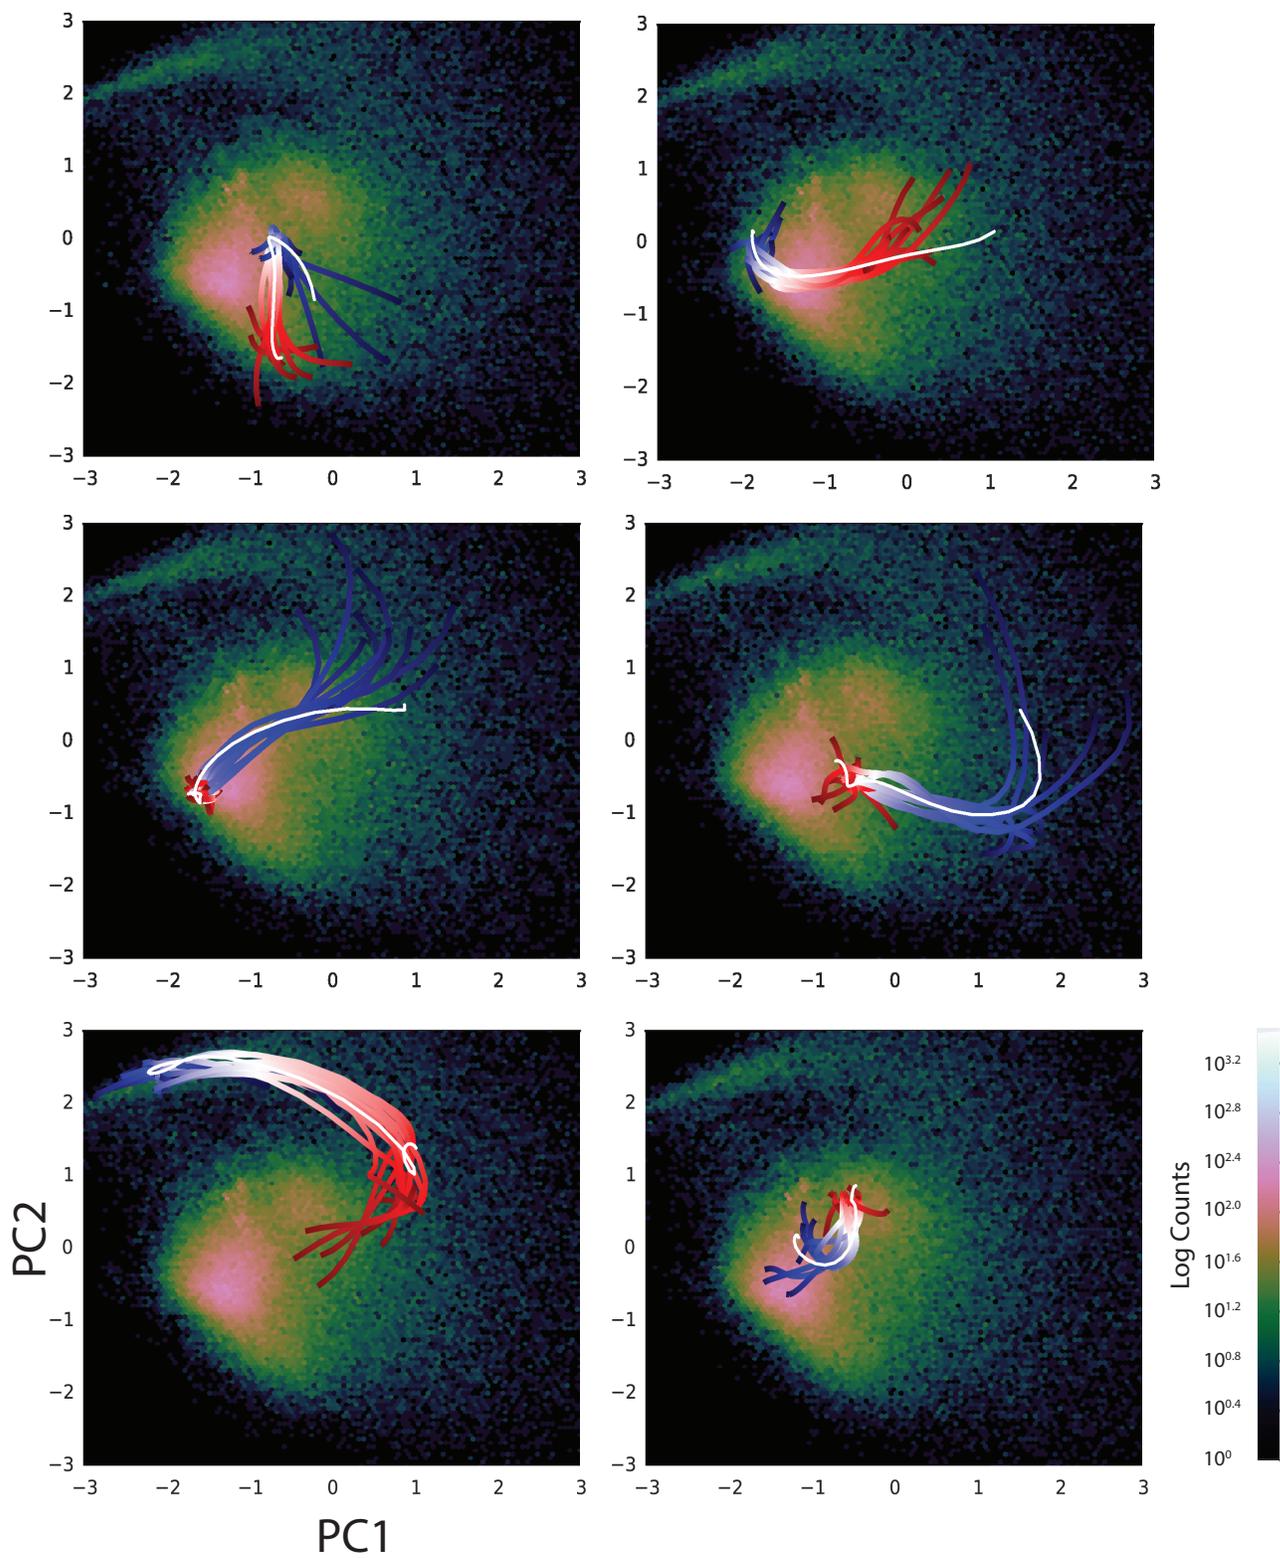
Wasserman, L. (2013). *All of Statistics* (Springer Science & Business Media).

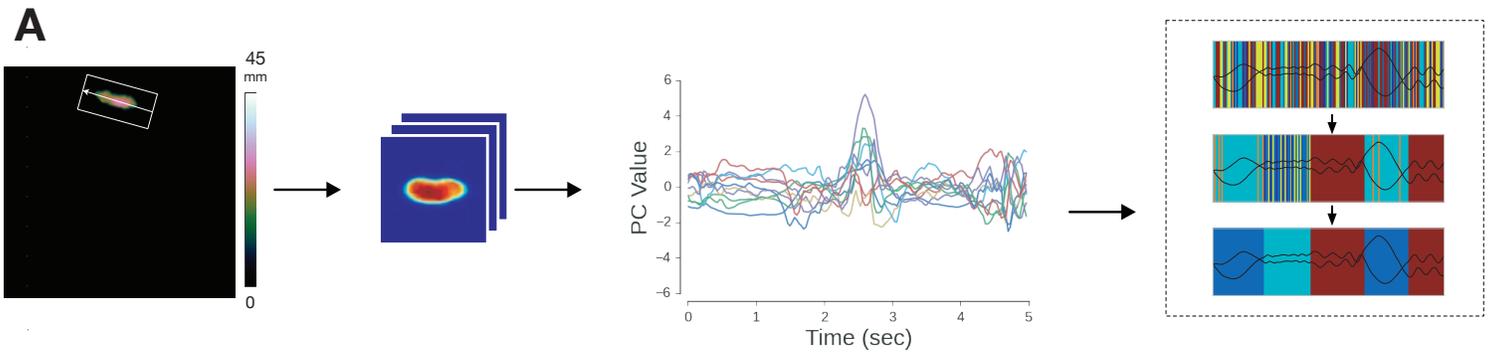
Welch, P. (1967). The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics* *15*, 70-73.

Wiener, N. (2013). *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications* (Martino Fine Books).









Preprocess Raw Imaging Data

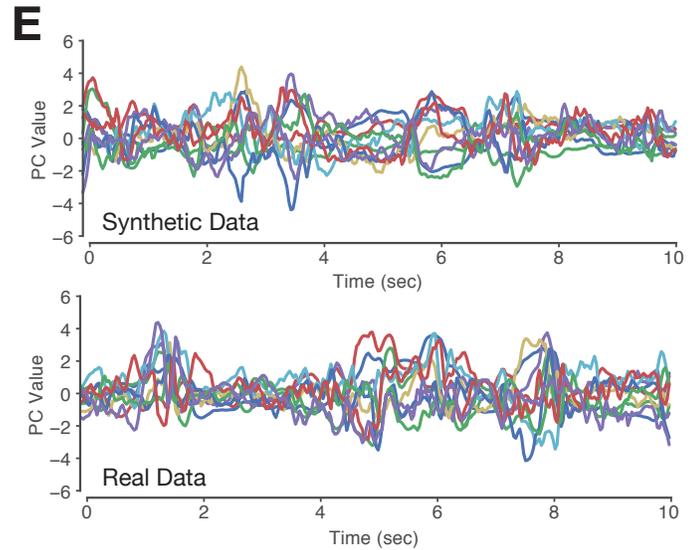
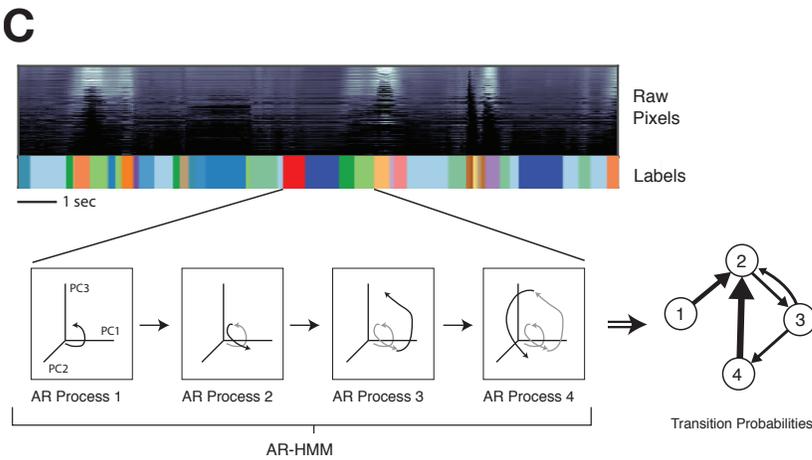
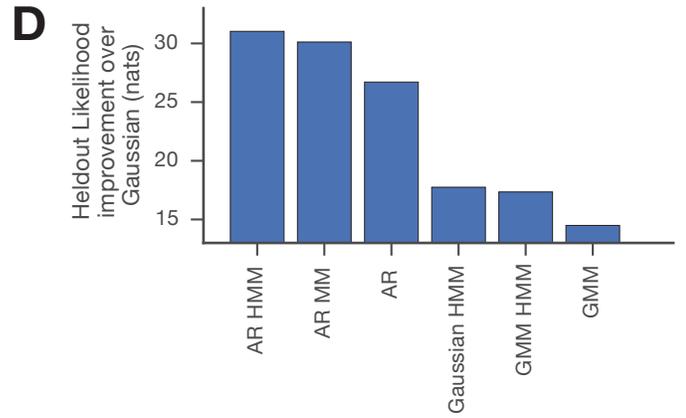
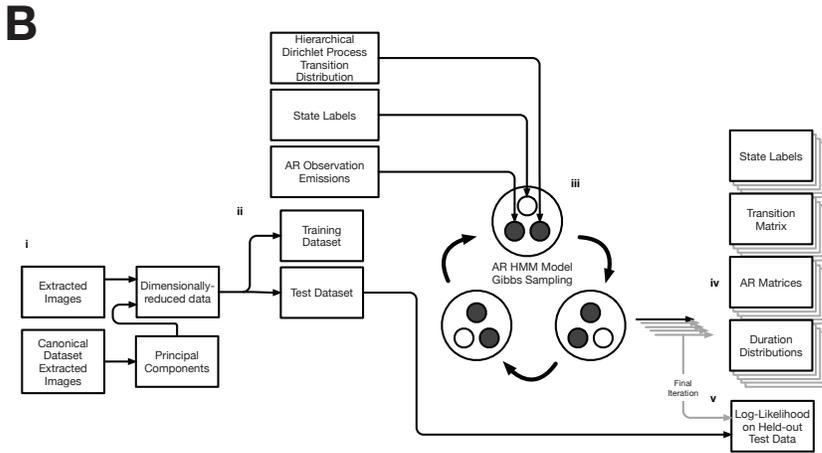
- Extract image of mouse
- Correct parallax artifacts
- Align mouse along axis of spine

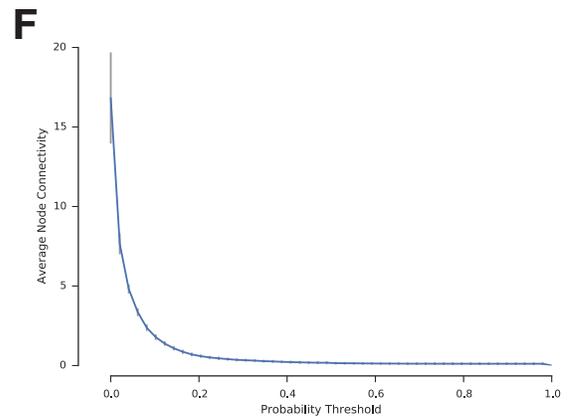
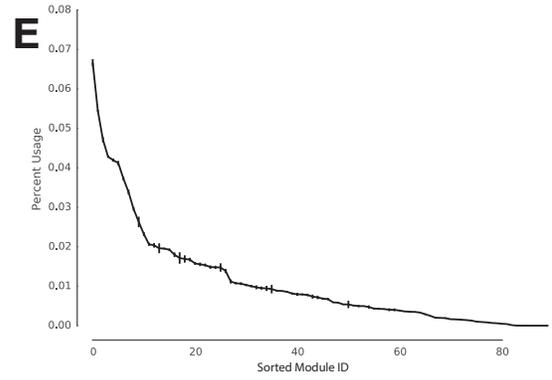
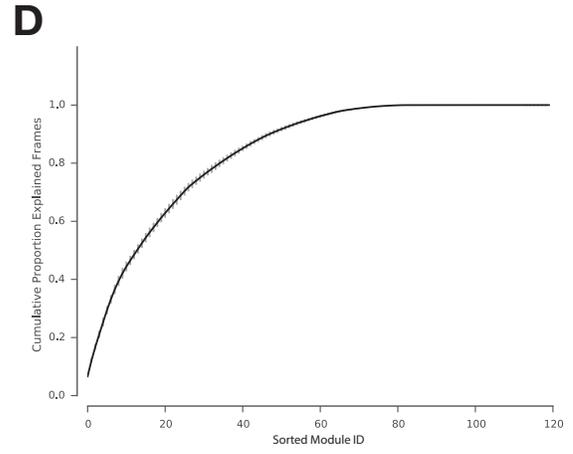
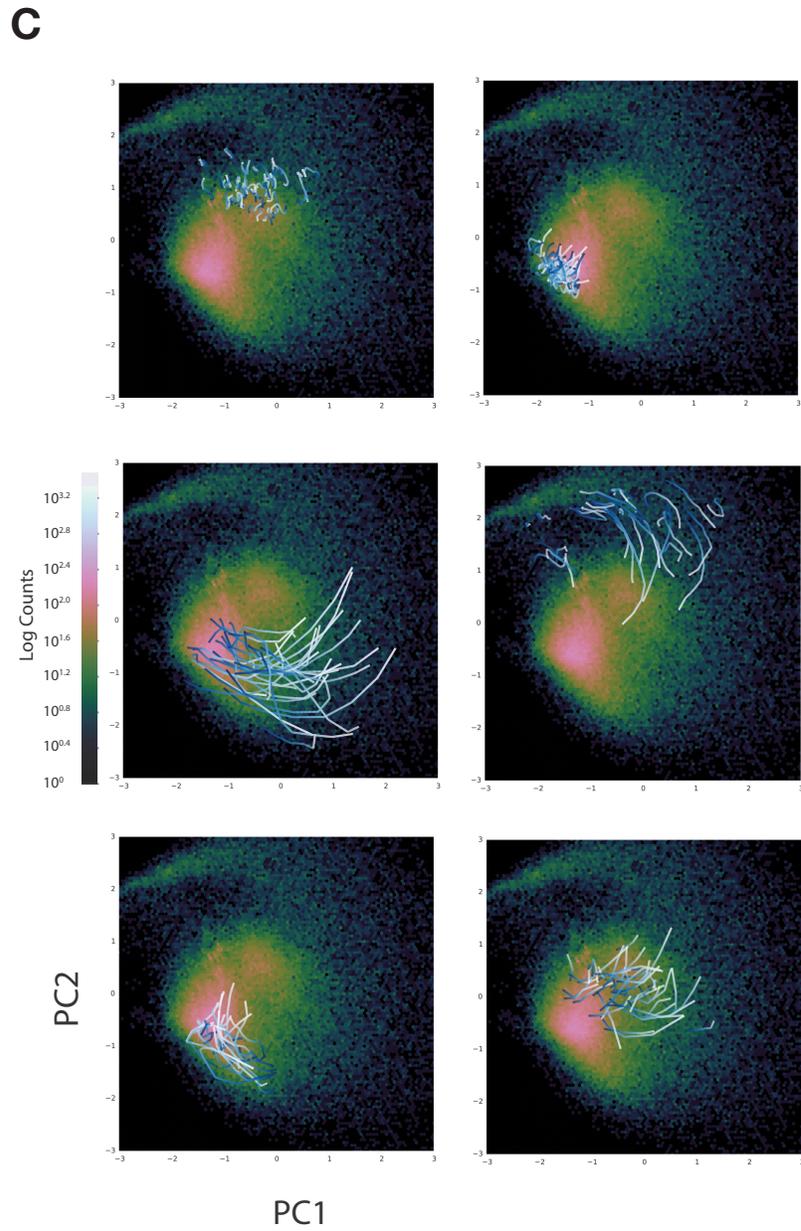
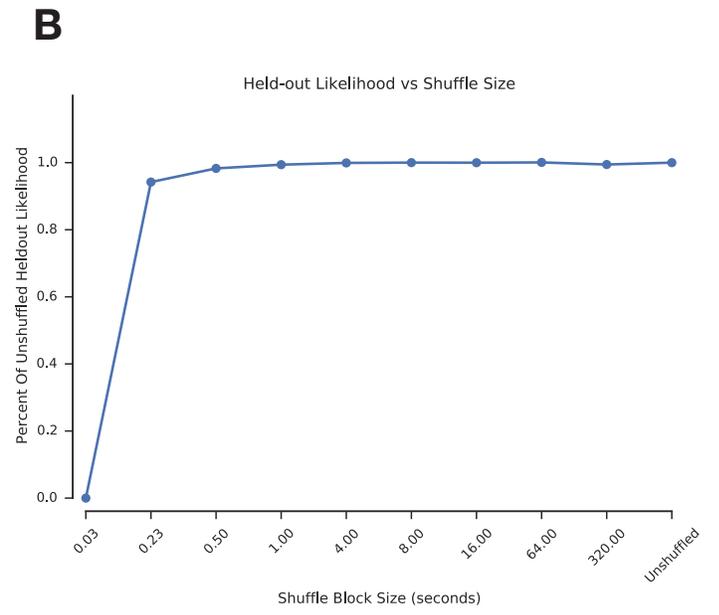
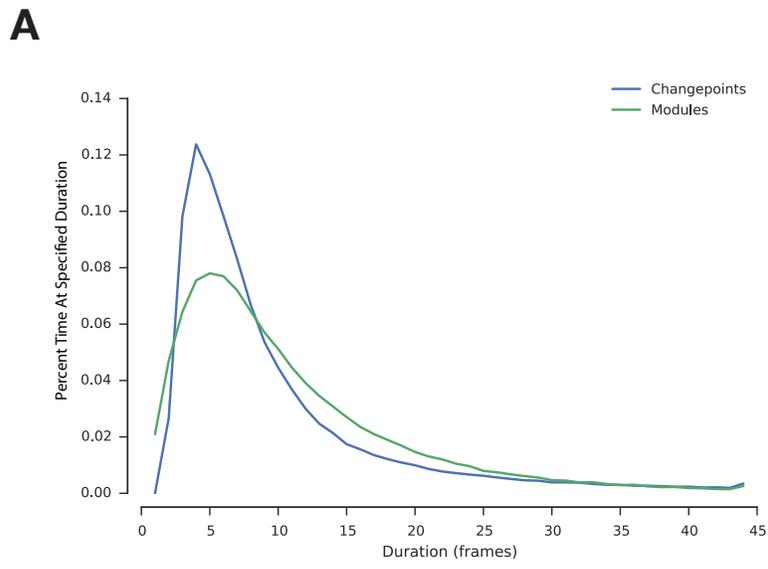
Dimensional Compression of 60x60 pixel data

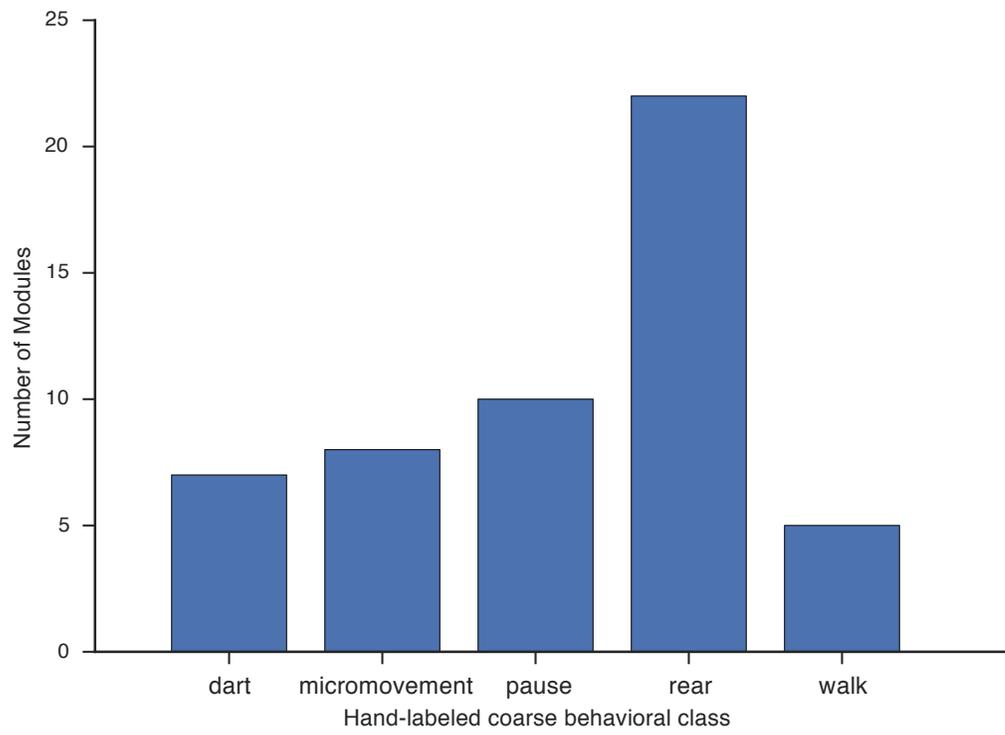
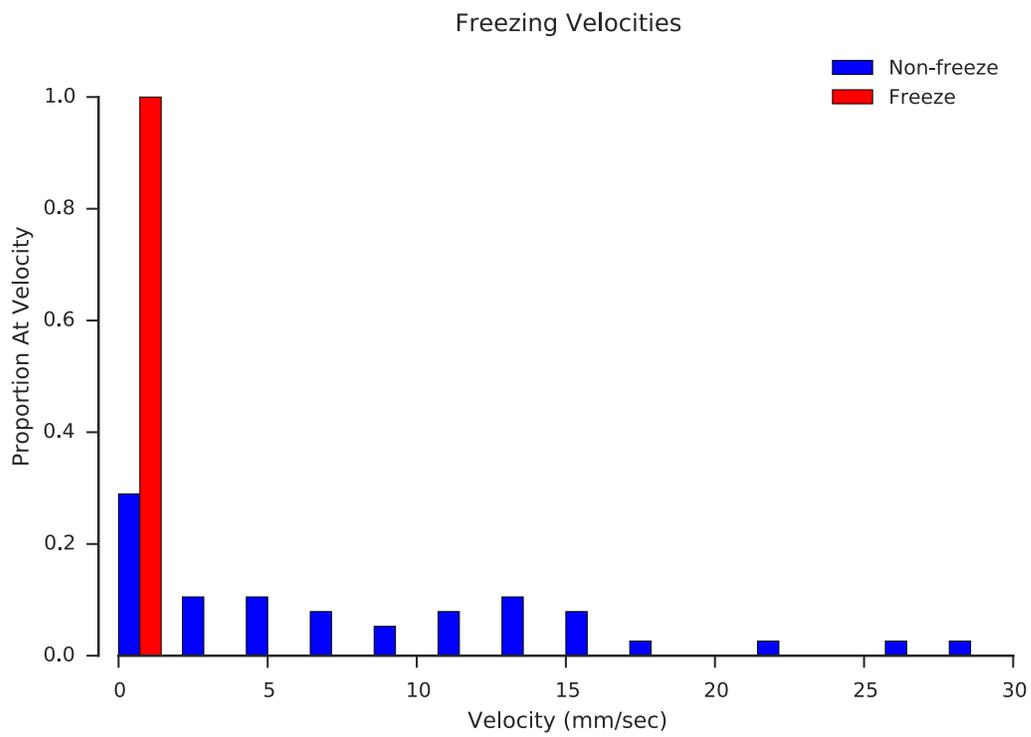
- Wavelet Transformation
- Principal components analysis
- Extract 10 PCs of pixel data

Submit PCA data to various models

Model Fitting with Gibbs Sampling to Identify Behavioral Modules And Transition Statistics

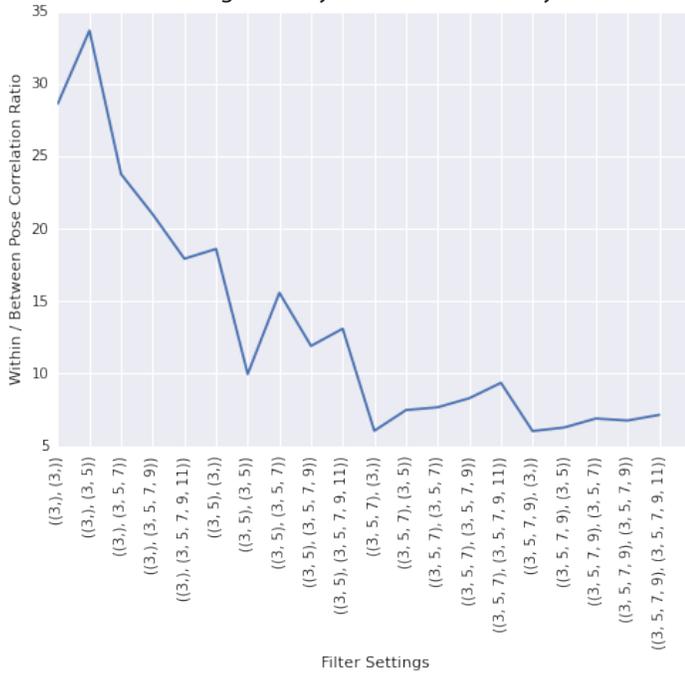




A**B**

A

Determining Filter Settings for Maximizing Ratio of Recognizability over Differentiability

**B**