



HAL
open science

A lexical approach for identifying behavioral action sequences

Gautam Reddy, Laura Desban, Hidenori Tanaka, Julian Roussel, Olivier Mirat, Claire Wyart

► **To cite this version:**

Gautam Reddy, Laura Desban, Hidenori Tanaka, Julian Roussel, Olivier Mirat, et al.. A lexical approach for identifying behavioral action sequences. 2021. hal-03334730

HAL Id: hal-03334730

<https://hal.science/hal-03334730>

Preprint submitted on 4 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A lexical approach for identifying behavioral action sequences

Gautam Reddy,^{1,*} Laura Desban,² Hidenori Tanaka,^{3,4}

Julian Roussel,² Olivier Mirat,² and Claire Wyart^{2,†}

¹*NSF-Simons Center for Mathematical & Statistical Analysis of Biology,
Harvard University, Cambridge, MA 02138*

²*Institut du Cerveau (ICM), Inserm U 1127,
CNRS UMR 7225, Sorbonne Université, Paris, France*

³*Physics & Informatics Laboratories,
NTT Research, Inc., East Palo Alto, CA 94303*

⁴*Department of Applied Physics, Stanford University, Stanford, CA 94305*

(Dated: August 27, 2020)

Abstract

Animals display characteristic behavioral patterns when performing a task, such as the spiraling of a soaring bird or the surge-and-cast of a male moth searching for a female. Identifying such conserved patterns occurring rarely in noisy behavioral data is key to understanding the behavioral response to a distributed stimulus in unrestrained animals. Existing models seek to describe the dynamics of behavior or segment individual locomotor episodes rather than to identify occasional, transient irregularities that make up the behavioral response. To fill this gap, we develop a lexical, hierarchical model of behavior. We designed an unsupervised algorithm called “BASS” to efficiently identify and segment conserved behavioral action sequences transiently occurring in long behavioral recordings. When applied to navigating larval zebrafish, BASS extracts a dictionary of remarkably long, non-Markovian sequences consisting of repeats and mixtures of slow forward and turn bouts. Applied to a novel chemotaxis assay, BASS uncovers conserved chemotactic strategies deployed by zebrafish to avoid aversive cues consisting of sequences of fast large-angle turns and burst swims. In a simulated dataset of soaring gliders climbing thermals, BASS finds the spiralling patterns characteristic of soaring behavior. In both cases, BASS succeeds in identifying action sequences that are highly conserved but transient in the behavior deployed by freely moving animals. BASS can be easily incorporated into the pipelines of existing behavioral analyses across diverse species, and even more broadly used as a generic algorithm for pattern recognition in low-dimensional sequential data.

I. INTRODUCTION

Animal behavior is extremely diverse and context-dependent. Yet, animals often exhibit certain characteristic patterns in their behavior, particularly when executing a task or in response to a stimulus. How can we identify such conserved elements from noisy behavioral data? One increasingly common computational approach to quantitatively describe behavior is to leverage recent developments in the automated tracking of postural dynamics [1–8]. These methods exploit clusters in low-dimensional embeddings of postural dynamics to describe behavior as a sequence of stereotyped elementary locomotor episodes drawn from

* gautam_nallamala@fas.harvard.edu

† claire.wyart@icm-institute.org

a probabilistic model. The resulting descriptions parallel language models, containing information about local dynamics in the form of a probabilistic syntax over individual locomotor episodes [9–19].

Rather than modeling the dynamics of an animal’s behavior or clustering individual elementary locomotor episodes, we propose an alternative approach, where we instead aim to identify conserved *sequences* of behavior. This approach is particularly well-suited for highlighting the differences in navigation across different environments or genetic variants, which reveal how the animal reacts to an environmental or genetic perturbation. Such differences are often quite subtle, making comparative analyses difficult to implement using existing dynamical models. The key difficulty lies in that the majority of behaviorally-relevant responses are transient and occur only a few times in the dataset. By focusing on capturing short time-scale dynamics, dynamical models miss low-copy-number, behaviorally-relevant patterns. Such stretches are lost in the noise and are difficult to pick out from a large dataset.

To give a simple example, consider a scenario where one is presented with a control ‘behavioral’ dataset consisting of a sequence of 100,000 fair coin tosses, and a treatment dataset which is otherwise statistically identical except for 50 sequences of 20 consecutive tails placed at random locations within the sequence. By eye, the sequences in the treatment dataset clearly stand out as abnormal, relevant stretches. On the other hand, a Markov model on heads and tails, for example, when fit to the treatment data may indeed show a statistically significant deviation in its transition matrix from the control, but does not point to the nature of the abnormal stretches or where to find them. Similarly, the process of navigation in a spatially-distributed sensory landscape is highly complex to infer. Despite being robust and conserved, the responses to gradients of attractive or aversive chemical [20], updraft [21, 22] or visual [23, 24] gradients are subtle. The moment at which the animal experiences and reacts to the cue will vary from one animal to the other and is difficult to infer. To identify conserved elements of responses, one then has to find transient behavioral patterns that are over-represented in the environment enriched with chemical, thermal or visual gradients relative to a uniform control.

To discover and construct dictionaries of conserved action sequences, referred to as ‘motifs’ [3], we developed a lexical model of behavior and the unsupervised method BASS (Behavioral Action Sequence Segmentation). Motifs correspond to recurring *sequences* of

elementary locomotor episodes of arbitrary length. A well-established motif for larval zebrafish corresponds to the sequence of locomotor episodes consisting of a J-turn followed by pursuit and capture swims deployed for hunting prey [17, 25, 26]. BASS considers sequences rather than individual locomotor episodes and has the following advantages: 1) BASS does not need to know the sensory information perceived by the animal, 2) BASS captures extended responses that last much longer than a typical locomotor episode, 3) BASS identifies conserved sequences despite them rarely occurring in the behavioral recording.

Our key generative assumption is that animals navigating in a particular environment draw their behavior sequentially from a context-dependent dictionary of motifs. Our goal then is to develop efficient, unsupervised methods to infer this dictionary based on noisy instantiations of the motifs observed in experiments.

If elementary locomotor episodes are represented by symbols, one straightforward approach to motif discovery is to simply enumerate over-represented sequences of n symbols (n -grams). However, the memory and computation time required for this approach increases exponentially with n . An alternative, more efficient approach is to maintain a set of possible sequences (in the form of a dictionary [27] or a suffix tree [28–31]) and add a new motif $\mathbf{m}_1\mathbf{m}_2$ to this set by concatenating two existing motifs \mathbf{m}_1 and \mathbf{m}_2 only if they are juxtaposed more often than chance. While the latter class of methods have enjoyed great success in bioinformatics [27, 29] and text processing [28, 31], the complexity of behavioral data poses an additional challenge. Data in bioinformatics and text processing consist of a well-defined sequence of letters (AGTC or the English alphabet respectively) with little variability in instantiations of a particular word (words are rarely misspelled). We identify three sources of variability in behavioral data: (1) Action pattern noise, which is the variability in instantiations of a particular motif template, (2) Locomotor episode noise, i.e., the variations in observed output, which may lead to a movement appearing as a similar one, and (3) Background variability due to rare and/or erratic movements. To make an analogy with speech learning [32], our task is similar to learning new words from *spoken* language (with no distinctive pauses separating the words) and given prior knowledge of phonology. Action pattern noise, in this analogy, corresponds (not exclusively) to stutters in speech, locomotor episode noise to substitutions of similar phonemes (for example, the aspirated /p^h/ and the unaspirated /p/), and background noise to the utterance of unique proper nouns or unusual sounds. To take into account these sources of variability, we generalize a model-

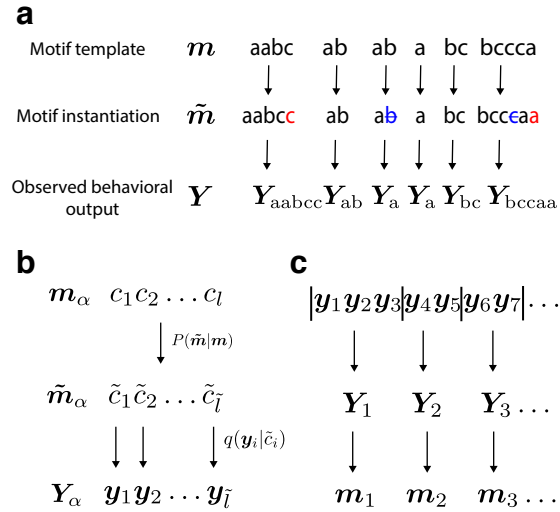


FIG. 1. The generative model from motifs to behavioral output. (a) Motif templates are fixed sequences of elementary locomotor episodes (labeled a, b and c in this example). The observed behavioral output is generated from motif templates drawn sequentially from a dictionary. An instantiation of a template may "mutate" by insertions (red) or deletions (blue), which then generates the observed output as shown in panel (b). (b) The generative process from a motif template $c_1 c_2 \dots c_l$ to instantiation $\tilde{c}_1 \tilde{c}_2 \dots \tilde{c}_l$ to observed output $\mathbf{y}_1 \mathbf{y}_2 \dots \mathbf{y}_l$. (c) The unsupervised inference procedure (BASS) first learns a dictionary of motifs and then segments (vertical bars) the observed behavioral output $\mathbf{y}_1, \mathbf{y}_2, \dots$ into the most likely sequence of motifs $\mathbf{m}_1, \mathbf{m}_2, \dots$ from the dictionary that generated it.

ing framework from bioinformatics ([27]) by introducing an additional two-level hierarchical model. The lower level maps observed behavioral data to a latent state space (similar to going from a Markov model to a Hidden Markov model) and the second level introduces a model for noisy instantiations of motifs (Figure 1a). Despite the model's complexity, we show that inference is tractable and that motifs can be efficiently extracted from datasets of sizes within reach of current experiments.

The larval zebrafish is an interesting vertebrate model organism to investigate the emergence of behavioral action sequences and how they are used to navigate in chemical gradients. In order to survive, five days old zebrafish larvae actively explore their environment avoiding toxic cues and searching for food using stereotypical locomotor episodes consisting of bouts of activity lasting few hundreds of milliseconds separated by distinct pauses [16, 17, 25, 33].

Their small size enables the recording of numerous larvae in parallel, leading to the collection of thousands of swim bouts in a few minutes. Using our lexical approach, we first investigate the behavioral action sequences, i.e., the stereotyped sequences of bout types that larval zebrafish use to spontaneously navigate their environment. Next, we take advantage of a novel chemotaxis assay in which larvae navigate in arenas with gradients of noxious stimuli (acidic pH) and effectively avoid aversive regions. The behavioral response that enables zebrafish larvae to avoid aversive environments is unknown. Examining global kinematic parameters reveals only minor differences, which makes identifying the chemotactic response challenging with classical approaches and thus makes for an appropriate benchmark for our approach.

We first develop the lexical model and the motif identification algorithm, BASS. We apply the algorithm to synthetic data and to datasets obtained from freely-exploring and chemotactic zebrafish larvae. By comparing the dictionaries in the two environments, we identify the sequences that larvae use to chemotax. Lastly, we apply BASS to synthetic datasets of a soaring glider executing a thermalling strategy and show that the algorithm successfully identifies characteristic spiralling patterns.

II. RESULTS

A. A lexical model of animal behavior

Much like language, we assume the behavior of an animal in a particular environment can be described by a sequence of motifs drawn from a dictionary \mathcal{D} , where each motif is a string of arbitrary length containing characters from an alphabet. Motifs are to be considered as *templates* for the generation of action sequences. Each of the K characters in the alphabet represent elementary locomotor episodes and correspond to the unique label of one of the K soft clusters that define the elementary locomotor episodes in postural space. The probability density $q(\mathbf{y}|c)$, in practice obtained through clustering, specifies the probability of observing \mathbf{y} when the animal executes the locomotor episode c . The implicit assumption here is the existence of well-defined elementary locomotor episodes, which has indeed been shown in a variety of systems including rodents, flies, worms and zebrafish larvae [11, 16, 33–35]. We may relax this assumption and instead use clustering as a tiling of postural space, which would manifest as additional noise and a larger alphabet.

Behavior is generated from motif templates, which are sequentially sampled independently and identically from a distribution $\{p_{\mathbf{m}}\}$ over the motifs in the dictionary and individual characters (Figure 1a). The inclusion of individual characters accounts for movements that are not part of any motif, for instance, rare behaviors and erratic movements. These movements constitute background noise that impair motif identification since a motif $\mathbf{m} = c_1 c_2 \dots c_l$ is detectable only if its likelihood is comparable to its constituent characters, $p_{\mathbf{m}} \gtrsim \prod_i p_{c_i}$. Given a sequence of motifs, the data is generated from each template \mathbf{m} according to the probability density $Q(\cdot|\mathbf{m})$ defined below, which is a central element of the model.

The probability, $Q(\mathbf{Y}_\alpha|\mathbf{m}_\alpha)$, of an observed output pattern $\mathbf{Y}_\alpha = \mathbf{y}_1 \mathbf{y}_2 \dots \mathbf{y}_l$ given a motif template $\mathbf{m}_\alpha = c_1 c_2 \dots c_l$ (Figure 1b) defines the behavioral output generated by \mathbf{m}_α . We introduce a model for 'pattern noise': intuitively, if a motif template is viewed as the averaged trajectory of a stochastic dynamical system traversing through a state space, our model for pattern noise corresponds to one where in a particular realization, the trajectory spends a longer or shorter duration at certain regions of state space, but does not deviate into distant regions of state space. In particular, in each instantiation, \mathbf{m}_α 'mutates' to $\tilde{\mathbf{m}} = \tilde{c}_1 \tilde{c}_2 \dots \tilde{c}_l$ with probability $P(\tilde{\mathbf{m}}|\mathbf{m}_\alpha)$. The output \mathbf{y}_i is drawn independently for each character in the mutated sequence from $q(\mathbf{y}_i|\tilde{c}_i)$. To quantify pattern noise, we fix the probability of error per character that results either in the deletion or duplication of that symbol. Note that locomotor episode noise is implicitly incorporated via soft clustering and is determined by the discriminability of neighboring states. We derive a recursive equation for the efficient calculation of $Q(\mathbf{Y}_\alpha|\mathbf{m}_\alpha)$ (see SI Appendix).

Performing inference on this model requires constructing the dictionary \mathcal{D} as well as estimating the motif probabilities $\{p_{\mathbf{m}}\}$. To build our dictionary, we use an iterative procedure generalized from ref. [27] to our latent space model, where we start from a dictionary with only single characters and progressively add new motifs based on how often existing motifs occur next to each other. In particular, we cycle between: (1) estimating $\{p_{\mathbf{m}}\}$ using maximum likelihood estimation (MLE), (2) expanding \mathcal{D} if certain pairs of motifs occur next to each other more often than you would expect from $\{p_{\mathbf{m}}\}$, (3) truncate shorter motifs from \mathcal{D} that are "explained away" by the addition of the longer motifs into the dictionary. We briefly expand on these three steps in the next paragraph; see SI Appendix for further details.

Given a behavioral dataset $\mathbf{Y} = \mathbf{y}_1\mathbf{y}_2 \dots \mathbf{y}_L$, the series of motif templates that generate it is unknown. For example, if $L = 3$, we have $\mathbf{Y} = \mathbf{y}_1\mathbf{y}_2\mathbf{y}_3$, whose likelihood is obtained by summing over all possible ways the dataset can be partitioned: $Q(\mathbf{y}_1)Q(\mathbf{y}_2)Q(\mathbf{y}_3) + Q(\mathbf{y}_1)Q(\mathbf{y}_2\mathbf{y}_3) + Q(\mathbf{y}_1\mathbf{y}_2)Q(\mathbf{y}_3) + Q(\mathbf{y}_1\mathbf{y}_2\mathbf{y}_3)$, where each marginal probability factor in each term is from an instantiation of a particular motif template. In general, the likelihood of \mathbf{Y} under our generative model is the sum over all possible partitionings $\{\pi\}$ of the dataset (of which there are 2^{L-1}) into observed data sequences $\{\mathbf{Y}_\alpha^\pi\}$, weighted by the likelihood of each partitioning:

$$P(\mathbf{Y}; \{p_{\mathbf{m}}\}) = \sum_{\pi} \prod_{\alpha=1}^{N(\pi)} Q(\mathbf{Y}_\alpha^\pi), \quad (1)$$

where the marginal probability is $Q(\mathbf{Y}_\alpha^\pi) = \sum_{\mathbf{m}} Q(\mathbf{Y}_\alpha^\pi | \mathbf{m}) p_{\mathbf{m}}$ and $N(\pi)$ is the total number of templates in partitioning π . We show (SI Appendix) that the MLE for $p_{\mathbf{m}}$ satisfies the implicit equation

$$p_{\mathbf{m}}^* \propto \sum_{\pi} \sum_{\alpha'=1}^{N(\pi)} p(\mathbf{m} | \mathbf{Y}_{\alpha'}^\pi) \prod_{\alpha=1}^{N(\pi)} Q(\mathbf{Y}_\alpha^\pi), \quad (2)$$

where $p(\mathbf{m} | \mathbf{Y}_\alpha^\pi)$ is the posterior probability of \mathbf{m} given the data and the pre-factor is determined from normalization. The sum over the posterior probabilities can be interpreted as an effective number of counts of \mathbf{m} in the partitioning π ; (2) can then be re-cast as $p_{\mathbf{m}}^* = \langle N_{\mathbf{m}} \rangle / \bar{N}$, where $\langle N_{\mathbf{m}} \rangle$ is the expected number of counts of \mathbf{m} over the ensemble of partitionings and $\bar{N} = \sum_{\mathbf{m}'} \langle N_{\mathbf{m}'} \rangle$ is the average number of partitions.

Given the implicit dependence of $p_{\mathbf{m}}^*$ within the large sum in (2), it is rather surprising that the MLE can be performed efficiently. To compute $p_{\mathbf{m}}^*$, it is useful to define the free energy, $F \equiv -\ln P(\mathbf{Y}; \{p_{\mathbf{m}}\})$, which is to be minimized. The gradients of F can be efficiently calculated using dynamic programming methods (SI Appendix), which allows for computation of $p_{\mathbf{m}}^*$ using standard gradient descent methods. Note that the number of counts is then $\langle N_{\mathbf{m}} \rangle = -p_{\mathbf{m}} \partial_{\mathbf{m}} F$. New motifs are added to the dictionary if they occur more often than expected by random concatenations of motifs already in the dictionary. The probability of a new motif \mathbf{m} being generated through all possible concatenations of smaller motifs in the dictionary, $\zeta(\mathbf{m})$, is compared to the empirical probability of \mathbf{m} , $-\zeta(\mathbf{m}) \partial_{\mathbf{m}} F / \bar{N}$. A standard likelihood ratio test yields a p -value and pairs below a p threshold (10^{-3}) are added to the dictionary. The results are not sensitive to this threshold.

The output of the algorithm is a dictionary of motifs and the number of times each motif occurs in the dataset, $\langle N_m \rangle$. Standard statistical tests on $\langle N_m \rangle$ can then be used in order to compare the relative abundance or rarity of a motif across datasets (for example, a control and a treatment dataset). Further, a Viterbi-like algorithm can be used to partition the data into the most likely partitioning (SI Appendix). An implementation of BASS is publicly available [36].

B. An illustration on synthetic data

To illustrate the generative process and the effectiveness of the method in identifying and segmenting motifs, we first apply it to a synthetically generated dataset. We assume individual data points are two-dimensional (representing a lower-dimensional embedding of postural dynamics) and are drawn from 7 distinct states (which make up the characters in our alphabet) with a Gaussian emission function as shown in Figure 2a. A dictionary of 50 motifs is constructed such that each motif has a mean length of five. Given the generated dictionary, the probability of each motif, p_m , is drawn and scaled with a parameter $1 - \epsilon_b$, where ϵ_b is the fraction of the dataset that is made up of individual characters. We use ϵ_b as a measure of ‘background noise’. Sequential data is sampled according to the lexical model, with ϵ_p as a measure of action pattern noise and locomotor episode noise μ , defined as the distance between neighboring clusters relative to the standard deviation of each cluster (Figure 2a). In the sample shown, we use $L = 40000$, $\epsilon_p = 0$, $\epsilon_b = 0.5$, $\mu = 3$.

On this dataset, the algorithm builds a dictionary containing 44 motifs with 11 false negatives and 6 false positives. Of the 11 false negatives (crosses in Figure 2b), 8 occur fewer than 25 times in the entire dataset. The three other false negatives (632, 631, 421325 in the inferred dictionary, see Figure 2a for cluster labels) were in fact closely related to the three most prominent motifs that were not identified (32,31,421335 in the true dictionary). The estimated probabilities of the true positive motifs match very well with their true probabilities (Figure 2b) despite significant background and locomotor episode noise. As a measure of performance, we compute the percentage of the sequence correctly segmented (i.e., both the state and the partition between motifs are correctly identified) by BASS in the most likely partitioning of the sequence. Figure 2c shows near-optimal performance for larger datasets, where optimal is defined with respect to the case when the true dictionary

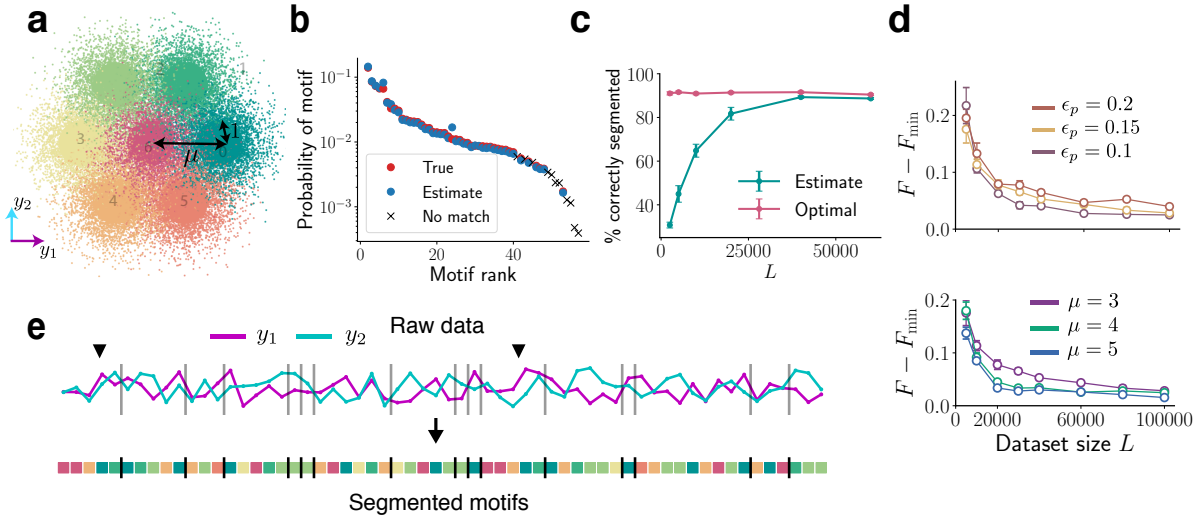


FIG. 2. BASS accurately identifies and segments motifs in noisy, synthetic data: (a) The seven clusters from which the two-dimensional data (along y_1, y_2) is drawn. (b) The true probabilities of the motifs (red dots) and probabilities estimated (blue dots) by our algorithm showing successful reconstruction of the dictionary. The crosses are low-probability motifs not identified by the algorithm (see main text). (c) The percentage of correct segmentations into motifs (cyan) with increasing dataset size. The optimal percentage when the true dictionary is known is shown in pink. In panels b,c, we use $L = 40000$, $\epsilon_p = 0$, $\epsilon_b = 0.5$, $\mu = 3$. (d) The difference in the negative log-likelihood per symbol after convergence when the true dictionary is unknown (F) and known (F_{\min}). Action pattern noise ϵ_p and locomotor episode noise μ are successfully integrated out with larger datasets. Top: $\mu = 3$, $p_d = 0.5$, Bottom: $\epsilon_p = 0.15$, $p_d = 0.5$, where p_d is the probability of an insertion in a motif instantiation. Error bars are s.e.m. (e) A snippet of the raw data sequence and the most likely partitioning into motifs found by the algorithm. The vertical bars delineate two successive motifs. The black arrows mark two instantiations of the same length-five motif.

is known. A snippet of the raw data is shown in Figure 2e along with the most likely partitioning into motifs from the learned dictionary. With larger datasets, the method robustly integrates out fluctuations due to significant action pattern and locomotor episode noise ϵ_p and μ (Figure 2d). BASS found no motifs in shuffled data.

We now apply BASS to larval zebrafish behavior in exploratory (pH neutral) and aversive (acidic) chemotaxis assays. An outline of our analysis pipeline is shown in Figure 3a.




FIG. 3. Analysis of larval zebrafish behavior exploring neutral and aversive environments. (a) Overview of the analysis pipeline. (b) A time series of the tail angle θ shows the discrete nature of locomotor episodes ('bouts') with the corresponding speed, change in heading and the summed tail angle calculated as the summed absolute amplitude of the tail angle. (c) Samples of the seven bout types identified using a Gaussian Mixture Model. In the superimposed images corresponding to a given bout type, the green and red dots correspond to the head position at the bout beginning and end respectively. Below each sample, the average tail angle θ is shown in solid color with 200 trajectories shown in grey.

C. A dictionary of action sequences for exploring larval zebrafish

Zebrafish larvae swim in short punctuated locomotor episodes called 'bouts' (duration $\text{mean} \pm \text{s.d} = 150 \pm 50$ ms) separated by longer periods of rest ($\text{mean} \pm \text{s.d} = 700 \pm 400$ ms) (Figure 3b). Larvae spontaneously explore their environment by performing mainly slow bouts occurring as forward swims and routine turns, often by repeating turns in the same direction [37], and rarely exhibit fast bouts such as burst swims or escapes [16, 33]. We collected a dataset of ≈ 85000 bouts from 171 fish swimming in elongated swim arenas.

A single bout is well-characterized by the fish's tail movement and other kinematic variables such as average speed and change in heading. From raw tracking data (SI Appendix), we use a six-dimensional parameterization \mathbf{y} for each bout, which includes the speed, the change in heading, the summed tail angle (summed absolute amplitude of the tail angle) and the first three principal components of the tail angle over time (SI Appendix). Based on this parameterization, bouts were categorized into different bout types using a Gaussian Mixture

Model (GMM). A GMM yields the probability density, $q(\mathbf{y}|c)$, for each bout type c , which serves as a statistical description of each bout type in terms of the means and covariances of the six variables. We clustered bouts into seven bout types (Figure 3c, SI Movie S1,S2), which correspond to two forward swims of different speeds (f , slow and F , fast), three turns based on the magnitude of change in heading (t , T and L , increasing angle), bursts (b) and an other (O) category. The O category contained a variety of different bouts that did not clearly fall into one class; these included O-bends, long turns and bursts, and improperly tracked bouts. The bout types are not sharply delineated; this is not an issue for the BASS algorithm since variability in \mathbf{y} is implicitly taken into account via $q(\mathbf{y}|c)$ as noted before.

Typical bout types are displayed in Figure 3c. Compared to previous categorizations performed on spontaneous exploration [16, 17, 25], our categories (except O) likely correspond to sub-divisions of forward swims, routine turns and burst swims. The f , t , T bout types typically correspond to the slow regime of locomotion enriched during basic exploration, while F , b and the rare O belong to the fast regime and occur overall less frequently when no aversive stimulus is applied [16, 33].

Zebrafish locomotor episodes in our conditions consist of seven bout types that make up the alphabet of our generative model. Sequences of consecutive bouts for each fish (≈ 500 bouts per fish) served as input to BASS. A coarse exploration of the pattern noise parameter ϵ_p and the probability of an insertion p_d using a held-out dataset yielded $\epsilon_p = 0.1$ and $p_d = 0.2$, which were used for the rest of our analysis (Figure S4). These numbers suggest noisy motif instantiation and a bias towards insertions (i.e., repeats). Notably, the greater held-out likelihood for $\epsilon_p > 0$ highlights the advantage of incorporating variability into our modeling framework.

The algorithm converged to a dictionary consisting of 66 motifs with similar results across trials and subsamples. The output of the algorithm is a dictionary consisting of the identified motifs and the expected number of occurrences of the motif, $\langle N_{\mathbf{m}} \rangle$, in the dataset. Intuitively, $\langle N_{\mathbf{m}} \rangle$ is the number of times a motif occurs after appropriately discounting its occurrences within a longer motif and taking into account pattern noise and locomotor episode noise. For example, a locomotor episode which is on the boundary between the f and t clusters (locomotor episode noise) is appropriately weighted as one half f and one half t . This example further highlights the importance of incorporating locomotor episode noise as a source of variability. Similarly, the observed sequence $ffff$ will also contribute to the

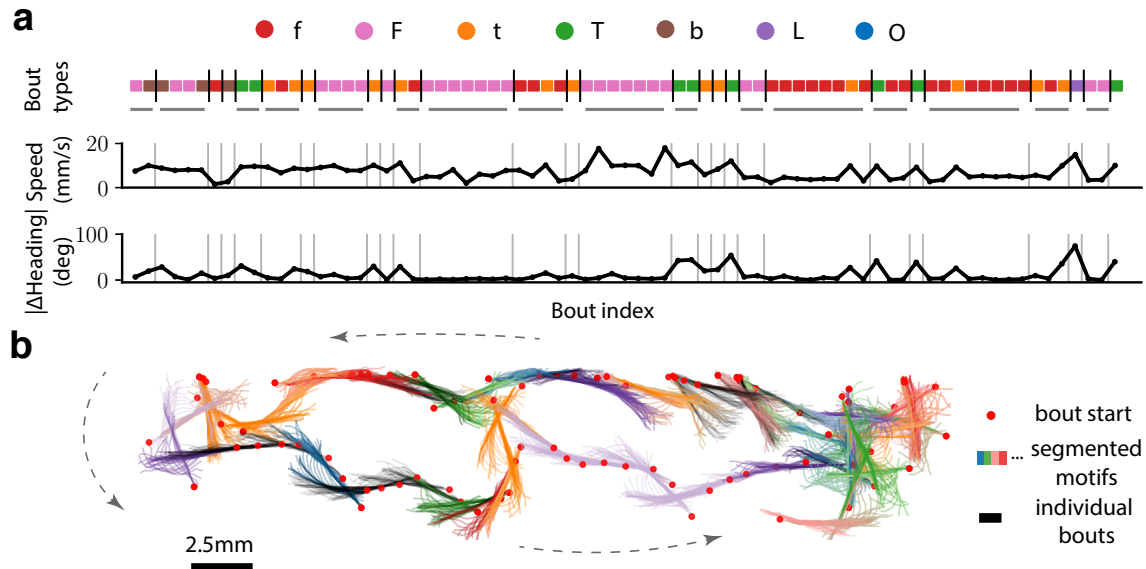


FIG. 4. Motifs identified by BASS make up a significant fraction of the dataset. (a) A sample sequence of 75 bouts from the exploratory data segmented (separated by vertical bars) into the most likely sequence of motifs from the learned dictionary. The corresponding speed and absolute change in heading are shown. Motifs longer than one locomotor episode are underlined in gray. (b) A sample trajectory consisting of 80 bouts (head position at the beginning of a bout is shown as a red dot) are segmented into motifs (head and tail at each frame are shown), where successive bouts from the same motif have the same color. The black-colored segments of the trajectory are motifs of length one i.e., single locomotor episodes.

number of counts of the motif fff since the extra f could be due to an insertion (pattern noise). The observed sequence $ffff$ will not contribute as one count for the motif $ffff$ and also contribute as two counts for the motif ff . The counts are instead shared between the motifs $ffff$ and ff based on the relative probability of the two motifs (implicitly using Bayes' rule) and the number of ways ff can appear in the observed sequence $ffff$. A subset of these motifs is shown in Table I (see also Table S1). In Figure 4a,b, we provide a typical sample sequence of bouts segmented into a sequence of motifs.

The dictionary reveals several surprising features. A significant fraction of the fish movements were made of motifs: motifs covered on average 78% of all bouts per fish with a standard deviation of 7.3% across fish. We found motifs as large as 14 bouts, and this could further expand in a particular realization due to insertions. In particular, f repeated 14

times occurred more than 500 times. While this may be explained by the large fraction of f , repeats were also found for T , F and b . Overall, the most enriched and common motifs correspond to repetitions of the same bout type, and typically occur 2-14 times in a row. Motifs containing mixtures of bouts included typically only 2 different bout types. Remarkably, we noticed throughout the list of enriched motifs that the two bout types belonging to a given motif, belonged either to low speed (mixtures of either f and t in fft or in $TfTf$) or to high speed (mixtures of F and b in $FbFb$ or $bbFb$), but never combined fast and slow locomotor episodes.

To quantify how unusual these sequences were under a Markov model, we compared the observed occurrence of the identified motifs to those predicted from the best-fit Hidden Markov Model (HMM). Our lexical model yielded a better fit compared to an HMM (difference in held-out free energy per bout of 0.12), and a significant portion of motifs deviated from Markovianity (Tables I,S1). Two aspects of the behavior likely lead to the observed non-Markovianity. First, while long repeats of the same bout type occur often, the distribution of the number of repeats has a heavy tail and decays much slower than a geometric distribution (Tables I,S1). Second, sequences with mixtures of two bout types such as $TfTf$ and fft are common; while the repeats emphasize (say) $f \rightarrow f$ transitions, the motifs with mixtures of bout types on the other hand emphasize $f \rightarrow t$ transitions, creating a tension between the two in a purely Markovian picture. One concern could be that the t bouts in the mixtures of f and t (e.g. fft) lie on the border between the f and t clusters, and that the apparent non-Markovianity is simply due to excessive coarse-graining of the elementary locomotor episodes. However, this issue is solved by our soft clustering procedure, which assigns the appropriate probability weight to locomotor episodes that lie on the boundary between two clusters.

To verify that the long chain of repeats were not an artifact due to our elongated well geometry, we applied a similar pipeline of bout categorization and motif identification on a previously published dataset [16] (see SI Appendix, Figure S3). The dataset consists of $\approx 120,000$ bouts obtained from 23 fish freely swimming in a square well (of side ~ 25 mm) under varying light intensities. Notably, the resulting dictionary also displays long chains of repeats and significant non-Markovianity albeit with a heavier emphasis on turns compared to forward swims (Table S2). We observe again in the motifs a dichotomy between the classical slow and fast locomotor episodes when bouts occur either via forward swims or

TABLE I. Motifs over-represented in the exploratory dataset. A subset of motifs occur (‘Observed’ column) more often than predicted by a first-order Markov model (the ‘Expected’ column). The p -value is obtained using a likelihood ratio test. Single-length motifs are not shown (see Figure 5c). See also Table S1. The rightmost column shows the percentage of unique fish (out of 171 total) which executed the motif at least once in the most likely partitioning of the data. Note that a motif may never appear in the most likely partitioning even though it has non-zero p_m (for eg. *tttt* below, which is partitioned into four individual *ts*)

Motifs	$-\log_{10} p$	Observed	Expected	% of fish
ffffffff	>300	1366	387	51% (87/171)
ffffffffffff	>300	510	50	38% (65/171)
ffffff	208.01	3234	1797	69% (118/171)
fff	42.33	9544	8327	93% (159/171)
FFFFFFF	28.23	311	153	53% (91/171)
fffftf	27.64	497	290	70% (119/171)
ftffff	25.07	495	297	66% (113/171)
ftfff	22.72	1125	824	76% (130/171)
ftff	21.12	1745	1377	84% (144/171)
ftf	18.5	2724	2289	57% (98/171)
ftf	13.96	4337	3859	92% (158/171)
TfT	11.12	722	554	68% (116/171)
FFFF	7.94	1428	1224	67% (115/171)
TfTf	7.28	346	254	58% (100/171)
tttt	6.7	256	181	0% (0/171)
TTTT	5.06	160	110	54% (92/171)
bb	3.87	924	1044	30% (52/171)
bbbb	3.21	115	82	21% (36/171)
FbFb	2.19	99	74	19% (32/171)

turns. Mixtures are composed either of slow or of fast locomotor episodes: slow turns and forward swims appear together in a set of sequences (*ffTf*, *TfTf*) while fast forward and burst

swims ($FbFb$, $bbFb$) appear together in other sequences. Remarkably, mixtures of *slow* bouts (forward swims or turns) and *fast* bouts (forward or burst swims) are conspicuously absent in both dictionaries.

D. Fish chemotax away from acidic pH using conserved sequences of fast bursts and large avoidance turns

Although few studies have reported that zebrafish respond to acute applications of chemicals in the surrounding water [38–40], the behavioral responses of freely-swimming zebrafish larvae navigating in chemical gradients of aversive or appetitive cues have not yet been investigated. We applied acid to the two ends of our extended arenas (of length 14 cm) forming a sharp gradient (SI Appendix); diffusive transport at the time scale of the experiment (ten minutes) is at most 1 cm and therefore is confined to the ends. Zebrafish larvae robustly performed chemotaxis and avoided the two extremities (Figure 5a) despite displaying only minor differences in kinematic parameters (Figure 5b). The distribution of bout types was similar for larval zebrafish navigating in acidic gradients. We noticed a small over-representation of certain bout types (Figure 5c), suggesting that fish perform more fast forward swims F , burst swims b , and large-angle avoidance turns O in response to the aversive gradients. Burst b and fast forward swims F occurred uniformly around the center of the well (Figure S2), while larval zebrafish executed more large-angle avoidance O bouts closer to the acidic gradient, suggesting that O bouts are potentially implicated in the response to noxious stimuli during aversive chemotaxis.

We next investigated whether the chemotaxis response extended beyond a single bout to action sequences composed of a sequence of bouts that were consistently observed across fish. We implemented a comparative approach of finding conserved sequences of actions that are highly over-represented in the aversive environment compared to exploration. We applied the BASS algorithm to the dataset from fish in the aversive environment ($\sim 66,000$ bouts from 135 fish). The resulting dictionary of motifs contained a total of 81 motifs, slightly larger than the one obtained from exploration (Table S3). The two dictionaries contain broad similarities: both contain long repeats of the same bout type and mixtures of t and f , yet contain important differences, particularly in the over-representation of mixtures of b, F and O, b in the aversive environment.

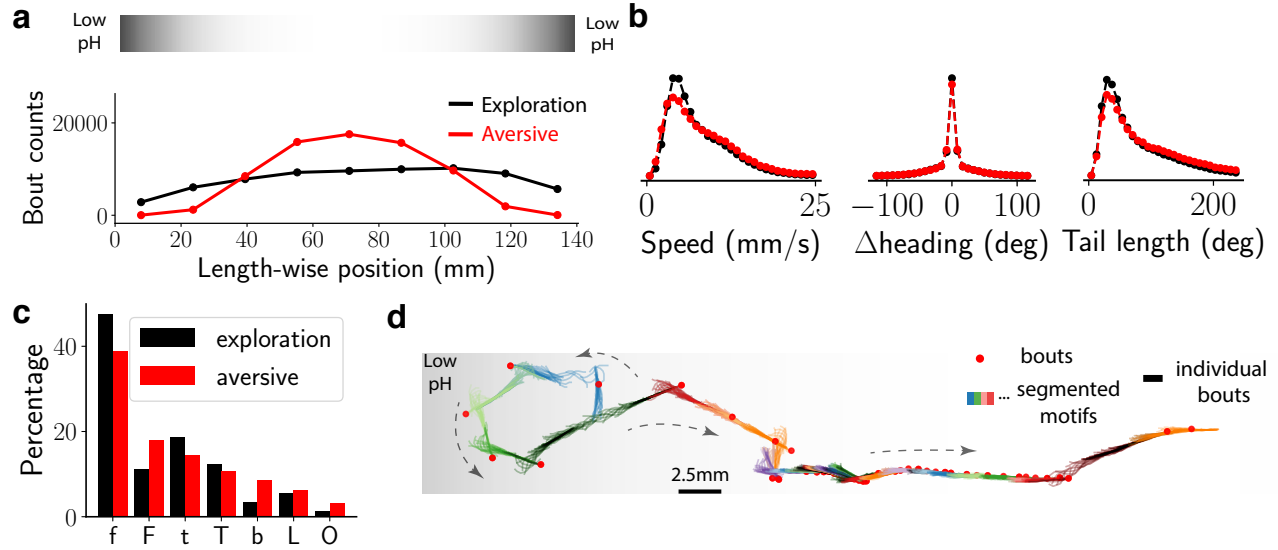


FIG. 5. Fish avoid an aversive environment using a transient chemotactic response. (a) Histogram of larvae positions along the well with and without the aversive (acidic) gradient, located at the ends of the well. The s.e.m is square-root of the counts, which is negligible. An illustration of the aversive gradient is shown above. (b) The distribution of speed, change in heading and the summed tail angle of all bouts during exploration (black) and in aversive environment (red). The difference in global kinematic parameters between the two environments is small. (c) The fraction of each bout type in exploratory and aversive environments, where a total of ≈ 85000 and ≈ 66000 bouts were collected respectively, shows an increase in fast bouts : b, F and O . (d) Localisation of bout types along the well with and without the aversive (acidic) gradient, (e) BASS segments a sequence of bouts from the aversive environment into a sequence of motifs. Shown here is a sample trajectory (as in Figure 4b) where the fish escapes from the aversive environment.

We examined over-represented sequences by comparing the relative occurrences of motifs in the exploratory and aversive environments. The two dictionaries were combined to obtain a total of 103 unique motifs and the expected number of occurrences of each motif, $\langle N_m \rangle$, for the two environments. We use $-\log_{10} p$ as a measure of over-representation, where p is obtained from a likelihood ratio test on $\langle N_m \rangle$ in the two environments. To calibrate this score, we first split the exploratory dataset into halves and computed the $-\log_{10} p$ for each motif; the threshold 15 was chosen for a false positive rate of 10%. To find conserved sequences that were executed by many fish and not just a few abnormal fish, we sub-sampled our dataset from the aversive environment to 80% its size ten times, performed the

TABLE II. Motifs consistently over-represented in the aversive chemotaxis assay. The rightmost column shows the percentage of unique fish (out of 135 total) which executed the motif at least once.

Motifs	$-\log_{10} p$	$\langle N_{\mathbf{m}} \rangle_{\text{aver}}$	$\langle N_{\mathbf{m}} \rangle_{\text{explo}}$	% of fish
fTff	90.84	101	5	64% (87/135)
OO	66.63	111	11	60% (81/135)
bb	55.84	210	55	82% (82/135)
bOOb	46.06	63	5	46% (46/135)
Ob	35.74	135	36	77% (77/135)
bO	34.6	140	39	82% (82/135)
bFbb	24.96	49	6	49% (49/135)
Obbb	24.07	56	9	49% (49/135)

comparison for each sub-sample, and chose only those motifs above threshold of 15 in *all* ten sub-samplings.

Table II shows the eight over-represented motifs conserved across fish. The eight motifs were executed by multiple fish (rightmost column of the table). All motifs except one (*fTff*) are mixtures of *b* and *O*. The bouts from the selected motifs (except *fTff*) are significantly over-represented close to the aversive gradient compared to the rest of the bouts in the aversive environment (Figure 6a), though no such selection was explicitly imposed *a priori* in our analysis. To verify that these motifs were indeed the ones involved in aversive chemotaxis, we computed the distance per bout the fish travels in the direction down the aversive gradient for bouts within motifs flagged as over-represented and the rest of the bouts. We find a highly significant bias for the over-represented bouts for swimming down the gradient (Figure 6b). Strikingly, when bouts from the flagged sequences that begin at the two extreme quarters of the well excluding *fTff* are considered, the bias greatly increases, even beyond the typical length-wise distance travelled per bout, indicating that these bouts are triggered by the sharp acidic gradient and are almost exclusively aimed down the gradient. In contrast, the motif *fTff*, which was also reliably observed across fish, was not directed away from the aversive gradients nor did it preferentially occur close to the gradient, suggesting that *fTff* may be induced by a global acidification of the swim arena but not as a




FIG. 6. Fish chemotaxis against an acidic gradient via sequences of fast bursts and large-angle avoidance turns. (a) Distributions of length-wise positions for the Control bouts (all bouts from aversive environment except the ones from sequences flagged as over-represented in Table II), *fTff* only and Flagged bouts (from the sequences in Table II) except *fTff* in red, orange and blue respectively. (b) The length-wise displacement travelled in a bout down the gradient for bouts tagged as Control, Flagged (as defined in (a)), *fTff* only, Control(ends) (all bouts from the two ends of well shown in (a)), Flagged(ends) (flagged but with *fTff* removed and in the ends of the well). For scale, the red, dashed line shows the mean length-wise distance per bout for unflagged bouts. Error bars are s.e.m. (c) The mean speed, change in heading and duration of the bouts (black, red error bars for s.d, s.e.m respectively) from *b* and *O* bout types that are part of the Flagged(ends) sequences from (b). (d) Superimposed images for four random samples of *b* and *O* bout types from the flagged sequences in (a). The green and red are the head positions at the beginning and end of the bout respectively. (e) Superimposed images for four random samples highlighting the Flagged(ends) sequences (blue dots), which include the three bouts before (green dots) and after (red dots) the flagged sequence. Note that the depicted gradient is illustrative.

direct response to the acidic gradient (see SI Movie S3 for a sample of the motif *fTff*).

Further examination of the bouts from the *b* and *O* bout types implicated in chemotaxis showed that both bout types were significantly longer and the fish swam faster compared to the unflagged bout types (Figure 6c, Figure S1). Inspection of the tail movements from this subset of *O* bouts showed that these bouts typically consisted of a large-angle avoidance turn followed by a long burst swim (see Figure 6d for examples). The induced sequences of

fast bouts composed of b and O near the edges of gradient are distinct from the recently described slow avoidance response to CO_2 [38]. The previous report of a lack of behavioral response to HCl $\text{pH} = 4.5$ ([38]) suggests that b and O may be elicited by more acidic pH .

E. Spiraling motifs in soaring trajectories

We note that BASS receives no information about the stimulus experienced by the animal. To confirm that BASS can extract relevant aspects of behavior even when the cues in the environment are unknown and the stimulus-response map is complex, we turn to simulations of a soaring glider using thermals (updrafts) to gain height in a convective environment [21, 22]. The purpose of this exercise is to present an example of how BASS could be applied to naturalistic behavioral data obtained in the field. We used a dataset of size within experimental reach and input variables that can be readily measured using existing instrumentation.

We simulated trajectories of a soaring glider executing a thermalling strategy in an environment that contains updrafts and downdrafts of varying strengths. The thermalling strategy was learned using reinforcement learning in the field [22]. The glider takes mechanical cues (i.e., torques and accelerations) as input and responds by modulating its bank angle. Past simulations in a turbulent convective environment have shown that a glider executing this thermalling strategy successfully gains height using thermals and exhibits spiraling trajectories similar to those displayed by birds [21]. A sample trajectory is shown in Figure 7a, which shows the glider’s meandering path before finding the updraft and the spiraling trajectory while ascending the thermal.

We simulated 75 episodes each (250 seconds per episode) of the glider executing the thermalling strategy and the glider executing a random policy (i.e., it increases, decreases or keeps the same bank angle with equal probability). Units are defined using realistic parameters for glider aerodynamics; for instance, the glider has an airspeed $\sim 10\text{m/s}$ and takes actions every 1.5 seconds (see refs [21, 22] for full details of the simulation). BASS receives as input solely the bank angle at every 1.5 seconds, which is divided into five states corresponding to bank angles of -30° , -15° , 0° , 15° , 30° .

BASS converged to dictionaries with 80 and 65 motifs for the thermalling and random policy datasets respectively ($\epsilon_p = 0.1, p_d = 0.5$). In Figure 7b,c we show examples of the

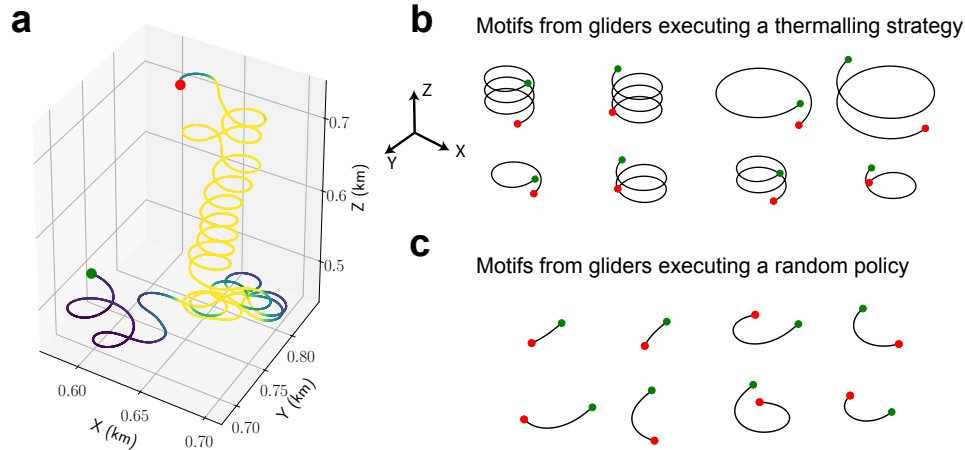


FIG. 7. BASS finds spiralling motifs in simulated trajectories of a soaring glider executing a thermalling strategy. (a) A sample trajectory showing the thermalling behavior of the soaring glider. The glider reacts to mechanical cues such as torques and accelerations. See [21, 22] for more details. Regions of downdraft and updraft (blue through yellow) are marked, which highlight the relatively disperse trajectories in downdrafts compared to the spiraling behavior in updrafts. (b) Motifs found by BASS that span the largest fraction of the thermalling dataset (i.e., episodes where the glider used a thermalling strategy). Most enriched motifs are those of a spiraling pattern. The green and red dots mark the start and end points. Note that the glider sinks without updrafts. (c) Same as (b) for motifs found in the dataset of a glider executing a random policy.

eight motifs that spanned that largest fraction in the two dataset (i.e., the motifs were sorted according to the number of times they occur in the dataset multiplied by the length of the motif). BASS indeed finds a consistent spiraling pattern in the thermalling dataset, with the longest two motifs of opposite helicity (first two in the top row of Figure 7b) lasting 24 actions, i.e., 36 seconds. Full circles with a lower bank angle and spiraling patterns of shorter length were also found. The motifs from the random policy dataset were roughly half as long (mean length ≈ 8 s) on average compared to the motifs from the thermalling dataset (mean length ≈ 16 s) and largely consist of short turns as shown in Figure 7c.

III. DISCUSSION

In this study, we develop a method to extract conserved sequences that rarely occur in noisy behavioral data from freely-moving animals. We present a lexical model of animal

behavior, where we view observed behavior as a composition of recurring motif templates drawn from a dictionary. We develop the BASS algorithm for performing inference on this model, which ultimately yields a dictionary of motifs that the animal performs in its designated environment. Applying the method on data from exploring zebrafish larvae revealed a long time-scale organization of bout sequences that cannot be explained in a Markovian model on single bouts. In a novel aversive chemotaxis task, we identified conserved sequences of bouts that the fish employ to avoid a noxious environment. When applied to simulated thermalling dataset, BASS succeeded in identifying the long spiraling patterns characteristic to soaring.

BASS is a useful tool to identify behavioral action sequences that are highly conserved despite being rare and transient in naturalistic settings when animals are freely moving. To summarize, BASS can be optimally used: 1) when the sensory information perceived by the animal is unknown / uncontrolled, 2) when the behavioral responses is complex and last much longer than a typical movement, 3) when the behavioral response consist of highly-conserved sequences that occur rarely in the behavioral recording.

We argue that our lexical model yields complementary insight to traditional dynamical models, and is better suited for comparative behavioral analyses, particularly for comparisons across animals in different environments and/or closely related genetic variants. The generative model of independent motifs is the simplest one and fits better to our behavioral data than a Markov model with a similar number of parameters, which are often used to depict quantitative ethograms.

Our model generalizes previous work on motif discovery from bioinformatics, machine learning and time series analysis by incorporating two important generative processes crucial for behavioral modeling that were not modeled previously: elevating motif templates to latent variables and introducing a data-generating process for elementary locomotor episodes. This generalization is necessary to take into account the significant variability in behavioral data, where clusters in postural space are not always well-defined and erratic movements are common. Note that the generative process for *motif templates* can also be viewed here as a particular hidden Markov model, where the $|\mathcal{D}| + 1$ hidden states at the topmost level are the motifs in the dictionary and the ‘background’. In this picture, the full generative model has three levels of hierarchy. Several extensions of the model are possible. Prior knowledge can be easily applied. For instance, priors on the distributions of motif lengths

or the distribution of frequencies can be introduced by weighting different partitions or a Dirichlet prior on the probabilities, respectively. In both cases, an MLE equation (or a MAP estimate in the latter case) similar to (2) can be derived. More complex, hierarchical models over motifs may be learned by noticing that the Markovian structure of the partitioning is compatible with structured variational approximations [41]. Further structure can be introduced into the ‘background’, which we have assumed is made of independently drawn characters, similar to those used in bioinformatics [42]. Compression methods [43–45] optimize an altogether different “coding” objective, which do not necessarily lead to meaningful motifs; for example, the two-symbol word ab could be identified as a motif simply because a and b occur often, even if a and b occur next to each other purely by chance.

In certain cases, behavior is better described as continuous motor output in contrast to the discrete bouts that comprise zebrafish locomotion (e.g the eigenworms that describe *C.elegans* dynamics [35]). In such cases, locally linear approximations similar to those in ref. [18] can be used. The approach there is to fit short windows of dynamics using linear models and then use hierarchical clustering to cluster models in model space into discrete states. A similar approach can be used to generalize BASS to continuous motor output. In particular, we may define the elementary locomotor episodes (via $q(\mathbf{y}|c)$) as a set of short continuous snippets (such as lines and curves) equipped with an appropriate probability measure to accommodate imperfect fits. Further, motifs from the extracted dictionaries can be hierarchically clustered to define broader behavioral classes of motifs. These generalizations will be explored in future work.

It is important to emphasize that BASS receives no explicit information about the stimulus experienced by the animal; the extracted motifs therefore do not contain information about the precise stimulus-response map of the animal, but can reveal relevant qualitative aspects of the animal’s behavior when the sensory information is not known. As illustrated here with our simulation, BASS may discover the spiraling of a soaring bird [46] with no reference to what stimulus triggers those responses. More generally, BASS is applicable when characteristic patterns of behavior are repeated on several occasions during a task. However, it is possible that a single stimulus-response map may yield a highly diverse set of outcomes based on diversity in input and memory of past inputs. In this case, inference is difficult for any algorithm since little statistical information is available without access to the input stream and other methods should be devised.

We show here that freely-swimming zebrafish larvae navigating in a gradient of aversive cues exhibit conserved sequences of fast bouts mixing burst and large-angle avoidance turns to avoid the aversive environment. Notably, these sequences make up only a very small fraction of the dataset ($\sim 0.2\%$), yet are successfully captured in our analysis.

Furthermore, we also discover motifs deployed by zebrafish larvae when exploring either rectangular arenas (this study) or confined square arenas [16]. We find that larval zebrafish repeat bout types belonging to the same speed regime (slow or fast). Long repeats of the same bout type, particularly forward and turn swims, occur frequently for ~ 3 -8 iterations. This observation is consistent with prior observations of repetition of turns in freely-swimming larvae [37]. In the baseline conditions when the environment contains no aversive cue, larvae mainly display motifs consisting of slow, forward bouts and turns. In contrast, when larval zebrafish are exposed to noxious stimuli, such as the acidic gradients that activate cutaneous sensory afferents, they respond using transient and conserved sequences of fast forward and/or turns in order to swim away. While an automated categorization of bouts can identify that large-angle avoidance turns O often occur in the acidic gradients (Figure S2), our method reveals a richer chemotactic response composed of fast locomotor episodes including burst swims and large angle turns. Remarkably, BASS led to the discovery of a very interesting phenomenon that is conserved across datasets from multiple labs operating with different zebrafish strains and different arenas. We show that larval zebrafish perform highly-specific sequences containing either only slow (during basic exploration) or only fast (during aversive chemotaxis) locomotor episodes. This observation suggests that the descending command signals sent to the spinal cord to elicit forward bouts [47] or turns [48] last over seconds to tens of seconds, possibly via sustained inputs to reticulospinal neurons in the hindbrain [47]. The dichotomy in terms of speed regime observed within conserved sequences of locomotor episodes may have important implications for the construction of neuronal networks sustaining foraging and avoidance.

BASS can be easily incorporated into existing behavioral analyses pipelines alongside the expanding repertoire of methods established for unsupervised behavioral clustering that deal with classifying individual locomotor episodes [1–5]. In its current form, our implementation can handle datasets of size $\lesssim 300,000$ bouts and dictionaries of size $\lesssim 500$, beyond which approximations for scalable inference have to be developed. Our approach further highlights the connection between behavioral modeling and genomics, where a wealth of algorithms

have been developed. Exploiting this analogy may lead to a fruitful exchange of techniques between the two seemingly-disparate fields. Finally, we remark that our method is an addition to a rapidly enlarging computational toolkit for extracting mechanistic answers to behavioral questions. BASS complements existing computational methods for unsupervised behavioral clustering, typically used to segment individual locomotor episodes, by finding conserved structure in the form of sequences of locomotor episodes.

IV. METHODS AND MATERIALS

The code for BASS and the scripts used to reproduce the figures are publicly available [36]. The repository includes the six-dimensional data used for bout categorization and BASS analysis. Due to its large size, the raw tracking data has not been uploaded and is available at request.

ACKNOWLEDGMENTS

We thank Monica Dicu and Antoine Arneau for fish care in the Phenoparc animal core facility of ICM, Maxime Kermarquer for use of the ICM cluster, Kristen D’Elia, Sophia Horowitz and Clara Besserer for trouble-shooting at the early stages of the chemotaxis experiments, Bing Brunton, Elena Rivas and Massimo Vergassola for useful comments, and Joao Marques and Michael Orger for sharing their zebrafish larvae dataset. This research was initiated during the summer school “Neural computation for sensory navigation” held in 2018 in the Kavli Institute of Theoretical Physics (University of California in Santa Barbara, USA) and was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958, NIH Grant No. R25GM067110, and the Gordon and Betty Moore Foundation Grant No. 2919.01. This work was also supported by a New York Stem Cell Foundation (NYSCF) Robertson Award 2016 Grant #NYSCF-R-NI39, the HFSP Program Grants #RGP0063/2018, the Fondation Schlumberger pour l’Education et la Recherche (FSER/2017) for C.W. The research leading to these results has received funding from the program “Investissements d’avenir” ANR-10- IAIHU-06 (Big Brain Theory ICM Program) and ANR-11-INBS-0011 – NeurATRIS: Translational Research Infrastructure for Biotherapies in Neurosciences. C.W. received support from the ERC-Proof Of Concept grant ERC-

POC-2018#825273 for the development of the tracking algorithm ZebraZoom implemented by O.M. (www.zebrazoom.org). G.R. was partly supported by the NSF-Simons Center for Mathematical & Statistical Analysis of Biology at Harvard (award number #1764269) and the Harvard Quantitative Biology Initiative.

SI APPENDIX

1. Animal care and ethics statement

Animal handling and behavioral assays on 6-7 days old larvae were carried out with the validation of the Institut du Cerveau (ICM) in agreement with the French Ministry for Research and Education and the French National Ethics Committee (Comité National de Réflexion Ethique sur l'Expérimentation Animale, APAFIS #2018071217081175) following the European Communities Council Directive (2010/63/EU). All zebrafish animals were reared and maintained at +28.5°C on a 14/10 hour light/dark cycle until 7 days post fertilization (dpf) when behavioral experiments were conducted.

2. Behavioral assays and recording

Wild-type AB zebrafish were reared at the density of 20 larvae per petri dish filled with 30 mL of E3 medium. At 4 dpf, the E3 medium was replaced to avoid waste accumulation. On the day of behavioral recording, ten minutes prior to the beginning of experiments, all larvae were transferred onto the LED illumination plate to allow habituation to the lighting of the setup. At the beginning of each trial, we positioned on the illumination plate the recording chamber containing 12 swim arenas (length 140 mm x width 10 mm x height 4 mm) filled with 4 mL of E3 medium. Each individual larva was then transferred into the central region of a swim arena delimited by a custom-made comb. Once all larvae were positioned, the comb was carefully removed to enable animals to explore freely the arena and the video recording was launched.

To monitor the navigation in gradients of pH, larvae were exposed to acidic solution at both ends of the swim arenas. To create this double acidic gradient, an acidic solution was prepared by diluting 1 M HCl (Merck, #109057) 1:6 into E3 medium (final pH 1.5). 100 μ L of this solution was applied at both ends of the swim arena before removing the comb

and launching the recording. Out of 12 swim arenas, 6 were filled with the double gradient of acidic pH and 6 were filled with E3 medium only. Before any behavioral recordings, we estimated the stability of the pH gradient generated by following the same application protocol and monitoring the diffusion of red food colorant over 10 minutes. These pilot experiments indicated that in our conditions, the acidic solution slowly diffused and formed a steep gradient about 20 mm from each end of the swim arena over the 10 min-long recording.

Ten min-long videos were recorded at 160 Hz with an exposure time of 1 ms, and a pixel size of 70 μm using a ViewWorks camera (Basler acA2040-180km) controlled by the Hiris software (R&D Vision, Nogent sur Marne, <http://www.rd-vision.com/r-d-vision-eng>).

3. Tracking

To track the head and tail positions of zebrafish larvae, we used the open-source software ZebraZoom (<https://zebrazoom.org/>) [33, 49]. The algorithm begins by locating all the wells and by extracting the background of the video. ZebraZoom first applies a series of actions to detect the animal in each well: i) contours of head and entire body are detected using active contours, ii) the center of the head is identified as the center of mass of the head contour and the tip of the tail is detected using both the curvature along the body contour and distance to the center of the head. The midline is then identified between the left and right borders of the body contour. For each animal, the difference in pixel intensity between subsequent frames enables the automated detection of bout start and end. Then, for each bout, the algorithm calculates the head position, head direction and the tail angle from which kinematic parameters are subsequently estimated: number of oscillations, instantaneous tail beat frequency, maximum amplitude for each tail bend, bout speed, bout duration, and bout distance. Tunable parameters in the tracking algorithm were optimized to detect small amplitude forward bouts occurring frequently during exploration. In order to validate our algorithm, we manually inspected validation videos where the head direction and tail position were superimposed on the raw image when a bout is detected, allowing to check both the tracking and bout detection quality.

4. Bout categorization

A segmented bout contains the positions of the head and ten tail segments for each frame. For each bout, we computed the average speed, the change in heading angle and the duration of the bout. The tail angle relative to head orientation of the seven tail segments farthest away from the head were computed and the first 16 frames of these 7 quantities were concatenated to form a 112-dimensional vector. Principal Component Analysis (PCA) on this vector yielded a projection onto 3 principal components that together explained 80% of the variance. For bout categorization, the average speed, the change in heading and the summed tail angle (integrated absolute value of tail amplitude) were combined with the first 3 tail angle principal components, which together served as the six-dimensional input (denoted \mathbf{y} throughout the paper) for the clustering algorithm. Including other kinematic variables such as tail maximum amplitude, inter-bout interval, duration, number of oscillations, angular speed or displacement did not change our bout categorization and were excluded for the rest of the analysis.

For clustering of bouts, we employed a Gaussian Mixture Model (GMM) with a full covariance matrix. We modified the standard expectation-maximization algorithm to fit multiple datasets simultaneously such that the emission functions i.e., the probability of the input given cluster label, $q(\mathbf{y}|c)$, was fixed for all datasets but the cluster weights differ across datasets. To do this, we first sub-sampled the datasets so that they all have the same size. The objective function to be maximized for the GMM model was modified to be the sum of the log-likelihoods for each dataset. We found that adding a term proportional to the Jensen-Shannon divergence between the cluster weights across the two datasets encouraged the algorithm to find cluster assignments whose weights differed the most across datasets. The number of clusters was set to seven by measuring the log-likelihood on a held-out dataset. The distributions of the speed, duration and change in heading for different bout types are shown in Figure S1.

We performed a similar procedure to analyze the data from Marques et al [16]. The dataset contained approximately 120,000 bouts from 23 zebrafish larvae at 6-7 dpf imaged individually at 700 frames per second at a resolution of 62 microns per pixel with different light intensities. From the raw head position data, for each bout we extracted the displacement and change in heading in the first 120 frames (170 ms). Tail angles for five

tail segments (3rd furthest away from the head to the 2nd closest) sub-sampled at half the frame rate were concatenated together to form a 300-dimensional vector. PCA yielded four principal components, which were combined with the displacement and change in heading to form the input to the GMM (Figure S3). More clusters simply yielded finer divisions of the clusters. In their paper, Marques et al use a density-based clustering algorithm to identify 13 different bout types. In this dataset, most bouts corresponded to AS, Slow 1, Slow 2, Routine Turns with a small but significant fraction of SATs and HATs (which together formed the first four bout types f , F , t and T) and a much smaller fraction of other bout types. Most of the rare bout types were not identified in our clustering and were part of the other (O) bout type.

5. Modeling

In this Section, we expand on the model presented in the main text. We first develop the model followed by the method to build the dictionary.

a. Generative model

Given a time series of length L , $\mathbf{Y} = \mathbf{y}_1\mathbf{y}_2\dots\mathbf{y}_L$, a clustering algorithm with K clusters (labeled from 1 to K) applied on each \mathbf{y}_i converts the dataset \mathbf{Y} into a sequence of probability vectors, $\mathbf{R} = \boldsymbol{\rho}_1\boldsymbol{\rho}_2\dots\boldsymbol{\rho}_L$, where each $\boldsymbol{\rho}_i$ is a K -dimensional probability vector whose components sum to one. To fix notation, the j th component of this vector, ρ_{ij} , is given by $\rho_{ij} \equiv q(\mathbf{y}_i|j)$, where $q(\cdot|j)$ is the probability density function associated to cluster label j . We use cluster label and characters of an alphabet interchangeably. Each character represents a well-defined locomotor episode; in the case of larval zebrafish, this is the category of a single bout. In the case of the soaring glider, the glider takes discrete actions of changing the bank angle every 1.5 seconds. For each action, the glider changes its bank angle linearly over 1.5 seconds for a total change of 0° or $\pm 15^\circ$ starting from one of five angles: $-30^\circ, -15^\circ, 0^\circ, 15^\circ, 30^\circ$. The bank angle at the beginning of the action is taken as the postural state. The appropriate choice of the timescale used to define a postural state is organism-dependent. We assume the timescale and the clustering scheme have been chosen appropriately based on prior knowledge about the organism. We do not address the details

of how to find an appropriate clustering scheme for each scenario; we refer to reviews on the topic [2, 4, 5].

We define a motif $\mathbf{m} = c_1c_2 \dots c_l$ as a string of characters (or cluster labels) of arbitrary length l (l is not a parameter), which belongs to a dictionary \mathcal{D} that contains all the motifs. Each motif in the dictionary and each character in the alphabet are assigned probabilities $p_{\mathbf{m}}$ which are to be inferred from data. Henceforth, we will use ‘templates’ to mean both motifs and characters. We assume the data is generated as follows: First, templates are generated sequentially, drawn independently and identically from the distribution $p_{\mathbf{m}}$. Second, in each instantiation, the template has a probability of ‘mutating’ to a closely-related pattern. The model for mutation of templates is described in the following paragraph. Third, data is generated from the mutated templates character-by-character from $q(\mathbf{y}|c)$. For instance, if the mutated template is given by $\tilde{\mathbf{m}} = \tilde{c}_1\tilde{c}_2 \dots \tilde{c}_{\tilde{l}}$, a data vector of length \tilde{l} , $\mathbf{y}_1\mathbf{y}_2 \dots \mathbf{y}_{\tilde{l}}$, corresponding to $\tilde{\mathbf{m}}$ is drawn from the distribution $\prod_i q(\mathbf{y}_i|\tilde{c}_i)$.

Now, we define the likelihood function for motifs $Q(\mathbf{Y}|\mathbf{m})$, where \mathbf{Y} is a sequence of data vectors $\mathbf{y}_1\mathbf{y}_2 \dots \mathbf{y}_{\tilde{l}}$ and $\mathbf{m} = c_1c_2 \dots c_l$. A character in a motif can be mutated in two ways: a deletion, where the character does not appear, or copy insertions, where the character is copied in-place a certain number of times. This is similar to the operations allowed in a profile HMM. We do not allow for transpositions of two characters or insertions of new characters that are not identical to the one already present. To give an example, a motif abb ($K = 2$ alphabet size) may mutate to ab or bb corresponding to one deletion, and $aabb$ or $abbb$ corresponding to one insertion, but cannot mutate to $abab$ via one insertion or deletion (see main text for an explanation for this choice). We assume each character m_i in the motif has a probability $\epsilon_p p_d$ of being deleted, probability $\epsilon_p p_{c,j}$ of being copied exactly j ($j > 0$) times and is neither deleted nor copied with probability $1 - \epsilon_p$. To normalize, we require $p_d + \sum_j p_{c,j} = 1$. We can write down a recursive formula for computing $Q(\mathbf{Y}|\mathbf{m})$. Intuitively, we start from the last character in the motif and assign it a certain number of data points depending on whether there is a deletion (0 points), j insertions ($j + 1$ points) or neither (1 point). After the assignment, the last character and all the data points assigned to it are removed, and the process is repeated until there are no data points left to be assigned. Denote $Q(\mathbf{y}_1\mathbf{y}_2 \dots \mathbf{y}_{k_1}|c_1c_2 \dots c_{k_2}) \equiv M(k_1, k_2)$, where $0 \leq k_1 \leq L$ and $0 \leq k_2 \leq l$. For

$k_1, k_2 > 0$, M can be computed recursively as

$$M(k_1, k_2) = (1 - \epsilon_p)q(\mathbf{y}_{k_1}|c_{k_2})M(k_1 - 1, k_2 - 1) + \epsilon_p p_d M(k_1, k_2 - 1) + \sum_{j=1}^{k_1} \epsilon_p p_{c,j} M(k_1 - j, k_2 - 1) \prod_{k'=0}^j q(\mathbf{y}_{k_1-k'}|c_{k_2}), \quad (3)$$

with boundary conditions $M(0, k_1) = 0$ for $k_2 > 0$ and $M(0, 0) = 1$. In the applications presented in this paper, we set $p_{c,1} = 1 - p_d$ non-zero and the rest zero (i.e., only one copy of a character is allowed).

b. Likelihood function and proof that normalization is unity

Given a dataset $\mathbf{Y} = \mathbf{y}_1 \mathbf{y}_2 \dots \mathbf{y}_L$, its likelihood under the model is given by the sum over all possible 2^{L-1} partitionings $\{\pi\}$ of the dataset, where *partition* α in *partitioning* π is denoted \mathbf{Y}_α^π . Each partitioning is weighted by its likelihood. We have

$$P(\mathbf{Y}; \{p_m\}) = Z'^{-1} \sum_{\pi} \prod_{\alpha=1}^{N(\pi)} Q(\mathbf{Y}_\alpha^\pi), \quad (4)$$

where the marginal probability is $Q(\mathbf{Y}_\alpha) = \sum_{\mathbf{m}} Q(\mathbf{Y}_\alpha|\mathbf{m})p_{\mathbf{m}}$, $N(\pi)$ is the total number of templates in partition π and Z' is the normalization constant, which is integrated over all possible \mathbf{Y} . The product of the marginals Q in the above equation is the likelihood of each partitioning. We give a simple (non-rigorous) argument for Z' being close to unity for large L . Intuitively, this is because the normalization should be one if not for boundary effects, which we show are negligible for large datasets. We have

$$Z'(L) = \int \mathcal{D}\mathbf{Y}_L P(\mathbf{Y}_L; \{p_m\}) \quad (5)$$

where the subscript is introduced to keep track of its length. For any l , conditioning on the first partition before l , we have the recursive formula

$$\begin{aligned} Z'(l) &= \sum_{i=1}^{l-1} \int \mathcal{D}\mathbf{Y}_i Q(\mathbf{Y}_i) Z'(l-i) \\ &= \sum_{\mathbf{m}} p_{\mathbf{m}} \sum_{i=1}^{l-1} \int \mathcal{D}\mathbf{Y}_i Q(\mathbf{Y}_i|\mathbf{m}) Z'(l-i) \end{aligned} \quad (6)$$

Consider the case when $l \ll L$ and we modify the sequence so that the first $l-1$ points in the sequence are replaced by 1s instead of the appropriate integrals over the marginals Q i.e.,

$Z'(i) = 1$ for $i < l$. Since $l \ll L$, we must have Z' for this modified sequence approximately equal to the original Z' , with corrections being $O(l/L)$. By the above recursive formula, for this modified sequence $Z'(l) = 1$ because of the normalization of $p_{\mathbf{m}}$ and $Q(\mathbf{Y}|\mathbf{m})$ and it follows that $Z'(l') = 1$ for any $l' > l$. An approach for a rigorous argument could be to find an upper bound for $|1 - Z'(i)|$ for $i < l \ll L$ and proceed similarly.

c. Maximum likelihood estimation of motif probabilities

Since $Z' \simeq 1$, $\{p_{\mathbf{m}}\}$ enter the equation only via $Q(\mathbf{Y}_{\alpha})$. We use a Lagrange multiplier λ to impose the normalization constraint $\sum p_{\mathbf{m}} = 1$ and maximize $P(\mathbf{Y}; \{p_{\mathbf{m}}\})$. Taking the derivative w.r.t $p_{\mathbf{m}}$ and equating to zero, we get

$$\lambda = \sum_{\pi} \sum_{\alpha'} Q(\mathbf{Y}_{\alpha'}^{\pi}|\mathbf{m}) \prod_{\alpha \neq \alpha'} Q(\mathbf{Y}_{\alpha}^{\pi}). \quad (7)$$

Multiplying on both sides by $p_{\mathbf{m}}$ and using Bayes' rule, $Q(\mathbf{Y}_{\alpha'}^{\pi}|\mathbf{m})p_{\mathbf{m}} = Q(\mathbf{Y}_{\alpha'}^{\pi})p(\mathbf{m}|\mathbf{Y}_{\alpha'}^{\pi})$, we get the implicit equation for the MLE,

$$\lambda p_{\mathbf{m}}^* = \sum_{\pi} \sum_{\alpha'=1}^{N(\pi)} p(\mathbf{m}|\mathbf{Y}_{\alpha'}^{\pi}) \prod_{\alpha=1}^{N(\pi)} Q(\mathbf{Y}_{\alpha}^{\pi}). \quad (8)$$

The sum over the posterior probabilities can be interpreted as an effective number of counts of \mathbf{m} in the partition π . We then have $p_{\mathbf{m}}^* = \langle N_{\mathbf{m}} \rangle / \sum_{\mathbf{m}'} \langle N_{\mathbf{m}'} \rangle$, where $\langle N_{\mathbf{m}} \rangle$ is the expected number of counts of \mathbf{m} over the ensemble of partitions.

d. Numerical optimization

To compute $p_{\mathbf{m}}$ numerically, we define the free energy, $F = -\ln P(\mathbf{Y}; \{p_{\mathbf{m}}\})$ and minimize it under the constraint $\sum_{\mathbf{m}} p_{\mathbf{m}} = 1$. We impose this constraint using the transformation $p_{\mathbf{m}} = \frac{e^{-\beta_{\mathbf{m}}}}{\sum_{\mathbf{m}'} e^{-\beta_{\mathbf{m}'}}}$, fixing one of the $\beta_{\mathbf{m}}$'s to zero and optimizing for the rest. We use the L-BFGS-B algorithm, which requires F and the gradient of F , whose calculation is described below. The Hessian can be computed exactly and Hessian-based methods applied, but we found that the gradient-based methods work equally well and are in fact slightly faster. Note that since $\langle N_{\mathbf{m}} \rangle = -p_{\mathbf{m}} \partial_{\mathbf{m}} F$, the gradient calculation can also be used to calculate the number of occurrences of a motif.

e. Free energy and gradient calculation

The calculation of F and $\partial_{\mathbf{m}}F$ is performed using transfer matrix-like techniques. The procedure is similar to the one described in ref [27] (see [50] for technical details). The key insight from ref [27] is that the partitioning of the data vector has a Markovian structure, which allows calculations similar to the forward-backward algorithm for Hidden Markov Models. The same structure is preserved in our model albeit with the added complication that the motifs themselves are latent variables and thus all possibilities need to be considered for each partition, which is taken into account in the marginal $Q(\mathbf{Y}_\alpha)$ in (4). The procedure described below holds for a single unbroken sequence. In practice, the dataset contains a collection of sequences from different sources and the procedure is easily adapted to the multiple sequence case. Below, we use a subscript $i : j$ to mean the quantity is evaluated for the sequence $\mathbf{y}_i \mathbf{y}_{i+1} \dots \mathbf{y}_j$. Let $Z_{i:j}$ denote $P(\mathbf{Y}_{i:j} | \{p_{\mathbf{m}}\})$. We can write the recursive formula

$$Z_{1:i} = \sum_{l=1}^m Q_{i-l+1:i} Z_{1:i-l}, \quad (9)$$

where the marginal probability in shorthand is $Q_{i-l+1:i} = \sum_{\mathbf{m}} Q_{i-l+1:i}^{\mathbf{m}} p_{\mathbf{m}}$ and $Q^{\mathbf{m}}$ is the motif likelihood function of \mathbf{m} , $Q(\cdot | \mathbf{m})$. Here we truncate the length of the partition to $m = 2l_{\max}$, where l_{\max} is the length of the longest motif since we allow at most one duplication per character.

In practice, the Q values become extremely small and there is underflow. We instead calculate the ratio $R_i = Z_{1:i} / Z_{1:i-1}$, for which we can write

$$R_i = Q_{i:i} + \sum_{l=2}^m Q_{i-l+1:i} \left(\prod_{k=i-l+1}^{i-1} R_k \right)^{-1}. \quad (10)$$

The free energy is $F = -\sum_{i=1}^L \ln R_i$. For the derivative, it is also useful to calculate the quantity $R'_i = Z_{i:L} / Z_{i+1:L}$, which is calculated as

$$R'_i = Q_{i:i} + \sum_{l=2}^m Q_{i+l-1:i} \left(\prod_{k=i+1}^{i+l-1} R'_k \right)^{-1} \quad (11)$$

The derivative $\partial_{\mathbf{m}}F$ is calculated using

$$\partial_{\mathbf{m}}F = -\sum_{i=1}^L G_{i-l+1,i} Q_{i-l+1:i}^{\mathbf{m}}, \quad (12)$$

where $G_{i-l+1,i} \equiv Z_{1:i-l}Z_{i+1:L}/Z_{1:L}$ is calculated using

$$G_{i-l+1,i} = \left(\prod_{k=1}^{i-l} R_k/R'_k \right) \left(\prod_{k=i-l+1}^i R'_k \right)^{-1}. \quad (13)$$

In practice, calculating $Q_{i-l+1:i}^m$ is the bottleneck, and we pre-calculate and store non-zero (above a threshold) values of $Q_{i-l+1:i}^m$ before the optimization for p_m is done. Finally, gradients in p_m are converted to gradients in β_m for optimization.

f. Dictionary expansion

As described in the main text, new motifs are added by considering all possible concatenations of pairs of motifs in the dictionary. Consider the motif $\mathbf{m}\mathbf{m}'$ obtained by two existing motifs \mathbf{m} and \mathbf{m}' . The expected number of counts in \mathbf{Y} of $\mathbf{m}\mathbf{m}'$ is calculated using the derivative of F :

$$\langle N_{\mathbf{m}\mathbf{m}'} \rangle = \zeta(\mathbf{m}\mathbf{m}') \sum_{i=1}^L G_{i-l+1,i} Q_{i-l+1:i}^{\mathbf{m}\mathbf{m}'}, \quad (14)$$

where $\zeta(\mathbf{m}\mathbf{m}')$ is the probability under the model of all possible partitionings of the concatenated motif $\mathbf{m}\mathbf{m}'$. For example, if the dictionary contains a , b and ab with probabilities p_a, p_b, p_{ab} , we have $\zeta(abab) = p_a p_b p_a p_b + p_a p_b p_{ab} + p_{ab} p_a p_b + p_{ab} p_{ab} = p_a^2 p_b^2 + 2p_a p_b p_{ab} + p_{ab}^2$. ζ can be computed efficiently using an equation similar to (9). $\langle N_{\mathbf{m}\mathbf{m}'} \rangle$ is compared to the model prediction $\bar{N}\zeta(\mathbf{m}\mathbf{m}') \simeq L\zeta(\mathbf{m}\mathbf{m}')/\bar{l}$, where \bar{l} is the mean motif length, $\bar{l} \equiv \sum_{\mathbf{m}} l_{\mathbf{m}} p_{\mathbf{m}}$. A likelihood ratio test yields a p -value and pairs below a threshold (we use 0.001) are added to the dictionary.

g. Dictionary truncation

For truncation, the set $\{p_{\mathbf{m}}\}$ is re-calculated. To prune clusters of similar motifs, we define a distance measure between a pair of motifs using the Jensen-Shannon (JS) divergence of the data distributions Q generated by the pair, calculated using Monte Carlo sampling i.e., the distance between two motifs is

$$d(\mathbf{m}, \mathbf{m}') = \frac{D_{\text{KL}}(Q(\cdot|\mathbf{m})||Q') + D_{\text{KL}}(Q(\cdot|\mathbf{m}')||Q')}{2 \ln 2}, \quad (15)$$

where $Q' = (Q(\cdot|\mathbf{m}) + Q(\cdot|\mathbf{m}'))/2$ and D_{KL} is the standard Kullback-Leibler divergence between two distributions. The factor $1/2 \ln 2$ restricts the range of d between 0 and 1. Motifs which have a distance less than a threshold (we use $J_{thr} = 0.15$, see Figure S4) are considered similar. We re-weight motif probabilities by adding the probabilities of similar motifs to the estimated one, keep the representative motif with the highest re-weighted probability in a cluster and discard the rest. Finally, $\{p_{\mathbf{m}}\}$ is again calculated and motifs which have low counts (we use 5) are discarded. Individual characters are not discarded from the dictionary using the above criteria until the last iteration. Note that pruning using the JS divergence is only performed when $\epsilon_p > 0$; the JS distance is always 1 otherwise.

We keep track of the free energy per symbol of the dataset after every iteration of the above three steps. We find that the free energy monotonically decreases in almost all cases and converges to a local optimum, though we have no formal proof as to why it is monotonic. The iterative procedure is stopped when the relative change in free energy is $< 0.1\%$ for two consecutive iterations. The procedure usually converges within 15 iterations.

h. Most likely sequence of motifs

The HMM structure of the generative model allows us to compute the most likely sequence of motifs that generate the dataset in a manner similar to the Viterbi algorithm. In particular, given the dataset \mathbf{Y} , we would like to find the most likely sequence $\mathbf{m}_1^*, \mathbf{m}_2^*, \dots$ defined as

$$\mathbf{m}_1^*, \mathbf{m}_2^*, \dots = \arg \max_{\mathbf{m}_1, \mathbf{m}_2, \dots} P(\mathbf{m}_1, \mathbf{m}_2, \dots | \mathbf{Y}) \quad (16)$$

$$= \arg \max_{\mathbf{m}_1, \mathbf{m}_2, \dots} P(\mathbf{Y} | \mathbf{m}_1, \mathbf{m}_2, \dots) \prod_i p_{\mathbf{m}_i} \quad (17)$$

The max over the quantity on the right hand side can be computed using dynamic programming; we define a value function for the probability of the most likely motif sequence $\max_{\mathbf{m}_1, \mathbf{m}_2, \dots} P(\mathbf{m}_1, \mathbf{m}_2, \dots | \mathbf{Y}_{1:k})$ for the length k sequence $\mathbf{Y}_{1:k}$ and iterate over the next most likely motif.

i. Generation of the synthetic dataset

The data is drawn from a Gaussian mixture model with 7 components in two dimensions as shown in Figure 2a of the main text. We sample a frequency distribution over the 7 characters in the alphabet from a symmetric Dirichlet distribution ($\alpha = 5$). We then construct a dictionary of 50 motifs, with each motif having a length drawn from a Poisson distribution with mean 5. The characters in each motif are drawn according to the frequency distribution over the alphabet. The probability of each motif, p_m , is drawn from a symmetric Dirichlet distribution ($\alpha = 1$) and scaled with a parameter $1 - \epsilon_b$, where ϵ_b is the fraction of the dataset that is made up of individual characters and specifies the ‘background noise’.

The data is sampled according to the generative process described previously. We use ϵ_p as a measure of action pattern noise, which is the probability of a deletion or duplication (which are set equal) per symbol in the motif as described previously. To quantify locomotor episode noise, we define the discriminability μ between clusters using the distance between neighboring clusters (each of which has standard deviation one) as shown in Figure 2a. For the numerical experiments, we set $\epsilon_b = 0.5$ fixed and vary ϵ_p and μ . In the example shown in main text, we set dataset size $L = 40000$, $\epsilon_p = 0$, $\mu = 3$.

6. Movie legends

Movie S1. Movie showing the raw video captured by the camera and the most likely bout categorization using the Gaussian Mixture Model. The data is from a single larva in the exploratory environment.

Movie S2. Movie showing the raw video captured by the camera and the most likely bout categorization using the Gaussian Mixture Model. The data is from a single larva in the aversive environment.

Movie S3. Movie showing a sample of the the abnormal motif $fTff$ (see main text).

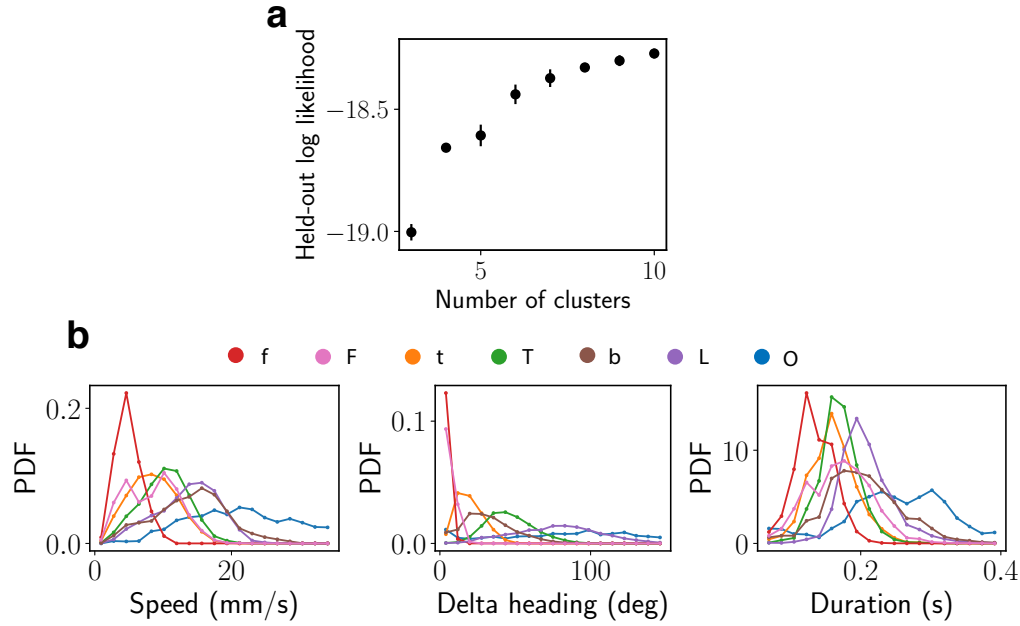


FIG. S1. Bout categorization. (a) The held-out log likelihood from fitting the Gaussian Mixture Model to the exploratory and aversive datasets vs the number of mixture components. We characterize each bout using its speed, change in heading, summed tail angle (integrated absolute tail angle over time) and the first three PC components of the tail angles. (b) The histograms of speed, change in heading and duration of bout for each bout type.

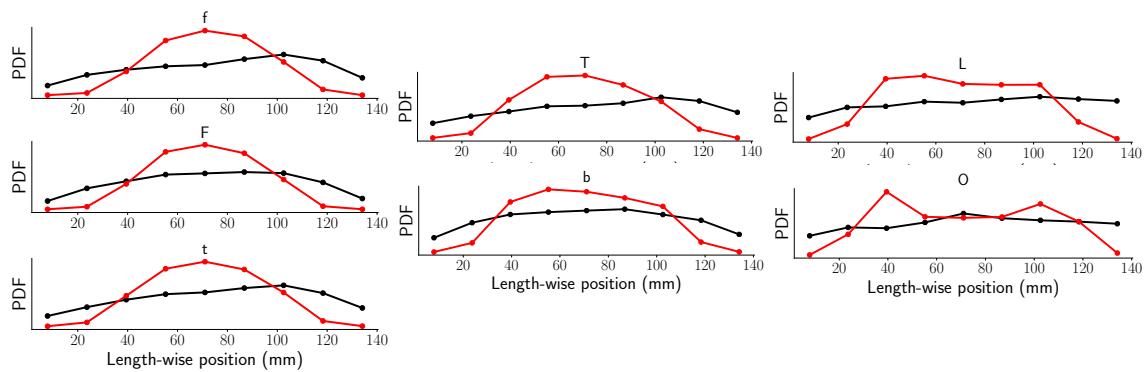


FIG. S2. Histogram of larvae positions along the well in exploratory (black) and aversive (red) datasets.

TABLE S1. Motifs discovered in exploratory data: The 25 motifs that deviate most from Markovianity are shown (as measured by a p -value with Markovianity as the null hypothesis, see main text).

Motif	$-\log_{10} p$	Observed	Expected
fffffff	inf	1366	387
fffffff	inf	510	50
ffffff	208.01	3234	1797
fff	42.33	9544	8327
FFFFFFF	28.23	311	153
fffffff	27.64	497	290
ftffff	25.07	495	297
ftfff	22.72	1125	824
ftff	21.12	1745	1377
ftf	18.5	2724	2289
ftf	13.96	4337	3859
TfT	11.12	722	554
FFFF	7.94	1428	1224
TfTf	7.28	346	254
tttt	6.7	256	181
TTTT	5.06	160	110
bb	3.87	924	1044
bbbb	3.21	115	82
FbFb	2.19	99	74
Fb	1.79	1300	1388
Tfff	1.74	1690	1788
bbFb	1.72	72	54
bFbF	1.72	100	78
bF	1.66	1374	1460
bFb	1.57	242	209

TABLE S2. Motifs discovered in the data from Marques et al: The 25 motifs that deviate most from Markovianity are shown (as measured by a p -value with Markovianity as the null hypothesis, see main text).

Motif	$-\log_{10} p$	Observed	Expected
TTTTTTT	142.48	580	161
TTTTTTTT	142.23	378	71
ffffff	132.95	883	338
ffffff	130.9	546	155
TTTTTT	128.43	915	364
fffff	127.1	1483	741
fTf	119.55	1896	1057
ffff	105.39	2569	1620
TTTTT	102.33	1515	826
fTf	96.1	1009	483
fTff	72.61	540	221
ttttt	64.91	475	192
fff	64.26	4586	3544
fffTf	59.68	306	101
TTTT	57.99	2604	1870
ttttttt	56.51	148	28
TfTf	55.98	355	133
FFFFFF	53.79	385	154
fffTff	51.81	184	46
FfF	50.8	1447	948
FFFFFFF	47.97	201	57
TfTf	47.27	206	61
fTffff	42.53	111	21
ttff	42.3	881	1352
tttt	42.11	1853	1326
ffF	42.09	2217	1637

TABLE S3. Motifs discovered in the data from chemotactic fish in the aversive environment: The 25 motifs that deviate most from Markovianity are shown (as measured by a p -value with Markovianity as the null hypothesis, see main text).

Motif	$-\log_{10} p$	Observed	Expected
fffffffffff	inf	497	33
fffffffffff	inf	751	93
ffffffffffff	inf	334	12
ffffff	278.21	1860	719
fffff	155.38	3266	1990
FFFFFFF	71.86	653	294
fff	44.26	6533	5506
FFFFFFFFFFFFFFF	43.87	59	4
ffftfff	17.45	290	166
fffftff	16.56	285	165
FFFF	14.89	2167	1821
ffftff	14.89	641	460
fttff	10.83	960	767
TfT	8.42	456	342
fTfT	6.54	184	123
bbbb	6.45	316	234
tttt	6.09	162	107
bFFFFFFF	5.94	87	49
FFFFFFFb	5.67	83	47
TTTT	5.3	139	92
ftff	4.87	1448	1290
FFFbFFFF	4.79	56	30
bFb	4.18	405	330
bbFb	3.81	148	107
Obbb	3.68	106	72

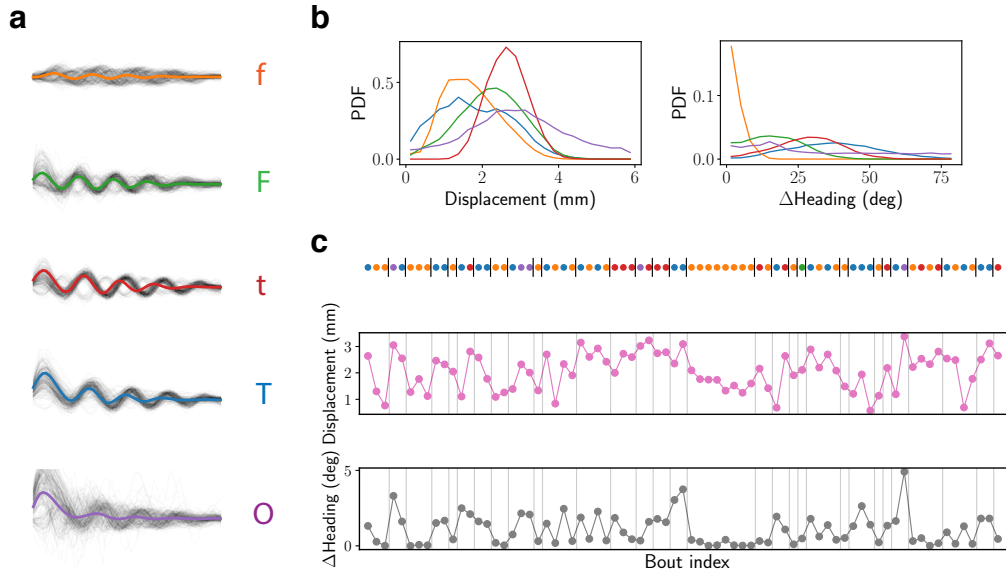


FIG. S3. Bout categorization and motif segmentation of the dataset from Marques et al. (a) The five bout types identified using a Gaussian Mixture Model. Each bout is characterized using its displacement, change in heading and the first four PC components of the tail angles. (b) The histograms of displacement and change in heading for the five bout types. (c) A sample from the dataset showing the most likely segmentation of the data into motifs from the learned dictionary.

-
- [1] D. J. Anderson and P. Perona, Toward a science of computational ethology, *Neuron* **84**, 18 (2014).
 - [2] A. Gomez-Marin, J. J. Paton, A. R. Kampff, R. M. Costa, and Z. F. Mainen, Big behavioral data: psychology, ethology and the foundations of neuroscience, *Nature neuroscience* **17**, 1455 (2014).
 - [3] S. R. Datta, D. J. Anderson, K. Branson, P. Perona, and A. Leifer, Computational neuroethology: A call to action, *Neuron* **104**, 11 (2019).
 - [4] A. E. Brown and B. De Bivort, Ethology as a physical science, *Nature Physics* **14**, 653 (2018).
 - [5] G. J. Berman, Measuring behavior across scales, *BMC biology* **16**, 23 (2018).
 - [6] T. D. Pereira, D. E. Aldarondo, L. Willmore, M. Kislin, S. S.-H. Wang, M. Murthy, and J. W. Shaevitz, Fast animal pose estimation using deep neural networks, *Nature methods* **16**, 117 (2019).




FIG. S4. Hyperparameter selection via held-out test sets. Pattern noise ϵ_p , the probability of a deletion p_d and the Jensen-Shannon divergence threshold for removing similar bouts, J_{thr} , were set by computing the free energy per bout on a randomly split held-out dataset of approximately 18000 bouts (training dataset size was 70000). The free energy per bout for each tested parameter triplet is computed by averaging over 18 trials (with a separate random split for each trial). The red square shows the parameters ($\epsilon_p = 0.1, p_d = 0.2, J_{thr} = 0.15$) chosen for the rest of our analysis.

- [7] A. Mathis, P. Mamidanna, K. M. Cury, T. Abe, V. N. Murthy, M. W. Mathis, and M. Bethge, DeepLabcut: markerless pose estimation of user-defined body parts with deep learning, *Nature neuroscience* **21**, 1281 (2018).
- [8] J. M. Graving, D. Chae, H. Naik, L. Li, B. Koger, B. R. Costelloe, and I. D. Couzin, Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning, *Elife* **8**, e47994 (2019).
- [9] G. J. Berman, W. Bialek, and J. W. Shaevitz, Predictability and hierarchy in drosophila behavior, *Proceedings of the National Academy of Sciences* **113**, 11943 (2016).
- [10] A. Sharma, R. Johnson, F. Engert, and S. Linderman, Point process latent variable models of larval zebrafish behavior, in *Advances in Neural Information Processing Systems* (2018) pp. 10919–10930.
- [11] A. B. Wiltschko, M. J. Johnson, G. Iurilli, R. E. Peterson, J. M. Katon, S. L. Pashkovski, V. E. Abraira, R. P. Adams, and S. R. Datta, Mapping sub-second structure in mouse behavior, *Neuron* **88**, 1121 (2015).
- [12] B. C. Daniels, W. S. Ryu, and I. Nemenman, Automated, predictive, and interpretable in-

- ference of *caenorhabditis elegans* escape dynamics, *Proceedings of the National Academy of Sciences* **116**, 7226 (2019).
- [13] G. J. Stephens, B. Johnson-Kerner, W. Bialek, and W. S. Ryu, Dimensionality and dynamics in the behavior of *c. elegans*, *PLoS computational biology* **4** (2008).
- [14] G. J. Stephens, M. B. De Mesquita, W. S. Ryu, and W. Bialek, Emergence of long timescales and stereotyped behaviors in *caenorhabditis elegans*, *Proceedings of the National Academy of Sciences* **108**, 7286 (2011).
- [15] A. J. Calhoun, J. W. Pillow, and M. Murthy, Unsupervised identification of the internal states that shape natural behavior, *Nature neuroscience* **22**, 2040 (2019).
- [16] J. C. Marques, S. Lackner, R. Félix, and M. B. Orger, Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering, *Current Biology* **28**, 181 (2018).
- [17] D. S. Mearns, J. C. Donovan, A. M. Fernandes, J. L. Semmelhack, and H. Baier, Deconstructing hunting behavior reveals a tightly coupled stimulus-response loop, *Current Biology* **30**, 54 (2020).
- [18] A. C. Costa, T. Ahamed, and G. J. Stephens, Adaptive, locally linear models of complex dynamics, *Proceedings of the National Academy of Sciences* **116**, 1501 (2019).
- [19] S. Gupta and A. Gomez-Marin, A context-free grammar for *caenorhabditis elegans* behavior, *BioRxiv* , 708891 (2019).
- [20] N. J. Plowes, T. Colella, R. A. Johnson, and B. Hölldobler, Chemical communication during foraging in the harvesting ants *messor pergandei* and *messor andrei*, *Journal of Comparative Physiology A* **200**, 129 (2014).
- [21] G. Reddy, A. Celani, T. J. Sejnowski, and M. Vergassola, Learning to soar in turbulent environments, *Proceedings of the National Academy of Sciences* **113**, E4877 (2016).
- [22] G. Reddy, J. Wong-Ng, A. Celani, T. J. Sejnowski, and M. Vergassola, Glider soaring via reinforcement learning in the field, *Nature* **562**, 236 (2018).
- [23] T. L. Warren, Y. M. Giraldo, and M. H. Dickinson, Celestial navigation in *drosophila*, *Journal of Experimental Biology* **222** (2019).
- [24] A. M. Kist and R. Portugues, Optomotor swimming in larval zebrafish is driven by global whole-field visual motion and local light-dark transitions, *Cell Reports* **29**, 659 (2019).
- [25] R. E. Johnson, S. Linderman, T. Panier, C. L. Wee, E. Song, K. J. Herrera, A. Miller, and F. Engert, Probabilistic models of larval zebrafish behavior reveal structure on many scales,

- Current Biology **30**, 70 (2020).
- [26] E. Gahtan, P. Tanger, and H. Baier, Visual prey capture in larval zebrafish is controlled by identified reticulospinal neurons downstream of the tectum, *Journal of Neuroscience* **25**, 9294 (2005).
- [27] H. J. Bussemaker, H. Li, and E. D. Siggia, Building a dictionary for genomes: identification of presumptive regulatory sites by statistical analysis, *Proceedings of the National Academy of Sciences* **97**, 10096 (2000).
- [28] D. Ron, Y. Singer, and N. Tishby, The power of amnesia: Learning probabilistic automata with variable memory length, *Machine learning* **25**, 117 (1996).
- [29] G. Bejerano and G. Yona, Variations on probabilistic suffix trees: statistical modeling and prediction of protein families, *Bioinformatics* **17**, 23 (2001).
- [30] J. E. Markowitz, E. Ivie, L. Kligler, and T. J. Gardner, Long-range order in canary song, *PLoS computational biology* **9** (2013).
- [31] D. Mochihashi and E. Sumita, The infinite markov model, in *Advances in neural information processing systems* (2008) pp. 1017–1024.
- [32] S. Goldwater, T. L. Griffiths, and M. Johnson, A bayesian framework for word segmentation: Exploring the effects of context, *Cognition* **112**, 21 (2009).
- [33] O. Mirat, J. R. Sternberg, K. E. Severi, and C. Wyart, Zebrazoom: an automated program for high-throughput behavioral analysis and categorization, *Frontiers in neural circuits* **7**, 107 (2013).
- [34] G. J. Berman, D. M. Choi, W. Bialek, and J. W. Shaevitz, Mapping the stereotyped behaviour of freely moving fruit flies, *Journal of The Royal Society Interface* **11**, 20140672 (2014).
- [35] A. E. Brown, E. I. Yemini, L. J. Grundy, T. Jucikas, and W. R. Schafer, A dictionary of behavioral motifs reveals clusters of genes affecting *caenorhabditis elegans* locomotion, *Proceedings of the National Academy of Sciences* **110**, 791 (2013).
- [36] G. Reddy, <https://github.com/greddy992/bass> (2020).
- [37] T. W. Dunn, Y. Mu, S. Narayan, O. Randlett, E. A. Naumann, C.-T. Yang, A. F. Schier, J. Freeman, F. Engert, and M. B. Ahrens, Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion, *Elife* **5**, e12741 (2016).
- [38] T. Koide, Y. Yabuki, and Y. Yoshihara, Terminal nerve *gnrh3* neurons mediate slow avoidance of carbon dioxide in larval zebrafish, *Cell reports* **22**, 1115 (2018).

- [39] W.-y. Chen, X.-l. Peng, Q.-s. Deng, M.-j. Chen, J.-l. Du, and B.-b. Zhang, Role of olfactorily responsive neurons in the right dorsal habenula–ventral interpeduncular nucleus pathway in food-seeking behaviors of larval zebrafish, *Neuroscience* **404**, 259 (2019).
- [40] R. Candelier, M. S. Murmu, S. A. Romano, A. Jouary, G. Debrégeas, and G. Sumbre, A microfluidic device to study neuronal and motor responses to acute chemical stimuli in zebrafish, *Scientific reports* **5**, 12196 (2015).
- [41] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, An introduction to variational methods for graphical models, *Machine learning* **37**, 183 (1999).
- [42] N. Rajewsky, M. Vergassola, U. Gaul, and E. D. Siggia, Computational detection of genomic cis-regulatory modules applied to body patterning in the early drosophila embryo, *BMC bioinformatics* **3**, 30 (2002).
- [43] C. G. Nevill-Manning and I. H. Witten, On-line and off-line heuristics for inferring hierarchies of repetitions in sequences, *Proceedings of the IEEE* **88**, 1745 (2000).
- [44] A. Gomez-Marin, G. J. Stephens, and A. E. Brown, Hierarchical compression of caenorhabditis elegans locomotion reveals phenotypic differences in the organization of behaviour, *Journal of The Royal Society Interface* **13**, 20160466 (2016).
- [45] M. Ghosh and J. Rihel, Hierarchical compression reveals sub-second to day-long structure in larval zebrafish behaviour, *bioRxiv* , 694471 (2019).
- [46] Z. Akos, M. Nagy, and T. Vicsek, Comparing bird and human soaring strategies, *Proceedings of the National Academy of Sciences* **105**, 4139 (2008).
- [47] A. Pujala and M. Koyama, Chronology-based architecture of descending circuits that underlie the development of locomotor repertoire after birth, *Elife* **8**, e42135 (2019).
- [48] K.-H. Huang, M. B. Ahrens, T. W. Dunn, and F. Engert, Spinal projection neurons control turning behaviors in zebrafish, *Current Biology* **23**, 1566 (2013).
- [49] O. Mirat, *Zebrazoom* (2020).
- [50] H. J. Bussemaker, H. Li, E. D. Siggia, *et al.*, Regulatory element detection using a probabilistic segmentation model., in *Ismb* (2000) pp. 67–74.