Unsupervised Bayesian Ising Approximation for revealing the neural dictionary in songbirds

Damián G. Hernández^{1,2}, Samuel J. Sober³, and Ilya Nemenman^{2,3,4}

¹Department of Medical Physics, Centro Atómico Bariloche and Instituto Balseiro, Bariloche 8400, Argentina

²Department of Physics, Emory University, Atlanta, GA 30322, USA

³Department of Biology, Emory University, Atlanta, GA 30322, USA

⁴Initiative in Theory and Modeling of Living Systems, Atlanta, GA 30322, USA

ABSTRACT

The problem of deciphering how low-level patterns (action potentials in the brain, amino acids in a protein, etc.) drive high-level biological features (sensorimotor behavior, enzymatic function) represents the central challenge of quantitative biology. The lack of general methods for doing so from the size of datasets that can be collected experimentally severely limits our understanding of the biological world. For example, in neuroscience, some sensory and motor codes have been shown to consist of precisely timed multi-spike patterns. However, the combinatorial complexity of such pattern codes have precluded development of methods for their comprehensive analysis. Thus, just as it is hard to predict a protein's function based on its sequence, we still do not understand how to accurately predict an organism's behavior based on neural activity. Here we derive a method for solving this class of problems. We demonstrate its utility in an application to neural data, detecting precisely timed spike patterns that code for specific motor behaviors in a songbird vocal system. Our method detects such codewords with an arbitrary number of spikes, does so from small data sets, and accounts for dependencies in occurrences of codewords. Detecting such dictionaries of important spike patterns – rather than merely identifying the timescale on which such patterns exist, as in some prior approaches – opens the door for understanding fine motor control and the neural bases of sensorimotor learning in animals. For example, for the first time, we identify differences in encoding motor exploration versus typical behavior. Crucially, our method can be used not only for analysis of neural systems, but also for understanding the structure of correlations in other biological and nonbiological datasets.

Introduction

One of the goals of modern high-throughput biology is to generate predictive models of interaction networks, from interactions among individual biological molecules¹ to the encoding of information by networks of neurons in the brain². To be able to make predictions about activity across networks requires one to accurately approximate—or build a *model* of—their joint probability distribution, such as the distribution of joint firing patterns in neural populations or the distribution of co-occurring mutations in proteins of the same family. To successfully generalize and to improve interpretability, models should contain as few as possible terms. Thus constructing a model requires detecting *relevant* features in the data: namely, the smallest possible set of spike patterns or nucleotide sequences that capture the most correlations among the network components. By analogy with the human language, where words are strongly correlated, co-occurring combinations of letters, we refer to the problem of detecting features that succinctly describe correlations in a data set as the problem of *dictionary reconstruction*, see Fig. 1.

In recent years, the problem of dictionary reconstruction has been addressed under different names for a variety of biological contexts³ including gene expression networks^{4,5}, protein structure, protein-protein interactions^{1,6–8}, the structure of regulatory DNA⁹, distribution of antibodies and pathogenic sequences^{10,11}, species abundance¹², and collective behaviors^{13–17}. The efforts to identify interactions in neural activity have been particularly plentiful^{2,18–21}. The diversity of biological applications notwithstanding, most of these attempts have relied on similar mathematical constructs, and most have suffered from the same limitations. First, unlike in classical statistics and traditional quantitative model building, where the number of observations, M , usually vastly exceeds the number of unknowns to be estimated, K , $K/M \ll 1$, modern biological data often has $M \gg 1$, but also $K/M \gg 1$. Indeed, because of, for example, protein allostery, recurrent connections within neural populations, or coupling to global stimuli, biological systems are rarely limited to local interactions only^{22–24}, so that the number of pairwise interactions among N variables is $K \sim N^2$, and the number of all higher order interactions among them is $K \sim 2^N$. Put differently, words in biological dictionaries can be of an arbitrary length, and spelling rules may involve many letters simultaneously, some of which are far away from each other. Because of this, reconstruction of biological dictionaries from data sets of realistic sizes requires assumptions and simplifications about the structure of possible biological correlations, and will not be possible by brute force. The second problem is that, as in human languages, biological dictionaries have redundancies: there are synonyms

Figure 1. The dictionary reconstruction problem. In many biological systems, such as understanding the neural code, identifying protein-DNA binding sites, or predicting 3-D protein structures, we need to infer dictionaries — the sets of statistically over- or under-represented features in the datasets, which we refer to as words in the dictionary. To do so, we represent the data as a matrix of binary activities of N biological units (spike/no spike, presence/absence of a mutation, etc.), and view M different experimental instantiations as samples from an underlying stationary probability distribution. We then use the uBIA method to identify the significant words in the data. Specifically, uBIA systematically searches for combinatorial activity patterns that are over- or under-represented compared to their expectation given the marginal activity of the individual units. If multiple similar patterns can (partially) explain the same statistical regularities in the data, they compete with each other for importance, resulting in an irreducible dictionary of significant codewords. In different biological problems, such dictionaries can represent neural control words, DNA binding motifs, or conserved patterns of amino acids that must be neighbors in the 3-D protein structure.

and words that share roots. For example, a set of gene expressions may be correlated not because the genes interact directly, but because they are related to some other genes that do. Similarly, a certain pattern of neural activity may be statistically over- or under-represented not on its own, but because it is a subset or a superset of another, more important, pattern. Identifying *irreducible words*—the root forms of biological dictionaries—is therefore harder than detecting all correlations while also being crucial to fully understanding biological systems. Altogether, these complications make it impossible to use off-the-shelf methods for reconstructing combinatorially complex dictionaries from datasets of realistic sizes.

In this work, we propose a new method for reconstructing complex biological dictionaries from relatively small datasets, as few as $M \sim 10^2 \dots 10^3$ samples of the joint activity. We impose no limitation on the structure of the words that can enter the dictionary — neither their length nor their rules of spelling — beyond the obvious limitation that (i) words that do not happen in the data cannot be detected, and (ii) that data contain few samples of many words, rather than of just a few that repeat many times. Additionally, we address the problem of irreducibility, making the inferred dictionaries compact, non-redundant, and easier to comprehend. The main realization that allows this progress is that instead of approximating the entire joint probability distribution of a system’s states and hence answering *how* specific significant words matter, we can focus on a more restricted, and hence simpler, question: *which* specific words contribute to the dictionary. We answer this question in the language of Bayesian inference and statistical mechanics by developing an unsupervised version of the Bayesian Ising Approximation²⁵ and by merging it with the *reliable interactions* model²⁶.

To illustrate the capabilities of our approach, we develop it in the context of a specific biological system: recordings from single neurons in brain area RA (the robust nucleus of the arcopallium) in the Bengalese finch, a songbird. Neurons communicate with each other using patterns of action potentials (spikes), which encode sensory information and motor commands, and hence behavior. Reconstructing the neural dictionary, and specifically detecting irreducible patterns of neural activity that correlate with (or “encode”) sensory stimuli or motor behaviors — which we hereafter call *codewords* — has been a key problem in computational neuroscience for decades¹⁸. It is now known that both in the sensory^{27–34} and in the motor systems^{35–38} the timing of neural action potentials (spikes) in multispike patterns, down to millisecond resolution, can contribute to the encoding of sensory or motor information. Such dictionaries that involve long sequences of neural activities (or incorporate multiple neurons) at high temporal resolution are both complex and severely undersampled. Specifically, even though the songbird datasets considered here are large by neurophysiological standards, they are too small for most existing

Figure 2. Quantification of the neural activity and the behavior. A spectrogram of a single song syllable in top-left corner shows the acoustic power (color scale) at different frequencies as a function of time. Each tick mark (bottom-left) represents one spike and each row represents one instantiation of the syllable. We analyze spikes produced in a 40 ms premotor window (red box) prior to the time when acoustic features were measured (red arrowhead). These spikes were binarized as 0 (no spike) or 1 (at least one spike) in 2 ms bins, totaling 20 time bins. The different acoustic features (pitch, amplitude, spectral entropy) were also binarized. For different analyses in this paper, 0/1 can denote the behaviors feature that is below/above or above/below its median, or as not belonging/belonging to a specific percentile interval. The inset shows the area RA within the song pathway, two synapses away from the vocal muscles, from which these data were recorded.

analysis approaches for modeling neural activity to detect the codewords. This motivates the general inference problem we address here.

Understanding the neural-motor dictionary answers important questions about vertebrate motor control which could not be addressed by previous techniques (see *Online Methods*). Specifically, while it is known that various features of the complex vocal behaviors are encoded by millisecond-scale firing patterns³⁵, here for the first time we can identify which specific patterns most strongly predict behavioral variations. Further, we show that dictionaries of individual neurons are rather large and quite variable, so that neurons speak different languages, which nonetheless share some universal features. Intriguingly, we detect that codewords that predict large, exploratory deviations in vocal acoustics are statistically different from those that predict typical behaviors. Collectively, these findings pave the way for development of future theories of the structure of these dictionaries, of how they are formed during development, how they adapt to different contexts, and how motor biophysics translates them into specific movements. More generally, these findings open gates for using our method in other biological domains, where reconstruction of feature dictionaries is equally important.

Results

The neuro-motor code reconstruction problem

Owing to their complex and tightly regulated vocal behavior and experimentally accessible nervous system, songbirds provide an ideal model system for investigating the neural dictionaries underlying complex motor behaviors^{39,40}. We recorded from individual neurons in the motor cortical area RA of Bengalese finches during spontaneous singing, while at the same

time quantifying acoustic feature of song, specifically the fundamental frequency ("pitch"), amplitude, and spectral entropy of individual vocal gestures, or "syllables", as described previously^{35,41,42}. The data sets are sufficiently large to allow reconstruction of the dictionaries: we have 49 data sets — spanning 4 birds, 30 neurons and sometimes multiple song syllables for each neuron — for which we observed at least 200 instances of the behavior and the associated neural activity, which we estimate below to be the lower threshold for a sufficient statistical power.

To represent analysis of this motor code as a dictionary reconstruction problem, we binarize the recorded spiking time series so that $\sigma_t = (0, 1)$ indicates the absence or presence of a spike in a time slice of $[(t-1)\Delta t, t\Delta t)$, see Fig. 2. Thus each time interval is represented by a binary variable, and interactions among these patterns are described by over-represented or under-represented sequences of zeros and ones in the data. Using a complementary information-theoretic analysis, Tang et al.³⁵ showed that the mutual information between the neural spike train and various features of song acoustics peaks at $\Delta t = 1 \dots 2$ ms. Thus studying precise timing codes means that we focus on $\Delta t = 2$ ms (our datasets are not large enough to explore smaller Δt) as discussed previously in³⁵. Detection of statistically significant codewords at this temporal resolution would both re-confirm that this neural motor code is timing based, consistent with previous analyses³⁵, as well as for the first time revealing the specific patterns that most strongly predict behavior. We focus on neural time series of length $T = 40$ ms duration preceding a certain acoustic syllable, which includes the approximate premotor delay with which neurons and muscles influence behavior³⁵. Thus the index t runs between 1 and $N = T/\Delta t = 20$.

Since we are interested in codewords that relate neural activity and behavior, we similarly binarize the motor output³⁵, denoting by 0 or 1 different binary characteristics of the behavior, such as pitch being either below or above its median value, or outside or inside a certain range, see Fig. 2. We treat the behavioral variable as the 0'th component of the neuro-behavioral activity network, which then has $N = 21$ binary variables, $\sigma = \{\sigma_i\}_0^N$. Finding the codewords and reconstructing the neural-behavioral dictionary is then equivalent to detecting significantly over- or under-represented patterns in the probability distribution $P(\sigma)$, and specifically those patterns that are overrepresented together with the behavioral bit. Note that $2^N = 2^{21} \approx 2 \cdot 10^6$, which is much greater than $M \sim 100 \dots 1000$ observations of the activity that we can record, illustrating the complexity of the problem. In fact, similar to the famous birthday problem, one expects a substantial number of repeating samples of the activity — and hence the ability to detect statistically over- and under-represented binary words — only when $M \sim \sqrt{2^N}$, which is what limits the statistical power of any dictionary reconstruction method.

The unsupervised BIA method (uBIA)

To reconstruct the neural-motor dictionaries and detect codewords that predict specific behaviors, we develop the unsupervised Bayesian Ising Approximation (uBIA) method based on the Bayesian Ising Approximation for detection of significant features in regression problems²⁵. For this, we write the probability distribution $P(\sigma)$ as

$$\log P(\sigma|\vec{\theta}) = -\log Z + \sum_{i=0}^N \theta_i \sigma_i + \sum_{j \geq i} \theta_{ij} \sigma_i \sigma_j + \sum_{k \geq j \geq i} \theta_{ijk} \sigma_i \sigma_j \sigma_k + \dots + \theta_{0\dots N} \sigma_0 \times \dots \times \sigma_N = -\log Z + \sum_{\mu} \theta_{\mu} \prod_{i \in V_{\mu}} \sigma_i, \quad (1)$$

where Z is the normalization coefficient⁴³. We use the notation such that V_{μ} is a nonempty subset of indexes $i \in [0, N]$, and $\mu = [1, 2^{N+1} - 1]$ is the subset number. Then $\{\theta_{\mu}\} = \vec{\theta}$ are coefficients in the log-linear model in front of the corresponding product of binary σ s. In other words, V_{μ} denotes a specific combination of the behavior and / or times when the neuron is active. If θ_{μ} is nonzero for a term $\prod_{i \in V_{\mu}} \sigma_i$, where $i = 0$ (the response variable) is in V_{μ} , then this specific spike word is correlated with the motor output, and is a significant codeword in the neural code, see Fig. 1. Finding nonzero θ_{μ} s is then equivalent to identifying *which* codewords matter and should be included in the dictionary in Fig. 1, and inferring the exact values of θ_{μ} tells *how* they matter.

A common alternative model of probability distributions uses $x = 2\sigma - 1 = \pm 1$ instead of $\sigma = (0, 1)$. A third order term coupling, for example, $\sigma_i \sigma_j \sigma_k$ represents a combination of first, second, and third order terms in the corresponding x s, and vice versa. Thus which words are codewords may depend on the parameterization used, but the longest codewords and nonoverlapping groups of codewords remain the same in both parameterizations. Our choice of σ vs x is for a practical reason: a codeword in the σ basis does not contribute to P unless *all* of its constituent bins are nonzero. Thus since spikes are rare, we do not need to consider contributions of very long words to the code.

We would like to investigate the neural dictionary systematically and to avoid pitfalls of other methods that arbitrarily truncate Eq. (1) at some order of interactions. For this, we define indicator variables $s_{\mu} = (0, 1)$, $\mu = 1, \dots, 2^{N+1} - 1$, which denote if a particular sequence of $\sigma_i = 1$, $i \in V_{\mu}$, and $\sigma_i = 0$, $i \notin V_{\mu}$, "matters" (is a putative word in the dictionary), i. e., it is either statistically significantly over- or under-represented in the data set compared to a null model (which we define later). In

other words, we rewrite

$$\log P(\boldsymbol{\sigma}|\vec{\theta}) = -\log Z + \sum_{\mu} \theta_{\mu} s_{\mu} \prod_{i \in V_{\mu}} \sigma_i. \quad (2)$$

We then choose a prior on θ_{μ} and on s_{μ} . As mentioned above, detecting words that are severely anomalously represented is easy, and it is not our focus. Instead, since many neurons control the muscles and hence the behavioral output, we assume that individual neuron can only have a very weak effect on the motor behavior. In other words, typically $|\theta_{\mu}| \ll 1$. We thus work in the strong regularization limit and impose priors

$$\mathcal{P}(\theta_{\mu}|s_{\mu} = 1) \propto \exp \left[-\frac{1}{2} \varepsilon (\theta_{\mu} - \theta_{\mu}^*)^2 \right], \quad \varepsilon \ll 1. \quad (3)$$

Note that the prior distribution $\mathcal{P}(\theta_{\mu}|s_{\mu} = 0)$ is irrelevant since, for $s_{\mu} = 0$, θ_{μ} does not contribute to $P(\boldsymbol{\sigma}|\vec{\theta}, \vec{s})$. We then choose $\bar{\theta}_{\mu}$ in a way such that the *a priori* and the empirical averages of individual σ_i s are equal, $\langle \sigma_i \rangle = \bar{\sigma}_i$ (we always use $\langle \dots \rangle$ and $\langle \dots \rangle$ to denote *a priori* and *a posteriori* expectations). This is equivalent to saying that we constrain our model to reproduce the firing rate of neurons and the frequency of the behavior. In typical problems, such marginal expectations are, indeed, well-known, and it is the higher order interaction terms, the complex words in the dictionary, that make the reconstruction hard. Finally, we choose $P(s_{\mu} = 1) = P(s_{\mu} = 0) = 0.5$, so that *a priori* a particular codeword has a fifty-fifty chance of being included in the neural dictionary.

Since we are only interested in whether a spike word is important in predicting the behavior and not how important it is, we integrate out all θ_{μ} having observed M samples of the $N + 1$ dimensional vector $\boldsymbol{\sigma}$. We perform this integration to the quadratic order in ε to get the posterior probability of indicator variables (see *Online Methods*):

$$P_{\varepsilon}(\mathbf{s}|\{\boldsymbol{\sigma}\}) = \frac{1}{\mathcal{Z}(\varepsilon)} \exp \left[\varepsilon \sum_{\mu} h_{\mu}(\{\boldsymbol{\sigma}\}) s_{\mu} + \varepsilon^2 \sum_{\mu\nu} J_{\mu\nu}(\{\boldsymbol{\sigma}\}) s_{\mu} s_{\nu} \right]. \quad (4)$$

This is of a familiar pairwise Ising form⁴⁴, with data-dependent magnetic fields h_{μ} and the couplings $J_{\mu\nu}$. Note that this Ising model has 2^{N+1} spins, replacing the model with $N + 1$ spins with higher order interactions in Eq. (1). This is naively a much harder problem. However, since most of the 2^{N+1} words do not appear in the actual data, and because of the ε^2 in front of the pairwise coupling term, evaluating posterior expectations $\langle s_{\mu} \rangle$ for all word that actually occur is relatively easy (see *Online Methods*). Knowing $\langle s_{\mu} \rangle$ allows us to tell which specific words should enter the neural-behavioral dictionary with a high *a posteriori* probability. To interpret Eq. (4), we notice (see details in *Online Methods*) that the linear terms $h_{\mu}(\{\boldsymbol{\sigma}\})$ bias the indicator variables: those words that over- or under-occur (are anomalously represented) in the data relative to their expected frequency in the null model that matches $\langle \sigma_i \rangle$ will have $h_{\mu} > 0$. In their turn, the couplings $J_{\mu\nu}$ are typically negative for those words that frequently co-occur, which includes words that at least partially overlap. Thus in this Ising model, words compete to explain certain correlations in data. Once a word has a large evidence (the h_{μ} term), it suppresses all other correlated words that explain the data in a weaker way, bringing the dictionary closer to the irreducible form. To verify that our analysis can, indeed, recover neural dictionaries and to set various adjustable parameters involved in the method, we successfully tested the approach on synthetic data that are statistically similar to our neural recordings (see *Online Methods*). In particular, for the data we have, we expect false discovery rate of < 1 codeword per neuron, but the fraction of the words that we discover depends on the (unknown) typical values of θ s.

Statistical properties of neural motor codes

Figure 3 illustrates the occurrence of two specific codewords found by uBIA that encode high-pitch renditions of syllables. Note that these codewords are, indeed, overrepresented together with the high pitch vocalizations. Analyzing if a particular word is correlated with an acoustic feature is, of course, not hard. However, detecting words that should be tested, without a multiple hypothesis testing significance penalty is nontrivial. Thus the power of uBIA comes from being able to systematically analyze abundances of *combinatorially many* such spike patterns, and further to identify which of them are *irreducibly* over- or under-represented. Figure 4 illustrates statistical properties of entire neural-behavior dictionaries discovered by uBIA for different songbird premotor neurons and for three features of the acoustic behavior. While we reconstruct the dictionaries that include all irreducible words, including those that have only anomalous firing patterns, here we primarily focus on codewords, that is the words that relate behavior to the neural activity. We do the analysis twice, first for behavior binarized as $\sigma_0 = 1$ for the above-median acoustic features, and then for the below-median acoustic features. We then combine the results. Note that the same pattern of spikes should not be simultaneously over- or under-represented when studying both the above and the below median codes, since the pattern cannot code for two mutually exclusive features. There were 0.7 such codewords on

Figure 3. Sample multispike codewords. A: Probability distribution an acoustic parameter (fundamental frequency, or pitch). For this analysis, we consider the output to be $\sigma_0 = 1$ when the pitch is above median (blue), and zero otherwise (red). B: Distribution of two sample codewords (a two-spike word in the top raster, and a three-spike word in the bottom raster) conditional on pitch. In each raster plot, a column represents 40 ms of the spiking activity preceding the syllable, with a grey tick denoting a spike. Every time a particular pattern is observed, its ticks are plotted in black. Note that these two spike words are codewords since they are overrepresented for above-median pitch (blue box) compared both to the null model based on the marginal expectation of individual spikes, and to the presence of the patterns in the low pitch region. Labels (a) and (b) identify these patterns in Fig. 4B.

Figure 4. Statistical properties of neural dictionaries. A: Sample neural-behavioral dictionaries for two neurons from two different birds (columns) and for three different acoustic features of the song (rows: pitch, amplitude, and the spectral entropy). The light gray curve in the background and the vertical axis corresponds to the probability of neural firing in each 2 ms bin (the firing rate). The rectangular ticks represents the timing of spikes in neural words that predict the acoustic features. For example, a two spike word with ticks at points $t = i, j$ corresponds to the probability that the word $\mu = (i, j)$ is a codeword for the acoustic feature with a probability statistically significantly higher than $1/2$. Codewords for high (low) output, i. e., $\sigma_0 = 1$ above (below) the median, are shown in blue (red). Full (empty) symbols correspond to over(under)-occurrence of the codeword-behavior combinations compared to the null model. Finally full (empty) black symbols represent words that over(under)-occur in the blue code and under(over)-occur in the red code. Words labeled (c)-(g) are also shown in (B). B: Frequency of occurrence of statistically significant codewords for different acoustic features in different neurons. Only first 200 codewords shown for clarity. Plotting conventions same as in (A), and letters label the same codewords as in (A) and in Fig. 3B. C: Proportion of m -spike codewords found in the dictionaries analyzed. An m -spike word corresponds to an $(m + 1)$ -dimensional word in the neural-behavioral dictionary. We note that about a third of all analyzed dictionaries are empty, so that those that have words in them typically have more than illustrated here. D: Mean number of significant codewords, averaged across all neurons and acoustic features. An average neuron has 5.6 codewords in our dataset, of which 3.1 code for the pitch, 2.5 for the amplitude, and 2.8 for the spectral entropy, with the number of words coding for pairs of features or for all three of them indicated by the overlap of rectangles in the Venn diagram. For comparison, our estimated false discovery rate is 0.3 words, so that only ~ 0.3 spurious words are expected to be discovered in each individual dictionary. We note that about a third of all analyzed dictionaries are empty, so that those that have words in them typically have more than illustrated here. E: Mean inter-spike interval (ISI) for the codewords (spike words that code for behavior) vs. all spike words that are significantly over- or under-represented, but do not code for behavior. Averages in each of the four analyzed birds are shown, illustrating that the ISI statistics of the coding and non-coding words are different, but the differences themselves vary across the birds. Star denotes 95% confidence. Other properties of the dictionaries (mean number of spikes in codewords, fraction of codewords shared by three vocal features, proportion of under/over-occurring codewords), do not differ statistically significantly across the birds.

average per dictionary. This is consistent with the expected false discovery rate of about 0.3 codewords per neuron for data sets of our size and statistical properties.

The most salient observation is that the inferred codewords consist of present or absent spikes in specific 2 ms time bins. This is consistent with previous analysis³⁵, which identified the same timescale for this dataset by analyzing the dependence of the mutual information between the activity and the behavior on the temporal resolution, but was unable to detect the specific words that carry the information. The second crucial observation is that most of codewords are composed of multiple spikes, representing an orthographically complex *pattern* timing code³⁷, in contrast to single spike timing codes, such as in⁴⁵. Large number of codewords of 2 or more spikes (and thus 3 or more features, including the behavior itself) suggests that analyzing these dictionaries with the widely-used lower order MaxEnt or GLM methods that typically focus on lower-order statistics (see *Online Methods*) would miss their significant components. Our third crucial observation is that very few sub-words / super-words pairs occur in the dictionaries, cf. Fig. 4 (e. g., the second codeword coding for entropy in neuron 2 in the panel A is a subword of the others). This indicates that uBIA fulfills its goal of rejecting multiple correlated explanations for the same data.

We quantify these observations as follows. In the 49 different datasets, the average size of a dictionary is 14. Of these words, on average 5.6 include the behavioral feature and hence are *codewords*, cf. Fig. 4(D). That there are so many specific temporally precise codewords suggests that the behaviorally-relevant spike timing patterns are the rule, rather than the exception, in this dataset. We found that 66% of codewords are unique to one of the three analyzed acoustic features. This further quantifies the observation that some neurons in RA are *selective* for specific acoustic features, as noted previously in⁴¹. Across all neurons and all acoustic features, only 15% of codewords consist of a single spike (or absence of spike), while 58%, 23%, and 4% consist of two, three, and four spikes respectively, cf. Fig. 4(C) (we are likely missing many long codewords, especially with small θ 's due to undersampling, see *Online methods*). This observation is consistent across all neurons and acoustic features, again indicating that coding by temporally precise spike patterns is a rule and not an exception.

At the same time, the observed dictionaries are quite variable across neurons and the production of particular song syllables. Codewords are built by stitching together multiple spikes or spike absences, and individual spikes occur at certain time points in the $(-40, 0)$ ms window with different probabilities in different neurons and syllables (i. e., the firing rate is both time and neuron dependent, cf. Fig. 4(A), grey lines). Codewords are likely to occur where the probability of seeing a spike in a bin is $\sim 50\%$, since these are the times that have more capacity to transmit information. Thus variability in firing rates as a function of time across neurons necessarily creates variability in the dictionaries across these neurons. Beyond this, we observe additional variability among the dictionaries that is *not* explained by the rate fluctuations. For example, we can differentiate one of the four birds from two of the others just by looking at the proportions of high-order codewords (an average of 0.21 bits in Jensen-Shannon divergence between the target bird and the rest), which means that we need around five independent samples/codewords to distinguish this bird from the others). Further, the mean inter-spike interval (ISI) for codewords is different from that of other words in the dictionaries, and this difference is also bird-dependent, see Fig. 4(E).

Verification of the inferred dictionaries

To show that the dictionaries we decoded are biologically (and not just statistically) significant, we verify whether the codewords can, in fact, be used to predict the behavioral features. For this, we built two logistic regression models that relate the neural activity to behavior. The first one uses the presence / absence of spikes in individual time bins and the second the presence / absence of the uBIA detected codewords as predictor variables (see *Online Methods*). Note that the individual spikes model is still a precise-timing model, which has 20 predictors (20 time bins, each 2 msec long), and hence one may expect it to predict better than the codewords model, which typically has many fewer predictors. To account for the possibility of overfitting, in all comparisons we test the predictive power of models using cross-validation. We emphasize that we do not expect either of the two models to capture an especially large fraction of the behavioral variation. Indeed, Tang et al.³⁵ have shown that, at 2 ms resolution, on average, there is only about 0.12 bits of information between the activity of an individual neuron and the behavioral features, and the assumption behind our entire approach is that none of individual predictors have strong effects. Further, a specific model, such as logistic regression, will likely recover even less predictive power from the data. With this, Supplementary Fig. 8 compares prediction between the two models, obtaining a significantly higher accuracy and a lower mean cross-entropy between the model and the data for the models that use codewords as predictors. In other words, the higher order, but smaller, dictionaries decoded by uBIA outperform larger, non-specific dictionaries in predicting behavior.

Dictionaries for exploratory vs. typical behaviors

Bengalese finches retain the ability to learn through their lifetimes, updating their vocalizations based on the sensorimotor feedback^{40,46,47}. A key element of this lifelong learning capacity is the precise regulation of vocal variability, which songbirds use to explore the space of possible motor outputs, cf. Fig. 5A,B. For example, male songbirds minimize variability when singing to females during courtship, but dramatically increase the range of variability in acoustic features such as pitch when singing alone^{48,49}. The variability is controlled by the activity of nucleus LMAN. Silencing or lesioning LMAN reduces the

Figure 5. Codes for vocal motor exploration. **A:** Distribution of syllable pitch relative to the mean for exploratory and performance behaviors (blue, intact birds, vs. grey, LMAN-lesioned animals, see main text). **B:** Ratio of the histograms in (A) evaluated in the quintiles of the exploratory (blue) distribution centered around [10%, 20%, ..., 90%] points. **C:** Total number of codewords when considering the vocal output as 1 if it belongs to a specific 20-percentile interval of the output distribution, and 0 otherwise. We observe that there are significantly more codewords for the exploratory behavior (tails of the distribution compared to the middle intervals). Notice that the shape of the curves parallels that in (B), suggesting that exploration drives the diversity of the codewords. **D:** Number of codewords when considering the vocal output as 1 if it belongs to a 33-percentile (non-overlapping) interval of the output distribution, and 0 otherwise. Here there are significantly more codewords when coding for high pitch. Further, the codewords found for each of the three intervals are mostly multi-spike (histograms show the distribution of the order of the codewords for each percentile interval). **E:** For codewords for the 33-percentile intervals, we compare the mean inter-spike intervals (ISIs). Codewords for high outputs (especially for pitch and spectral entropy) have a significantly larger mean ISI. **F:** We illustrate dictionaries of two neurons for the medium and the high spectral entropy ranges. Notice that the high entropy range has significantly more codewords.

acoustic variance of undirected song (Fig. 5A) to a level approximately equal to that of female-directed song^{50,51}. Using uBIA, we can ask for the first time whether the statistics of codewords controlling the exploratory vs. the baseline range of motor variability are different. To do this, we analyze the statistics of codewords representing different parts of the pitch distribution. First, we define the output as $\sigma_0 = 1$ if the behavior belongs to a specific 20-percentile interval ([0 – 20], [10 – 30], ..., [80 – 100]) and compare the dictionaries that code for behavior in each of the intervals. We find that there are significantly more codewords for exploratory behaviors (percentile intervals farthest from the median, cf. Fig. 5C). This holds true for different features of the vocal output, though the results are only statistically significant if pooled over all features. To improve statistical power by increasing the number of trials in each acoustic interval, we also consider a division of the output into three equal intervals: low, medium, and high. In this case, there are still more codewords for the high exploratory pitch, and the dictionaries for each of the intervals are still multispike, cf. Fig. 5D. We further observe that the codewords themselves are different for the three percentile groups: the mean ISI of high pitch, amplitude, and spectral entropy codewords is higher, with the largest effect coming from the pitch and the spectral entropy, cf. Fig. 5E. Examples of typical and exploratory dictionaries are illustrated in Fig. 5F.

These findings challenge common accounts of motor variability, in songbirds and other systems, that motor exploration is induced by adding random spiking variations to a baseline motor program. Rather, the over-abundance of codewords in the exploratory flanks of the acoustic distributions indicates that the mapping between the neural activity and the behavior is more reliable than in the bulk of the behavioral activity: multiple renditions of the same neural command result in the same behaviors more frequently, making it easier to detect the codewords. One possibility is that the motor system is less biomechanically noisy for large behavioral deviations. This is unlikely due to the tremendous variation in the acoustic structure (pitch, etc.) of different song syllables within and across animals^{47,52}, which indicates that songbirds can produce a wide range of sounds and that particular pitches (i.e., those at at one syllable's exploratory tail) are not intrinsically different or harder for an animal to produce. Similarly, songbirds can dramatically modify syllable pitch in response to manipulations of auditory feedback^{40,47}. A more likely explanation for the greater prevalence of codewords in the exploratory tails is that the nervous system drives motor exploration by selectively introducing particular patterns into motor commands that are specifically chosen for their reliable neural-to-motor mapping. This would result in a more accurate deliberate exploration and evaluation of the sensory feedback signal, which, in turn, is likely to be useful during sensorimotor learning⁵³.

Discussion

In this work, we developed the unsupervised Bayesian Ising Approximation as a new method for reconstructing biological dictionaries — the sets of anomalously represented joint activities of multiple components of biological systems. Inferring these dictionaries directly from data is a key problem in many fields of modern data-rich biological and complex systems research including systems biology, immunology, collective animal behavior, and population genetics. Our approach addresses crucial shortcomings that so far have limited applicability of other methods. First, it does not limit the possible dictionaries, either by considering words of only limited length or of a pre-defined structure. Instead we perform a systematic analysis through all possible words that occur in the data sample. Second, it promotes construction of irreducible dictionaries, de-emphasizing related, co-occurring words. Further, uBIA does not make assumptions about the linear structure of dependencies unlike various linear methods.

To illustrate capabilities of the method, we applied it to analysis of motor activity in cortical area RA in a songbird. We were able to infer statistically significant codewords from large-dimensional probability distributions (2^{21} possible different words) with relatively small data sets ($\sim 10^2 \dots 10^3$ samples). We verified that the codewords are biologically meaningful, in the sense that they predict behavioral features better than alternative approaches. Importantly, most of words in hundreds of dictionaries that we reconstructed were more complex than is usually considered, involving multiple spikes in precisely timed patterns. The multi-spike, precisely timed nature of the codes was universal across individuals, neurons, and acoustic features, while details of the codes (e.g., specific codewords and their number) showed tremendous variability.

Further, we identified codewords that correlate with three different acoustic features of the behavior (pitch, amplitude, and spectral entropy), and different percentile ranges for each of these acoustic features. Across many of these analyses, various statistics of codewords predicting exploratory vs. typical behaviors were different. Specifically, the exploratory dictionaries contained more codewords than the dictionaries for typical behavior, suggesting that the exploratory spiking patterns are more consistently able to evoke particular behaviors. This is surprising since the exploratory behavior is usually viewed as being noisier than the typical one. Crucially, exploration is a fundamental aspect of sensorimotor learning^{40,46,54}, and it has been argued that large deviations in behaviors are crucial to explaining the observed learning phenomenology⁵³. However, the neural basis for controlling exploration vs. typical performance is not well understood. Intriguingly, vocal motor exploration in songbirds is driven by the output of a cortical-basal ganglia-thalamo-cortical circuit, and lesions of the output nucleus of this circuit (area LMAN) abolishes the exploratory (larger) pitch deviations^{50,51}. Our findings therefore suggest that the careful selection of the spike patterns most consistently able to drive behavior may be a key function of basal ganglia circuits.

While the identified codewords are statistically significant, and we show that they can predict the behavior better than larger, but non-specific features of the neural activity, a crucial future test of our findings will be in establishing their *causal* rather than merely correlative nature by means of stimulating neurons with patterns of pulses mimicking the identified codewords³⁶. This will be facilitated by the speed of our method, which can reconstruct dictionaries in real time on a laptop computer. Additional future work can explore how population-level dictionaries are built from the activity of individual neurons, how the dictionaries develop and are modified in development, and whether the structure of dictionaries as a whole can be predicted from various biomechanical and information-theoretic optimization principles. Finally, one needs to understand how these dictionaries are implemented in the recurrent dynamics of neural networks in animals' brains.

Online Methods

Overview of neural decoding methods

For many different experimental systems, it has been possible to measure the information content of spike trains^{34–36}, but the question of decoding — which spike patterns carry this information? — has turned out to be a harder one. Multiple approaches have been used to address this problem and to reconstruct neural dictionaries, whether in the context of sensory or motor systems, starting with linear decoding methods⁴⁵. All have fallen a bit short, especially in the context of motor codes, where an animal is free to perform any one of many behaviors it wishes, and hence statistics are usually poor. A leading method is Generalized Linear Models (GLMs)^{19,55,56}, which encode the rate of spike generation from a certain neuron at a certain time as a nonlinear function of a linear combination of past stimuli (sensory systems) or of future motor behavior (motor systems) and the past spiking activity of a neuron and its presynaptic partners. GLM approaches can detect the importance of the timing of individual spikes and sometimes interspike intervals for information encoding, but generalizations to detect importance of higher order spiking patterns are not yet well established. Another common approach is based on maximum entropy (MaxEnt) models^{2,57,58}. These replace the true distribution of the data with the least constrained (i. e., maximum entropy) approximation consistent with low-order, well-sampled correlation functions of the distribution. The approach is computationally intensive, especially when higher order correlations are constrained by data. At the same time, to approximate empirical distributions well, a large number of such constraints is often required. This needs very large datasets, especially if one is interested in relating the neural activity to the external (behavioral or sensory) signals. Such large datasets are hard to obtain in the motor control setting.

More recently, feed-forward and recurrent artificial neural network approaches have been used to decode large-scale neural activity^{59,60}, but these have focused primarily on neural firing rates over large (tens of milliseconds) temporal windows. As a result, to date, there have not been successful attempts to reconstruct neural dictionaries from data, which would (i) resolve spike timing in words of the dictionary to a high temporal resolution, (ii) be comprehensive and orthographically complex, not limiting the words to just single spikes or pairs of spikes, and (iii) discount correlations among spiking words to produce irreducible dictionaries that only detect those codewords that cannot be explained away by correlations with other words in the dictionary.

Details of the Bayesian Ising Approximation approach

To obtain Eq. (4) for the posterior probability of including a word into the dictionary, we start with

$$P(\mathbf{s}|\boldsymbol{\sigma}) \propto p(\boldsymbol{\sigma}|\mathbf{s}) = \int d\boldsymbol{\theta} p(\{\boldsymbol{\sigma}|\boldsymbol{\theta}, \mathbf{s}\}) \prod_{\mu|s_{\mu}=1} p_{\epsilon}(\theta_{\mu}|s_{\mu}). \quad (5)$$

Now we make two approximations. First, we evaluate the integral in Eq. (5) using the saddle point approximation around the peak of the *prior*, $\boldsymbol{\theta}^*$. This is a low signal-to-noise limit, and it is different from most high signal-to-noise approaches that analyze the saddle around the peak of the *posterior*. Second, we do all calculations as a Taylor series in the small parameter ϵ (see below on the choice of ϵ). Both approximations are facets of the same strong regularization assumption, which insists that most coupling constants θ_{μ} are small. Following Fisher and Mehta²⁵, we obtain Eq. (4), where the magnetic fields (biases) h_{μ} and the exchange interactions $J_{\mu\nu}$ are

$$h_{\mu}(\boldsymbol{\sigma}) = \frac{1}{2} \left[\frac{\partial^2 \mathcal{L}}{\partial \theta_{\mu}^2} + \left(\frac{\partial \mathcal{L}}{\partial \theta_{\mu}} \right)^2 \right] \bigg|_{\boldsymbol{\theta}^*}, \quad (6)$$

$$J_{\mu\nu}(\boldsymbol{\sigma}) = \frac{1}{4} \left[\frac{\partial^2 \mathcal{L}}{\partial \theta_{\mu} \partial \theta_{\nu}} \left(\frac{\partial^2 \mathcal{L}}{\partial \theta_{\mu} \partial \theta_{\nu}} + 2 \frac{\partial \mathcal{L}}{\partial \theta_{\mu}} \frac{\partial \mathcal{L}}{\partial \theta_{\nu}} \right) \right] \bigg|_{\boldsymbol{\theta}^*},$$

where $\mathcal{L} = \log P(\boldsymbol{\sigma}|\boldsymbol{\theta})$ is the log-likelihood (see Fig. 6 for a geometric interpretation of the field h_{μ}). Plugging in the model of the probability distribution, Eq. (1), we get for the fields and the exchange interactions

$$h_{\mu} = \frac{M^2}{2} \left[(\overline{\boldsymbol{\sigma}}_{\mu} - \langle \boldsymbol{\sigma}_{\mu} \rangle)^2 - \frac{\text{var}(\boldsymbol{\sigma}_{\mu})}{M} \right], \quad (7)$$

$$J_{\mu\nu} = \frac{M^2}{4} \text{cov}(\boldsymbol{\sigma}_{\mu}, \boldsymbol{\sigma}_{\nu}) \left[\text{cov}(\boldsymbol{\sigma}_{\mu}, \boldsymbol{\sigma}_{\nu}) - 2M(\overline{\boldsymbol{\sigma}}_{\mu} - \langle \boldsymbol{\sigma}_{\mu} \rangle)(\overline{\boldsymbol{\sigma}}_{\nu} - \langle \boldsymbol{\sigma}_{\nu} \rangle) \right]. \quad (8)$$

Here, to simplify the notation, we defined $\boldsymbol{\sigma}_{\mu} \equiv \prod_{i \in V_{\mu}} \sigma_i$. Further, angular brackets, cov, and var denote the *a priori* expectations, covariances, and variances of frequencies of words in the null model, which matches frequency of occurrence of each individual σ_i (probability of firing in every time bin for the songbird data). Similarly, overlines denote the empirical counts or correlations between co-occurrences of words in the observed data. Specifically, denoting by n_{μ} the marginal frequencies of the word V_{μ} in the data, these expectations and frequencies are defined as follows:

$$\overline{\boldsymbol{\sigma}}_{\mu} = \frac{1}{M} \sum_{m=1}^M \left(\prod_{i \in V_{\mu}} \sigma_i^{m} \right) = \frac{n_{\mu}}{M}, \quad (9)$$

$$\langle \boldsymbol{\sigma}_{\mu} \rangle = \prod_{i \in V_{\mu}} \langle \sigma_i \rangle = \prod_{i \in V_{\mu}} \frac{n_i}{M}, \quad (10)$$

$$\text{var}(\boldsymbol{\sigma}_{\mu}) = \langle \boldsymbol{\sigma}_{\mu}^2 \rangle - \langle \boldsymbol{\sigma}_{\mu} \rangle^2 = \langle \boldsymbol{\sigma}_{\mu} \rangle (1 - \langle \boldsymbol{\sigma}_{\mu} \rangle) = \prod_{i \in V_{\mu}} \frac{n_i}{M} \left(1 - \prod_{i \in V_{\mu}} \frac{n_i}{M} \right), \quad (11)$$

$$\text{cov}(\boldsymbol{\sigma}_{\mu}, \boldsymbol{\sigma}_{\nu}) = \langle \boldsymbol{\sigma}_{\mu} \boldsymbol{\sigma}_{\nu} \rangle - \langle \boldsymbol{\sigma}_{\mu} \rangle \langle \boldsymbol{\sigma}_{\nu} \rangle = \prod_{k \in V_{\mu} \cup V_{\nu}} \frac{n_k}{M} - \left(\prod_{i \in V_{\mu}} \frac{n_i}{M} \right) \left(\prod_{j \in V_{\nu}} \frac{n_j}{M} \right), \quad (12)$$

To derive these equations, note that $\sigma_i^2 = \sigma_i$. Note also that $\text{cov}(\boldsymbol{\sigma}_{\mu}, \boldsymbol{\sigma}_{\nu}) = 0$ if the intersection of V_{μ} and V_{ν} is empty.

Equation (7) has a straightforward interpretation, which we illustrate in Fig. 6. Specifically, if the difference between the *a priori* expected frequency and the empirical frequency of a word is statistically significantly nonzero (compared to the *a*

Figure 6. Geometric interpretation of the fields h_μ in the uBIA method, in relation to the log-likelihood function $\mathcal{L}(\{\sigma\}|\theta)$. The uBIA method makes an approximate guess (dashed line in left panel) of how much in log-likelihood $\Delta\tilde{\mathcal{L}}_\mu$ we would win by fitting a parameter θ_μ , and how far in parameter space we would need to go, $\Delta\tilde{\theta}_\mu$ (see left panel). The sign of the field only depends on the improvement in log-likelihood, being positive beyond a threshold (inclusion of a word). This complexity penalization comes from the Bayesian approach in this strong regularized regime. On the other hand, the farther we go in parameter space, the smaller in absolute value the field becomes (see right panel).

priori standard error), then the corresponding word is anomalously represented. It does not matter whether the word is over- or under-represented: in either case, if the frequency deviated from the expectation, then the field h_μ is positive, biasing the indicator s_μ towards 1, and hence towards inclusion of the word in the dictionary. If the frequency is as expected, then the field is negative, and the indicator is biased towards 0, excluding the word from the dictionary. Note that as M increases, the standard error goes down, and the field generally increases, allowing us to consider more words. The sign of θ_μ would determine whether the word is over- or underrepresented. However, estimating the exact value of θ_μ from small datasets is often impossible and is not our goal, even though, in Fig. 2, we denote words as under- or over-represented by whether their empirical frequency is smaller or larger than the *a priori* expectation. Thus in some aspects, our approach is similar to the previous work⁶¹, where multi-neuronal patterns are found by comparing empirical firing probabilities to expectations. However, we do this comprehensively for *all* patterns that occur in data, and we account for reducibility of the dictionaries (also see below).

The exchange interactions $J_{\mu\nu}$ are also interpretable. As explained above, correlations among words are a serious problem. Indeed, for example, a word $\sigma_0\sigma_i\sigma_j = 1$ may occur too frequently simply because its sub-words $\sigma_0\sigma_i = 1$, $\sigma_i\sigma_j = 1$, or $\sigma_0\sigma_j = 1$ are common. Alternatively, the word may be frequent because, in its turn, it is a sub-word of a larger common word, for example, $\sigma_0\sigma_i\sigma_j\sigma_k = 1$. In GLMs, resolving these overlaps requires imposing sparsity or other additional constraints. In contrast, the couplings $J_{\mu\nu}$ address this problem for uBIA naturally and computationally efficiently. Notice that because of the factor 2 in Eq. (8), the exchange interactions are predominantly negative if one expects the two studied words to be correlated, and if they co-occur in the empirical data as much as they are expected to co-occur in the null model because of the overlaps in their composition, V_μ and V_ν . Negative $J_{\mu\nu}$ s implement a “winner-take-all” mechanism, where statistical anomalies in data that can be explained, in principle, by many different θ_μ s are attributed predominantly to one such θ_μ that explains them the best. On the other hand, the exchange interactions are positive if one expects correlations between the words *a priori*, but does not observe them. Thus, in principle, a word can be included in the dictionary even at zero field h_μ .

Knowing the coefficients h_μ and $J_{\mu\nu}$, one can numerically estimate $\langle s_\mu \rangle$, the posterior expectation for including a word V_μ in the dictionary. Generally, finding such marginal expectations from the joint distribution in disordered systems is a hard problem. However, here $h_\mu \propto \varepsilon$ and $J_{\mu\nu} \propto \varepsilon^2$, so that the fields and the interactions create small perturbations around the “total ignorance” solution, $\langle s_\mu \rangle = 1/2$ (this is a manifestation of our general assumption that none of the words is very easy to detect). Therefore, we calculate the marginal expectation using fast mean field techniques⁶². We use the *naïve* mean field approximation, which is given by self-consistent equations for the posterior expectations in terms of the magnetizations $m_\mu = 2\langle s_\mu \rangle - 1$,

$$\tanh^{-1}(m_\mu(\varepsilon)) = \frac{\varepsilon}{2} \left[h_\mu + \varepsilon \sum_{\nu} J_{\mu\nu} + \frac{\varepsilon}{2} \sum_{\nu} J_{\mu\nu} m_\nu(\varepsilon) \right] \quad (13)$$

$$= \frac{1}{2} [\varepsilon h_\mu + \varepsilon^2 h_\mu^{\text{eff}}(\varepsilon)], \quad (14)$$

so that interactions among spins are encapsulated in an effective field $\varepsilon h_\mu^{\text{eff}}$. We solve Eq. (13) iteratively²⁵, by increasing ε from 0—that is, from the total ignorance $\langle s_\mu \rangle = 1/2$ or $m_\mu = 0$ —and up to the limiting value ε_{max} in steps of $\delta\varepsilon = M^{-1}/20$. This limiting value ε_{max} is determined by the two approximations involved in the strong regularization assumption. First, the

saddle point approximation around the peak of the prior in Eq. (5) implies that the characteristic width of the prior should be smaller than that of the likelihood, $\varepsilon \leq \varepsilon_1 = 1/M$. Second, the Taylor series up to second order in ε for the posterior of the indicator variables implies that the quadratic corrections should not be larger than the linear terms. Within the mean field approximation, this means that $\langle |h_\mu| \rangle_\mu \geq \langle |\varepsilon h_\mu^{\text{eff}}(\varepsilon)| \rangle_\mu$, which is saturated at some ε_2 (notice that, in contrast to our usual notation, the averages here are over the indices, and not the data). Thus overall we take $\varepsilon_{\text{max}} = \min\{\varepsilon_1, \varepsilon_2\}$.

Additionally we have used the TAP equations⁶², instead of Eq. (13) to calculate magnetizations. These are more accurate since they account for how a spin affects itself through its couplings with the other spins. However, corrections due to this more complicated method were observed to be negligible in our strong regularized regime, since they were of higher order in $\varepsilon \ll 1$. Thus all results that we report are based on the mean field estimation.

Effect of absent words

Of the exponentially many possible words, majority will not happen in a realistic data set. In particular, this includes most of long words. At the same time, *a priori* expectations for the frequency of such words, Eq. (10), decrease exponentially fast with the word length. Thus the fields, Eq. (7), for the words that do not occur are small, and the posterior expectation for including these words in the dictionary is $\langle s_\mu \rangle \approx 1/2$, so that we do not need to analyze them explicitly. A bit more complicated is the fact that all words affect each other's probability to be included in the dictionary through the exchange couplings $J_{\mu\nu}$, so that, in principle, the sum in the mean field equations, Eq. (14), is over exponentially many terms. Here we show that the effect of non-occurring words on the interaction terms is exponentially small in N , as long as the empirical averages $n_i/M \ll 1$.

To illustrate this, we start with the probabilities $p(\sigma_i = 1) = p_i$ of a single variable i being active. We then define the average such probability $q = N^{-1} \sum_i p_i$. Without the loss of generality, we assume $q < 1/2$, and otherwise we rename $\sigma_i \rightarrow 1 - \sigma_i$. Denoting a long word of a high order k that does not occur in the data as σ_ω , we have $n_\omega = 0$. Then the corresponding field is

$$h_\omega = \frac{M^2}{2} \left[(0 - \langle \sigma_\omega \rangle)^2 - \frac{\text{var}(\sigma_\omega)}{M} \right] \quad (15)$$

$$\sim -\frac{M}{4} q^k [1 - q^k M] \sim -\frac{M}{4} q^k. \quad (16)$$

Here we consider as *high order* words those, for which $q^k M \ll 1$ (in general, $\langle \sigma_\omega \rangle M = \langle n_\omega \rangle \ll 1$, which happens for $k \sim 4 \dots 5$ for our datasets). Then the magnetization is

$$m_\omega(k) \simeq \tanh\left(\frac{\varepsilon}{2} h_\omega\right) \sim -\frac{\varepsilon M}{8} q^k. \quad (17)$$

This illustrates our first assertion that none of these non-occurring words will be included in the dictionary. However, as a group, they may still have an effect on words of lower orders. To estimate this effect, for a word σ_μ of a low order k_0 , we calculate the effective field $\tilde{h}_\mu^{\text{eff}}$, which all of the non-occurring words σ_ω have on it. First we notice that, if V_μ and V_ω do not overlap, then their covariance is zero, and $J_{\mu\omega} = 0$. That is, only high order words that overlap with V_μ can contribute to $\tilde{h}_\mu^{\text{eff}}$. Since $\text{cov}(\sigma_\mu, \sigma_\omega) \sim q^k (1 - q^{k_0})$, the couplings are

$$J_{\mu\omega}(k) \sim \frac{M^2}{4} q^{2k} (1 - q^{k_0}) \times \mathcal{O}(1). \quad (18)$$

Using Eq. (13), this gives for the typical effective field that absent words have on the word μ

$$\tilde{h}_\mu^{\text{eff}}(\{\omega\} \rightarrow \mu) \simeq \varepsilon \sum_{k \gtrsim k_0}^N \mathcal{N}(k) J_{\mu\omega}(k) \left(1 + \frac{\varepsilon}{2} m_\omega(k)\right), \quad (19)$$

where the number of words of order k that overlap with σ_μ and can affect it is given by the combinatorial coefficient $\mathcal{N}(k) \simeq \binom{N-k_0}{k-k_0}$. This has a very sharp peak at $k = (N + k_0)/2$, where $\mathcal{N} \simeq 2^{N-k_0}$. We can approximate the sum in Eq. (19) as the argument of the sum evaluated at this peak $k = (N + k_0)/2$, obtaining an effective field coming from high order words

$$\tilde{h}_\mu^{\text{eff}}(\{\omega\} \rightarrow \mu) \sim \varepsilon 2^{N-k_0} \frac{M^2}{4} q^{N+k_0} (1 - q^{k_0}) \left[1 - \frac{\varepsilon^2 M}{16} q^{\frac{N+k_0}{2}}\right] \quad (20)$$

$$\sim \frac{\varepsilon M^2}{4} \left(\frac{q}{1/2}\right)^N \frac{q^{k_0}}{2} (1 - q^{k_0}) \quad (21)$$

$$\propto \left(\frac{q}{1/2}\right)^N. \quad (22)$$

In other words, even the combined effect of all higher order absent words is small if the average frequency of individual letters is smaller than $1/2$. We thus can disregard all non-occurring words in the mean field equations.

We stress that, for this to hold, the average of the binary variables σ_i must be small, $q = N^{-1} \sum_i p(\sigma_i = 1) < 1/2$. In our songbird dataset, this condition was fulfilled with $q \sim 0.2$. However, in 4% of cases the probability to have a spike in a certain time bin was $p_i > 1/2$. Thus to stay on the safe side, we performed additional analyses by redefining variables as $\sigma_i \rightarrow 1 - \sigma_i$ if the presence of a spike in a bin was $> 50\%$. In other words, in such cases, we defined the absence of the spike as 1 and the presence as 0. For our datasets, the findings did not change with this redefinition.

This previous analysis does not imply that absent words of high order are irrelevant — it only says that they cannot be detected with the available datasets. In the numerical implementation of the method, we filter out long absent words ω such that $\langle \sigma_\omega \rangle M = \langle n_\omega \rangle < 0.02$, with this cutoff determined by Eqs. (15-17), so that, for these words, $h_\omega \ll 1$. These words get assigned $1/2$ as the posterior probability of inclusion in the dictionary, and their contribution to the mean field equations is neglected. In contrast, if a word ω is absent but $\langle n_\omega \rangle \geq 0.02$, we include them in the analysis, Eq. (4). Such words may turn out to be relevant code words, especially if they happen a lot less frequently than expected *a priori*.

Synthetic data for testing and fine-tuning the method

To set the free parameters of our approach and to quantify its ability to reconstruct dictionaries, we test it on synthetic data sets that are similar to the songbird data, analyzed in this work. We use the log-linear model, Eq. (1), as a generative model for binary correlated observables σ with $N = 20$. We choose the individual biases in the generative model from a normal distribution, $\theta_i \sim \mathcal{N}(-0.7, 0.1^2)$, which matches the observed probability of a spike in a bin in the bird data. That is, $p(\sigma_i = 1) \simeq [1 + \exp(-2\theta_i)]^{-1} \sim q \sim 0.2$. Then we select which binary variables interact. We allow interactions of 2nd, 3rd, and 4th order, with an equal number of interactions per order. For different tests, we choose the interaction strengths from (a) the sum of two Gaussian distributions, one with a positive mean and the other with a negative one, $\text{mean}(\theta_\mu) = \pm 0.5$, $\text{std}(\theta_\mu) = 0.1$, and (b) from one Gaussian distribution centered at zero with $\text{std}(\theta_\mu) = 0.5$. Both choices reflect our strong regularization assumption, so that effects of individual variables on each other are weak, and a state of one variable does not determine the state of the others, and hence does not “freeze” the system. We are specifically interested in performance of the algorithm in the case where θ are distributed as the sum of Gaussians. On the one hand, this tests how the algorithm deals with data that are atypical within its own assumptions. On the other hand, this choice ensures that there are very few values of θ that are statistically indistinguishable from zero, making it easier to quantify findings of the algorithm as either true or false. We have additionally tested other distributions of θ , but no new qualitatively different behaviors were observed. Finally, for both types of distributions of θ , we also varied the density of interactions α , from $\alpha = 2$ to $\alpha = 4$, which spans the interaction densities of tree-like and 2D lattice-like networks.

Next, we generate M samples from these random probability distribution and we apply our pipeline to reconstruct the dictionary. We test on 400 distributions from each family. As the first step, we discard high-order words absent in the data using a threshold on the expected number of occurrences $\langle \sigma_\mu \rangle M = \langle n_\mu \rangle < 0.02$, as explained above. Next, we select $N_{\max}$ words that have the highest (absolute) values in magnetic field (we have tested $N_{\max} = 200, 500, 2000, 5000$, and finally use 500 after not observing differences). To decide which of these remaining words are to be included in the dictionary, we build the Ising model on the indicator variables, Eq. (4), with its corresponding magnetizations m_μ given by the mean field equations. We start from an inverse regularization strength of $\varepsilon = 0$ and then increase ε in steps of $\delta\varepsilon = 1/(20M)$, up to $\varepsilon_{\max} = \min\{\varepsilon_1, \varepsilon_2\}$, as detailed above. We then set the full false discovery rate n_{false} — the number of dictionary words that we allow ourselves to identify falsely from the fully reshuffled data, which must have zero words in it. Note that we do reshuffling while keeping the observed frequency of individual variables n_i constant. For the value of n_{false} , we identify the significance threshold for the magnetization, which allow this many false words to be accepted in the shuffled data, on average (from the same 400 distributions). Finally, we select as dictionary words those that have their marginal magnetizations m_μ above a significance threshold $m_\mu > m(n_{\text{false}})$.

We then repeat the analysis for different values of n_{false} , seeking the thresholds that minimize false negatives (exclusions of true words from the dictionary) and false positives (inclusion of spurious words into the dictionary) simultaneously. To measure these, we use two metrics: precision and recall. First, precision measures the fraction of the words included in the dictionary that are true, i. e., have a nonzero θ_μ in the generative model. Second, recall measures the fraction of the words in the generative model with $\theta_\mu \neq 0$ that were included in the dictionary. Results of the analysis are shown in Fig. 7B. Since data set sizes are relatively small, we do not expect to detect all words, especially in the case where θ_μ are allowed to be close to 0 in the generative model (Gaussian distributed). Thus we emphasize precision over recall in setting parameters of the algorithm: we are willing to not include words in a dictionary, but those words that we include should have a large probability of being true words in the underlying model. Our most crucial observation is that the precision-recall curves are remarkably stable with the changing density of interactions. Recall is smaller when interactions coefficients are taken from a Gaussian centered at zero. However, one could argue that missing words with very small θ_μ should not be considered a mistake: they are not significant words in the studied dictionary. We observe that by keeping the full false discovery rate n_{false} below 0.5 (only about half a

Figure 7. Results of the synthetic data analysis. Performance on synthetic data as a function of the density of interactions α , the distributions $p(\theta)$ for the strength of interactions, and the number of samples $M \in \{200, \dots, 1600\}$ (logarithmically spaced). The first and the second columns correspond to precision-recall curves for the different density of interactions (significant words) per variable, $\alpha \in \{2, 4\}$, within the true generative model. The top and the bottom rows corresponds to the interaction strengths θ selected from the sum of two Gaussian distributions, or a single Gaussian, as described in the text. For the first two columns, we vary the significance threshold in marginal magnetization $m(n_{\text{false}})$, such that the full false discovery rate on the shuffled data $n_{\text{false}} \in [0.005, 40]$. In the third column we show the value of n_{false} that corresponds to the precision of 80% as a function of M , so that the precision is larger than 80% in the shaded region. This region is quite large and overlaps considerably for the four cases analyzed, illustrating robustness of the method.

Figure 8. Prediction accuracy with uBIA dictionaries. We compare prediction of the behavior using logistic regression models that have as features (i) neural activity in all the time bins at 2 ms resolution versus (ii) only the detected relevant codewords. **A:** Scatter plot of accuracy of models of both types, evaluated using 2-fold cross-validation. Inset shows that the difference between the prediction is significant with $p < 0.01$ according to the paired t-test. **B:** Scatter plots of the mean cross-entropy between the data and the models for the two model classes. Inset: Even though the models that use the codewords are simpler (have fewer terms), they are able to predict better (with lower cross-entropy) according to the paired t-test.

word detected falsely, on average, in shuffled data), we can reach a precision as high as 80%, extracting 20% – 30% of the codewords depending on the number of samples, the distribution of θ , and α . We are thus confident that our method produces dictionaries, in which a substantial fraction of words correspond to true words in the data.

Testing the predictive power of the uBIA dictionaries

In this section, we test whether the codewords found in data from songbird premotor neurons can be used to predict the subsequent behavior. We compare two logistic regression models: one that uses the activity in the 20 time bins to predict the behavior and another that only uses as features the activity of the few relevant codewords, usually far fewer than 20. The features corresponding to the codewords are binary, and they are only active when all the time bins of such words are active. This means that the model using the time bins is more complex, as it already has all the information that the codewords model has and more, though it does not account for combinatorial effects of combining spikes into patterns. In order to properly test the predictive power between these two models with different complexity we perform 2-fold cross-validation, using a log-likelihood loss function. As is common in these cases, an L2 penalty is included to help the convergence to a solution (the models were implemented with the Classify function from Mathematica, whose optimization is done by the LBFGS algorithm). As shown by Tang et al.³⁵, not all neurons in our dataset are timing neurons, or even code for the behavior at all. Thus we restrict the comparison to those cases that have at least 4 codewords (27 case in total, with 10 codewords on average). Both of the logistic regression models have the following structure

$$p(y = 1 | \mathbf{z}, \boldsymbol{\beta}) = \frac{1}{1 + \exp(-\beta_0 - \sum_i \beta_i z_i)}, \quad (23)$$

where y corresponds to the behavior, and the features correspond to the time bins in one case ($z_i = x_i$) and to the codewords in the other ($z_i = \prod_{j \in V_i} x_j$), while β_i are the coefficients of the model. The loss function used is the log-likelihood with the L2 penalty,

$$\mathcal{L} = \sum_{m=1}^M \log p(y^{(m)} | \mathbf{z}^{(m)}, \boldsymbol{\beta}) - \frac{\lambda}{2} \sum_i \beta_i^2, \quad (24)$$

where M is the number of samples, and λ is the regularization strength. In our analysis, as different datasets have different number of samples, we show the results for the mean cross-entropy over the test data, which correspond to the normalized log-likelihood.

Tang et al.³⁵ showed that individual neurons on average carry around 0.12 bits at a 2 ms scale. So for both models we expect the prediction accuracy to be barely above chance, especially since we are focusing on a particular prediction model (a logistic

regression), and may be missing predictive features not easily incorporated in it. Figure 8a shows the scatter plot of accuracy in the 27 analyzed datasets, plotting the prediction using the time bins activity on the horizontal axis versus prediction using only the codewords activity on the vertical one. We observe that the models based on codewords are consistently better than the ones using all the 20 time bins, and the difference is significant (inset). We additionally evaluate the quality of prediction using the mean cross-entropy between the model and the data. Figure 8b shows that the models with the codewords have lower mean cross-entropies and thus generalize better (see Inset).

Software Implementation

The software implementation of uBIA is available from <https://github.com/dghernandez/decomotor>.

Data

The data used in this work is available from https://figshare.com/articles/Songbird_premotor_dictionaries/10315844.

References

1. Marks, D. S. *et al.* Protein 3d structure computed from evolutionary sequence variation. *PLOS One* **6**, e28766 (2011).
2. Schneidman, E., Berry, M. J., Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
3. Natale, J. L., Hofmann, D., Hernández, D. G. & Nemenman, I. Reverse-engineering biological networks from large data sets. In Munsky, B., Hlavacek, W. S. & Tsimring, L. (eds.) *Quantitative Biology: Theory, Computational Methods, and Models*, 213–246 (MIT Press, Cambridge, MA, 2018).
4. Margolin, A. A. *et al.* Aracne: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinforma.* **7**, S7 (2006).
5. Lezon, T. R., Banavar, J. R., Cieplak, M., Maritan, A. & Fedoroff, N. V. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proc. Natl. Acad. Sci. (USA)* **103**, 19033–19038 (2006).
6. Morcos, F. *et al.* Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc. Natl. Acad. Sci. (USA)* **108**, E1293–E1301 (2011).
7. Bitbol, A.-F., Dwyer, R. S., Colwell, L. J. & Wingreen, N. S. Inferring interaction partners from protein sequences. *Proc. Natl. Acad. Sci. (USA)* **113**, 12180–12185 (2016).
8. Halabi, N., Rivoire, O., Leibler, S. & Ranganathan, R. Protein sectors: evolutionary units of three-dimensional structure. *Cell* **138**, 774–786 (2009).
9. Otwinowski, J. & Nemenman, I. Genotype to phenotype mapping and the fitness landscape of the e. coli lac promoter. *PloS One* **8**, e61570 (2013).
10. Mora, T., Walczak, A. M., Bialek, W. & Callan, C. G. Maximum entropy models for antibody diversity. *Proc. Natl. Acad. Sci. (USA)* **107**, 5405–5410 (2010).
11. Ferguson, A. L. *et al.* Computational prediction of broadly neutralizing hiv-1 antibody epitopes from neutralization activity data. *PLOS One* **8**, e80562 (2013).
12. Tikhonov, M., Leach, R. W. & Wingreen, N. S. Interpreting 16s metagenomic data without clustering to achieve sub-otu resolution. *The ISME J* **9**, 68 (2015).
13. Bialek, W. *et al.* Statistical mechanics for natural flocks of birds. *Proc. Natl. Acad. Sci. (USA)* **109**, 4786–4791 (2012).
14. Couzin, I. D. & Krause, J. *Self-organization and collective behavior in vertebrates* (Academic Press, 2003).
15. Lukeman, R., Li, Y.-X. & Edelstein-Keshet, L. Inferring individual rules from collective behavior. *Proc. Natl. Acad. Sci. (USA)* **107**, 12576–12580 (2010).
16. Kelley, D. H. & Ouellette, N. T. Emergent dynamics of laboratory insect swarms. *Sci. Reports* **3**, 1073 (2013).
17. Pérez-Escudero, A. & De Polavieja, G. G. Collective animal behavior from bayesian estimation and probability matching. *PLOS Comput. Bio.* **7**, e1002282 (2011).
18. Stevens, C. & Zador, A. Neural Coding: The enigma of the brain. *Curr. Biol.* **5**, 1370–1371 (1995).

19. Pillow, J. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
20. Bassett, D. S. & Sporns, O. Network neuroscience. *Nat. Neurosci.* **20**, 353 (2017).
21. Williams, A. H. *et al.* Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping. *BioRxiv* 661165 (2019).
22. Schwab, D. J., Nemenman, I. & Mehta, P. Zipf’s law and criticality in multivariate data without fine-tuning. *Phys. Rev. Lett.* **113**, 068102 (2014).
23. Merchan, L. & Nemenman, I. On the sufficiency of pairwise interactions in maximum entropy models of networks. *J. Stat. Phys.* **162**, 1294–1308 (2016).
24. Nemenman, I. Renormalizing complex models: It is hard without landau! *J. Club Condens. Matter Phys.* (2017).
25. Fisher, C. & Mehta, P. Bayesian feature selection for high-dimensional linear regression via the ising approximation with applications to genomics. *Bioinformatics* btv037 (2015).
26. Ganmor, E., Segev, R. & Schneidman, E. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc Natl Acad Sci (USA)* **108**, 9679–9684 (2011).
27. Berry, M., Warland, D. K. & Meister, M. The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. (USA)* **94**, 5411–5416 (1997).
28. Strong, S., Koberle, R., de Ruyter van Steveninck, R. & Bialek, W. Entropy and information in neural spike trains. *Phys. Rev. Lett.* **80**, 197 (1998).
29. Reinagel, P. & Reid, R. C. Temporal coding of visual information in the thalamus. *J. Neurosci.* **20**, 5392–5400 (2000).
30. Arabzadeh, E., Panzeri, S. & Diamond, M. E. Deciphering the spike train of a sensory neuron: counts and temporal patterns in the rat whisker pathway. *J. Neurosci.* **26**, 9216–9226 (2006).
31. Rokem, A. *et al.* Spike-timing precision underlies the coding efficiency of auditory receptor neurons. *J. Neurophys.* **95**, 2541–2552 (2006).
32. Nemenman, I., Lewen, G., Bialek, W. & de Ruyter van Steveninck, R. Neural coding of natural stimuli: information at sub-millisecond resolution. *PLOS Comput. Bio.* **4**, e1000025 (2008).
33. Lawhern, V., Nikonov, A. A., Wu, W. & Contreras, R. J. Spike rate and spike timing contributions to coding taste quality information in rat periphery. *Front. Integr. Neurosci.* **5**, 18 (2011).
34. Fairhall, A., Shea-Brown, E. & Barreiro, A. Information theoretic approaches to understanding circuit function. *Curr. Opin. Neurobiol.* **22**, 653–659 (2012).
35. Tang, C., Chehayeb, D., Srivastava, K., Nemenman, I. & Sober, S. Millisecond-scale motor encoding in a cortical vocal area. *PLOS Biol.* **12**, e1002018 (2014).
36. Srivastava, K. *et al.* Motor control by precisely timed spike patterns. *Proc. Natl. Acad. Sci. (USA)* **114**, 1171–1176 (2017).
37. Sober, S. J., Sponberg, S., Nemenman, I. & Ting, L. H. Millisecond spike timing codes for motor control. *Trends Neurosci.* **41**, 644–648 (2018).
38. Putney, J., Conn, R. & Sponberg, S. Timing is (almost) everything in a comprehensive, spike-resolved flight motor program. *bioRxiv* (2019). <https://www.biorxiv.org/content/early/2019/05/03/602961.full.pdf>.
39. Doupe, A. & Kuhl, P. Birdsong and human speech: Common themes and mechanisms. *Ann. Rev. Neurosci.* **22**, 567 (1999).
40. Kuebrich, B. & Sober, S. Variations on a theme: Songbirds, variability, and sensorimotor error correction. *Neuroscience* **296**, 48–54 (2015).
41. Sober, S., Wohlgenuth, M. & Brainard, M. Central contributions to acoustic variation in birdsong. *J. Neurosci.* **28**, 10370–10379 (2008).
42. Wohlgenuth, M. J., Sober, S. J. & Brainard, M. S. Linked control of syllable sequence and phonology in birdsong. *J. Neurosci.* **30**, 12936–12949 (2010).
43. Amari, S.-I. Information geometry on hierarchy of probability distributions. *Inf. Theory, IEEE Transactions on* **47**, 1701–1711 (2001).
44. Thompson, C. J. *Mathematical statistical mechanics* (Princeton University Press, 2015).
45. Bialek, W., Rieke, F., de Ruyter van Steveninck, R. & Warland, D. Reading a neural code. *Science* **252**, 1854–1857 (1991).

46. Kelly, C. W. & J, S. S. A simple computational principle predicts vocal adaptation dynamics across age and error size. *Front. Integr. Neurosci.* **8**, 75 (2014).
47. Sober, S. J. & Brainard, M. S. Adult birdsong is actively maintained by error correction. *Nat. Neurosci.* **12**, 927 (2009).
48. Hessler, N. A. & Doupe, A. J. Social context modulates singing-related neural activity in the songbird forebrain. *Nat. Neurosci.* **2**, 209–211 (1999).
49. Woolley, S. C., Rajan, R., Joshua, M. & Doupe, A. J. Emergence of context-dependent variability across a basal ganglia network. *Neuron* **82**, 208–223 (2014).
50. Kao, M., Doupe, A. & Brainard, M. Contributions of an avian basal ganglia–forebrain circuit to real-time modulation of song. *Nature* **433**, 638 (2005).
51. Olveczky, B. P., Andalman, A. S. & Fee, M. S. Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLOS Biol.* **3**, e153 (2005).
52. Elemans, C. P. H. *et al.* Universal mechanisms of sound production and control in birds and mammals. *Nature Commun.* **6**, 8978 (2015).
53. Zhou, B., Hofmann, D., Pinkovitzky, I., Sober, S. & Nemenman, I. Chance, long tails, and inference: a non-gaussian, bayesian theory of vocal learning in songbirds. *Proc. Natl. Acad. Sci. (USA)* **115**, E8538–E8546 (2018).
54. Tumer, E. C. & Brainard, M. S. Performance variability enables adaptive plasticity of ‘crystallized’ adult birdsong. *Nature* **450**, 1240–1244 (2007).
55. Paninski, L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Comput. Neural Syst.* **15**, 243–262 (2004).
56. Gerwinn, S., Macke, J. & Bethge, M. Bayesian inference for generalized linear models for spiking neurons. *Front. Comput. Neurosci.* **4** (2010).
57. Granot-Atedgi, E., Tkacik, G., Segev, R. & Schneidman, E. Stimulus-dependent maximum entropy models of neural population codes. *PLOS Comput. Bio.* **9** (2013).
58. Savin, C. & Tkačik, G. Maximum entropy models as a tool for building precise neural controls. *Curr. Opin. Neurobiol.* **46**, 120–126 (2017).
59. Pandarinath, C. *et al.* Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. methods* **1** (2018).
60. Glaser, J., Chowdhury, R., Perich, M., Miller, L. & Körding, K. Machine learning for neural decoding. *arXiv.org* (2017). [1708.00909v2](https://arxiv.org/abs/1708.00909v2).
61. Schnitzer, M. J. & Meister, M. Multineuronal firing patterns in the signal from eye to brain. *Neuron* **37**, 499–511 (2003).
62. Oppen, M. & Saad, D. *Advanced mean field methods: Theory and practice* (MIT press, 2001).

Acknowledgements

We thank David Hoffman and Pankaj Mehta for valuable discussions. This work was supported in part by NIH Grants R01-EB022872, R01-NS084844, and R01-NS099375, and NSF grant BCS-1822677. IN acknowledges hospitality of the Kavli Institute for Theoretical Physics, supported in part by NSF Grant PHY-1748958, NIH Grant R25GM067110, and the Gordon and Betty Moore Foundation Grant 2919.01. IN and SJS further acknowledge hospitality of the Aspen Center for Physics, which is supported by NSF grant PHY-1607611.

Author contributions statement

DGH, SJS, and IN conceived the project, performed the research, and wrote the paper. DGH wrote the software implementation of uBIA.

Additional information

Competing interests The authors declare no competing interests.