The choice-wide behavioral association study: data-driven identification of interpretable behavioral components

David B. Kastner[1,5], Nicole O. Yokota[1], Christina Y. Lee[1], Cristofer Holobetz[1], Ali Chaudhry[1], Greer Williams[1], Jane Ton[1], Mihika Mehra[1], Viktor Kharazia[1], Dilworth Y. Parkinson[2], Joseph P. Romano[3], Peter Dayan[4]

[1]*Department of Psychiatry and Behavioral Sciences, University of California, San Francisco, CA 94143, USA*

[2]*Advanced Light Source, Lawrence Berkeley National Labs, Berkeley, CA, 94720, USA*

[3]*Department of Statistics, Stanford University, Stanford, CA 94305, USA*

[4]*Max Planck Institute for Biological Cybernetics, Tübingen 72076, Germany*

[5]Lead Contact

Correspondence: David B. Kastner: david.kastner2@ucsf.edu

Behavior contains rich structure across many timescales. However, there is a dearth of methods to identify relevant components, especially over the longer periods required for learning and decision-making. Inspired by the goals and techniques of genome-wide association studies, we present a data-driven method—the choice-wide behavioral association study: CBAS—that systematically identifies such behavioral features. CBAS extracts sequences of actions or choices that either significantly differ between groups or correlate with a covariate of interest, using a powerful resampling-based method to correct for multiple comparisons. We illustrate CBAS through application to tasks performed by flies, rats and humans, showing that it provides unexpected and interpretable information in each case.

INTRODUCTION

Understanding how behavior differs between or among groups of humans or other animals is critical for generating and testing hypotheses about the functional role of genes, neural circuits, and regions of the brain. It is also essential for characterizing neurologic and psychiatric dysfunction[1]. However, behavior is highly complex, evolving over multiple timescales and exhibiting substantial path dependencies due to individual experience[2–6]. To deal with these complexities, it is increasingly popular to collect large datasets through automated behavioral paradigms, and to analyze them using computational and machine learning methods[7–16]. There are two main computational approaches[17,18]: top-down/model-dependent, and bottom-up/data-driven. Model-dependent methods are substantially more prevalent in the analysis of behavior; however, data-driven methods dominate in other areas with rich and complex findings and afford the possibility of addressing some of the many limitations of top-down approaches.

In model-dependent analyses, behavioral data, such as choices in a decision-making task, are processed under the specific assumptions of a hypothesis or model. If the model or hypothesis is correct, this is highly efficient, since large volumes of data can be reduced to a handful of semantically meaningful parameters, such as learning rates or differential sensitivity to rewards or punishments. By finding specific parameters whose values differ significantly between groups, or correlate with facets of those groups, model-dependent analyses conflate identification with interpretation. That is both their strength and their weakness. If the model is correct, it allows for the difference or correlation to be directly interpreted in terms of the parameters concerned; however, if the model is incorrect or limited, confirmation bias[19] imposes an incorrect interpretation on the difference or correlation. That is, even when substantial effort is put into building multiple alternative models, and comparing them in a statistically rigorous manner, it remains possible that the best fitting model nevertheless fails to characterize the behavior properly. This limits the interpretability of the model or the interrogation of hypotheses. In the worst cases, model and hypothesis-based analyses offer post-hoc explanations for any difference that can be found in a behavioral dataset[20].

Data-driven analyses start from the other end, characterizing behavior without relying on parametric assumptions, and making very few assumptions about how the data are generated. Some of these approaches are unsupervised, for instance finding clusters in behavioral space and comparing them between groups[12,21,22]; others are more supervised,

directly looking for discriminative differences between populations or individuals[7–10]. By not trying to shoehorn data into a limited set of parameters, these methods should be more sensitive; however, comparisons between groups pose severe statistical challenges, because of the complexity and dimensionality of behavioral datasets; and interpretation remains a challenge.

Modern machine learning methods[23–25] can sit between model-dependent and data-driven camps. They have the particular benefit of providing useful lower bounds for how well hypothesis-driven models should fit but often struggle with convincing interpretability and data efficiency.

Here, we present the choice-wide behavioral association study (CBAS), a novel data-driven analysis method designed to identify sequences of choices (or other discrete behavioral features) that differ between or among groups. CBAS has two components: 1) breaking down behavior into a comprehensive language that allows comparisons to be made between two groups or to assess correlation with a covariate of interest; and 2) using rigorous, resampling-based, corrections to account for the resulting large number of comparisons and to maintain statistical power despite statistical complexity in the data. CBAS breaks the conflation between detection and interpretation of model-dependent analyses because it is designed simply to identify differences or find correlations with a covariate. The presence (or absence) of that significance is interpretation-free. However, by design, the output of CBAS provides a rich and rigorous substrate for interpreting behavior, for testing the adequacy of models and hypotheses concerning the source of behavioral differences, and our method provides a general data-driven framework to advance our resulting understanding.

RESULTS

Choice as a common discretization for behavior

The development of CBAS was motivated by the powerful data-driven approaches of genome-wide association studies (GWAS), whole exome sequencing (WES) and whole genome sequencing (WGS)[26]. Prior to the wide-scale implementation of these methods, candidate gene studies that attempted to associate candidate genes with phenotypes were often underpowered and failed to replicate[27–29]. It has been argued that many data-driven behavioral analyses suffer the same problems[30]. GWAS/WES/WGS famously solved the problem of statistical reliability by employing rigorous multiple-comparisons correction. These corrections depend on the discrete representation afforded by base-pairs[31,32]. CBAS follows an equivalent statistical path, and also requires a similarly appropriate, and potentially general, discretization, but for behavior. In many cases, choices provide just such a discretization (Fig 1a).

We apply CBAS to three tasks in three diverse species—flies, humans, and rats—to show the breadth of applicability of CBAS. For simplicity and clarity, we describe the motivation and development of CBAS using one of the tasks and then show its application to the other tasks below. The task that we will initially use to describe CBAS involves drosophila walking on a y-maze (Fig 1a; left). The flies are tracked as they spontaneously make the choice of going to either the left or right arm after leaving a current arm[31]. This left/right discretization of the task has enabled many conclusions to be drawn about the genetic nature of individual variability[31–36]. Here, we compare the choices of two outbred strains of fly, Cambridge-A (CA) and w1118,

from a publicly available dataset[31]. The w1118 strain is the background strain for many transgenic flies[31]. Analyses in the original paper[31] hypothesized a potential source of the difference between the groups: handedness, or the preference for left or right turns, by the individual flies; however, there are many different ways in which that might manifest, and the rich dataset combined with CBAS allows multiple possibilities to be evaluated.

CBAS evaluates all sequences of choices up to a set length

Using choice as the basis of the comparison for the behavioral analysis requires an additional consideration beyond what is done for the genome with base-pairs. For GWAS/WES/WGS, an individual base pair can be a meaningful unit of information (although this is only a partial story[37,38]) that can be compared between subjects. However, isolated choices are rarely meaningful units by themselves, as the choices that precede and follow them can change their interpretation. CBAS does not just evaluate the occurrence of individual choices, but, rather, the occurrence of all sequences up to a certain, user-defined, length. In general, the longer the sequences, the more data and computational resources will be necessary, but the richer the aspects of behavior that can be encompassed.

For the fly spontaneous alternation behavior, we applied CBAS to all sequences of choices of turns up to a length of 10. For each sequence of choices, all instances of that sequence up to a criterion of 250 choices are found in the dataset, and the rate of usage of that sequence in the two populations is used as the point of comparison at the basis of CBAS (Fig 1b shows one particular example sequence).

Correcting for the many comparisons inherent in evaluating the occurrence of such sequences is complicated by the fact that their incidence can be highly correlated. Standard methods to correct statistically for multiple comparisons (e.g. Bonferroni, Holm, Benjamini-Hochberg, Benjamini-Yekutieli) are underpowered for correlated data. Instead, we use an approach based on resampling to correct for the multiple comparisons—called the Romano-Wolf method. This approach retains power in the face of correlations[39]. In practice, we set the significance threshold according to median control of the false discovery proportion at 5%. This control is comparable to 5% control of the false discovery rate, which provides mean control of the false discovery proportion[40].

Using the Romano-Wolf method for resampling-based multiple comparisons correction[39,40], CBAS identifies many sequences of choices that differ significantly between the two strains of flies in the spontaneous alternation task (1,605/2,046 sequences; Fig 1c). We are also able to determine the sample size required to provide ~80% power to detect a difference at this level of statistical certainty, using subsampling of the data (Fig S1).

CBAS identifies interpretable differences for fly spontaneous alternation

CBAS identifies many sequences that significantly differ between the CA and w1118 fly lines. Upon inspection of all of the significantly different sequences, a clear, interpretable, contrast is apparent between those sequences that occur more in the CA line and those that occur more in the w1118 line (Fig 2a). The CA line utilizes sequences with extended numbers of the same turn in a row, whereas the w1118 line utilizes sequences with more frequent changes in the direction of turning. Therefore, CBAS not only identifies that there is a difference

between the two fly lines, but it also provides information to support an interpretation as to the nature of the difference.

The paper from which this open-source data comes from highlights the handedness of the individual flies, whereby individual flies show a propensity towards either right or left turns[31]. Such an observation creates a straightforward testable hypothesis about the nature of the difference between the two lines, i.e., that the lines differ in the extremes of their handedness bias, with individual CA flies hypothesized to being more extreme in their left/right turn bias than the w1118 lines. This hypothesis can be formalized by randomly generating virtual twin subjects by sampling turns according to the left/right turn biases of each individual fly. CBAS can compare this turn bias model of the lines, through generating an expectation for what its output would look like if the lines solely differed based on the left/right turn bias (Fig 2b). The CBAS results from the data are inconsistent with this hypothesis (judged by the numbers of sequences that are significantly different in two transparent metrics).

Instead of modeling the choices of the individual flies as being generated by a left/right bias on each turn, we could instead generate the turns based on the bias of the individual flies to persist in their direction of turning (i.e., turning in the same direction as their previous turn). If we generate virtual twin subjects by sampling turn directions according to these one-step persistence biases of each individual fly—persistence rate—the resulting CBAS output is closer to the data than the left/right turn bias but still is inconsistent overall (Fig 2b). However, if the choices of virtual twin flies are generated by using the distributions of persistence for each of the number of turns in a row from each individual fly—persistence distribution—then the CBAS of that model almost perfectly matches the CBAS output of the data according to our metrics (Fig 2b).

Graceful decay of CBAS output with decreasing group size

We exploited the very large numbers of flies in the spontaneous alternation dataset to understand how the output of CBAS behaves with different sample sizes. This determines our ability to make conclusions and derive interpretations from fewer subjects. We first evaluate the number of significant sequences identified from smaller group sizes as a fraction of the total number of significant sequences identified with the full dataset (Fig 3a). As expected, with smaller group sizes, we find fewer overall significantly different sequences; however, across the range of group sizes evaluated, the median fraction of sequences was larger than the fraction of the population being used in the smaller group CBAS (Fig 3a). This indicates that at smaller group sizes, the proportion of sequences identified by CBAS for these subjects grows faster than the proportion of subjects in the CBAS. This provides a rapid increase in the amount of information provided by CBAS as the sample size increases.

We next evaluate the fraction of the sequences identified by CBAS for the smaller sized groups that are not identified in the CBAS on the full dataset (which we might interpret as false positives). Since we control the false discovery proportion and not the familywise error rate, we do not expect this value to be zero; however, it was consistently small across all the group sizes evaluated (Fig 3b). The medians across the different repeats of the same sample sizes across the different sample sizes are all less than 2%.

Finally, we evaluated the similarity of the sequences identified by CBAS with nonoverlapping sets of subjects with the same group size (Fig 3c). At smaller groups sizes (40 per group), this overlap could be quite low (median 24.4%), even though there was sufficient evidence to discriminate the strains (Fig. S1). This means that the specific sequences identified by CBAS can be quite variable from one experiment to the next, especially at smaller group sizes. Given that, we sought to understand if the structure of the sequences comported with the conclusion from the full dataset—that the CA line uses more of the same turns in a row and the w1118 line more frequent changes in direction of turn. Even though the overlap between CBAS on different groups might not be that high (Fig. 3c), the sequences identified still follow the same structure as the full dataset in regard to number of the same turns in a row (seen across the repeats for each group size in Fig 3d). As the group size increases, the output of CBAS becomes more consistent, and more similar to the results from the full dataset (Fig 3e), licensing more generalized conclusions.

CBAS can also correlate sequences of choices with a covariate of interest

The second task to which we apply CBAS is the two-step task, developed to test the interplay between model-based and model-free (reinforcement) learning algorithms[41] (Fig 1a; middle). In this task, subjects choose between pairs of images at two different stages. The image choice at the first stage stochastically governs the pair of images from which the subject can choose at the second stage, and which enjoy drifting individual probabilities of delivering reward when picked. Model-dependent analyses of variants of this task (which make inferences based on the stochasticity in the transition between first and second stages) have led to conclusions about the algorithms used in different brains regions[41,42] and differences that underlie psychiatric symptoms[43], notably compulsivity. However, the interpretations are not without controversy[44], and the complexity of behavior makes it difficult to evaluate ways in which the model might be missing features in the data.

Data for the two-step task come from a large population of human subjects in the open-source dataset from Gillan et al.[43]. We applied CBAS to all sequences of up to 4 choices long over the course of the 200 trials—400 total choices—of the task. Following Akam et al.[42] we collapse image 3 and 4 into set A and image 5 and 6 into set B because the specific images within the set are not relevant for the critical decision at the first stage. Choices that are rewarded are treated as distinct from those that are not rewarded, as those choices can lead to different behaviors. That means that the sequence 2**A**1 is different from 2a1, where the upper-case emboldened letter indicates that a reward was received.

In building the two-step task dataset, the authors sought to examine how behavior varied along continuous trait dimensions. They did this by performing factor analysis on the answers to a series of psychiatric symptom questionnaires and showing that the way subjects performed that task was correlated with one of the factors (factor 2) that was associated with compulsive behavior and intrusive thoughts (CBIT)[43]. Therefore, we extended the CBAS so that it could also evaluate such correlations. We applied CBAS to the correlation between the sequence usage for each subject and their CBIT score (Fig 1b), finding 31/408 sequences to be significant (Fig 1c).

CBAS provides evidence for testing model-dependent analyses of the two-step task

To understand the output of CBAS for the two-step task, we review the expectation associated with the hypothesis-driven analysis of this experiment. The two-step task was designed to evaluate the interplay between model-based and model-free decision making[41] based on the stochastically rare transitions (from image 1 to set B; and image 2 to set A). Model-based decision making develops an understanding of the structure of the world (which is the model) and makes choices based on that understanding. Model-free decision making makes choices based on reinforced past successful actions, without developing an understanding of the structure of the world. A way to see the difference between these two decision-making schemes is when a learner chooses image 2 at the first step, makes an unlikely stochastic transition to set A, and gains a reward at this second step ('2**A**' in our language). A regular model-free learner will reinforce the choice of image 2 (along with the image chosen at A) and will therefore be more likely to choose image 2 on the next trial. By contrast, a model-based learner will be more likely to choose image 1 on the next opportunity, based on the observation that this will make the transition to set A more likely (because of the common/rare transition structure in the environment)[41].

We evaluated all of the sequences that CBAS found as either significantly positively or negatively correlated with CBIT (Fig 4a). CBAS identified a positive correlation with many sequences involving being rewarded at A and then selecting image 2 or being rewarded at B and then selecting image 1. In particular, the structures 1**A**2 or 2**B**1 occur commonly in the positively correlated sequences. These behavioral motifs are inconsistent with either model-free or model-based learning, as the subjects get rewarded after choosing from the common side (A from 1 or B from 2) but then select the image that will rarely bring them back to the previously rewarded side (2 from A or 1 from B). The model-based decision maker would return to 1 as that is the side that would commonly go to set A where there was the reward. The model-free decision maker would also return to 1 as that was the side that was reinforced as being connected to the recent reward at A.

We next used CBAS to test whether the combined model-based/model free reinforcement learning (MB/MF RL) model used by Gillan et al.[43] which has parameters describing model-based and model-free decision making and the propensity to employ each, correctly captures the data. To do this, we generated virtual twin choices for each subject based on fits of each subject to this model and ran CBAS on the model output (Fig 4b). This led to a positive correlation between CBIT and the model inconsistent sequences with the choices 1**A**2 and 2**B**1, like the human data. However, there was the important difference in that CBAS identified a positive correlation with the unrewarded versions of those units (1a2 and 2b1) that was not present in the human data.

CBAS therefore identified an important failing in the MB/MF RL model. We can explain this by noting that this model has a parameter for "stickiness"—the propensity to continue to choose the same of the first images (1 or 2) in a row. Decreased stickiness would also comparatively favor sequences that change the first image, including 1**A**2, 1a2, 2**B**1 and 2b1. Indeed, in our fits, along with the negative correlation between the propensity to employ model-based choices and CBIT ($\rho = -0.08; p = 1.7x10^{-3}$) reported by Gillan et al.[43], there was

263 also a negative correlation between the stickiness parameter and CBIT ($\rho = -0.11$; $p =$
264 $1.6x10^{-5}$).

265 As an alternative, we generated behavior based on just model-inconsistent choices. That
266 is, we created a virtual twin of each subject s from the MB/MF RL model by hybridizing the rate
267 with which that subject performed model-inconsistent choices with choices that were
268 otherwise generated according to the fit parameters for a single, reference subject s*. The
269 reference subject s*, was randomly chosen for each of the 50 repeats. That is, we used the
270 MB/MF RL model based on s* to generate data, but every time the model produced 1**A** or 2**B**
271 we probabilistically selected the next choice according to the overall model-inconsistent rate of
272 subject s (but then continued to generate choices as if nothing else had changed). This ensures
273 that the majority of choices came from a subject with a single CBIT value, and only the choices
274 following 1**A** or 2**B** were manipulated (19% on average). The CBAS results from this model-
275 variant were much closer to the CBAS output of the data in the metric reported (shown in the
276 'Model inconsistent' results in Fig. 4b). Furthermore, when we fit the output of the model-
277 inconsistent variant of the model to the MB/MF RL model, we found that a positive correlation
278 between model-inconsistent choices and CBIT alone can appear as a negative correlation
279 between the model-based parameter of the model and CBIT (Fig 4c).

280 **CBAS identifies a hippocampal dependent strategy for spatial alternation learning**

281 The third, and final task we evaluated with CBAS is spatial alternation behavior in rats.
282 For this task, to get reward, the rats must alternate between pairs of arms of a track whilst
283 visiting a different arm of the track in between (Fig 1a; right). Spatial alternation is a common
284 behavioral paradigm for phenotyping and neurophysiology, and the discretization of the
285 behavior into the arms chosen by the animals forms the basis of many of the conclusions from
286 these studies[45–49]. However, our recent work calls into question the assumptions and
287 hypotheses that motivate the standard analysis for spatial alternation behavior[50,51]. Such
288 limitations could even be general across other tasks used to evaluate memory[52]. Furthermore,
289 even though we have developed many variants of reinforcement learning agents to fit
290 individual behavior, these agents learned in ways that were very evidently different from the
291 animals[51], limiting their use for phenotyping. Here, we use CBAS to evaluate the role of the
292 hippocampus in this task by comparing rats with lesions to their hippocampus to rats that
293 underwent control surgery.

294 For the spatial alternation task, we applied CBAS to all sequences of arm choices up to 6
295 choices long over the first 800 arm choices of the rats. As with the two-step task, rewarded
296 choices are treated differently from unrewarded choices. As with the fly spontaneous
297 alternation dataset, CBAS compared the rate of usage of each sequence between the
298 populations of hippocampally-lesioned and control rats (Fig 1b; right). 409/24,342 sequences
299 differed significantly (Fig 1c).

300 CBAS identified many sequences discriminating the fly strains. However, the
301 dimensionality of the task is quite low, so it is fairly straightforward to structure all of the
302 sequences to develop an interpretation (Fig 2). By contrast, CBAS did not identify that many
303 sequences in the two-step task that significantly correlated with CBIT, so the extra complexity
304 and dimensionality of this task did not prevent us from directly evaluating all the sequences to

305 develop an interpretation (Fig 4). For the case of rat spatial alternation, CBAS identified a good
306 number of sequences, and the task is complex and high dimensional, making it a challenge to
307 directly develop an interpretation of the output of CBAS.

308 Therefore, we developed an analysis method that takes advantage of the structure of
309 the sequences at the output of CBAS to help inform an interpretation of the spatial alternation
310 behavior. CBAS evaluates all sequences up to a certain length, and so sequences with shared
311 and overlapping structure might all be significant. We therefore defined the notion of a
312 "complete" target sequence (Fig S2) as one with the property that all sufficiently prevalent
313 longer sequences that contain it are found by CBAS to differ significantly in the same way as the
314 target (e.g., all are found to be greater in the control than the lesion rats). Effectively, a
315 "complete" sequence has the property that no matter what happens before or after that
316 sequence, the series of choices remain significantly different between the groups. Such
317 sequences are a powerful way to simplify the complex behavior.

318 We find "complete" sequences that the control rats do more than the lesion rats and
319 that the control rats do less than the lesion rats (Fig 5a). Inspection of these sequences provides
320 two axes to compare all of the significant sequences. The "complete" sequences that the
321 control rats use more than the lesion rats show systematic exploration of the arms by largely
322 continuing in the same direction and transitioning mostly to neighboring arms. These features
323 are much less the case in the "complete" sequences that the control rats do less than the lesion
324 rats. We used these two axes to compare all the significant sequences (Fig 5b). Consistent with
325 the finding for "complete" sequences, sequences that the control rats use more than the lesion
326 rats tend towards higher directional inertia—the propensity to continue in the same direction—
327 with smaller overall transition distance.

328 The hypothesis generated by the "complete" sequences and verified across all of the
329 significant sequences, motivates a relatively straightforward metric by which we can evaluate
330 all of the choices of the rats for comparison between the two groups. We quantified the
331 transition normalized directional inertia in the control and lesion rats. This metric measures, for
332 each choice, whether that choice continues in the same direction as the preceding choice and
333 normalizes it by the number of arms transitioned. For example, choosing arm 3 after coming
334 from 2 and 1 (sequence 1-2-3) has a transition normalized directional inertia of 1, choosing arm
335 4 after coming from 2 and 1 (sequences 1-2-4) has a transition normalized directional inertia of
336 0.5, and choosing arm 1 after coming from 2 and 1 (sequence 1-2-1) has a transition normalized
337 directional inertia of 0. Lesion and control rats show a clear and significant difference in their
338 transition normalized directional inertia (Fig 5c).

339 Although "complete" sequences were not used to interpret the fly and human CBAS
340 outputs, evaluating them would lead to identical conclusions (Fig S3).

341 **CBAS generated hypothesis replicates for spatial alternation**

342 The fact that the transition normalized directional inertia is significantly different
343 between the control and lesion rats, is partly a circular claim. We motivated the metric based
344 on the sequences identified by CBAS as differing between the groups. It did not have to be the
345 case that the metric would differ across all the choices of the animals (as opposed to just the

significant sequences), but it should not be overly surprising that it does. Therefore, we sought to replicate our understanding in a separate dataset of hippocampal lesion and control rats.

Even with a smaller sample size, CBAS identifies differences due to the hippocampal lesioning (Fig 6a), consistent with the power calculation on the initial dataset (Fig S1e&f). The sequences identified by CBAS separate along the two axes identified by the "complete" sequences from the initial dataset (Fig 6b). This comportment of the structure of the sequences with a smaller sample size is consistent with what was found with subsampling the fly data (Fig 3d). Consistent with the hypothesis from the initial dataset, the groups significantly differ in their transition normalized directional inertia (Fig 6c). Indeed, we even replicated specific sequences across the two datasets (Fig 6d). Interestingly, a much larger fraction of the sequences that the control rats use more than the lesion rats replicate compared to the sequences the control rats use less than the lesion rats. This is consistent with the overall hypothesis that the hippocampus enables a learning strategy that systematically engages with the arms of the track through moving in the same direction and transitioning to nearby arms. Without a hippocampus that systematic behavior goes away, exposing less systematic choices, that inevitably will be less likely to replicate between samples.

DISCUSSSION

We have presented a general data-driven behavioral analysis method, CBAS, for identifying interpretable behavioral components that is grounded in sequences of choices made by subjects. There has been significant progress, in recent years, in tracking and analyzing, short timescale actions of subjects[7–16]; however, we lack methods for generally analyzing and interpreting sequences of these actions or choices and decisions of subjects that unfold over a longer period than moment to moment movements. CBAS provides just such a method and is applicable across a wide array of species and different behavioral paradigms (Fig 1). It can be used to test models and hypotheses (Fig 2&4), and, for instance, using generalizable simple principles about the relationship between sequences, it can generate hypotheses in complex behaviors where reliable computational understanding has yet to emerge (Fig 5&6). Through taking advantage of large-scale data collection and rigorous statistical methods, CBAS has the potential to transform our use of behavior in a comparable way to how GWAS/WES/WGS changed the paradigm for genomic studies.

Somewhat surprisingly, in all tasks, CBAS provides information and generates interpretations that are either inconsistent with, or unexpected from, the prior understanding of these data and tasks. For the flies, as mentioned above, the dataset came from a paper characterizing the existence of handedness in these flies, as manifested by extremes in individual preferences for left or right turns[31]. The output of CBAS on these data allowed for the testing of generative models to explain that handedness. The most straightforward model that one could postulate based on the hypothesis of handedness is not consistent with the data (Fig 2b); whereas a model drawing upon the distribution of individual flies to persist in the same direction for a range of the same turn in a row was almost indistinguishable from the data according to our metrics. Therefore, CBAS generates testable hypotheses as to underlying computational structure that these flies draw upon as they perform this behavior.

For the two-step task, the CBAS results provide an alternative interpretation for the relationship between the task and individual CBIT levels. Previously, these data were interpreted to mean that individuals with more CBIT are less likely to draw upon model-based planning[43]. The CBAS results indicate that the relationship with CBIT rests more on the increase of model-inconsistent choices with CBIT (Fig 4a&b), and that that increase alone could appear as a negative correlation with model-based planning, even if there was in reality not one (Fig 4c). What an increase in model-inconsistent choices means in the context CBIT is something that should be explored in the future.

The result that the hippocampus is involved in spatial alternation behavior is not a new one. There have been prior reports of changes to spatial alternation behavior in the context of hippocampal lesioning[45] and manipulation[48,53]. However, all of the prior work occurred from the prevailing perspective of spatial alternation exclusively reflecting memory allocation, something we have recently shown to be inconsistent with the way in which the rats learn[50,51]. Guided by the results of CBAS we find that the hippocampus is responsible for implementing a specific strategy (or set of strategies) whereby the animals systematically move across the track and transition between closer arms as they learn. This strategy is far more consistent with the hippocampus implementing a structuring of the task space as described in recent models of hippocampal function[54,55], than with the typical memory function assumed in the prior studies of spatial alternation behavior.

These three tasks were just the first ones to which we applied CBAS; they were not selected from a larger set of tasks and used for illustrative purposes. We strongly anticipate that there will be comparable benefits to applying CBAS and other data-driven methods to a broad array of datasets and behavioral tasks to provide a stronger foundation for our understanding of behavior. The ability of CBAS to identify differences or correlations in a dataset is directly built into its function. However, the ability to interpret the differences or correlations found remains nontrivial. We have provided examples of ways to extract interpretable information from the sequences output by CBAS and have developed a general analysis ("complete" sequences) to assist in that process. However, there remains room for future work to assist in going from the identification of a difference or correlation by CBAS to understanding and interpreting what those differences and correlations mean. Of course this is in no way a unique problem to CBAS. Genomics methods (GWAS/WES/WGS) similarly require secondary and follow up analyses.

Considerations for the application of CBAS to datasets

To apply CBAS to a dataset, there are some decisions that need to be made. At a fundamental level, CBAS requires a discretized description of the behavior. This was motivated by the common use of choices to describe behavior (Fig 1a). Furthermore, there is evidence that, at a rather fundamental level, behavior occurs through discrete choices[12,56]. CBAS is not currently capable of working on continuous data. The discretization informs the language that is used to make up the sequences that form the basis of the statistical evaluation within CBAS. In deciding on the language there is a range of biases that can be imposed with different decisions. For example, with the two-step task, we combined the second image choice into sets A or B. This prevents CBAS from differentiating a difference in choices at that level. Based on the literature we deemed that bias to be justified, but it is a bias, nonetheless.

A next decision to make is the length of sequence up to which CBAS will evaluate. The longer the sequences, the more information CBAS can provide, but the more data and computational resources that will be necessary. If the language is large, i.e. there are many things that can happen in the behavior at any point in time, and the transition matrix between one choice and another is not sparse, then computational resources might not be feasible to calculate the statistics underlying CBAS, especially if longer sequences are evaluated. Additionally, it is necessary to specify a criterion for the total number of choices per animal or session that are decomposed to form the sequences that are compared. Although it is possible just to include all the data, this might introduce bias if the groups have different total numbers of choices overall. Conversely, one might try and study aspects of learning or the acquisition of competence by applying CBAS to restricted portions of the behavior. It would also be possible to use a discrete language that characterizes aspects of behavior that unfold over longer timescales (for instance, defining a discrete language using unsupervised clustering over windows of choices), or indeed to employ a multi-scale description.

Finally, CBAS either needs to compare groups or correlate with a covariate of interest. Currently CBAS is not able to just evaluate a single group of animals by itself in an unsupervised manner without a covariate of interest. Future work will focus on relaxing this limitation.

In sum, we offer CBAS as a new data-driven analysis technique that extracts statistically rigorous differences in behavioral sequences between or among groups. We show that such sequence distinctions can be revealing about the underlying information processing functions that vary and can circumvent biases that plague model-dependent analysis methods. We illustrated the properties and promise of CBAS in three different datasets in three different species.

Methods

Romano-Wolf resampling based multiple comparisons correction: We follow the terminology and description laid out in Clarke et al.[57] to describe the Romano-Wolf multiple comparison correction. First, we describe the way the method corrects for the familywise error rate (FWER) and then explain how the procedure is extended to provide median control of the false discovery proportion. FWER control at a level of α means that across all comparisons there is a α percent chance of having at least one false positive rejection of a null hypothesis. The Romano-Wolf procedure provides FWER control through resampling the data[39]. We describe the case of testing a total of S hypotheses.

It is not generally known if the Romano-Wolf procedure controls for type III, or directional, errors. Type III errors are errors in the sign, or direction, of the conclusion. For example, if a statistical test provided information to reject the null hypothesis $\theta_1 = \theta_2$, and you then concluded that $\theta_1 > \theta_2$, when in fact $\theta_1 < \theta_2$. Therefore, instead of running a single two-tailed test for each sequence, we run two one-tailed tests for each sequence. As a consequence, the total number of hypotheses tested, S, is twice the total number of sequences being compared. For CBAS those hypotheses take one of two forms: i) the rate of each sequence (r_s) is the same between two groups $\Delta r_s = 0$, or ii) that there is no correlation (ρ_s) between each sequence and a covariate of interest, $\rho_s = 0$. For the one-tailed versions of each

hypothesis, we ask if $r_{s_1} - r_{s_2} > 0$ and $r_{s_2} - r_{s_1} > 0$ for the comparison CBAS, where r_{s_N} is the rate of sequence s for group N, or if $\rho_s > 0$ and $\rho_s < 0$ for the correlational CBAS.

The first step in the procedure is to create a studentized test statistic for each hypothesis. The studentization is different based on whether the CBAS is comparing two groups or calculating a correlation. In the case where two groups are being compared:

$$t_s = \frac{\Delta r_s}{\sigma_s} \tag{1}$$

where σ_s is the standard error of Δr_s, which we calculate by combining the standard error of the mean of the rate for each group using error propagation, i.e. $\sigma_s = \sqrt{{\sigma_{s_1}}^2 + {\sigma_{s_2}}^2}$, where σ_{s_N} is the standard error of the occurrence rate of sequence s for group N.

In the case where the correlation is being calculated the studentized test statistic is[58]:

$$t_s = \frac{\sqrt{n}\hat{\rho}_n}{\hat{\tau}_n} \tag{2}$$

where:

$$\hat{\rho}_s = \frac{\sum X_i Y_i - n\bar{X}_s\bar{Y}}{\sqrt{\sum(X_i - \bar{X}_n)^2 \sum(Y_i - \bar{Y})^2}} \tag{3}$$

$$\hat{\tau}_s = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(X_i - \bar{X}_s)^2(Y_i - \bar{Y})^2}}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(X_i - \bar{X}_s)^2}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(Y_i - \bar{Y})^2}} \tag{4}$$

For eq. 2, 3, and 4, n is the number of subjects for which the correlation is being calculated, X_i and Y_i are the values of the metrics being correlated (in our case, the sequence count and CBIT score for each individual respectively), $\bar{X}_s$ is the mean sequence count for the specific sequence being considered, and $\bar{Y}$ is the mean of the covariate of interest (CBIT score, which is the same for any sequence).

When two groups are being compared, we resample from the entire population with replacement (separately for each group) and build up a null distribution by bootstrapping M times. The test statistic from the mth bootstrap sample for $m = 1, \ldots M$ is:

$$t_s^{*,m} = \frac{\Delta r_s^{*,m}}{\sigma_s^{*,m}} \tag{5}$$

where, $\Delta r_s^{*,m}$ is the difference in the rate of each sequence whilst resampling, with replacement, from the entire population, ignoring the group labels. $\sigma_s^{*,m}$ is the accompanying

bootstrap standard error of the difference in means. The resampled group sizes are the same as the two groups of interest.

In the case where the correlation is being calculated, the test statistic based on the mth bootstrap sample is:

$$t_s^{*,m} = \frac{\sqrt{n}\hat{\rho}_s^{*,m}}{\hat{\tau}_s^{*,m}} \tag{6}$$

where, $\hat{\rho}_s^{*,m}$ is the correlations of each sequence whilst permuting the relationship between the metrics being calculated across the entire population. $\hat{\tau}_s^{*,m}$ is the accompanying normalization for the permuted relationship. The permuted group size is the same as the original.

For the main CBAS calculations, we used a value of $M = 10{,}000$. When we run many repeats of different models or for power calculations, we use a value of $M = 1{,}000$. Importantly, for each individual resampling, m, the same resampled set is used for all sequences.

The test statistics, and their accompanying estimators are ordered from largest to smallest values. This creates an $M \times S$ matrix where each column contains all the estimators of the test statistics. The first column contains the estimators from the largest test statistic, the second column contains the estimators from the second largest, etc.

To define the distribution for which each test statistic is compared, which then determines the adjusted p-value, the following algorithm is used. The first sequence considered is the one with the maximum test statistic, t_s. Its comparison distribution is defined as the maximum value within each row of the matrix of estimators of the test statistic:

$$\max(t^{*,m}) = \max\{t_1^{*,m}, \dots, t_S^{*,m}\} \tag{7}$$

which provides a total of M values, $t^{*,m}$ (there is no longer an association with s, because these values can come from a resampling of any of the sequences). Using those M values, the adjusted p-value is calculated as follows:

$$p_s^{adj} = \frac{\#\{\max(t^{*,m}) \geq t_s\} + 1}{M + 1} \tag{8}$$

After calculating p_s^{adj} for the first sequence, the column with the test statistic estimators generated from the first sequence is removed from the matrix. This now leaves a matrix that is $M \times (S - 1)$. The above procedure is then used to calculate p_s^{adj} for the sequence with the second largest test statistic, and then its column of test statistic estimators is removed, etc until all p_s^{adj} have been calculated. Following the algorithm described in Clarke et al.[57], we enforce monotonicity of the p-values by resetting the p-value for each sequence:

$$p_s^{adj} = \max\{p_s^{adj}, p_{s-1}^{adj}\} \tag{9}$$

This is done prior to calculating the p-value for the next sequence.

To control the false discovery proportion (i.e., control the number of false positives divided by the total number of hypotheses rejected), the idea of k-FWER is introduced[40]. For control of the FWER, k is equal to 1, and that leads to an α percent chance that there is at least 1 false positive among all hypotheses rejected (for FWER control $\alpha = 0.05$ is most commonly used). If k equals 2, then there is an α percent chance that there are at least 2 false positives among all hypotheses rejected. Therefore, to get control of the false discovery proportion we need to find the k that provides the proportion of interest given the number of hypotheses rejected. So, if we want a false discovery proportion, γ, of 0.05, we need $k \sim 0.05 \times \text{number of hypotheses rejected}$.

Romano and Wolf also derived an algorithm to do just that. The algorithm is as follows. Start with $k = 1$. Apply the k-FWER procedure, and note the total number of hypotheses rejected, N. If $N < \frac{k}{\gamma} - 1$, stop and you have identified k. Otherwise, increase k by 1, and repeat[40]. The way you determine k-FWER is in eq. 7, instead of taking the maximum value in each row, you take the kth largest value. Finally, to get median control of the false discovery proportion, $\alpha = 0.5$. This means that 50% of the time you will get a value greater than γ, and 50% of the time you will get a value less than γ, leading to median control of the false discovery proportion[40]. This is a similar decision to what is done when calculating the false discovery rate (FDR) with Benjamini-Hochberg or Benjamini-Yekutieli, except the false discovery rate controls the mean of the false discovery proportion instead of the median.

Lastly to avoid confusion around what the statistical control is we label the adjusted p-value with the k-FWER control (i.e. median control of the false discovery proportion) as ζ instead of p, comparable to the use of q with FDR control.

Power estimation for sample size calculation with CBAS. To estimate the statistical power of CBAS for a given sample size (Fig S1) we resampled the dataset without replacement and ran CBAS to determine the number of significant sequences. For each sample size we performed 20 repeats. We also ran a CBAS comparing each group to itself, with 20 repeats for each group; or correlating the sequence counts with a randomly generated set of CBIT scores drawn from the same distribution as the actual CBIT scores (Fig S1), with 40 repeats. Then for each sample size, the power is estimated by identifying the fraction of largest 20 number of significant sequences that came from the CBAS run on the resampled data (as opposed to the comparisons comparing groups to themselves or randomly generated CBIT).

Power estimation for sequence usage for "complete" sequences. Populations the same size as the datasets were bootstrapped (sampled with replacement). 1,000 CBAS were run on different bootstrapped samples. The smallest sequence usage rate for which there was 80% power to detect a difference was defined as the smallest usage rate from the sequences that were significant on 80% of the CBAS runs.

Spontaneous alternation in flies

Data. Data was downloaded from: https://lab.debivort.org/neuronal-control-of-locomotor-handedness/. The first 250 turns for each fly were used. Flies that did not do at least 250 turn

were excluded from the analysis. 466/1,225 CA and 565/1372 w1118 flies did not reach the 250-turn criterion.

Two-step task in humans

Data. Data was downloaded from: https://osf.io/usdgt/overview. All subjects in the dataset completed all 200 trials/400 choices. Therefore, all subjects were included in the analysis.

MB/MF RL model. The model was implemented following the description from Gillan et al.[43]. The model combines units for model-based and model-free control, as well as stickiness. For each trial, t, there are choices at two stages, $c_{1,t}$ and $c_{2,t}$. The choice at stage 1 leads stochastically to a stage 2 state, s_t. Reward, r_t, occurs following the second stage choice based on dynamic probabilities for each image at the second stage. $r_t = 1$ indicates reward and $r_t = -1$ indicates no reward. Stage 1 choices depend upon what is learned at stage 2; therefore, stage 2 learning is described first.

At stage 2, the model learns a value function, $Q_t^{\text{stage2}}(s, c)$, where s is state. The reward received at each trial leads to updating Q according to: $Q_{t+1}^{\text{stage2}}\left(s_t, c_{2,t}\right) = (1 - \alpha)Q_t^{\text{sttge2}}\left(s_t, c_{2,t}\right) + r_t$, where α is the learning rate parameter. The probability to make a specific stage 2 choice is governed by a softmax: $P\left(c_{2,t} = c\right) \propto \exp\left(\beta^{\text{stage2}} Q_t^{\text{stage2}}(s_t, c)\right)$, which is normalized over both options of c, and β^{stage2} is the inverse temperature parameter.

Stage 1 choices are determined by weighting the value predictions of the different components of the model about the stage 2 value of each stage 1 choice. Model-based values are dictated by the learned values of the stage 2 state, maximized over the two actions: $Q_t^{MB} = \max_{c_2}\left(Q_t^{\text{stage2}}(s, c_2)\right)$, where s is the stage 2 state predominantly produced by stage 1 choice c. Model-free values are dictated by two learning rules: $Q_{t+1}^{MF_0}(c_1, t) = (1 - \alpha)Q_t^{MF_0} + Q_t^{\text{stage2}}(s_t c_{2,t})$ and $Q_{t+1}^{MF_1}(c_1, t) = (1 - \alpha)Q_{t+1}^{MF_1} + r_t$. Stage 1 choice probabilities are also governed by a softmax: $P\left(c_{1,t} = c\right) \propto \exp\left(\beta^{MB} Q_t^{MB}(c) + \beta^{MF_0} Q_t^{MF_0}(c) + \beta^{MF_1} Q_t^{MF_1}(c) + \beta^{\text{stick}} I(c = c_{1,t-1})\right)$, where the β values are the temperature parameters for each component, and $I(c = c_{1,t-1})$ is 1 for the choice that repeats the one made on the previous trial and 0 for the other choice.

Following each trial, the value estimates in all Q for all unchosen actions and unvisited states are decayed multiplicatively by $(1 - \alpha)$.

Simulated annealing was used to fit the model, initializing with random parameters. At least 10 different fits were run for each set of parameters and the fit parameters leading to the largest log-likelihood was chosen. The data contained instances where a subject did not choose either of the images presented. Those choices were not included in the likelihood for the fitting. However, the lack of choice was coded as its own symbol and analyzed within CBAS for correlations with CBIT.

Spatial alternation in rats

Animals: All experiments on rats were conducted in accordance with University of California San Francisco Institutional Animal Care and Use Committee and US National Institutes of Health guidelines. Rats were fed standard rat chow (LabDiet 5001). To motivate the rats to perform

599 the task, reward was sweetened evaporated milk: 25 g of sugar per can (354 ml) of evaporated
600 milk (Carnation). The rats were food restricted to up to 85% of their basal body weight.

601 **Surgery:** Anesthesia was induced using ketamine, xylazine, atropine, and isoflurane. The
602 animals underwent surgery to either lesion the dorsal and intermediate hippocampus or as
603 control for the lesioning. For the lesion rats, craniotomies were drilled bilaterally to allow for
604 access to all 24 injection sites (Table 1) and the injection needle (33 gauge) was inserted
605 through the dura. For each lesion site, the needle was lowered 0.1 mm below the target
606 coordinates and then raised back to the target prior to injection. 120 or 110 nl of NMDA (20
607 mg/ml) was injected at each site at a rate of 0.1 µl/min, for males and females, respectively.
608 The needle was left at the site for 2 minutes after finishing the injection. For the rats that
609 underwent control surgeries, everything was the same as the lesion rats, except no NMDA was
610 injected, and the needle did not dwell at the hypothetical injection site. All animals that
611 underwent the lesion surgery received Diazepam 10 mg/kg intraperitoneally post-operatively.

612 Hippocampal lesioning was confirmed via micro-CT (Fig S4) collected at the Advanced
613 Light Source at Lawrence Berkeley National Lab. Brains were processed and imaged, and scans
614 were processed as described in Kastner et al.[59]. CT images were aligned using manual
615 registration in Slicer-3D (http://www.slicer.org)[60,61]. Male and female brains were separately
616 registered.

617 **Spatial alternation behavior:** The automated behavior system for spatial alternation behavior
618 was previously described[51]. There are different symbols on each arm of the track serving as
619 proximal cues, and there are distal cues distinguishing the different walls of the room.
620 Pneumatic pistons (Clippard) open and close the doors. Python scripts, run through Trodes
621 (Spike Gadgets), control the logic of the automated system. The reward wells contain an
622 infrared beam adjacent to the reward spigot. The automated system uses the breakage of that
623 infrared beam to progress through the logic of the behavior. In addition to the infrared beam
624 and the spigot to deliver the reward, each reward well has an associated white light LED.

625 Each cohort of rats is divided into groups of four (or three) animals. The same groups
626 were maintained throughout the duration of the experiment. Within a group, a given rat is
627 always placed in the same rest box, and the four rats of a group serially perform the behavior.
628 The rats have multiple sessions on the track each day. Prior to beginning the spatial alternation
629 task, the rats experience multiple days and sessions where they get rewarded at any arm that
630 they visit (provided it is not an immediate repeat). During this period of the behavior, the
631 duration of a session is defined by a fixed number of rewards, or a fixed amount of time on the
632 track (15 minutes), whichever came first. During the alternation task the duration of a session
633 was defined either by a fixed number of center arm visits and at least one subsequent visit to
634 any other arm, or a fixed amount of time on the track (15 minutes), whichever came first.

635 The algorithm underlying the spatial alternation task is such that three arms on the track
636 have the potential for reward: arms 2, 3, and 4 have the potential to be rewarded, and arms 1,
637 5, and 6 do not. Of those three potentially rewarded arms we refer to the middle of the three
638 arms as the center arm (arm 3) and the other two arms as the outer arms (arms 2 and 4).
639 Reward is delivered at the center arms if and only if: 1) the immediately preceding arm whose
640 reward well infrared beam was broken was not the center arm. Reward was delivered at the

outer two arms if and only if: 1) the immediately preceding arm whose reward well infrared beam was broken was the center arm, and 2) prior to breaking the infrared beam at the center arm, the most recently broken outer arm infrared beam was not the currently broken outer arm infrared beam. The one exception to the outer arm rules was at the beginning of a session, if no outer arm infrared beam was broken prior to the first infrared beam break at the center arm, then only the first condition had to be met.

For the running of the behavior, the infrared beam break determined an arm visit; however, the rats sometimes go down an arm, get very close to the reward wells, but do not break the infrared beam. Therefore, for all the analyses described for the rats, an arm choice is defined as when a rat gets close to a reward well. These times were extracted from a video recording of the behavior. This proximity-based definition of an arm visit added additional arm visits to those defined by the infrared beam breaks, and none of them could ever be rewarded, nor alter the logic of the underlying algorithm. However, because of the non-Markovian nature of the reward contingency, they could affect the rewards provided for subsequent choices.

For the initial dataset, a total of 46 rats that underwent control surgery (22 males, 24 females) and 45 rats that underwent hippocampal lesion surgery (23 males, 22 females) were run on the spatial alternation task. 6 of the rats that underwent hippocampal lesion surgery showed no evidence of lesioning on CT and were therefore not included in the analysis. For the replication dataset, a total of 9 rats that underwent control surgery (all males) and 11 rats that underwent hippocampal lesion surgery (all males) were run on the spatial alternation task. For the replication dataset, all of the rats that underwent hippocampal lesion surgery showed evidence of lesioning on CT. Rats were housed together prior to their undergoing surgery. They were single housed for post-operative recovery and through food restriction and the behavior.

The first 800 choices during the spatial alternation behavior for each rat were used. Only 1 rats did not reach the criterion. That control, female, rat from the initial dataset, completed 618 choices, and was included in the analysis.

Data and Code availability. Rat spatial alternation data generated for this work and code used to calculate CBAS will be posted to Github upon publication.

Acknowledgements. We thank Claire Gillan and Benjamin de Bivort for making their data publicly available. We thank Thomas Akam, Nathanial Daw, and Adam Frank for very helpful comments on the manuscript. This research used resources of the Advanced Light Source, a DOE Office of Science User Facility under contract no. DE-AC02-05CH11231. **Funding:** This work was supported by grants from the Jane Coffin Childs Memorial Fund for Medical Research (D.B.K.), the UCSF Physician Scientist Scholars Program (D.B.K.), an NIH R25 (R25MH060482) (D.B.K), a Simons Foundation Autism Research Initiative grant (899599) (D.B.K.), an NIMH K08 Award (K08MH139983) (D.B.K.), the Max Planck Society (P.D.) and the Humboldt Foundation (P.D.). **Author Contributions.** D.B.K. and P.D. designed the study, D.B.K., J.P.R., and P.D. developed the analysis method, D.B.K., N.O.Y., C.Y.L., C.H., J.T., A.C., M.M., V.K. and G.W. collected the data, D.L.P. provided expertise and resources. D.B.K., N.O.Y., C.Y.L., and M.M. analyzed the data, D.B.K and P.D. wrote the manuscript, and D.B.K., J.P.R. and P.D. edited the manuscript. **Declaration of interests.** The authors declare no competing interests.

682 **Figure Legends**

683 **Figure 1. Choice-Wide Behavioral Association Study (CBAS) is applicable across diverse tasks**
684 **and species.** (**a**) The starting point for the three behavioral tasks we consider is to decompose
685 the actions of the subjects into a series of discrete choices. Left: Fly y-maze (top), and example
686 snippet of the choices of left (red) and right (blue) turns (bottom) of an individual fly. Data from
687 Buchanan et al.[31]. Middle: Two-step task, performed by human subjects (top), and an example
688 snippet of choices of an individual subject (bottom). The colors correspond to the different
689 images chosen (bottom). Data from Gillan et al.[43]. Right: Spatial alternation behavior,
690 performed by rats (top), and an example snippet of choices of an individual rat (bottom). Colors
691 correspond to the arm chosen by the rat. The rats are rewarded if they visit arms 2, 3, and 4 by
692 alternating visits to arms 2 and 4 after visiting arm 3 (arrows). The flies were not rewarded for
693 their choices; for the middle and right panels, filled circles indicate that reward was received.
694 (**b**) CBAS breaks down the task into all sequences of choices up to a defined total length and
695 compares the usage rate of the sequences between two groups (flies and rats) or correlates the
696 usage rate of each sequence with a covariate of interest (humans). In all cases, CBAS controls
697 statistically for the large number of sequence comparisons. Left: The occurrence of a single
698 sequence of 10 turns in the two fly strains during the first 250 turns of the spontaneous
699 alternation behavior. The rows show the occurrence of that sequence of turns for all individual
700 flies from each strain. For this example sequence, the CA strain has a rate of sequence usage of
701 9.9 per animal. The w1118 strain has a rate of sequence usage of just 1.5 per animal. This
702 difference in rate is significant with $\zeta = 1x10^{-4}$, corrected for the multiple comparisons.
703 Middle: Correlation between the participants' CBIT scores and their usage of the particular
704 sequence 1**A**2, which means choosing image 1, then making a choice in set A and getting
705 rewarded and then choosing image 2. The usage of this sequence shows a correlation of 0.12
706 ($\zeta < 2.8x10^{-3}$) with the CBIT score of each participant. Dotted line shows the linear fit to the
707 data. Right: The occurrence of a single sequence of arm choices 5654**3** (unrewarded for the first
708 5 choices but then rewarded for the move into the central arm of the rules, arm 3) in the two
709 groups of rats (control and hippocampal lesioned) during the first 800 arm choices. The rows
710 show the occurrence of that sequence of choices for all individual rats from each group. Control
711 rats use this sequence at a rate of sequence usage of 1.8 per animal. Hippocampal lesioned rats
712 use this sequence at a rate of sequence usage of 0.6 per animal. This difference in rate is
713 significant with $\zeta = 0.047$. (**c**) CBAS Manhattan plot displays the p-value for each sequence
714 evaluated in the 3 datasets. The sequences are ordered based on the number of choices each
715 involves (shown by colors, up to the total length considered), and they are displayed on a log
716 scale to make all sequence lengths visible. Within a given sequence length, the sequences are
717 ordered based on frequency of occurrence in the entire dataset. The horizontal dotted line
718 indicates the significance threshold of γ=5% control of the median false discovery proportion
719 ($\zeta = 0.5$). A total of 10,000 resamples were used to calculate the p-values, making the
720 maximum value on the plot $-\log_{10}\left(\frac{1}{10{,}001}\right) = 4.00004$. The location of the example sequences
721 shown in **b** are indicated by arrows.

722 **Figure 2. CBAS applied to fly spontaneous alternation produces interpretable differences**
723 **between fly strains.** CBAS was applied to flies tracked as they spontaneously made left and

right turns on a y-maze. (**a**) All sequences up to length 10 identified as significantly different between the CA and w1118 strains are displayed with left turns indicated in red and rights turns indicated in blue. Each row displays a sequence. The top set of rows shows the sequences CA uses more than w1118 and the bottom set of rows are the sequences w1118 use more than CA. Sequences are ordered in columns based on persistence, i.e., the maximum number of the same turns in a row in the sequence. For example, the sequence displayed in figure 1b, RLRLRLLLLL, is in the column labeled 5 (indicated with an arrow) because there are 5 left turns in a row, and that is the maximum number of the same turns in a row in the sequence. (**b**) The number of significant sequences as defined by CBAS is displayed for the data (black line and circles) and three different models (colors and box and whisker displays) that might characterize the turns of the individual flies. The turn bias model (orange) randomly generates each turn with a bias based on the individual fly's overall right turn bias over the first 250 turns. The persistence rate model (green) randomly generates each turn with a bias based on the individual fly's overall rate of continuing in the same direction. The persistence distribution model (purple) randomly generates a sequence of the same direction of turns based on the individual fly's probability distribution of continuing for any number of turns in a row. Top plots organize the sequences based on the maximum number of the same turns in row in the sequence (as in panel **a**). Bottom plots bin the same set of sequences based upon the fraction of right turns that occur in the sequence (30% for the example sequence in figure 1b). The left column quantifies the sequences that CA use more than w1118 and the right column the sequences w1118 use more than CA (or the equivalent model versions of each line). For the models, CBAS was run on 100 repeats of all three models, tailored to the choices of each individual fly. The center line of the box displays the median of the model output, the top and bottom lines of the box show the 75th and 25th quartiles, respectively, and the end of the whiskers show the full range of the model outputs.

Figure 3. Scaling of CBAS with sample size. For all panels in this figure, 40 repeats of smaller sample sizes were generated from the fly data by selecting at random subjects within the two groups (without replacement). (**a**) Each CBAS run on the smaller sample size identified some number of significant sequences. Those sequences were compared to the sequences identified in the CBAS on the full dataset (with 759 CA and 807 w1118 flies), and the graph shows the ratios of the numbers of sequences identified in each instance of the smaller sample size that were also found in the full dataset to the total number of sequences identified in the full dataset. As a comparison, the ratio of the number of flies in the smaller CBAS to the total number of flies in the dataset is plotted in the maroon horizontal lines. (**b**) The ratio of the number of significant sequences identified in each instance of the smaller sample size that were not also identified in the full dataset to the total number of significant sequences identified in that instance of the smaller sample size. (**c**) In creating the smaller samples sizes, 20 paired sets of animals were generated such that the two groups of flies within a paired set had no overlapping individuals. Plotted are the number of sequences that were found to be significant in both groups of flies within a paired set of nonoverlapping flies over the average number of significant sequences found between the two paired set of flies. (**d**) For each instance of each sample size, all of the sequences were categorized based on the maximum number of turns in the same direction and the fraction of significant sequences within those categories are plotted

767 to allow for an easier comparison across groups with very different numbers of sequences, i.e.,
768 there are only 2 sequences that have a maximum of 10 of the same turns in a row, whereas
769 there are over 600 with a maximum of 3 of the same turns in a row. The bottom row of this plot
770 is from the full dataset. (**e**) The Euclidean distance between each row from panel **d** and the full
771 dataset row is plotted. Colors correspond to the sample sizes as shown in panel **d**. For panels **a**
772 – **c** the center line of the box displays the median of the resamples, the top and bottom lines of
773 the box show the 75th and 25th quartiles, respectively, and the end of the whiskers show the full
774 range of the resampled outputs.

775 **Figure 4. CBAS helps interpret the correlation between CBIT and the two step task.** CBAS was
776 applied to humans performing the two-step task with sequences of choices up to length 4. (**a**)
777 All sequences whose usage was positively (left) and negatively (right) correlated with the
778 individual subject CBIT factor score are displayed. Colors correspond to the images and image
779 sets chosen as indicated by the diagram of the task (right). Upper case bold letters indicate that
780 a reward was received. The majority of the sequences positively correlated with CBIT include **B**1
781 and/or **A**2, and so the first two columns are arranged to center these motifs. For illustrative
782 purposes, the two sequences **B**1**A**2 and **A**2**B**1 are duplicated and listed in both columns. (**b**) The
783 number of significant sequences as defined by CBAS is displayed for the data (black line and
784 circles) and two different models of the task (colors and box and whisker displays). Sequences
785 are organized based upon their containing the subunit indicated on the x-axis. MB/MF RL model
786 (yellow) is the model-based/model-free RL model from Gillan et al.[43]. Each subject was fit to the
787 model, and their parameters were used to generate choices for CBAS. Model inconsistent
788 choices (grey) were generated by using the fit parameters from the MB/MF RL model for a
789 single subject, but whenever the sequence 1A or 2B was chosen, the subsequent choice was
790 probabilistically determined based on the "model inconsistent" rate of the individual subject.
791 The "model inconsistent" rate for a subject was their rate of choosing a 2 following 1A or a 1
792 following 2B. For each CBAS repeat of the model inconsistent choices a different subject's fit
793 parameters were used. For the models, CBAS was run on 100 repeats. The center line of the
794 box displays the median of the model repeats, the top and bottom lines of the box show the
795 75th and 25th quartiles, respectively, and the end of the whiskers show the full range of the odel
796 repeats. (**c**) Data generated from the sequences of choices matched to the model inconsistent
797 rate of each subject (as described above and shown in grey in panel **b**) was fit to the MB/MF RL
798 model and the p-value of the correlation of the model based parameter of the model fit with
799 CBIT is displayed as function of the correlation value. Dotted horizontal lines show $p = 0.05$, and
800 $p = 0.0083$. The latter number is the Bonferroni corrected p-value to control for the 6
801 parameters of the model. Bonferroni correction is used instead of Romano-Wolf, as that would
802 be the most typical, if not conservative, correction used. A total of 50 model repeats was fit,
803 86% and 56% had p-values <0.05 and 0.0083, respectively.

804 **Figure 5. CBAS identifies hippocampal based learning strategy for spatial alternation task in**
805 **rats.** CBAS was applied to rats performing the spatial alternation task with sequences of choices
806 up to length 6. (**a**) "Complete" sequences that the control rats use more than hippocampal
807 lesion (left) and that the lesion rats use more than control (right). "Complete" sequences are
808 ones that that always significantly different between the groups (in the same direction) no
809 matter what choices occur before or after up to the maximum length of 6 (see Fig. S2).

Emboldened numbers indicate that choice was rewarded. (**b**) The "complete" sequences differ between the groups in two broad ways: the consistency of the direction of travel (greater for controls); and the total absolute distance covered (so 21 would contribute a distance of 1 and 14 would contribute a distance of 3). This plot organizes all the sequences (i.e., not just the "complete" sequences) according to these two features (the former defining the rows, with NA for sequences with fewer than three choices for which this cannot be calculated; the latter the columns). For example, sequence **3**2**1**2**3**, is located in the row segment labeled 2 as two of the choices in the sequence continue in the same direction: **3**-**2**-1 (the 2-1 is in the same direction as 3-2) and 1-2-**3** (the 2-3 is in the same direction as 1-2), and is located in the column segment labeled 4, as all of the choices are 1 arm away from the previous choice. Sequences that are 1 choice long have a total transition number of 0. Each sequence in the plot is displayed along the tiny rows ordered by length, and, for completeness, the choices are colored based on the color of the arm numbers in the task diagram. (**c**) Given that the sequences that occur more in the control than the lesion rats tend to have more choices that continue in the same direction and smaller total transition numbers, we captured these feature of the data using the metric of transition normalized directional inertia, which measures the tendency to continue in the same direction divided by the distance of the transition. Transition normalized directional inertia was smoothed and averaged (±sem) across all rats in the control (black) and lesion (brown) groups and shown for the first 800 choices of the spatial alternation task. Control rats show significantly larger transition normalized directional inertia summed across the 800 choices of the spatial alternation task, $p = 4.0 \times 10^{-4}$ (permutation test).

Figure 6. Replication of the CBAS generated hypothesis in the rat spatial alternation task. CBAS was applied to a separate set of control and hippocampal lesion rats performing the spatial alternation task. (**a**) CBAS Manhattan plot displaying the p-value for each sequence compared between the two groups. Plot structured as described in Fig. 1c. (**b**) Significant sequences structured by the two metrics, number of choices in the same direction and total transition number, as described in Fig. 5b. (**c**) Transition normalized directional inertia as the two groups of rats in the replication dataset learn the spatial alternation task. Control rats show significantly larger transition normalized direction inertia summed across the 800 choices of the spatial alternation task, $p = 2.2 \times 10^{-3}$ (permutation test). (**d**) All of the significant sequences that CBAS identifies as occurring more in the control than lesion rats (two left columns) and as occurring less in the control than lesion rats (rightmost column). 65.4% (88.2% of for control > lesion, 22.2% of for control < lesion) of the sequences CBAS identifies as significant in this dataset were also significant in the initial dataset (indicated by boxes).

Supplementary Figure 1. Sample size calculation for CBAS (**a**) Flies: Left: the number of significant sequences when randomly resampling the populations without replacement and calculating CBAS on the smaller sample sizes. Middle: the number of significant sequences when comparing smaller samples sizes of the CA strain to itself, with nonoverlapping individuals in each group. Right: the number of significant sequences when comparing smaller samples sizes of the w1118 strain to itself, with nonoverlapping individuals in each group. (**b**) Power estimate for different sized groups of flies in each strain (see methods). (**c**) Humans: Left: the number of sequences significantly correlated with CBIT when randomly resampling the population separately without replacement and calculating CBAS on the smaller sample sizes. Right: the number of significant sequences when randomly resampling the population without replacement and comparing it to randomly generated CBIT scores drawn from a distribution imputed from the original CBIT scores and calculating CBAS on the smaller sample sizes. (**d**) Power estimate for different sized groups of human subjects (see methods). (**e**) Rats: Left: the number of significant sequences when randomly resampling the populations without replacement and calculating CBAS on the smaller sample sizes. Middle: the number of significant sequences when comparing smaller samples sizes of the control rats to itself, with nonoverlapping individuals in each group. Right: the number of significant sequences when comparing smaller samples sizes of the hippocampal lesion rats to itself, with nonoverlapping individuals in each group. (**f**) Power estimate for different sized groups of rats (see methods). For **a**, **c**, and **e** data points are overlayed by box and whisker plots. The center line of the box displays the median of the data, the top and bottom lines of the box show the 25th and 75th quartiles, respectively, and the end of the whiskers show the full range of the data. Additionally, the horizontal dotted lines for the left-most plot indicates the total number of significant sequences CBAS find in the entire dataset. For **b**, **d**, and **f** horizontal dotted line shows a value of 80%.

Supplementary Figure 2. "Complete" sequence calculation. (**a**) A sequence unit is evaluated, and all sequences in the dataset that contain that unit are found. For the example shown, the sequence unit is: **3**2**1**2**3** (indicated by an arrow). As usual, emboldened numbers indicate that that choice was rewarded. Sequences are colored as in Fig 5 and 6. The second column lists the rate at which each sequence was utilized by control animals (as this sequence occurs more in control than lesion rats). The rate in the control population above which there is 80% power to significantly detect a sequence difference was determined to be 0.174. The third column lists whether sequences occurred above the power level. The fourth column indicates which sequences were determined to be significant, and in what direction. A "complete" sequence is therefore defined as a sequence for which all longer sequences that contain that sequence of interest, and which occur above a rate that allows for 80% power for differentiation, are found by CBAS to differ significantly in the same direction as the sequence of interest (e.g., are all found to be more prevalent in the control, than the lesion, rats).

Supplementary Figure 3. "Complete" sequences for fly and human tasks. All the sequences found to be complete for the fly task (**a**) and the human task (**b**) are shown. There is no "complete" sequence that is negatively correlated with CBIT for the human two-step task. For **b** bold and upper-case letters indicate that reward was received for that choice.

886 **Supplementary Figure 4. Micro-CT used to confirm hippocampal lesioning.** Example virtual
887 slices, displaying two different parts of the hippocampus (**a**, anterior, and **b**, posterior), from
888 the micro CTs of a control and lesioned brain for males (top) and females (bottom).
889 Additionally, the averages for the virtual slices for the control males (n = 22) and females (n =
890 18) and lesioned males (n = 24) and females (n = 21) are displayed. The blur in the averaging is
891 due to imperfect registration across the different brains, as only a linear transformation was
892 used for the registration. Note that hippocampal cell layers remain well defined and intact in
893 the average images for the control, but largely absent in the lesioned average images, as
894 revealed by the irregular darker areas with no tissue in the location of the hippocampus. The
895 scale of the images are the same between example control and lesion images, and between
896 average control and lesion images, but the scale of the average images have an increased gain
897 compared to the example images to better highlight the features in the average images.

AP (mm)	ML (mm)	DV (mm)
-2.8	±1.4	-3.0
-3.3	±2.4	-3.0
-4.1	±1.8	-2.8
-4.1	±3.4	-2.8
-4.8	±2.0	-2.8
-4.8	±4.2	-3.1
-5.5	±2.6	-3.0
-5.5	±3.6	-2.9
-5.5	±5.0	-3.5
-5.5	±5.0	-5.5
-6.2	±4.0	-3.4
-6.2	±5.4	-4.4

898 **Table 1:** Stereotaxic coordinates of NMDA infusions to produce lesions within the hippocampal
899 formation. Adapted from Kim & Frank[45]. The coordinates are given for a Long Evans rat skull
900 which is leveled so that bregma and lambda lie in the same horizontal plane. AP:
901 anteroposterior, ML: mediolateral, DV: dorsoventral. The AP and ML coordinates are measured
902 from bregma. DV coordinates are measured from the dural surface above the brain.

903 1. Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., and Poeppel, D. (2017).
904 Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron *93*, 480–490.
905 https://doi.org/10.1016/j.neuron.2016.12.041.

906 2. Bialek, W. (2022). On the dimensionality of behavior. P Natl Acad Sci Usa *119*, e2021860119.
907 https://doi.org/10.1073/pnas.2021860119.

908 3. Meister, M. (2022). Learning, fast and slow. Curr Opin Neurobiol *75*, 102555.
909 https://doi.org/10.1016/j.conb.2022.102555.

910 4. Gomez-Marin, A., and Ghazanfar, A.A. (2019). The Life of Behavior. Neuron *104*, 25–36.
911 https://doi.org/10.1016/j.neuron.2019.09.017.

912 5. Trillmich, F., Günther, A., Müller, C., Reinhold, K., and Sachser, N. (2015). New perspectives in
913 behavioural development: adaptive shaping of behaviour over a lifetime? Front Zool *12*, S1.
914 https://doi.org/10.1186/1742-9994-12-s1-s1.

915 6. Dayan, P., Roiser, J.P., and Viding, E. (2018). The first steps on long marches: the costs of
916 active observation. In Rethinking Biopsychosocial Psychiatry. Rethinking Biopsychosocial
917 Psychiatry., W. Davies, R. Roache, and J. Savulescu, eds. (Rethinking Biopsychosocial
918 Psychiatry.).

919 7. Hession, L.E., Sabnis, G.S., Churchill, G.A., and Kumar, V. (2022). A machine-vision-based
920 frailty index for mice. Nat. Aging *2*, 756–766. https://doi.org/10.1038/s43587-022-00266-0.

921 8. Geuther, B.Q., Peer, A., He, H., Sabnis, G., Philip, V.M., and Kumar, V. (2021). Action detection
922 using a neural network elucidates the genetics of mouse grooming behavior. eLife *10*, e63207.
923 https://doi.org/10.7554/elife.63207.

924 9. Geuther, B., Chen, M., Galante, R.J., Han, O., Lian, J., George, J., Pack, A.I., and Kumar, V.
925 (2021). High-throughput visual assessment of sleep stages in mice using machine learning.
926 SLEEP *45*, zsab260. https://doi.org/10.1093/sleep/zsab260.

927 10. Sheppard, K., Gardin, J., Sabnis, G.S., Peer, A., Darrell, M., Deats, S., Geuther, B., Lutz, C.M.,
928 and Kumar, V. (2022). Stride-level analysis of mouse open field behavior using deep-learning-
929 based pose estimation. Cell Rep. *38*, 110231. https://doi.org/10.1016/j.celrep.2021.110231.

930 11. Weinreb, C., Pearl, J., Lin, S., Osman, M.A.M., Zhang, L., Annapragada, S., Conlin, E.,
931 Hoffman, R., Makowska, S., Gillis, W.F., et al. (2023). Keypoint-MoSeq: parsing behavior by
932 linking point tracking to pose dynamics. bioRxiv, 2023.03.16.532307.
933 https://doi.org/10.1101/2023.03.16.532307.

12. Wiltschko, A.B., Johnson, M.J., Iurilli, G., Peterson, R.E., Katon, J.M., Pashkovski, S.L., Abraira, V.E., Adams, R.P., and Datta, S.R. (2015). Mapping Sub-Second Structure in Mouse Behavior. Neuron *88*, 1121–1135. https://doi.org/10.1016/j.neuron.2015.11.031.

13. Pereira, T.D., Aldarondo, D.E., Willmore, L., Kislin, M., Wang, S.S.-H., Murthy, M., and Shaevitz, J.W. (2019). Fast animal pose estimation using deep neural networks. Nat. Methods *16*, 117–125. https://doi.org/10.1038/s41592-018-0234-5.

14. Pereira, T.D., Tabris, N., Matsliah, A., Turner, D.M., Li, J., Ravindranath, S., Papadoyannis, E.S., Normand, E., Deutsch, D.S., Wang, Z.Y., et al. (2022). SLEAP: A deep learning system for multi-animal pose tracking. Nat. Methods *19*, 486–495. https://doi.org/10.1038/s41592-022-01426-1.

15. Marshall, J.D., Aldarondo, D.E., Dunn, T.W., Wang, W.L., Berman, G.J., and Ölveczky, B.P. (2020). Continuous Whole-Body 3D Kinematic Recordings across the Rodent Behavioral Repertoire. Neuron *109*, 420-437.e8. https://doi.org/10.1016/j.neuron.2020.11.016.

16. Dunn, T.W., Marshall, J.D., Severson, K.S., Aldarondo, D.E., Hildebrand, D.G.C., Chettih, S.N., Wang, W.L., Gellis, A.J., Carlson, D.E., Aronov, D., et al. (2021). Geometric deep learning enables 3D kinematic profiling across species and environments. Nat. Methods *18*, 564–573. https://doi.org/10.1038/s41592-021-01106-6.

17. Huys, Q.J.M., Maia, T.V., and Frank, M.J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. Nat Neurosci *19*, 404–413. https://doi.org/10.1038/nn.4238.

18. Kumar, P., Dayan, P., and Wolfers, T. (2024). From Complexity to Precision—Charting Decision-Making Through Normative Modeling. JAMA Psychiatry *81*. https://doi.org/10.1001/jamapsychiatry.2023.4611.

19. Nickerson, R.S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. Rev Gen Psychol *2*, 175–220. https://doi.org/10.1037//1089-2680.2.2.175.

20. Kerr, N.L. (1998). HARKing: Hypothesizing After the Results are Known. Pers Soc Psychol Rev *2*, 196–217. https://doi.org/10.1207/s15327957pspr0203_4.

21. Wiltschko, A.B., Tsukahara, T., Zeine, A., Anyoha, R., Gillis, W.F., Markowitz, J.E., Peterson, R.E., Katon, J., Johnson, M.J., and Datta, S.R. (2020). Revealing the structure of pharmacobehavioral space through Motion Sequencing. Nat. Neurosci. *23*, 1433–1443. https://doi.org/10.1038/s41593-020-00706-3.

22. Gschwind, T., Zeine, A., Raikov, I., Markowitz, J.E., Gillis, W.F., Felong, S., Isom, L.L., Datta, S.R., and Soltesz, I. (2023). Hidden behavioral fingerprints in epilepsy. Neuron *111*, 1440-1452.e5. https://doi.org/10.1016/j.neuron.2023.02.003.

968 23. Dezfouli, A., Griffiths, K., Ramos, F., Dayan, P., and Balleine, B.W. (2019). Models that learn
969 how humans learn: The case of decision-making and its disorders. PLoS Comput. Biol. *15*,
970 e1006903. https://doi.org/10.1371/journal.pcbi.1006903.

971 24. Eckstein, M.K., Summerfield, C., Daw, N.D., and Miller, K.J. (2023). Predictive and
972 Interpretable: Combining Artificial Neural Networks and Classic Cognitive Models to
973 Understand Human Learning and Decision Making. bioRxiv, 2023.05.17.541226.
974 https://doi.org/10.1101/2023.05.17.541226.

975 25. Castro, P.S., Tomasev, N., Anand, A., Sharma, N., Mohanta, R., Dev, A., Perlin, K., Jain, S.,
976 Levin, K., Éltető, N., et al. (2025). Discovering Symbolic Cognitive Models from Human and
977 Animal Behavior. bioRxiv, 2025.02.05.636732. https://doi.org/10.1101/2025.02.05.636732.

978 26. Uffelmann, E., Huang, Q.Q., Munung, N.S., Vries, J. de, Okada, Y., Martin, A.R., Martin, H.C.,
979 Lappalainen, T., and Posthuma, D. (2021). Genome-wide association studies. Nat Rev Methods
980 Primers *1*, 59. https://doi.org/10.1038/s43586-021-00056-9.

981 27. Farrell, M.S., Werge, T., Sklar, P., Owen, M.J., Ophoff, R.A., O'Donovan, M.C., Corvin, A.,
982 Cichon, S., and Sullivan, P.F. (2015). Evaluating historical candidate genes for schizophrenia. Mol
983 Psychiatr *20*, 555–562. https://doi.org/10.1038/mp.2015.16.

984 28. Border, R., Johnson, E.C., Evans, L.M., Smolen, A., Berley, N., Sullivan, P.F., and Keller, M.C.
985 (2019). No Support for Historical Candidate Gene or Candidate Gene-by-Interaction Hypotheses
986 for Major Depression Across Multiple Large Samples. Am J Psychiat *176*, 376–387.
987 https://doi.org/10.1176/appi.ajp.2018.18070881.

988 29. Duncan, L.E., Ostacher, M., and Ballon, J. (2019). How genome-wide association studies
989 (GWAS) made traditional candidate gene studies obsolete. Neuropsychopharmacol *44*, 1518–
990 1523. https://doi.org/10.1038/s41386-019-0389-5.

991 30. Ioannidis, J.P.A. (2005). Why Most Published Research Findings Are False. Plos Med *2*, e124.
992 https://doi.org/10.1371/journal.pmed.0020124.

993 31. Buchanan, S.M., Kain, J.S., and Bivort, B.L. de (2015). Neuronal control of locomotor
994 handedness in Drosophila. Proc National Acad Sci *112*, 6700–6705.
995 https://doi.org/10.1073/pnas.1500804112.

996 32. Skutt-Kakaria, K., Reimers, P., Currier, T.A., Werkhoven, Z., and Bivort, B.L. de (2019). A
997 neural circuit basis for context-modulation of individual locomotor behavior. bioRxiv, 797126.
998 https://doi.org/10.1101/797126.

999 33. Ayroles, J.F., Buchanan, S.M., O'Leary, C., Skutt-Kakaria, K., Grenier, J.K., Clark, A.G., Hartl,
1000 D.L., and Bivort, B.L. de (2015). Behavioral idiosyncrasy reveals genetic control of phenotypic
1001 variability. Proc National Acad Sci *112*, 6706–6711. https://doi.org/10.1073/pnas.1503830112.

1002 34. Werkhoven, Z., Bravin, A., Skutt-Kakaria, K., Reimers, P., Pallares, L.F., Ayroles, J., and Bivort,
1003 B.L. de (2021). The structure of behavioral variation within a genotype. eLife *10*, e64988.
1004 https://doi.org/10.7554/elife.64988.

1005 35. Akhund-Zade, J., Ho, S., O'Leary, C., and Bivort, B. de (2019). The effect of environmental
1006 enrichment on behavioral variability depends on genotype, behavior, and type of enrichment. J.
1007 Exp. Biol. *222*, jeb202234. https://doi.org/10.1242/jeb.202234.

1008 36. Bivort, B. de, Buchanan, S., Skutt-Kakaria, K., Gajda, E., Ayroles, J., O'Leary, C., Reimers, P.,
1009 Akhund-Zade, J., Senft, R., Maloney, R., et al. (2022). Precise Quantification of Behavioral
1010 Individuality From 80 Million Decisions Across 183,000 Flies. Front. Behav. Neurosci. *16*,
1011 836626. https://doi.org/10.3389/fnbeh.2022.836626.

1012 37. Cui, T., Mekkaoui, K.E., Reinvall, J., Havulinna, A.S., Marttinen, P., and Kaski, S. (2022).
1013 Gene–gene interaction detection with deep learning. Commun. Biol. *5*, 1238.
1014 https://doi.org/10.1038/s42003-022-04186-y.

1015 38. Fang, G., Wang, W., Paunic, V., Heydari, H., Costanzo, M., Liu, X., Liu, X., VanderSluis, B.,
1016 Oately, B., Steinbach, M., et al. (2019). Discovering genetic interactions bridging pathways in
1017 genome-wide association studies. Nat. Commun. *10*, 4274. https://doi.org/10.1038/s41467-
1018 019-12131-7.

1019 39. Romano, J.P., and Wolf, M. (2005). Exact and Approximate Stepdown Methods for Multiple
1020 Hypothesis Testing. J Am Stat Assoc *100*, 94–108.
1021 https://doi.org/10.1198/016214504000000539.

1022 40. Romano, J.P., and Wolf, M. (2007). Control of generalized error rates in multiple testing.
1023 Ann Statistics *35*, 1378–1408. https://doi.org/10.1214/009053606000001622.

1024 41. Daw, N.D., Gershman, S.J., Seymour, B., Dayan, P., and Dolan, R.J. (2011). Model-Based
1025 Influences on Humans' Choices and Striatal Prediction Errors. Neuron *69*, 1204–1215.
1026 https://doi.org/10.1016/j.neuron.2011.02.027.

1027 42. Akam, T., Rodrigues-Vaz, I., Marcelo, I., Zhang, X., Pereira, M., Oliveira, R.F., Dayan, P., and
1028 Costa, R.M. (2020). The Anterior Cingulate Cortex Predicts Future States to Mediate Model-
1029 Based Action Selection. Neuron *109*, 149-163.e7.
1030 https://doi.org/10.1016/j.neuron.2020.10.013.

1031 43. Gillan, C.M., Kosinski, M., Whelan, R., Phelps, E.A., and Daw, N.D. (2016). Characterizing a
1032 psychiatric symptom dimension related to deficits in goal-directed control. eLife *5*, e11305.
1033 https://doi.org/10.7554/elife.11305.

1034 44. Silva, C.F. da, and Hare, T.A. (2020). Humans primarily use model-based inference in the
1035 two-stage task. Nat. Hum. Behav. *4*, 1053–1066. https://doi.org/10.1038/s41562-020-0905-y.

45. Kim, S.M., and Frank, L.M. (2009). Hippocampal Lesions Impair Rapid Learning of a Continuous Spatial Alternation Task. Plos One *4*, e5494. https://doi.org/10.1371/journal.pone.0005494.

46. Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR Activity Predicts Correct Decisions during the Initial Learning of an Alternation Task. Neuron *77*, 1163–1173. https://doi.org/10.1016/j.neuron.2013.01.027.

47. Kay, K., Chung, J.E., Sosa, M., Schor, J.S., Karlsson, M.P., Larkin, M.C., Liu, D.F., and Frank, L.M. (2020). Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. Cell *180*, 552-567.e25. https://doi.org/10.1016/j.cell.2020.01.014.

48. Fernández-Ruiz, A., Oliva, A., Oliveira, E.F. de, Rocha-Almeida, F., Tingley, D., and Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. Science *364*, 1082–1086. https://doi.org/10.1126/science.aax0758.

49. Sigurdsson, T., Stark, K.L., Karayiorgou, M., Gogos, J.A., and Gordon, J.A. (2010). Impaired hippocampal–prefrontal synchrony in a genetic mouse model of schizophrenia. Nature *464*, 763–767. https://doi.org/10.1038/nature08855.

50. Kastner, D.B., Gillespie, A.K., Dayan, P., and Frank, L.M. (2020). Memory alone does not account for the speed of learning of a simple spatial alternation task in rats. J Neurosci *40*, JN-RM-0972-20. https://doi.org/10.1523/jneurosci.0972-20.2020.

51. Kastner, D.B., Miller, E.A., Yang, Z., Roumis, D.K., Liu, D.F., Frank, L.M., and Dayan, P. (2022). Spatial preferences account for inter-animal variability during the continual learning of a dynamic cognitive task. Cell Reports *39*, 110708–110708. https://doi.org/10.1016/j.celrep.2022.110708.

52. Bennett, D., Nakamura, J., Vinnakota, C., Sokolenko, E., Nithianantharajah, J., Buuse, M. van den, Jones, N.C., Sundram, S., and Hill, R. (2023). Mouse Behavior on the Trial-Unique Nonmatching-to-Location (TUNL) Touchscreen Task Reflects a Mixture of Distinct Working Memory Codes and Response Biases. J. Neurosci. *43*, 5693–5709. https://doi.org/10.1523/jneurosci.2101-22.2023.

53. Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. Science *336*, 1454–1458. https://doi.org/10.1126/science.1217230.

54. Whittington, J.C.R., Muller, T.H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T.E.J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. Cell *183*, 1249-1263.e23. https://doi.org/10.1016/j.cell.2020.10.024.

1070 55. Benna, M.K., and Fusi, S. (2021). Place cells may simply be memory cells: Memory
1071 compression leads to spatial tuning and history dependence. Proc. Natl. Acad. Sci. *118*,
1072 e2018422118. https://doi.org/10.1073/pnas.2018422118.

1073 56. Markowitz, J.E., Gillis, W.F., Jay, M., Wood, J., Harris, R.W., Cieszkowski, R., Scott, R., Brann,
1074 D., Koveal, D., Kula, T., et al. (2023). Spontaneous behaviour is structured by reinforcement
1075 without explicit reward. Nature *614*, 108–117. https://doi.org/10.1038/s41586-022-05611-2.

1076 57. Clarke, D., Romano, J.P., and Wolf, M. (2020). The Romano-Wolf Multiple Hypothesis
1077 Correction in Stata. Ssrn Electron J. https://doi.org/10.2139/ssrn.3513687.

1078 58. DiCiccio, C.J., and Romano, J.P. (2017). Robust Permutation Tests For Correlation And
1079 Regression Coefficients. J. Am. Stat. Assoc. *112*, 1211–1220.
1080 https://doi.org/10.1080/01621459.2016.1202117.

1081 59. Kastner, D.B., Kharazia, V., Nevers, R., Astudillo-Maya, D.A., Williams, G.M., Yang, Z., Smyth,
1082 C., Santina, L.D., Parkinson, D.Y., and Frank, L.M. (2020). Scalable method for micro-CT analysis
1083 enables large scale quantitative characterization of brain lesions and implants. Sci Rep *10*,
1084 20851. https://doi.org/10.1038/s41598-020-77796-3.

1085 60. Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C.,
1086 Jennings, D., Fennessy, F., Sonka, M., et al. (2012). 3D Slicer as an image computing platform for
1087 the Quantitative Imaging Network. Magn Reson Imaging *30*, 1323–1341.
1088 https://doi.org/10.1016/j.mri.2012.05.001.

1089 61. Kikinis, R., Pieper, S.D., and Vosburgh, K.G. (2013). Intraoperative Imaging and Image-
1090 Guided Therapy. In Intraoperative Imaging and Image-Guided Therapy Intraoperative Imaging
1091 and Image-Guided Therapy. (Intraoperative Imaging and Image-Guided Therapy), pp. 277–289.
1092 https://doi.org/10.1007/978-1-4614-7657-3_19.

Figure 1

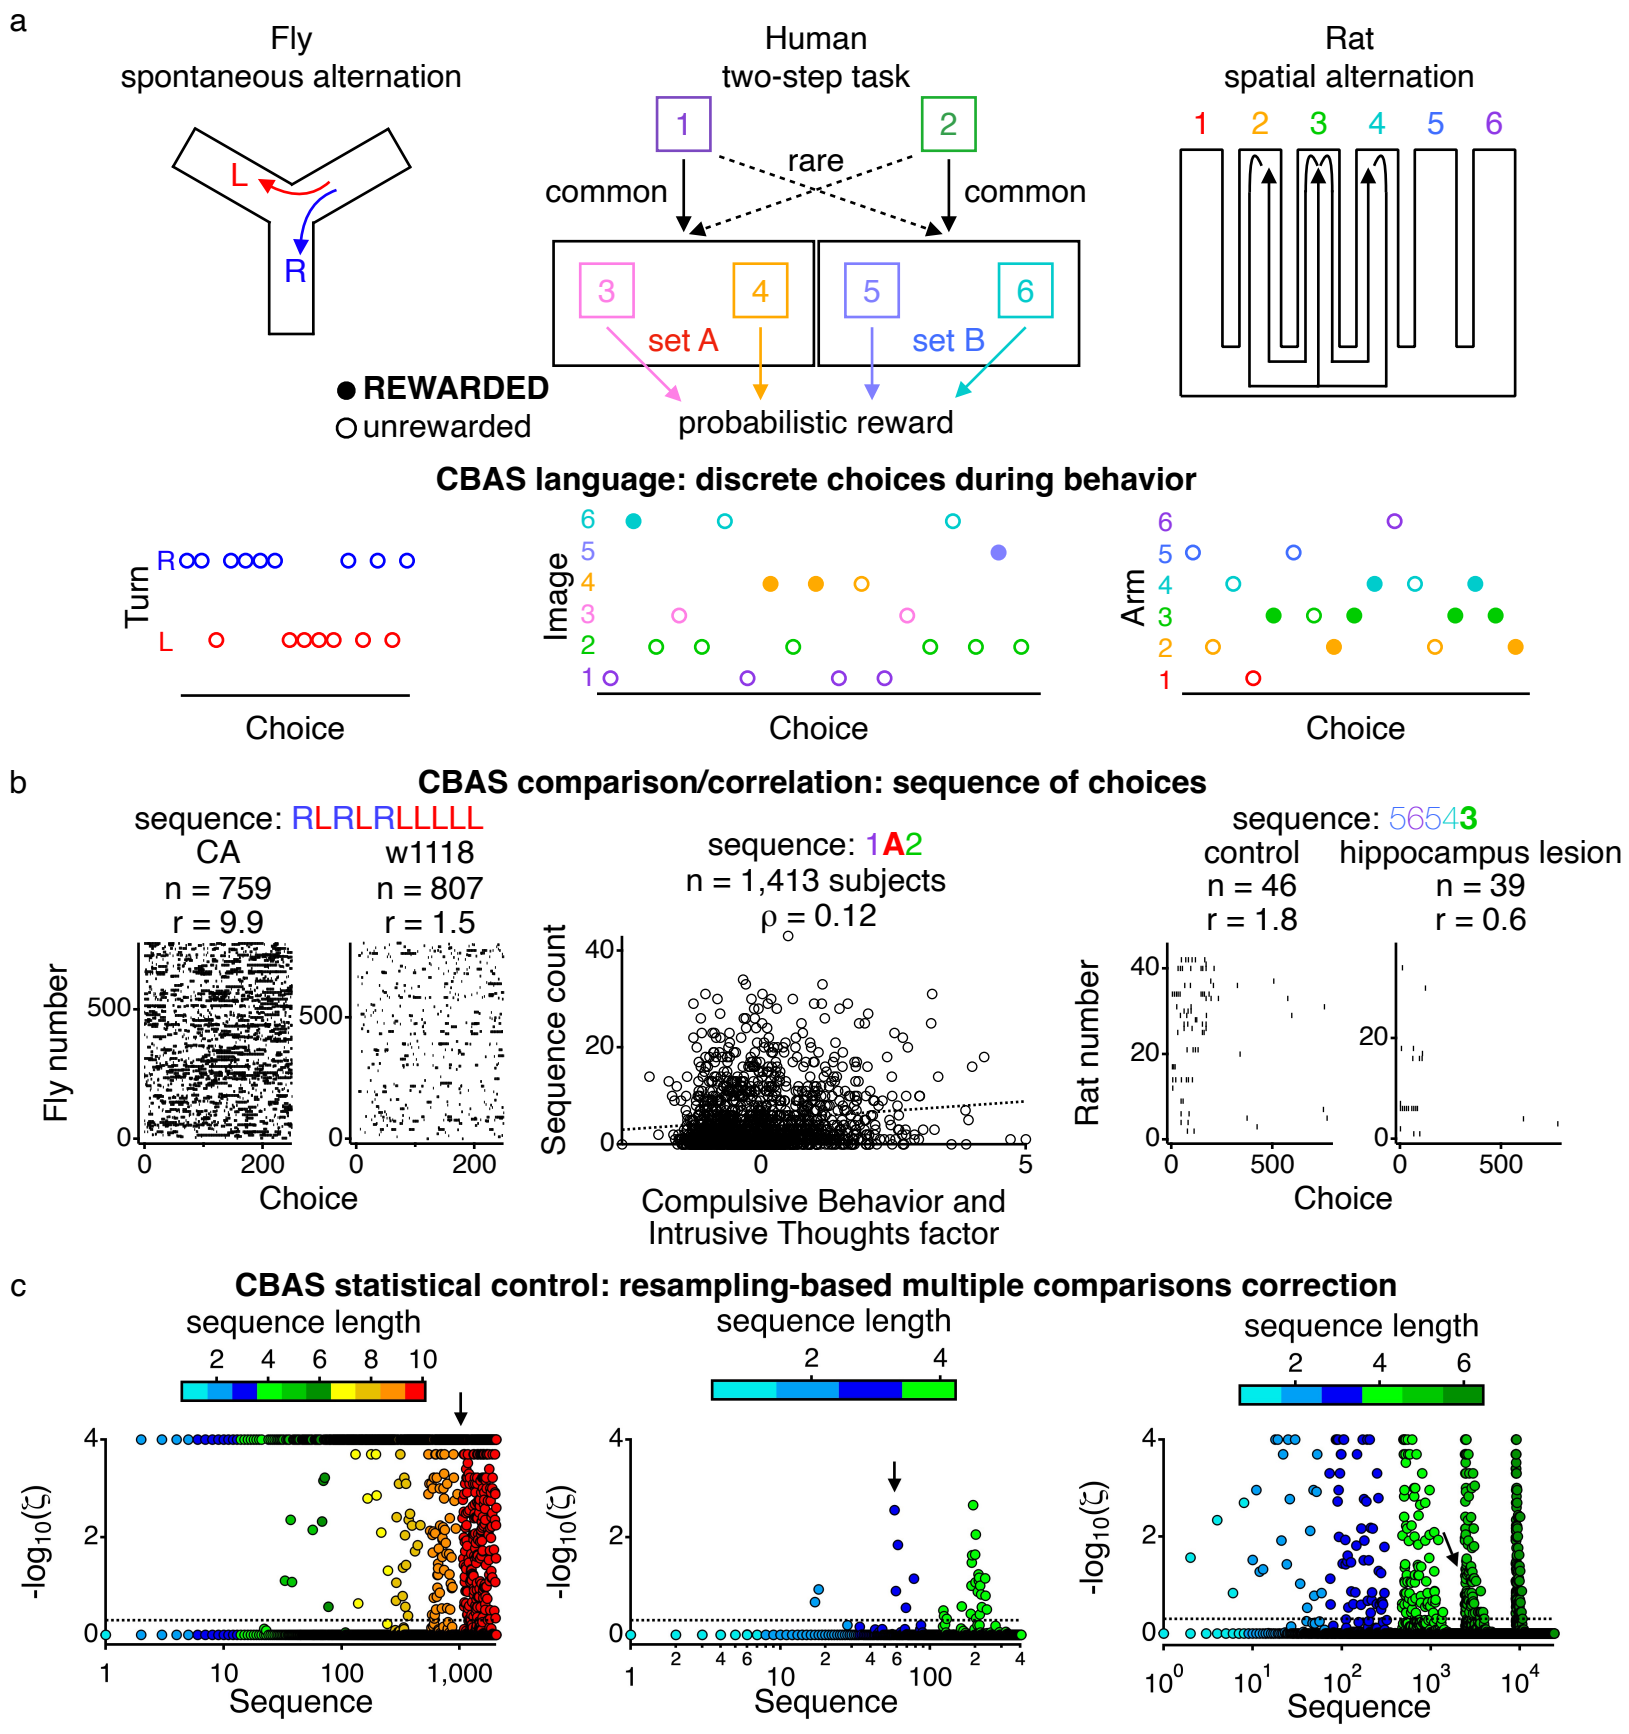
a
Fly
spontaneous alternation
L
R
Human
two-step task
1
2
rare
common
common
3
4
5
6
set A
set B
REWARDED
unrewarded
probabilistic reward
Rat
spatial alternation
1 2 3 4 5 6
CBAS language: discrete choices during behavior
Turn
R
L
Choice
Image
Choice
Arm
Choice
b
CBAS comparison/correlation: sequence of choices
sequence: RLRLRLLLLL
CA
n = 759
r = 9.9
w1118
n = 807
r = 1.5
Fly number
Choice
sequence: 1A2
n = 1,413 subjects
ρ = 0.12
Sequence count
Compulsive Behavior and
Intrusive Thoughts factor
sequence: 56543
control
n = 46
r = 1.8
hippocampus lesion
n = 39
r = 0.6
Rat number
Choice
c
CBAS statistical control: resampling-based multiple comparisons correction
sequence length
-log10(ζ)
Sequence
sequence length
-log10(ζ)
Sequence
sequence length
-log10(ζ)
Sequence

Figure 2

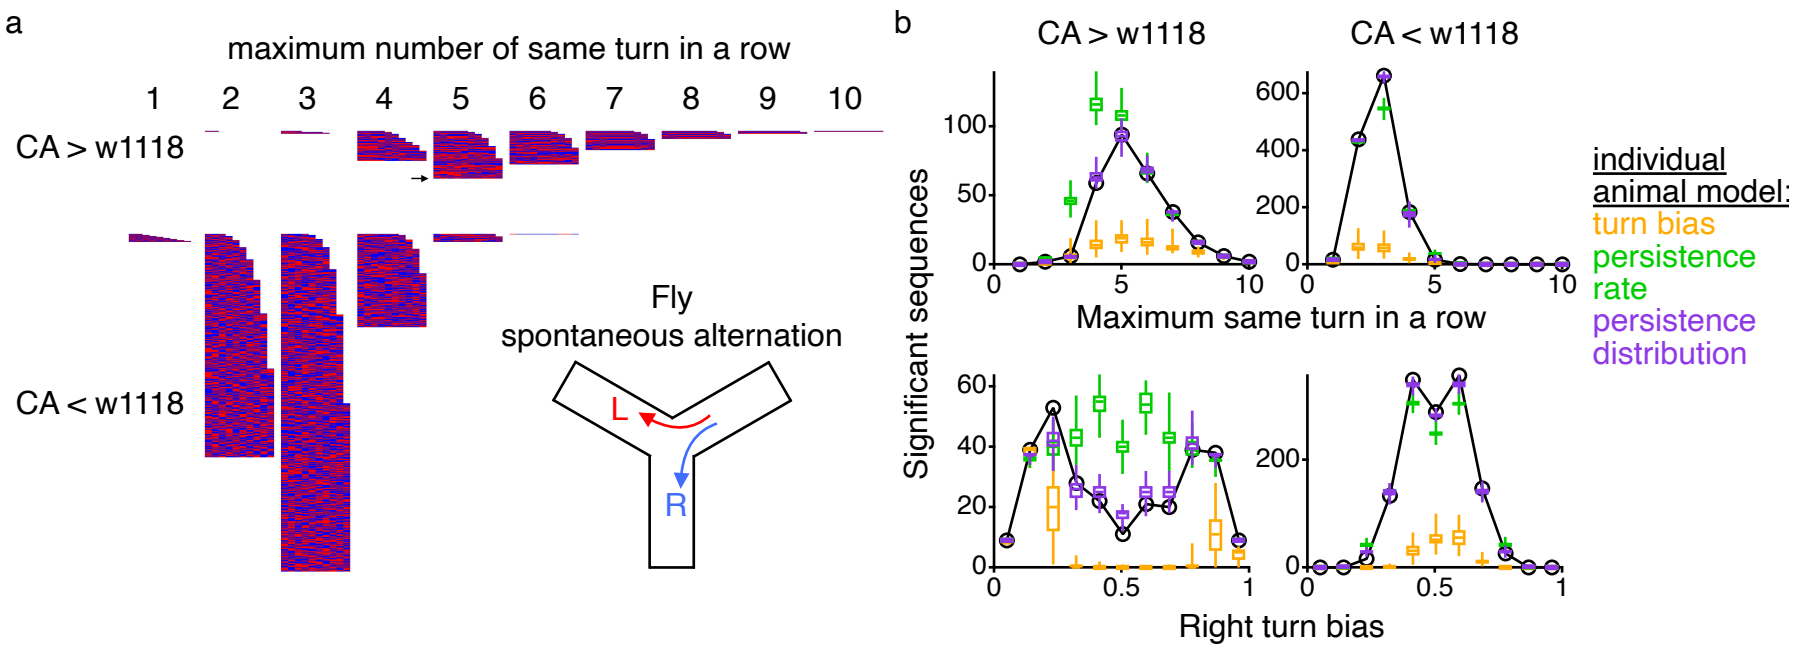
a
maximum number of same turn in a row
1
2
3
4
5
6
7
8
9
10
CA > w1118
CA < w1118
Fly
spontaneous alternation
L
R
b
CA > w1118
CA < w1118
Significant sequences
Maximum same turn in a row
Right turn bias
individual
animal model:
turn bias
persistence
rate
persistence
distribution

Figure 3

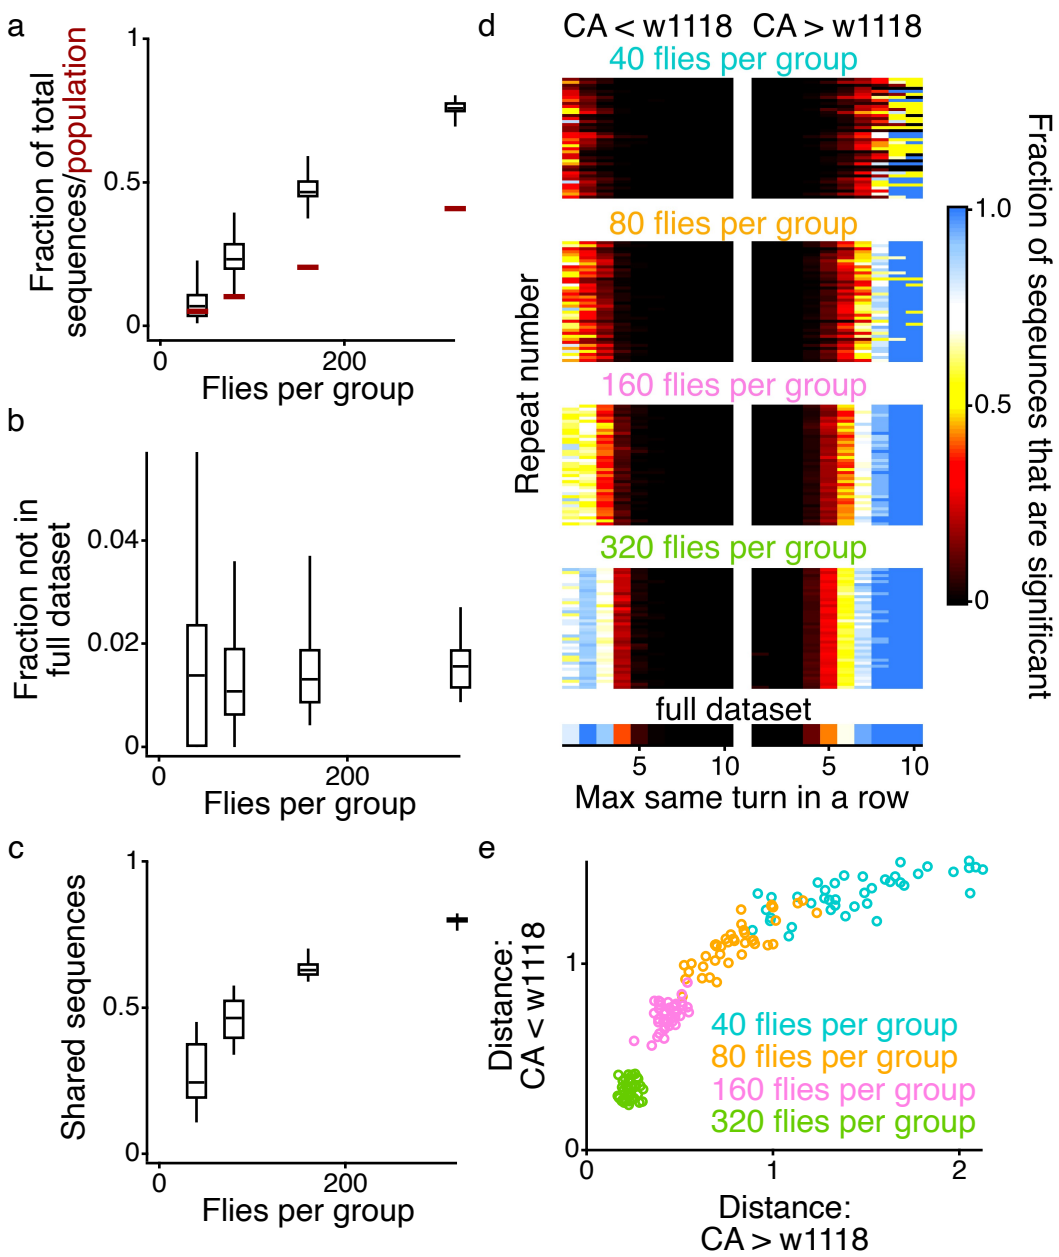
a
Fraction of total sequences/population
Flies per group
b
Fraction not in full dataset
Flies per group
c
Shared sequences
Flies per group
d
CA < w1118
CA > w1118
40 flies per group
80 flies per group
160 flies per group
320 flies per group
full dataset
Repeat number
Max same turn in a row
Fraction of seqeunces that are significant
e
Distance: CA < w1118
Distance: CA > w1118
40 flies per group
80 flies per group
160 flies per group
320 flies per group

Figure 4

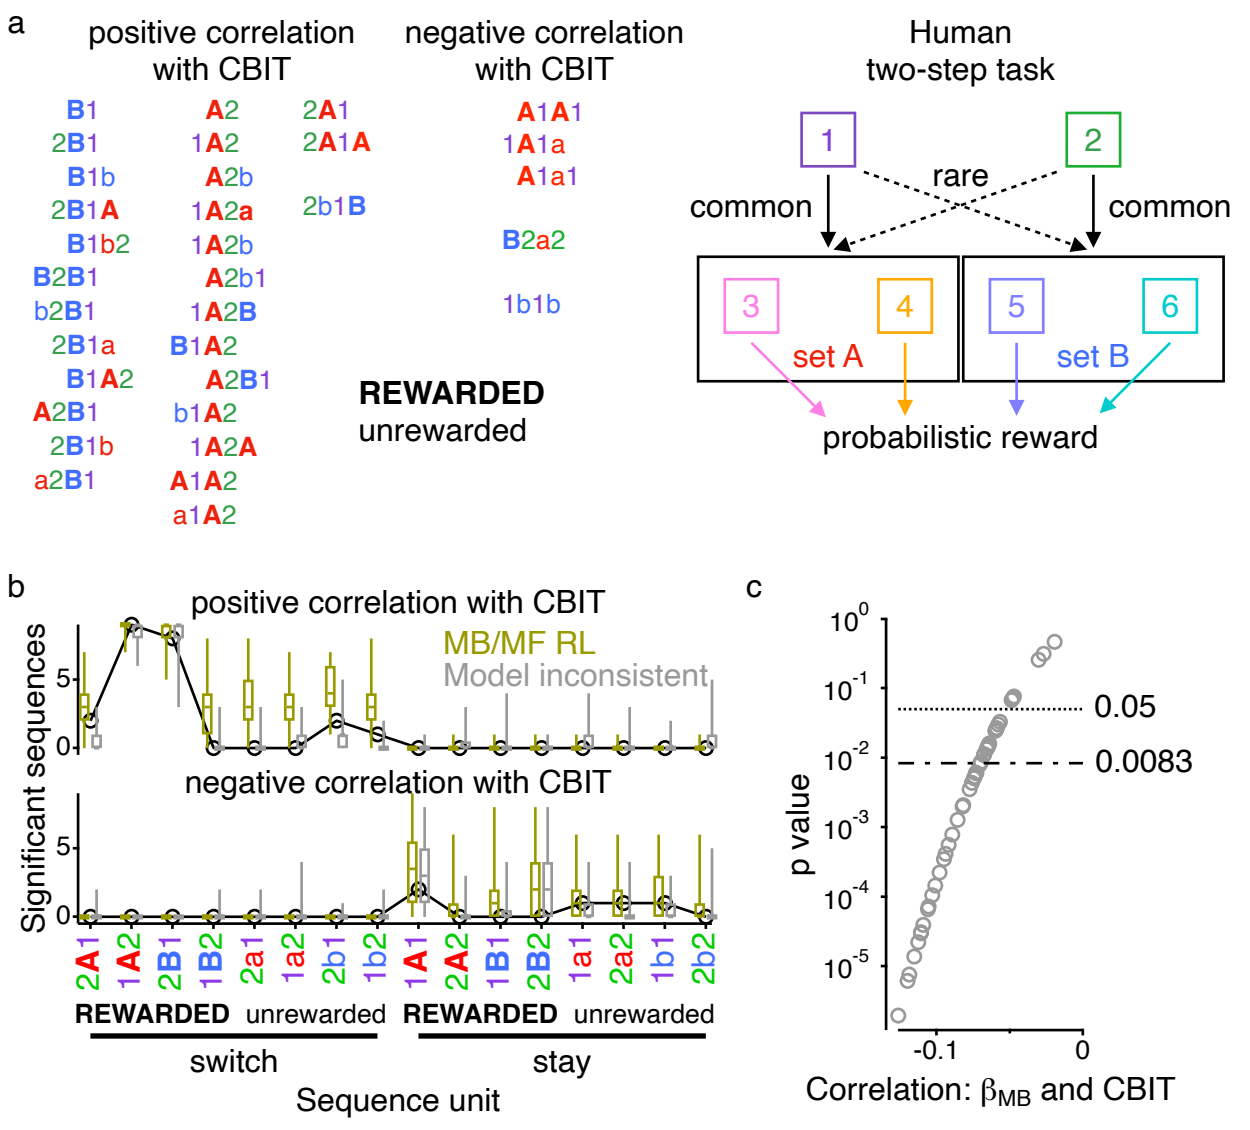
a
positive correlation with CBIT
negative correlation with CBIT
Human two-step task
B1
2B1
B1b
2B1A
B1b2
B2B1
b2B1
2B1a
B1A2
A2B1
2B1b
a2B1
A2
1A2
A2b
1A2a
1A2b
A2b1
1A2B
B1A2
A2B1
b1A2
1A2A
A1A2
a1A2
2A1
2A1A
2b1B
A1A1
1A1a
A1a1
B2a2
1b1b
REWARDED
unrewarded
1
2
rare
common
common
3
4
5
6
set A
set B
probabilistic reward
b
positive correlation with CBIT
MB/MF RL
Model inconsistent
negative correlation with CBIT
Significant sequences
2A1
1A2
2B1
1B2
2a1
1a2
2b1
1b2
1A1
2A2
1B1
2B2
1a1
2a2
1b1
2b2
REWARDED unrewarded REWARDED unrewarded
switch
stay
Sequence unit
c
p value
0.05
0.0083
-0.1
0
Correlation: βMB and CBIT

Figure 5

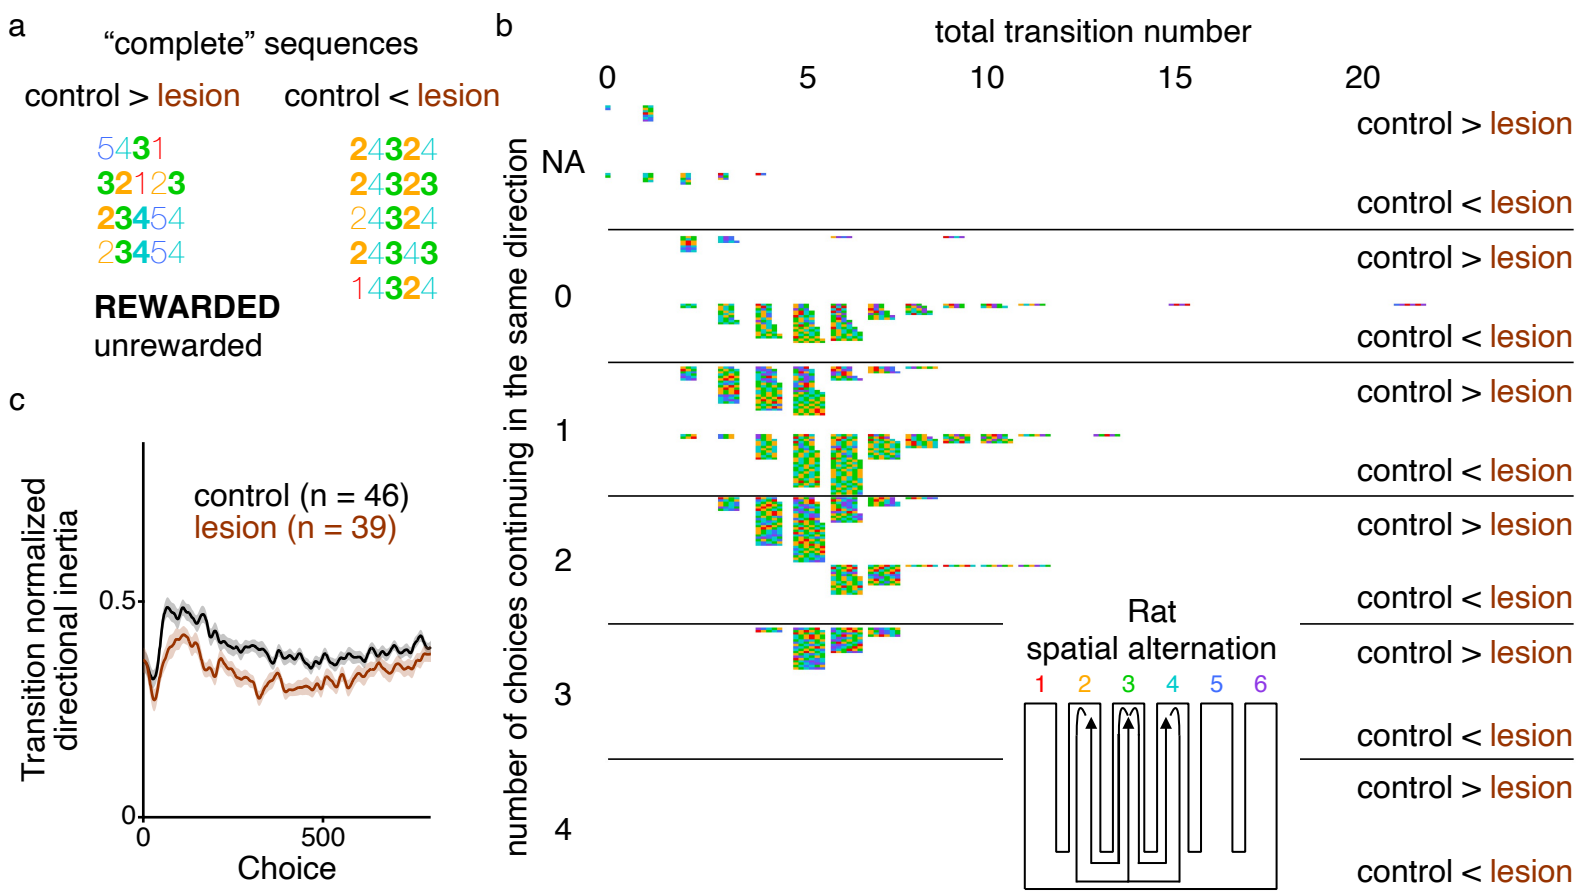
a
"complete" sequences
control > lesion
control < lesion
5431
32123
23454
23454
REWARDED
unrewarded
24324
24323
24324
24343
14324
b
total transition number
0
5
10
15
20
number of choices continuing in the same direction
NA
0
1
2
3
4
control > lesion
control < lesion
Rat
spatial alternation
1
2
3
4
5
6
c
Transition normalized directional inertia
control (n = 46)
lesion (n = 39)
0.5
0
0
500
Choice

Figure 6

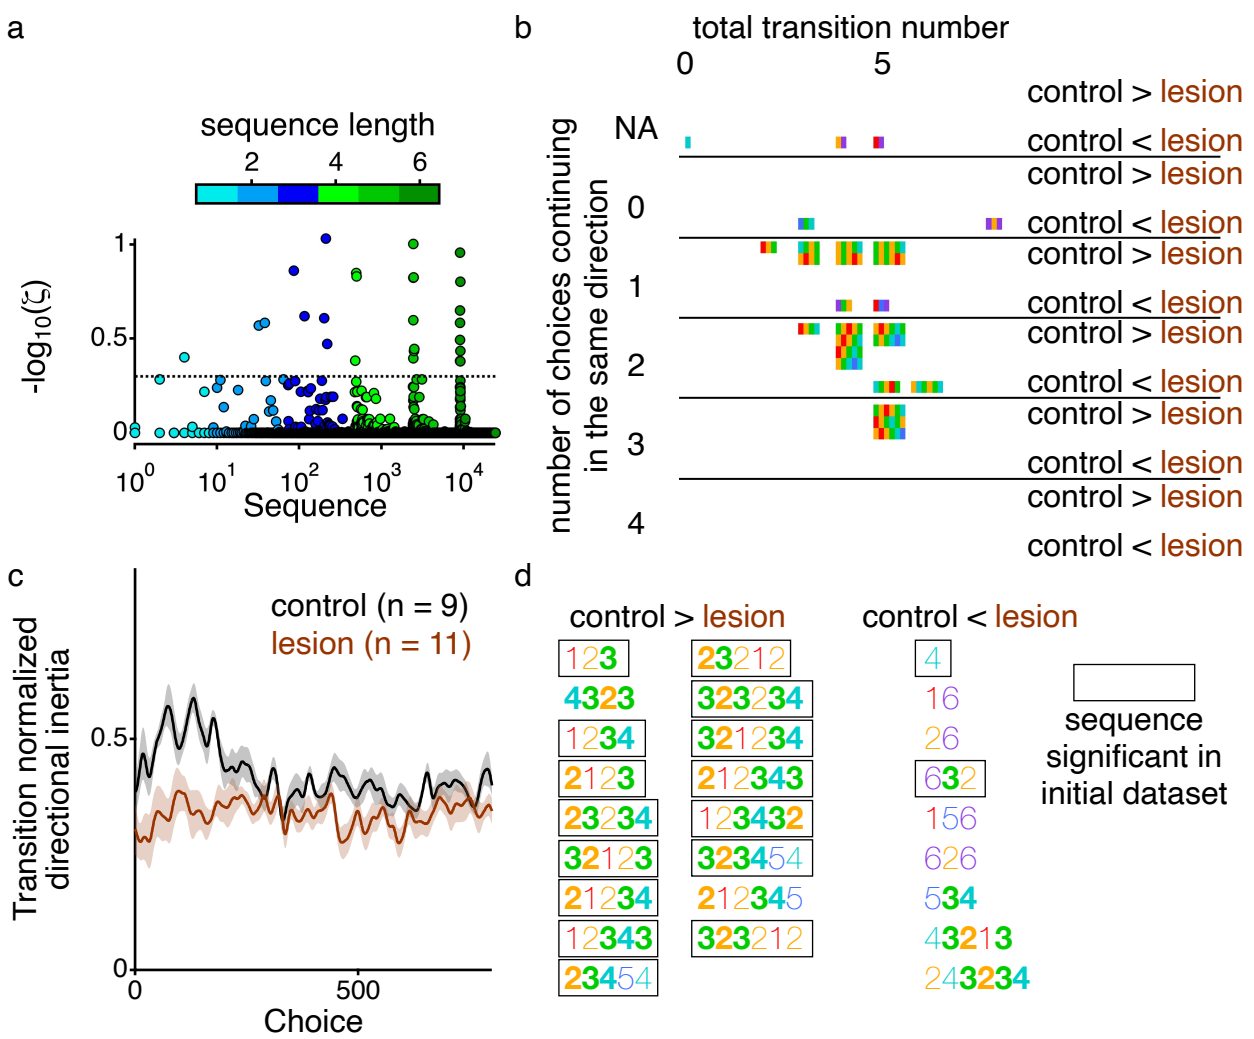
a
sequence length
2
4
6
-log10(ζ)
1
0.5
0
10^0
10^1
10^2
10^3
10^4
Sequence
b
total transition number
0
5
number of choices continuing in the same direction
NA
0
1
2
3
4
control > lesion
control < lesion
control > lesion
control < lesion
control > lesion
control < lesion
control > lesion
control < lesion
control > lesion
control < lesion
c
control (n = 9)
lesion (n = 11)
Transition normalized directional inertia
0.5
0
0
500
Choice
d
control > lesion
control < lesion
123
4323
1234
2123
23234
32123
21234
12343
23454
23212
323234
321234
212343
123432
323454
212345
323212
4
16
26
632
156
626
534
43213
243234
sequence significant in initial dataset

Supplementary Figure 1

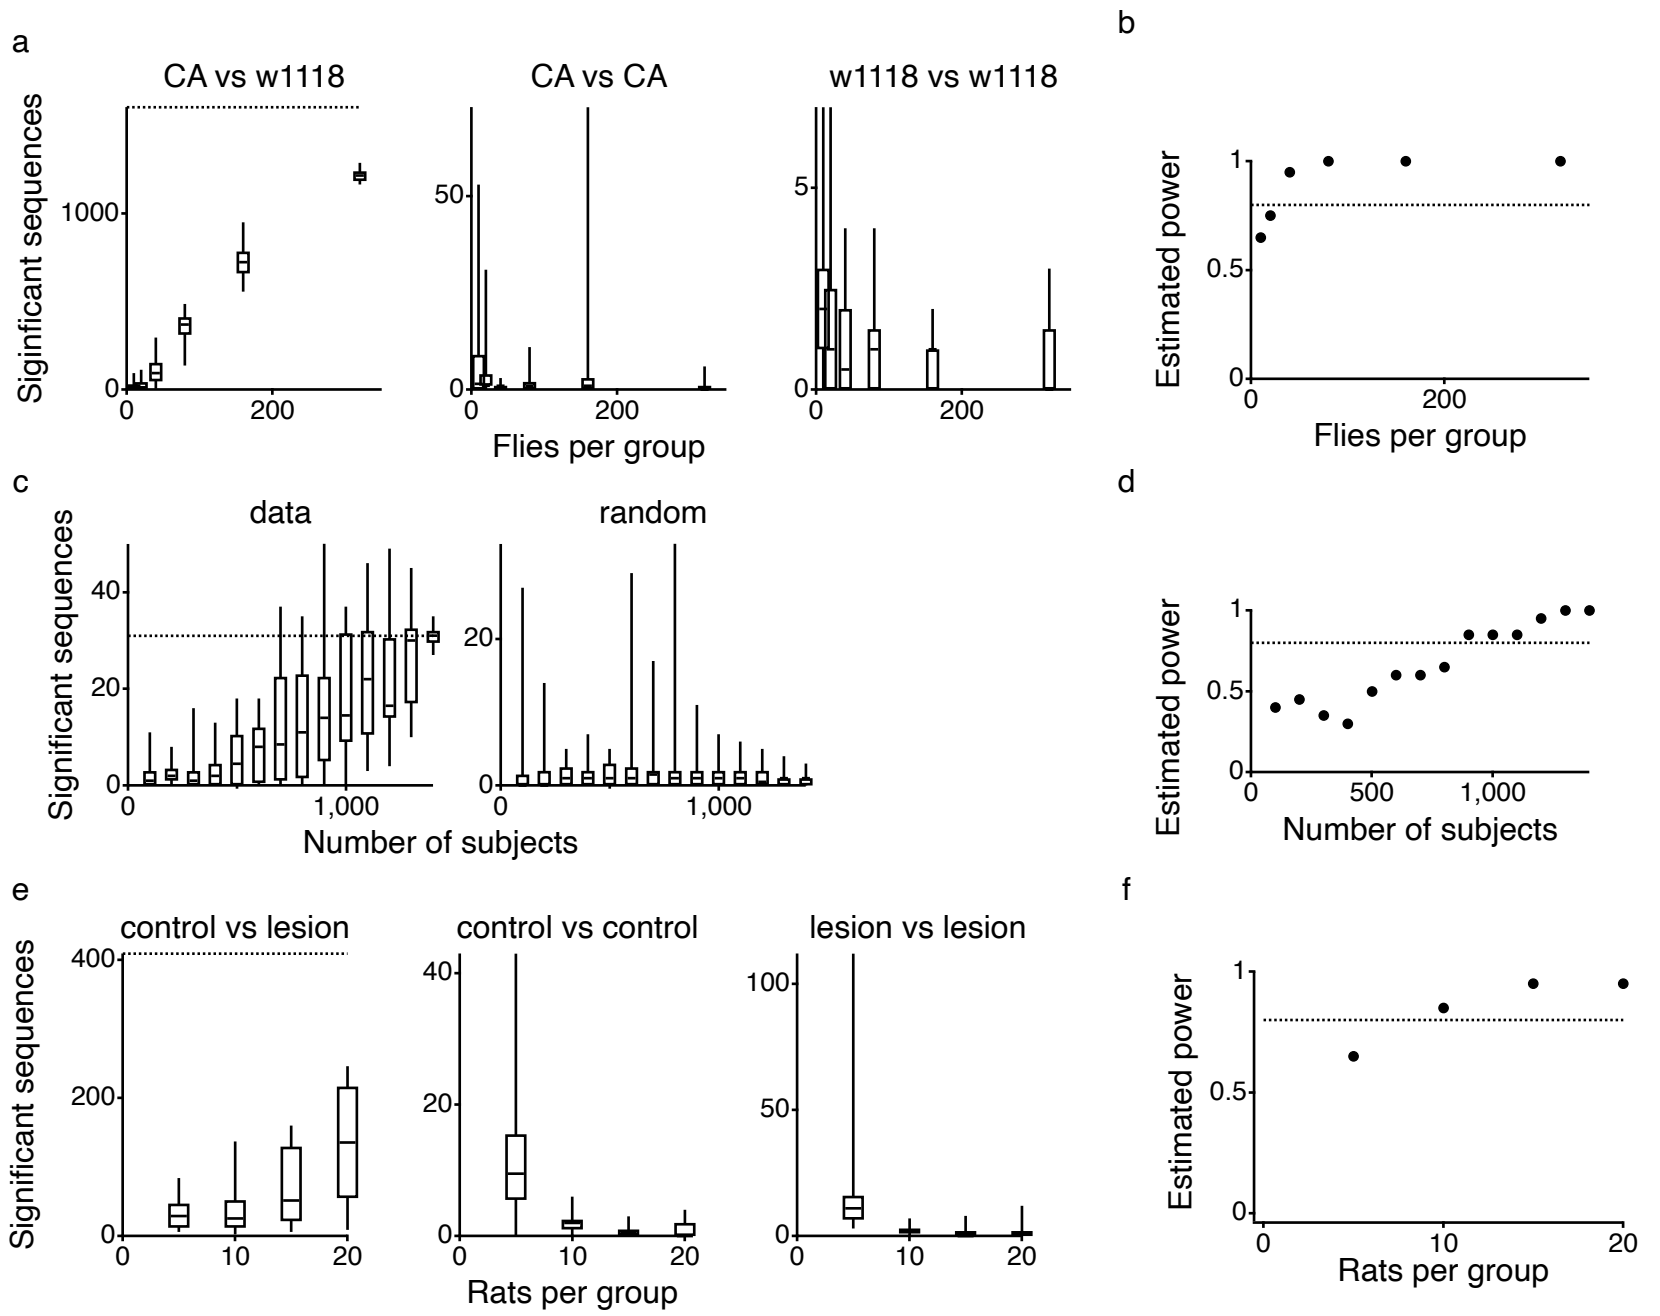

Supplementary Figure 2

a

	control rat sequence usage rate	lesion rat sequence usage rate	sufficient power	significance
32123**1**	0.04	0	no	
32123**2**	1.67	0.85	yes	control > lesion
32123**4**	0	0.05	no	
32123**4**	7.11	3.00	yes	control > lesion
32123**5**	0.11	0	no	
32123**6**	0.07	0.03	no	
→ **3**2123	9.00	3.92	yes	control > lesion
1**3**2123	0	0	no	
2**3**2123	0	0	no	
4**3**2123	3.61	1.56	yes	control > lesion
43**2**123	3.57	1.64	yes	control > lesion
5**3**2123	0.80	0.15	yes	control > lesion
6**3**2123	0.15	0.10	no	

REWARDED
unrewarded

Supplementary Figure 3

a

"complete" sequences

CA > w1118

LLLLLLL
RRRRRRRR
LLLLLLRR
LLLLRRRR
RLRRRRRR
RRRRLRRRR
LLLRLLLLL
RRRRRRRLL
LLLLRLLLL
LLLRRRRRR
RRRRRRLLL
RRRRRRLRL
RLLLRRRRR
LRLRRRRRR
RRRRRLRLR
LLLLLLRLR
RLRLLLLLL
LLLLLRLRL
RLLRLLLLL
LLLLLRLRR
LLRLRRRRR
LLLRLRRRR
RRRRLLLLR
LLLLRLRRR
LRLRLLLLL
LLLLLRRRL
LRRRRLLLL
LLLLLRRLR

CA < w1118

LRLLRLR
RLLRLLR
LRRRLLLR
RLRLRRRL
LRRLLLRL
RLRRLLLR
RLLRRRRL
LRRRLRLR
LRLLRRRL
RLLRLLLR
RLLRRRLR
RLRLRRLL
RLLLRLRL
RLRRRLRL
RLLRRLLR
RLRRRLLR
LRRLRLLR
LRRLLRLR
RRLLLRRL
RRLLLRLR
RLLLRLLR
LRRLLRRL
LRLRRRLR
RLRLRRLR
LRLRRLRR
RLLRRLRL
RLRLLLRL
RLRRLLRL
LRRRLLRL
LRLLLRRL
RLRRLRRL
LRLLRRLR
RLRRLRLR
LRLRLRRL
LRRLRLRL
LRLRRLLR
LLRRLRLR
RLLLRRLR
RLLRLRRL
LRRLRRLR
RLLRLRLR
RLLLRRRRL
LLRLLRRLL
RRLLRRLRR
RRLRLLRRR
RLLRRRLLL
LLRRRRLLR
LRRRRLLLR
LRRRLRRRL
RLRRLLRRR
LLRRLLLRR
RLLLLRRRL
LRLRLLLLR
RRLLLRRRL
LRRLRRRRL
RLRLLLLRL

Fly
spontaneous alternation

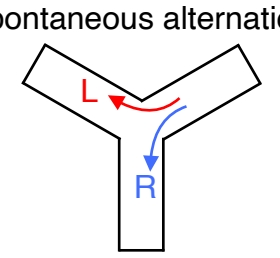

b

"complete" sequences
positive correlation
with CBIT

1A2
2B1
2A1

REWARDED
unrewarded

Human
two-step task

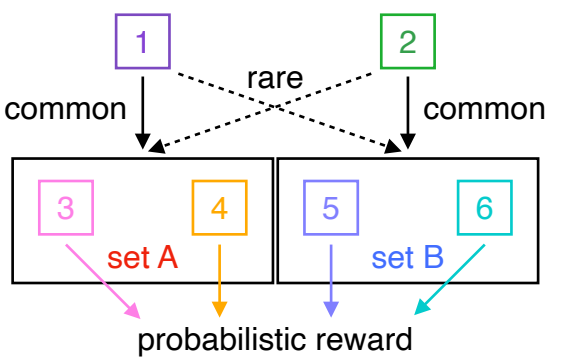
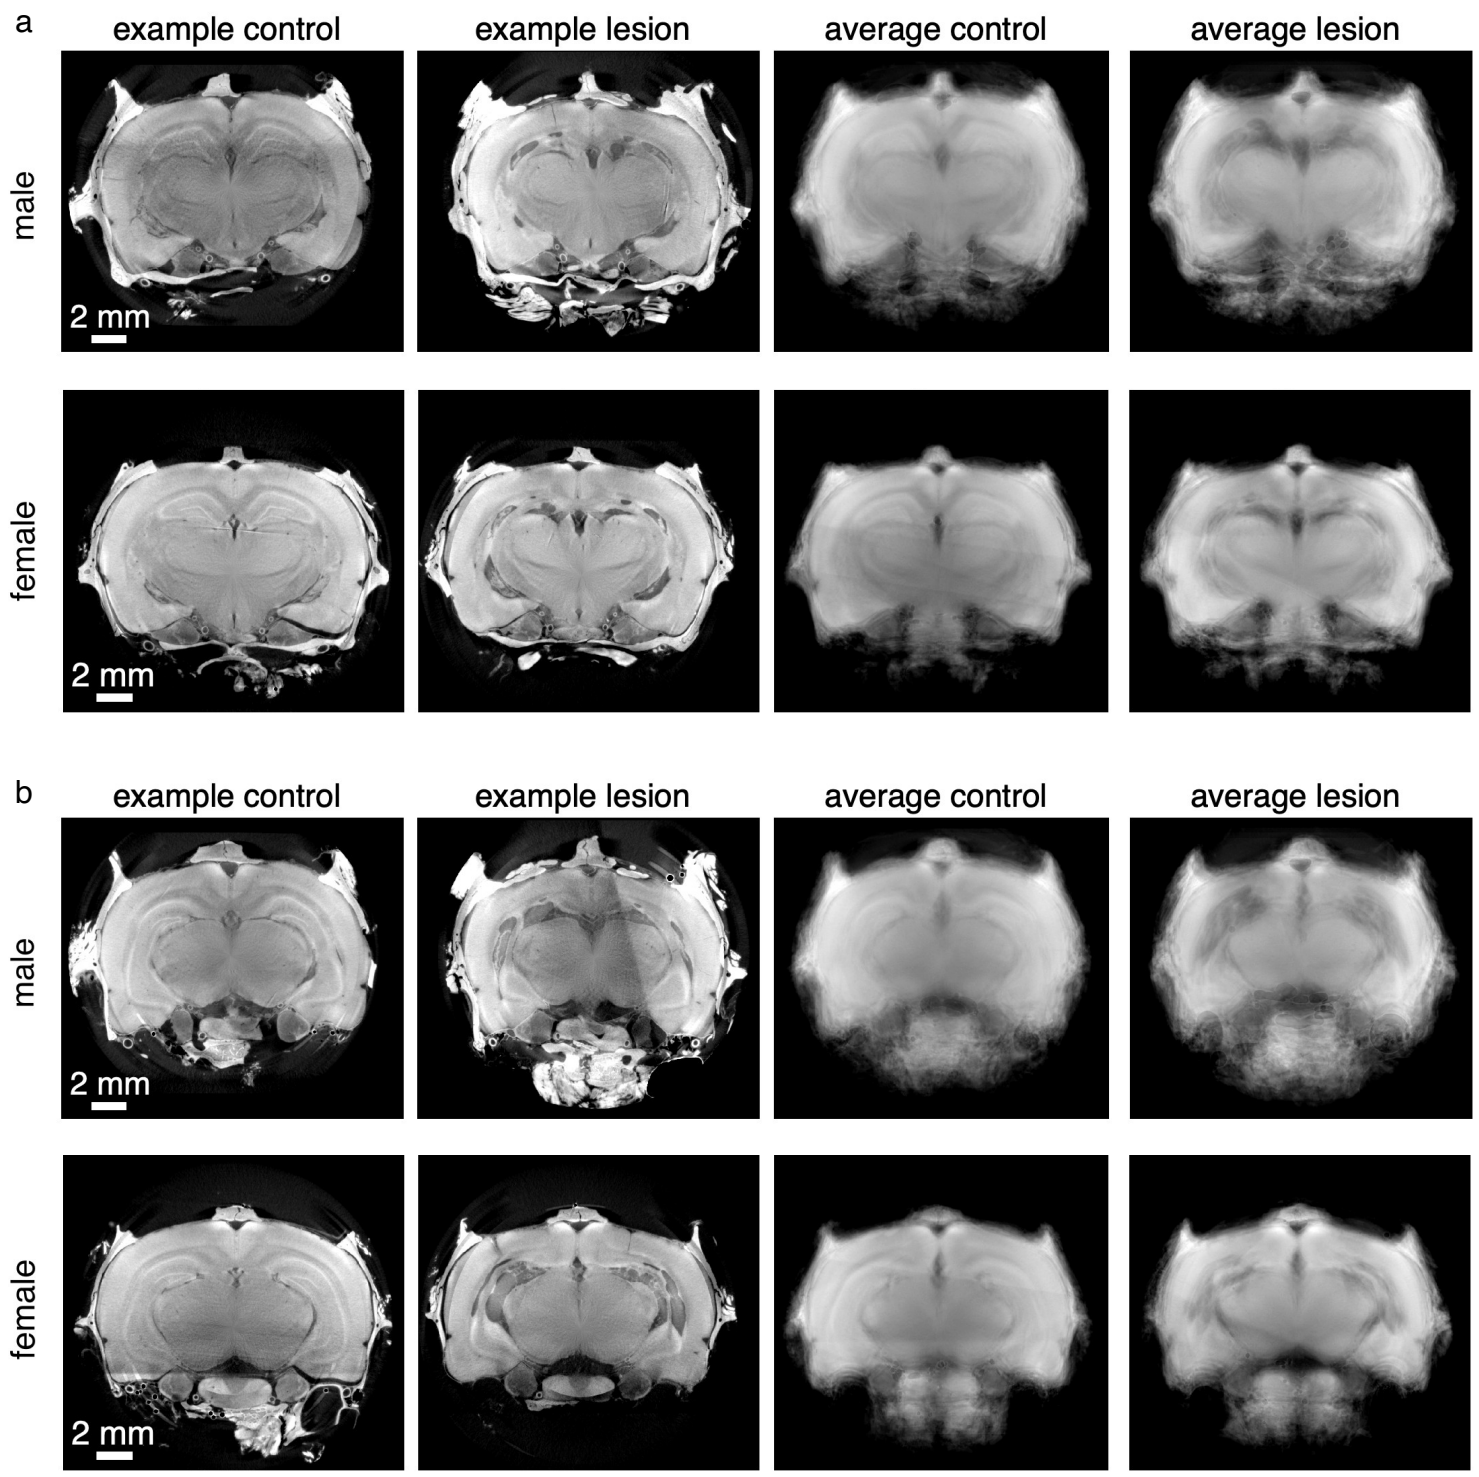
Supplementary Figure 4