Title: An unbiased decoding of ventral striatal oscillations in a rat model of binge eating: Finding the balance between model complexity and performance.

Authors: Lucas Dwiel[1], Michael A. Connerney[1], Alan I. Green[1,2,3], Jibran Y. Khokhar[1,3] and Wilder T. Doucette[1,3]

[1]Department of Psychiatry, Geisel School of Medicine at Dartmouth

[2]Department of Molecular and Systems Biology, Geisel School of Medicine at Dartmouth

[3]The Dartmouth Clinical and Translational Science Institute, Dartmouth College.

Corresponding Author: Wilder T. Doucette, MD, PhD

1 Medical Center Drive, Lebanon, NH 03756

Wilder.t.doucette@hitchcock.org

Tel: 603-650-7549

Abstract

Binge eating is a disruptive and treatment-resistant behavior that complicates the management of obesity and comorbid psychiatric disorders. The distributed network that governs feeding behavior, particularly of palatable food, has a nexus point within the ventral striatum (VS). This important structure integrates information from energy homeostasis circuits of the hypothalamus and brain stem, cognitive control regions of the prefrontal cortex and central nodes of learning and memory. The VS is thus an ideal target for neuromodulation-based interventions (e.g., deep brain stimulation) for the treatment of appetitive disorders. Here, we extracted information about feeding behavior from VS oscillations in a rat model of binge eating — a critical step in the development of responsive neuromodulation systems for appetitive disorders. Local field potentials (LFPs) were recorded from bilateral nucleus accumbens (NAc) core and shell (regions within the VS) during binge sessions across varying conditions. LFP features (n=58) were used as predictors in statistical (logistic regression) and machine learning (lasso) models to predict the type and quantity of food consumed, classify rat behavior (feeding or not feeding) in real time, and predict the onset of future feeding bouts. The complexity of real-time models was then manipulated to determine the impact on model stability and performance. Performance of models was quantified (effect size, *d*) by comparison to the performance of models built from permuted data. Models were able to predict the amount of food eaten at baseline (mean absolute error, MAE = 2.6 ± 0.08 gm; *d* = 0.43), the increase in binge size following food deprivation (MAE = 0.8% ± 0.01%; *d* = 1.1) and the type of food eaten (accuracy = 59 ± 1%; *d* = 0.55). Other models were able to predict whether initiation of feeding was imminent (area under the curve, AUC = 0.81 ± 0.03; *d* = 2.68) up to 42.5 seconds before feeding began and classify current behavior as feeding or not-feeding (AUC = 0.87 ± 0.01; *d* = 3.27). For real-time models, an optimal balance between model complexity and performance was achieved using 3 LFP features with the largest contributions from features in the alpha and

high gamma frequencies. Similar models could be incorporated into algorithms used to improve treatment outcomes in adaptive neuromodulation systems for disorders of appetitive behavior.

Introduction

The ventral striatum (VS) contains sub-regions (e.g., the nucleus accumbens [NAc] core and shell) that serve as nodes of convergence between cognitive control regions of the prefrontal cortex and central nodes of learning and memory (e.g., amygdala, hippocampus and the ventral tegmental area), the output of which influences both goal-directed and habitual behavior [1-3]. These nodes of convergence within the networks governing appetitive behavior have been proposed as possible targets for neuromodulation-based treatments for a variety of disorders including substance use disorders, eating disorders and obesity [4-7]. Targeting the VS for eating disorders and obesity is further supported by the known connectivity of the NAc with energy homeostasis circuits within the hypothalamus and brain stem [8, 9]. Further, activity within these VS sub-regions are modulated by both central and peripheral ligands signaling anorexigenic and orexigenic states [10-12]. This knowledge has led to both pre-clinical and clinical investigations using deep brain stimulation (DBS) targeted to sub-regions of the VS and most have demonstrated a capacity to alter appetitive behavior [13-18]. These initial studies support the treatment potential of targeting the VS but variable treatment outcomes combined with the known risks of the intervention have limited more widespread use. Fortunately, a subsequent wave of investigation aims to bring advances pioneered in Parkinson's disease and epilepsy to appetitive disorders by using relevant neural markers to provide feedback to trigger or modify ongoing stimulation (closed-loop and adaptive DBS systems) to improve treatment outcomes [19]. Such validated neural markers are already advancing the treatment of epilepsy (closed-loop DBS) and Parkinson's disease (adaptive DBS), while work to discover and validate similar markers in other neuropsychiatric disorders is ongoing [20-24].

To date, local field potential (LFP) oscillations recorded from the electrode arrays used to deliver stimulation have served as a rich source of neural systems-level information to control

stimulator output in advanced DBS systems. Thus, it is important to determine what information can be decoded from LFP oscillations recorded within the VS, as it is the most likely target for neuromodulation-based therapies for appetitive behaviors. Existing evidence suggests that connectivity between VS nodes and interconnected brain regions may be dynamically modulated during behavior. For example, the amygdala appears to be functionally connected to VS nodes when gamma-50 oscillations are prominent, however during gamma-80 oscillations the frontal cortical areas are more coherent with VS oscillations [25]. Such dynamic changes in oscillatory activity have been hypothesized to play a role in synchronizing information transfer through cortico-striatal loops [26, 27], and entrainment of single-unit firing to these oscillations suggests they are indeed relevant to striatal processing [28-31]. Supporting this premise, using a mouse model of binge eating, Wu *et al.* demonstrated that changes in LFP power spectra within the delta frequency range correlated with the initiation of a binge eating epoch; further, they used this LFP feature in a closed-loop system to reduce binge size [19]. The model used by Wu *et al.* is a good example of a low complexity population-based model (1 LFP feature built from 6 animals) demonstrating the trade-off of maximizing model simplicity at the expense of model performance (69% sensitivity and 64% specificity) [19].

Currently, there is insufficient evidence examining how differences in model complexity relate to model performance. When building predictive models for closed-loop and adaptive systems, understanding this trade-off (complexity vs. performance) is vital. This study determined how altering model complexity changed performance in classifying binge eating behavior in a rat model. Unbiased statistical and machine learning approaches were used to decode LFP recordings from bilateral NAc core and shell to predict imminent binge eating epochs and classify current behavior. Further, we used a machine learning approach to determine if VS oscillations recorded during binge eating sessions contained information that could predict the quantity of food consumed, the increase in binge size following food deprivation, or the type of food being consumed (low or high palatability). This work provides an

assessment of the quality and quantity of information that can be extracted from oscillations recorded in the VS related to binge-like feeding. Most importantly, this work helps to characterize the relationship between how predictive models are built (complex vs. simple; individualized vs. population-based) and the performance of those models across individuals, time, and variation of internal states. We expect these fundamental relationships, while demonstrated here in a rat model of binge eating, to generalize to other appetitive behaviors and neuropsychiatric conditions, thus providing a framework for the development of more effective models for closed-loop and adaptive neuromodulation-based treatments.

Methods and Materials

Animals

12 Male Sprague-Dawley rats were purchased from Charles River (Shrewsbury, MA) at 60 days of age and individually housed on a reverse 12 hour light/dark schedule with food and water *ad libitum*. Following habituation to the animal facility (1 week), but before beginning the binge eating paradigm, rats were implanted with custom electrode arrays targeted to the NAc core and shell bilaterally, according to the following coordinates relative to bregma: 1.6 mm anterior; ±1 and ±2.5 mm lateral; and 7.6 mm ventral. To induce “binge-like” feeding, a modification of previously published limited access protocols was used [32-34]. A powdered high-sugar, high-fat diet (“sweet-fat chow”), which contained 19% protein, 36.2% carbohydrates, and 44.8% fat by calories and 4.6 kcal/g (*Teklad Diets 06415*, South Easton, MA), was provided to the rats within the recording chambers for 2 hours [32]. A powdered diet was used to reduce chewing artifacts in the LFP signal. All experiments were carried out in accordance with the NIH Guide for the Care and Use of Laboratory Animals (NIH Publications No. 80-23) revised in 1996 and approved by the Institutional Animal Care and Use Committee at Dartmouth College.

Feeding behaviors and dynamics

Following completion of a DBS intervention published separately [35] and a 2 week washout period, animals re-established a baseline of binge-like feeding within recording

chambers. Once a stable baseline of palatable sweet-fat chow intake was established, baseline binge sessions (Base) were recorded (LFP and video). Binge sessions were also recorded following 24 and 48 hours of food deprivation (Dep24 and Dep48) with intervening periods (minimum 24 hours) of *ad libitum* access to house chow. To obtain a corresponding 2 hour session with consumption of enough, less palatable, house chow, animals underwent 24 hours of food deprivation prior to a 2 hour session (Chow) with access to powdered house chow. Videos of binge sessions were manually scored for feeding, approach (movement to food), and rest intervals by blinded scorers. The food was weighed before and after all recording sessions to accurately measure the amount of food consumed during the session with conversion to kilocalories for direct comparison of food types.

Since the binge sessions had minor variations in length, timestamps were normalized to session length to compare the temporal dynamics of feeding behavior across animals. Behavioral markers were assigned using the hand-scored behavioral start and stop times, then grouped to generate behavioral timelines by condition (e.g. baseline), and used to calculate the percentage of animals feeding at each normalized timestamp. For visualization, we divided these percentages into ten bins corresponding to 10% of the length of a binge session.

Trials were classified as pre-feeding if they occurred in the 45 seconds prior to feeding initiation and were not themselves within another feeding interval. All trials that were neither feeding nor pre-feeding were classified “not-feeding.” To determine if pre-feeding trials would be heavily influenced by approach behavior, we plotted the cumulative probability of approach trials through pre-feeding time and 100% of approach behavior was completed by 12.5 seconds before feeding initiation.

Local field potential (LFP) recording and processing

Rats were tethered in 18”X12”X24” chambers through a commutator to a Plexon data acquisition system while time-synchronized video was recorded (*Plexon,* Plano, TX) for offline analysis. We used standard LFP signal processing methods for estimating power spectral

densities (PSDs) and coherence to characterize activity recorded within and between NAc sub-regions at established frequency ranges from the rodent literature: delta = 1-4 Hz; theta = 5-10 Hz; alpha = 11-14 Hz; beta = 15-30 Hz; low gamma = 45-65 Hz; and high gamma = 70-90 Hz [36-38]. This produced 60 LFP features: 24 measures of power (6 frequency bands X 4 channels) and 36 measures of coherence (6 frequency bands X 6 channel combinations). In order to estimate low frequency components of our signals, the data were segmented into 5 second trials with each assigned a behavioral classification. The data were used either as a single trial or were averaged across trials of the same behavior type within a recording session. All signal processing was done using custom code written for Matlab R2017a (see Supplemental Methods for more details).

Assessing for potential chewing artifacts

To reduce the possibility of mechanical noise contaminating the LFP signal with non-neuronal signal we regressed the difference in feeding voracity (grams consumed per second spent eating) from baseline to Dep24 and Dep48 against the corresponding change in LFP features. To be conservative we removed any feature that had a significant, but uncorrected, p-value from the linear correlations (2/60 features were removed; Supplemental Figure S1).

Data preparation and model building

Two methods of model building were used: logistic regression and lasso. Logistic regressions were implemented using the Matlab function *fitglm*. The lasso algorithm applied a penalty term that reduced the data dimensionality to identify the subset of LFP features whose variation correlated with aspects of feeding behavior. We used the Matlab package *Glmnet* [39] to implement lasso with 100 iterations of k-fold cross-validation (see Supplemental Methods).

Linking ventral striatal activity to individual heterogeneity in feeding behavior

For predicting the individual heterogeneity in feeding behavior we used LFP features averaged across the recording as explanatory variables and the feeding aspect of interest as outcome variables (i.e. baseline food intake, change in binge size with food deprivation, and

food type). The sample sizes were relatively small (n = 24 for both baseline food intake and change in binge size from baseline; n = 21 for food type). Since there were an insufficient number of samples to construct a naïve test-set, the performance of the model was obtained from repeated 5-fold cross-validation to produce a performance distribution. Performances were reported with 95% confidence intervals and calculated as mean absolute error (MAE) for baseline food intake and change in binge size from baseline and as accuracy (%) for food type. The performance distributions were compared to distributions representing by-chance performance created by shuffling the assignment of explanatory and outcome variables using Monte Carlo sampling. The actual and permuted distributions were compared using the Mann-Whitney *U* test and converting the *U* test statistic into an r-family effect size and then a Cohen's *d* (see Supplemental Methods).

Linking ventral striatal activity to real-time feeding behavior

When predicting current or imminent behavior, LFP features from each non-overlapping 5-second trial were treated as observations with their corresponding behavioral classification. Importantly, no 5-second trial was given more than one of the following behavioral classifications: feeding, pre-feeding, and not-feeding.

For these predictions, the sample sizes were much larger allowing us to create independent training and test sets. For predicting current behavior, feeding vs. not-feeding, we removed pre-feeding trials (45 seconds prior to feeding initiation). We then randomly segregated ~20% of the data to be used in the naïve test set such that the original ratio of feeding to not-feeding trials was maintained. We used the adaptive synthetic sampling approach for imbalanced learning [40] to impute feeding trials from the remaining ~80% of data to obtain 250 feeding trials and these were combined with a randomly selected subset of 250 not-feeding trials. This process of splitting the data into test and training sets and imputing was then

repeated 20 times for each condition (Base, Dep24, Dep48, and Chow) to minimize variability introduced through random sub-sampling.

For predicting imminent feeding behavior we divided all pre-feeding trials (0 to 5 seconds before feeding) into a test and training set. The 80% of pre-feeding trials were imputed using ADASYN and matched with the same number of not-feeding trials to create training sets with 125 pre-feeding and 125 not-feeding trials. The 20% of pre-feeding trials in the test set were combined with enough not-feeding trials to match the original ratio of pre-feeding to not-feeding trials. Once the model was trained, its performance was tested on all pre-feeding trials that occurred up to 45 seconds before feeding. In total, seven sets of pre-feeding trials were created with enough not-feeding trials to match the initial distribution of pre-feeding to not-feeding trials across animals at each 5-second interval.

Model performances were measured by applying models to a naïve test set, calculating the classification probabilities of the new data, constructing a receiver operator characteristic (ROC) curve, calculating the area under the ROC curve (AUC), and reporting with 95% confidence intervals. Again, Monte Carlo sampling was used to shuffle the assignment of explanatory and outcome variables in order to assess by-chance accuracy. Effect size between the distribution of actual performances and permuted performances was calculated by converting the *U* test statistic into a Cohen's *d*.

Visualizing LFP feature changes around feeding epochs

Features of interest were extracted from all recordings around feeding epochs from 62.5 seconds before feeding to 32.5 seconds into feeding, and from 12.5 seconds before feeding ended to 52.5 seconds after feeding ended (only feeding trials 45 seconds or longer were used). For times outside of feeding, only data that did not overlap with other feeding epochs were used. The data from all of the baseline recordings were used to calculate the average LFP

feature value in each 5 second trial and were plotted through time around feeding epochs with one standard deviation.

Relating model complexity to performance

The complexity of models was adjusted in three ways: 1) the number of animals included; 2) the number of LFP features included; and 3) the number of conditions represented (Base, Dep24, Dep48, and Chow).

1) Individuals: To determine how many individuals were required to create a population-based model that generalized well to other animals, all permutations from 1/12 to 11/12 animals were used to build models. Population models (>1 animal used to train) were tested on the left out individual(s). To test the performance of models built from a single individual, training sets of 500 trials were used ensuring that the models had attained stable performance and were comparable to population-based models (see *Learning curves* below). The models built from individual animals were also tested on the left-out data from that individual to assess how an individualized model performed compared to a population-based model.

2) LFP features: The performance of the logistic model built from all 58 LFP features not correlated with potential chewing noise was considered the “gold standard.” The performance of this model was compared to those using fewer LFP features. The lasso algorithm reduced the number of features used to ~40 features. To test the lower bounds of simplicity, the performance of models built from all possible permutations of up to three features (logistic regression): 58 monads, 1,653 dyads, and 30,856 triads were evaluated. “Top-tier” models were models with performances that were not significantly different from the best performing model (from monad, dyad and triads). Significance was assessed using two-sample t-tests followed by the Bonferroni correction.

3) Conditions: Lastly, when assessing the importance of including data from multiple conditions in the model training data, 4000 trials from each state (500 trials from the 8 animals with data from each condition) were used. This approach allowed for evaluation of population-

based models or individualized models. Models were trained using data sets from just baseline or all four conditions, but in both cases more trials were used than needed to achieve stable model performance as determined by learning curves.

To compare performance between models with different manipulations of the above domains, two-sample t-tests were used with two tails followed by the Bonferroni correction for multiple comparisons.

Learning curves

To find the minimum number of trials required to achieve stable model performance, four learning curves were constructed: population, weighted population, individual, and pre-feeding (Figure S2). Because feeding trials were relatively rare (on average only 18% of trials) and rare event detection can lead to over-fit models or majority classifiers, we first tested two methods on the population data for dealing with rare events: imputation and weighting. For imputation we used the ADASYN algorithm [40] and we based weighting on the ratio of feeding to not-feeding trials (if there were twice as many not-feeding trials then the feeding trials had double the weight). For the imputed and weighted models, the number of trials in the training set was slowly increased from 6 to 500 per animal (72 to 6,000 trials for all 12 animals). The imputed models achieved stability sooner than the weighted models, so when needed we implemented ADASYN (Figure S2). The individual models were tested from 30 to 500 trials per animal and for the pre-feeding models from 4 to 350 trials. To determine the sample size at which model performance became stable, each sample size performance was compared with the performance of the largest model with correction for multiple comparisons. For the individual models the largest minimum sample size seen across the 12 animals was used for all individual models. (See Supplemental Figure S2)

Verification of electrode placement

At the conclusion of the experiment rats were euthanized. Brains were removed, prepared for cryostat sectioning, mounted on slides, and stained (thionine) for histologic analysis of electrode placement [17].

Results

Controlling for chewing noise

By quantifying the percent change in voracity (mg of food consumed per second) from baseline to 24- and 48-hour food deprived conditions and regressing that change against the percent change in every LFP feature between the same recordings we found 2 features out of 60 that had significant uncorrected p-values. To be conservative, we removed both features from all analyses to minimize the chances of chewing noise appearing in our models (Supplemental Figure S1).

Decoding amount of food consumed and food type from VS oscillations

The bimodal distribution of feeding activity during the baseline condition was not maintained across deprived conditions (Figure 1A) with more animals consistently feeding during the beginning of the recording. Only when the animals were deprived for 48 hours and given palatable food did they eat significantly more calories than in baseline ($t(19) = -4.25$, $p < 0.01$; Figure 1B). Although the feeding dynamics of animals given house chow after 24-hour food deprivation look similar to those of deprived animals given palatable food, during the Chow condition the animals ate significantly fewer calories than both the Dep24 and Dep48 conditions ($t(19) = 3.73$, $p < 0.01$; and $t(16) = 5.82$, $p < 0.001$; Figure 1B).

Using the brain activity recorded during baseline binge sessions (Figure 1C) models were built that outperformed permuted data in predicting the amount of food eaten at baseline (MAE = 2.6 ± 0.08 gm; $d = 0.43$; Figure 1D). Next, changes in brain activity from Base to Dep24 and Dep48 were used to predict changes in the quantity of food consumed in the corresponding sessions (MAE = 0.8 ± 0.01%; $d = 1.1$; Figure 1E). Lastly, brain activity was used to predict the

type of food being eaten following 24 hours of food deprivation — house chow vs. sweet-fat chow (accuracy = 59 ± 1%; *d* = 0.55; Figure 1F).

Predicting imminent feeding

LFP features from the trials immediately preceding feeding epochs in the baseline sessions were used to build models (Figure 2A) that predicted whether initiation of feeding was imminent (AUC = 0.81 ± 0.03; *d* = 2.68; Figure 2B). When these models were applied backwards through time, they were able to successfully differentiate trials as preceding feeding (pre-feeding — centered up to 42.5 seconds before feeding began) from all other trials outside of feeding trials (Figure 2C). These models were not merely detecting approach behaviors since only one third of the Pre-Feeding trials had overlap with approach behavior and none of the trials overlapped with approaches past -12.5 seconds from the initiation of feeding (Figure 2D).

The identity of the LFP features that were most frequently used by lasso to predict imminent feeding were plotted over time to investigate the temporal link between changes in these LFP features and feeding behavior. Both a power (shell left high gamma; Figure 2E) and coherence feature (core left - core right high gamma; Figure 2F) tended to transition from the non-feeding average towards the average level observed during feeding in the pre-feeding interval. After fluctuating around the feeding mean for that LFP feature, these features then began transitioning back to their non-feeding levels.

Relating model complexity to performance using models classifying feeding vs. not-feeding

To explore model complexity, three domains of the data used for model building were manipulated: number of individuals, number of LFP features, and number of conditions (Base, Dep24, Dep48, and Chow) (Figure 3A). Altering which data were used to train the models affected the tuning and performance of the models which was assessed using the same naïve test data sets which allowed for a comparison of performance.

Manipulation of the number of animals in population-based models

A population-based model was built from and tested on all animals and was able to classify feeding trials from not-feeding trials significantly better than would be expected by chance alone (AUC = 0.87 ± 0.01; *d* = 3.27; Figure 3Bi). All models were built using 80% of each individual's data from baseline binge sessions (Figure 2A) and the models performance was tested on the left-out 20%. When a leave-one-out (LOO) approach was used to capture the performance of population-based models on data from individuals not used in the model training data set, the performance decreased but remained better than the by-chance permuted models (AUC = 0.78 ± 0.06; *d* = 3.13; Figure 3Bi). Since population models were successful when built with the test animal left out, the first domain that was manipulated was the number of animals used to train the feeding vs. not-feeding models (Figure 3Bi Inset). This manipulation helped determine how many animals were necessary to create population models. After using 12 animals (All), performance dropped when using 11 and testing on the left-out animal (LOO) and continued to decrease down to using only one animal and testing on left-out animal data (Figure 3C).

Individualized vs. population-based models

The above data indicate that if enough individuals are used to train population-based models, then models will generalize and perform well on novel individuals. It is possible however that a model built from an individual will be able to perform better for that individual compared to a population based model. To test this, individualized models were built by training on 80% of a single animal's data and the models were then tested on that animal's left-out 20%. These models performed better than the permuted versions (AUC = 0.89 ± 0.06; *d* = 3.13; Figure 3Bii) and performed better than the population-based models tested on the same left-out 20% in 9/12 individuals when all animals were used to build the population model (Figure 3Biii All) and 11/12 individuals when all but the test individual was used to train the population-based model (Figure 3Biii LOO).

Manipulation of the number of LFP features

Next, the number of LFP features being used to train the models was manipulated (Figure 4A Inset). When classifying baseline feeding vs. not-feeding, models built from a single feature performed better than chance (Figure 4A), and although we found that increasing the number of features to two or three increased performance these simple models never achieved the performance of either a lasso (~40 features) or full logistic (58 features) model (Figure 4A). Further, when these simpler models were applied to testing data from other conditions, (Dep24, Dep48, Chow) the performance dropped considerably with the more complex triadic models providing the highest performance stability across conditions (Figure 4B). Further, the features in dyad and triad models as well as those used consistently by lasso were not all top performing features as monads, suggesting redundancy between LFP features in predicting feeding behavior (Supplemental Figure S3).

Manipulation of the conditions used to train and test models: impact on individual and population-based model performance

Given that individualized models tended to outperform population-based models when trained and tested in the baseline condition (Figure 3B) the performance of these models were also tested on data from the other three conditions. When the individual and population-based models from Figure 2 were tested on data from the other three conditions the population-based models significantly outperformed the individualized models (Dep24 $t(26) = -4.03$, $p < 0.01$; Dep48 $t(26) = -3.53$, $p < 0.05$; and Chow $t(26) = -3.63$, $p < 0.01$; Figure 5A). To determine if the performance of models could be improved across the different conditions if models were built using training data from all four conditions, new individualized and population-based models were created (Figure 5A Inset). In this set of models the individualized models significantly outperformed population-based models (Base $t(26) = 3.56$ and Chow $t(26) = 3.52$, $p < 0.05$; Dep24 $t(26) = 3.94$, $p < 0.01$; and Dep48 $t(26) = 9.69$, $p <0.001$; Figure 5B). In summary, the

population-based models built from the baseline condition generalized better across the other conditions compared to individualized models. However, if Individualized models were trained on data from all conditions then they outperformed population-based models in classifying feeding vs. not-feeding trials across conditions.

Manipulation of the number of LFP features used to build models from all conditions

Finally, the above findings were integrated by determining the minimum number of LFP features required in an individualized model built from all conditions to match the performance of the corresponding logistic model that uses all 58 LFP features (Figure 6 Inset). By using the performance of the 58 feature logistic models as the gold standard, we could determine how many features were needed to achieve a comparable performance when tested on each of the four conditions. Testing on Base and Chow data revealed that only two features were required to attain a performance that was not significantly different from the gold standard logistic regression (Figure 6A and D) and on Dep48 data three features are required (Figure 6C). The Dep24 condition required the 48 features used in the lasso model (Figure 6B) to perform at a comparable level to the logistic, although the minimum number could be as low as 4 features. Unfortunately, testing all possible combinations of LFP features from 4 to 48 was beyond the computational capacity of this study.

When we built and tested models using all conditions we found that there were many models which performed equally well (Table S1). Despite the flexibility in the exact features used in these models we did find that features within the alpha and high gamma ranges were the most common contributor to the top performing models (Table 1). For the simpler models, monads and dyads, most predictors were power features, while in the more complex triads coherence became equally represented (Table S1).

Discussion

This work demonstrates that ventral striatal oscillations can be decoded to predict aspects of feeding behavior such as the amount and type of food consumed. Further, these

oscillations can be decoded in real-time to predict imminent or current feeding in a rat model of binge eating that is stable across conditions. Lastly, by manipulating the complexity of data used to train models and then assess how these manipulations alter model performance this work creates a theoretical framework to guide the construction and implementation of predictive algorithms used in closed-loop and adaptive neuromodulation systems. Specifically, this work suggests that devices could be pre-loaded with a model built from a group of individuals utilizing a handful of oscillatory features and provide fair performance. However, to achieve optimal performance in a specific individual, additional data would need to be collected from that individual over time and across time (to capture variations in internal states) to create a personalized model.

Predicting aspects of feeding

The binge eating model used in this study resulted in variation in the size of binges between individuals (as previously reported [41]) and also in the amount that binge sizes increased following food deprivation. That this variation could be partially decoded from LFP oscillations recorded in the VS is not surprising given the connectivity of these VS regions (e.g., the nucleus accumbens) to the distributed circuit that controls feeding [42, 43]. Further, it is feasible that information about the type of food being consumed (high vs. low palatability) could be extracted from VS oscillation given that ‘liking’ of food [44, 45] and nutritive/caloric value [46] have signaling correlates within the nucleus accumbens. Overall, the capacity of VS oscillations to predict aspects of feeding behavior supports the role of VS regions within a distributed feeding circuit. Differences in the accuracy of the predictions across models may reflect the differential relevance of ventral striatal activity in regulating these dimensions of feeding behavior. For example, the variation between individuals in their baseline binge size is poorly predicted and thus, the neural determinants of this set point may be encoded within other nodes of the distributed feeding circuit. In contrast, the relative increase in binge size with food deprivation could more accurately be predicted by VS oscillations compared to baseline binge

size and this could reflect a more prominent role of the ventral striatum in binge size increase with food deprivation.

Population-based vs. individualized models

The presented data suggest that models perform best when they are built and tested within a given individual and condition and that performance drops when these individualized models are tested in other conditions even within the same animal. However, this can be mitigated by utilizing one of two approaches to eliminate the drop in performance across conditions: use data from multiple conditions within a given individual; or use data from multiple individuals within a single condition or from across multiple conditions for even better performance. Following either of these strategies will increase the probability of selecting and weighting LFP features most likely to generalize across animals and conditions with stable performance. Clinically, it is important that models are able to perform well across a heterogeneous population as well as the variable internal brain states that an individual will transition through across time. Therefore one or more of these approaches should be implemented when training models for use in algorithms that will control closed-loop or adaptive neuromodulation systems.

Machine learning and deep brain stimulation

Over the past few years there has been increasing interest in applying machine learning and other computational and statistical learning methods to guide the selection and implementation of psychiatric treatments. This work can be categorized into three domains: predicting treatment outcomes, defining treatment parameters, and actively optimizing treatment through time. Treatment outcomes have been successfully predicted using both structural and functional connectivity measures in Parkinsonian patients [47] as well as oscillatory activity in a rodent model of binge eating [35]. In a high throughput manner, machine learning has also been utilized to comb through the vast space of DBS parameters and medication combinations [48] and finding the set of DBS parameters that lead to the largest desired change in brain activity

[49]. Lastly, there is a push towards the development of adaptive and closed-loop systems which utilize machine learning to, for example, monitor stimulation evoked dopamine signaling and adjust stimulation parameters to optimize treatment [50] or to determine the time and brain state in which delivering stimulation has the largest effect [51]. With the exception of the work by Kumar et al. (2016), all of these implementations of machine learning used population-based datasets, i.e. they pooled data across subjects. For some predictions, using only individual datasets is impractical, such as determining the optimal target for stimulation [35]. However, when tuning stimulation parameters or trying to determine when stimulation would be most effective the data presented here suggest that using data from each individual would produce the best models as long as that data is sampled from across time and ideally represent different internal states (e.g., stress/anxiety or mood).

Number of predictors required for stable performance

This work highlights an important trade-off between model complexity (simple — using a single LFP feature) and model performance (the performance of the model across individuals and conditions). This study suggests that single feature models are unlikely to perform with enough accuracy or stability to be clinically useful. However, this work also shows that a relatively simple model (~3 features) can approach the performance and stability of high dimensional models (58 features). Further, these simple models overwhelmingly relied upon alpha and high gamma features while delta, the basis of the Wu model, only became used in the more complex triadic models [19]. To accomplish this, a series of statistical and computational methods were used to restrict the potential pool of predictive LFP features to generate a relatively simple model that was capable of performing with high accuracy and stability across conditions. This work suggests that enough information pertaining to a given behavior or symptom may be decoded from a few oscillatory features, thus limiting the

computational and power requirements needed to implement stable and accurate algorithms in closed-loop or adaptive neuromodulation systems.

A predictable pattern of brain activity precedes the initiation of feeding

As shown above (Figure 2), VS oscillations are able to classify times when an animal is about to feed (up to 45 seconds before feeding starts) from any other time-point outside of feeding. However, models built to classify feeding from not-feeding actually outperform the pre-feeding models in identifying the pre-feeding trials. This suggests that the same LFP features are used in both model types but are more clearly delineated (i.e. higher signal to noise ratio) between feeding and not-feeding. This finding also suggests that there is likely a “ramping” of the differences in LFP features in the transition leading up to feeding.

When this hypothesis was examined more closely (Figure 2E and F), a ramping phenomenon was observed in features that did much better classifying pre-feeding from not-feeding when trained on feeding vs. not-feeding data. These features were also top contributors in models classifying when feeding was occurring. For example, shell left high gamma power begins to approach its feeding level of activity before feeding initiates, stabilizes during feeding, and returns toward not-feeding levels after feeding stops. In other cases (Figure S4) there is a pattern of decreased variance around feeding which has been seen in the literature in response to visual stimuli and rewards [52], as a marker of stimulus perception/detection [53, 54], been correlated with behavioral performance [55, 56] and motor preparation/initiation [57, 58]. When the timing of the ramping events across the different LFP features is considered, it is apparent that they occur on different timescales. This likely underpins the accuracy of models predicting imminent feeding improving as the time gets closer to feeding initiation and more feature values have developed a more significant level of divergence (Figure 2C).

Confounds and noise

Given that motor initiation and preparation are known to manifest as decreased neural variability [57, 58] there is a possibility that the decrease in LFP signal variance around feeding

was merely a byproduct of the approach behavior of the animal. It is also possible that the success of our pre-feeding models was really detecting approach behavior. However, by 12.5 seconds before feeding initiation, none of the data contained any approach behavior (Figure 2D), while our models were able to continue to classify brain activity out to 42 seconds before feeding initiation as belonging to the pre-feeding period. Further, when models were built from feeding data during which there is no approach behavior they actually performed better in differentiating pre-feeding from not-feeding brain activity. Together, these data suggest that although approach behavior may also exhibit a decrease in neural variability or be detectable in its own right, these are not influencing the models described here.

Similarly, there is a concern that without electromyography of the masseters we are unable to completely remove the possibility of chewing noise influencing the models built from feeding data. However, if the success of our feeding models at differentiating feeding from other behaviors was merely reflecting the ability to decode mechanical chewing noise, then the feeding models should only be able to detect behaviors that have chewing noise. Instead, the feeding models are able to also detect pre-feeding activity which is not contaminated by chewing and therefore the LFP features that the model is using are unlikely to merely reflect chewing noise.

VS oscillations as a source of information for closed-loop and adaptive focal stimulation

Appetitive disorders such as binge eating have been associated with dysregulation of the brain reward circuit including the ventral striatum. Understanding the nature of network pathology is critical to the development of treatment that can therapeutically alter the network activity to ameliorate the problematic behavior or symptom. Neuromodulation-based interventions that use network activity to trigger and modify stimulation parameters have proven effective in the treatment of illnesses such as Parkinson's disease [24]. More recently, DBS for epilepsy has incorporated recording electrodes in order to detect pre-seizure

electrophysiological activity and use stimulation to prevent seizure generalization [23]. Future treatments of appetitive disorders will similarly utilize meaningful electrophysiology features in order to provide feedback for triggering and modification of stimulation parameters to optimize therapeutic efficacy and reduce the side effect profile [20]. A recent study demonstrated the theoretical feasibility of using a closed-loop system to trigger stimulation and decrease binge size in a mouse model of binge eating [19]. As noted in the introduction, the Halpern model was not very accurate due to only one LFP feature being used, but the success of the intervention highlights another point to consider when determining model complexity vs. model performance tradeoffs: what model performance is required to achieve a meaningful treatment outcome?

Conclusion

These findings suggest that imminent or ongoing feeding behavior can be decoded from ventral striatal oscillations. Further, information can be extracted from VS oscillations that correlates with the quantity and type of food being consumed. Together these findings suggest that VS oscillations are a rich source of information related to feeding behavior that could be used as real-time feedback for neuromodulation-based treatments or as neural markers for treatment development.

Funding and Disclosure

This work was supported by funds from the Department of Psychiatry at the Geisel School of Medicine at Dartmouth (AG), the Hitchcock Foundation (WD) and an LRP grant from the NIH NCATS (WD). By way of disclosure, over the past three years, Dr. Green has received research grants from Alkermes, Novartis and Janssen. He has served as an (uncompensated) consultant to Otsuka and Alkermes, and as a member of a Data Monitoring Board for Lilly. The other authors do not have any conflicts to disclose

References

[1] Berridge KC, Ho CY, Richard JM, DiFeliceantonio AG. The tempted brain eats: pleasure and desire circuits in obesity and eating disorders. Brain Res;1350:43-64.
[2] Ferrario CR, Labouebe G, Liu S, Nieh EH, Routh VH, Xu S, et al. Homeostasis Meets Motivation in the Battle to Control Food Intake. J Neurosci 2016;36(45):11469-81.
[3] Volkow ND, Wang GJ, Fowler JS, Tomasi D, Baler R. Food and drug reward: overlapping circuits in human obesity and addiction. Curr Top Behav Neurosci 2012;11:1-24.
[4] Muller UJ, Voges J, Steiner J, Galazky I, Heinze HJ, Moller M, et al. Deep brain stimulation of the nucleus accumbens for the treatment of addiction. Ann N Y Acad Sci 2013;1282:119-28.
[5] Pierce RC, Vassoler FM. Deep brain stimulation for the treatment of addiction: basic and clinical studies and potential mechanisms of action. Psychopharmacology (Berl) 2013;229(3):487-91.
[6] Kumar R, Simpson CV, Froelich CA, Baughman BC, Gienapp AJ, Sillay KA. Obesity and deep brain stimulation: an overview. Ann Neurosci 2015;22(3):181-8.
[7] Luigjes J, van den Brink W, Feenstra M, van den Munckhof P, Schuurman PR, Schippers R, et al. Deep brain stimulation in addiction: a review of potential brain targets. Mol Psychiatry 2012;17(6):572-83.
[8] Carus-Cadavieco M, Gorbati M, Ye L, Bender F, van der Veldt S, Kosse C, et al. Gamma oscillations organize top-down signalling to hypothalamus and enable food seeking. Nature 2017;542(7640):232-6.
[9] Leigh SJ, Morris MJ. The role of reward circuitry and food addiction in the obesity epidemic: An update. Biol Psychol 2016.
[10] Berendse HW, Groenewegen HJ, Lohman AH. Compartmental distribution of ventral striatal neurons projecting to the mesencephalon in the rat. J Neurosci 1992;12(6):2079-103.
[11] Zahm DS, Brog JS. On the significance of subterritories in the "accumbens" part of the rat ventral striatum. Neuroscience 1992;50(4):751-67.
[12] Groenewegen HJ, Wright CI, Beijer AV. The nucleus accumbens: gateway for limbic structures to reach the motor system? Prog Brain Res 1996;107:485-511.
[13] Vassoler FM, Schmidt HD, Gerard ME, Famous KR, Ciraulo DA, Kornetsky C, et al. Deep brain stimulation of the nucleus accumbens shell attenuates cocaine priming-induced reinstatement of drug seeking in rats. J Neurosci 2008;28(35):8735-9.
[14] Liu HY, Jin J, Tang JS, Sun WX, Jia H, Yang XP, et al. Chronic deep brain stimulation in the rat nucleus accumbens and its effect on morphine reinforcement. Addict Biol 2008;13(1):40-6.
[15] Knapp CM, Tozier L, Pak A, Ciraulo DA, Kornetsky C. Deep brain stimulation of the nucleus accumbens reduces ethanol consumption in rats. Pharmacol Biochem Behav 2009;92(3):474-9.
[16] Halpern CH, Tekriwal A, Santollo J, Keating JG, Wolf JA, Daniels D, et al. Amelioration of binge eating by nucleus accumbens shell deep brain stimulation in mice involves D2 receptor modulation. J Neurosci 2013;33(17):7122-9.
[17] Doucette WT, Khokhar JY, Green AI. Nucleus accumbens deep brain stimulation in a rat model of binge eating. Transl Psychiatry 2015;5:e695.
[18] Nangunoori RK, Tomycz ND, Oh MY, Whiting DM. Deep Brain Stimulation for Obesity: From a Theoretical Framework to Practical Application. Neural Plast 2016;2016:7971460.
[19] Wu H, Miller KJ, Blumenfeld Z, Williams NR, Ravikumar VK, Lee KE, et al. Closing the loop on impulsivity via nucleus accumbens delta-band activity in mice and man. Proc Natl Acad Sci U S A 2017.
[20] Deeb W, Giordano JJ, Rossi PJ, Mogilner AY, Gunduz A, Judy JW, et al. Proceedings of the Fourth Annual Deep Brain Stimulation Think Tank: A Review of Emerging Issues and Technologies. Front Integr Neurosci 2016;10:38.

[21] Rosin B, Slovik M, Mitelman R, Rivlin-Etzion M, Haber SN, Israel Z, et al. Closed-loop deep brain stimulation is superior in ameliorating parkinsonism. Neuron 2011;72(2):370-84.
[22] Beuter A, Lefaucheur JP, Modolo J. Closed-loop cortical neuromodulation in Parkinson's disease: An alternative to deep brain stimulation? Clin Neurophysiol 2014;125(5):874-85.
[23] Morrell MJ, Group RNSSiES. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology 2011;77(13):1295-304.
[24] Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, et al. Adaptive deep brain stimulation in advanced Parkinson disease. Ann Neurol 2013;74(3):449-57.
[25] van der Meer MA, Kalenscher T, Lansink CS, Pennartz CM, Berke JD, Redish AD. Integrating early results on ventral striatal gamma oscillations in the rat. Front Neurosci 2010;4:300.
[26] Fujisawa S, Buzsaki G. A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. Neuron 2011;72(1):153-65.
[27] Gatev P, Darbin O, Wichmann T. Oscillations in the basal ganglia under normal conditions and in movement disorders. Mov Disord 2006;21(10):1566-77.
[28] Sharott A, Moll CK, Engler G, Denker M, Grun S, Engel AK. Different subtypes of striatal neurons are selectively modulated by cortical oscillations. J Neurosci 2009;29(14):4571-85.
[29] van der Meer MA, Redish AD. Low and High Gamma Oscillations in Rat Ventral Striatum have Distinct Relationships to Behavior, Reward, and Spiking Activity on a Learned Spatial Decision Task. Front Integr Neurosci 2009;3:9.
[30] Tort AB, Komorowski RW, Manns JR, Kopell NJ, Eichenbaum H. Theta-gamma coupling increases during the learning of item-context associations. Proc Natl Acad Sci U S A 2009;106(49):20942-7.
[31] Howe MW, Atallah HE, McCool A, Gibson DJ, Graybiel AM. Habit learning is associated with major shifts in frequencies of oscillatory activity and synchronized spike firing in striatum. Proc Natl Acad Sci U S A 2011;108(40):16801-6.
[32] Berner LA, Avena NM, Hoebel BG. Bingeing, self-restriction, and increased body weight in rats with limited access to a sweet-fat diet. Obesity (Silver Spring) 2008;16(9):1998-2002.
[33] Corwin RL. Binge-type eating induced by limited access in rats does not require energy restriction on the previous day. Appetite 2004;42(2):139-42.
[34] Corwin RL, Buda-Levin A. Behavioral models of binge-type eating. Physiol Behav 2004;82(1):123-30.
[35] Doucette WT, Dwiel L, Boyce JE, Simon AS, Khokhar JY, Green AI. Machine learning based classification of deep brain stimulation outcomes in a rat model of binge eating using ventral striatal oscillations. 2018.
[36] Cohen MX, Axmacher N, Lenartz D, Elger CE, Sturm V, Schlaepfer TE. Nuclei accumbens phase synchrony predicts decision-making reversals following negative feedback. J Neurosci 2009;29(23):7591-8.
[37] McCracken CB, Grace AA. Nucleus accumbens deep brain stimulation produces region-specific alterations in local field potential oscillations and evoked responses in vivo. J Neurosci 2009;29(16):5354-63.
[38] Catanese J, Carmichael JE, van der Meer MA. Low- and high-gamma oscillations deviate in opposite directions from zero-phase synchrony in the limbic corticostriatal loop. J Neurophysiol 2016;116(1):5-17.
[39] Friedman J, Hastie T, Tibshirani R. Regularization Paths for Generalized Linear Models via Coordinate Descent. J Stat Softw 2010;33(1):1-22.
[40] He H, Bai Y, Garcia EA, Li S. ADASYN: Adaptive Synthetic Sampling Approach for Imbalanced Learning. Proceedings of the International Joint Conference on Neural Networks; June 20082008. p. 1322-8.
[41] Boggiano MM, Artiga AI, Pritchett CE, Chandler-Laney PC, Smith ML, Eldridge AJ. High intake of palatable food predicts binge-eating independent of susceptibility to obesity: an animal

model of lean vs obese binge-eating and obesity with and without binge-eating. Int J Obes (Lond) 2007;31(9):1357-67.
[42] Castro DC, Cole SL, Berridge KC. Lateral hypothalamus, nucleus accumbens, and ventral pallidum roles in eating and hunger: interactions between homeostatic and reward circuitry. Front Syst Neurosci 2015;9:90.
[43] Urstadt KR, Stanley BG. Direct hypothalamic and indirect trans-pallidal, trans-thalamic, or trans-septal control of accumbens signaling and their roles in food intake. Front Syst Neurosci 2015;9:8.
[44] Shin AC, Pistell PJ, Phifer CB, Berthoud HR. Reversible suppression of food reward behavior by chronic mu-opioid receptor antagonism in the nucleus accumbens. Neuroscience 2010;170(2):580-8.
[45] Pecina S, Berridge KC. Opioid site in nucleus accumbens shell mediates eating and hedonic 'liking' for food: map based on microinjection Fos plumes. Brain Res 2000;863(1-2):71-86.
[46] de Araujo IE, Oliveira-Maia AJ, Sotnikova TD, Gainetdinov RR, Caron MG, Nicolelis MA, et al. Food reward in the absence of taste receptor signaling. Neuron 2008;57(6):930-41.
[47] Horn A, Neumann WJ, Degen K, Schneider GH, Kuhn AA. Toward an electrophysiological "sweet spot" for deep brain stimulation in the subthalamic nucleus. Hum Brain Mapp 2017.
[48] Shamir RR, Dolber T, Noecker AM, Walter BL, McIntyre CC. Machine Learning Approach to Optimizing Combined Stimulation and Medication Therapies for Parkinson's Disease. Brain Stimul 2015;8(6):1025-32.
[49] Heldman DA, Pulliam CL, Urrea Mendoza E, Gartner M, Giuffrida JP, Montgomery EB, Jr., et al. Computer-Guided Deep Brain Stimulation Programming for Parkinson's Disease. Neuromodulation 2016;19(2):127-32.
[50] Grahn PJ, Mallory GW, Khurram OU, Berry BM, Hachmann JT, Bieber AJ, et al. A neurochemical closed-loop controller for deep brain stimulation: toward individualized smart neuromodulation therapies. Front Neurosci 2014;8:169.
[51] Kumar SS, Wulfing J, Okujeni S, Boedecker J, Riedmiller M, Egert U. Autonomous Optimization of Targeted Stimulation of Neuronal Networks. PLoS Comput Biol 2016;12(8):e1005054.
[52] Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 2010;13(3):369-78.
[53] Schurger A, Pereira F, Treisman A, Cohen JD. Reproducibility distinguishes conscious from nonconscious neural representations. Science 2010;327(5961):97-9.
[54] Schurger A, Sarigiannidis I, Naccache L, Sitt JD, Dehaene S. Cortical activity is more stable when sensory stimuli are consciously perceived. Proc Natl Acad Sci U S A 2015;112(16):E2083-92.
[55] Ledberg A, Montagnini A, Coppola R, Bressler SL. Reduced variability of ongoing and evoked cortical activity leads to improved behavioral performance. PLoS One 2012;7(8):e43166.
[56] Arazi A, Censor N, Dinstein I. Neural Variability Quenching Predicts Individual Perceptual Abilities. J Neurosci 2017;37(1):97-109.
[57] Saberi-Moghadam S, Ferrari-Toniolo S, Ferraina S, Caminiti R, Battaglia-Mayer A. Modulation of Neural Variability in Premotor, Motor, and Posterior Parietal Cortex during Change of Motor Intention. J Neurosci 2016;36(16):4614-23.
[58] Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. Neural variability in premotor cortex provides a signature of motor preparation. J Neurosci 2006;26(14):3697-712.

Figures and Legends

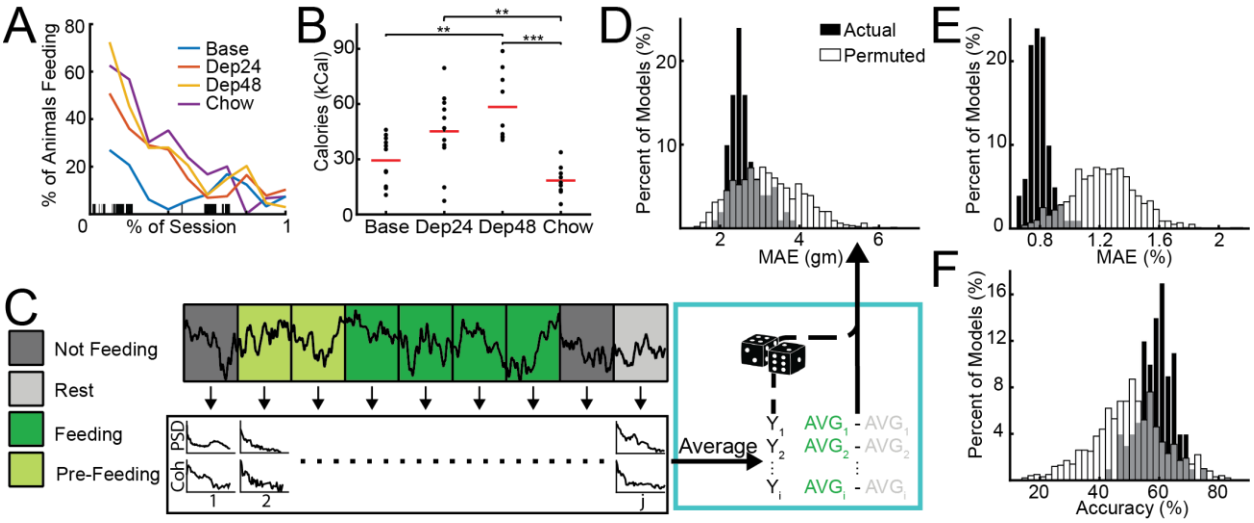

Figure 1. Predicting <u>aspects</u> of feeding behavior. **A.** Temporal dynamics of feeding behavior characterized on a population level across the four groups: baseline binge-like feeding of palatable chow (Base); binge-like feeding of palatable chow after 24-hours of food deprivation (Dep24) and 48-hours of food deprivation (Dep48); and limited access to house chow after 24-hours of food deprivation (Chow). Session lengths were normalized and the percent of animals feeding was smoothed for visualization. **B.** Kilocalories consumed across conditions. More calories were consumed in Dep48 (n = 8) than Base (n = 12), and in both Dep24 (n = 12) and Dep48 than Chow (n = 8). Red bar indicates group average. * $p < 0.05$, ** $p < 0.01$ **C.** Overview of analysis flow. Recordings were broken into non-overlapping 5 second trials and were assigned to one of 4 categories: not-feeding, rest, feeding, and pre-feeding. Power and coherence were calculated from each trial and averaged together according to category. Average brain activity of each animal during feeding was then normalized by activity during rest and used as predictor variables with behavioral metrics from that animal used as outcome variables. Outcome variables were then permuted to create distributions of 'by-chance' performances. **D-E.** Performance of models in predicting behavioral metrics; black distributions are performances of models using actual outcome variable assignment and white distributions are performances of models using permuted outcome variables. Distribution statistics reported as mean ± 95% confidence interval and effect size between Actual and Permuted given as Cohen's d. **D.** Mean absolute error (MAE) of predicting baseline food consumption in grams (gm). Actual = 2.6 ± 0.08 gm; Permuted = 3.2 ± 0.05 gm; d = 0.43. **E.** MAE of predicting percent change (%) in food consumption in Dep24 and Dep48 compared to Base (%): Actual = 0.8 ± 0.01%; Permuted = 1.2 ± 0.01%; d = 1.1. **F.** Percent accuracy (%) in predicting if house chow or palatable food was being consumed: Actual: 59 ± 1%; Permuted: 50 ± 1%; d = 0.52.

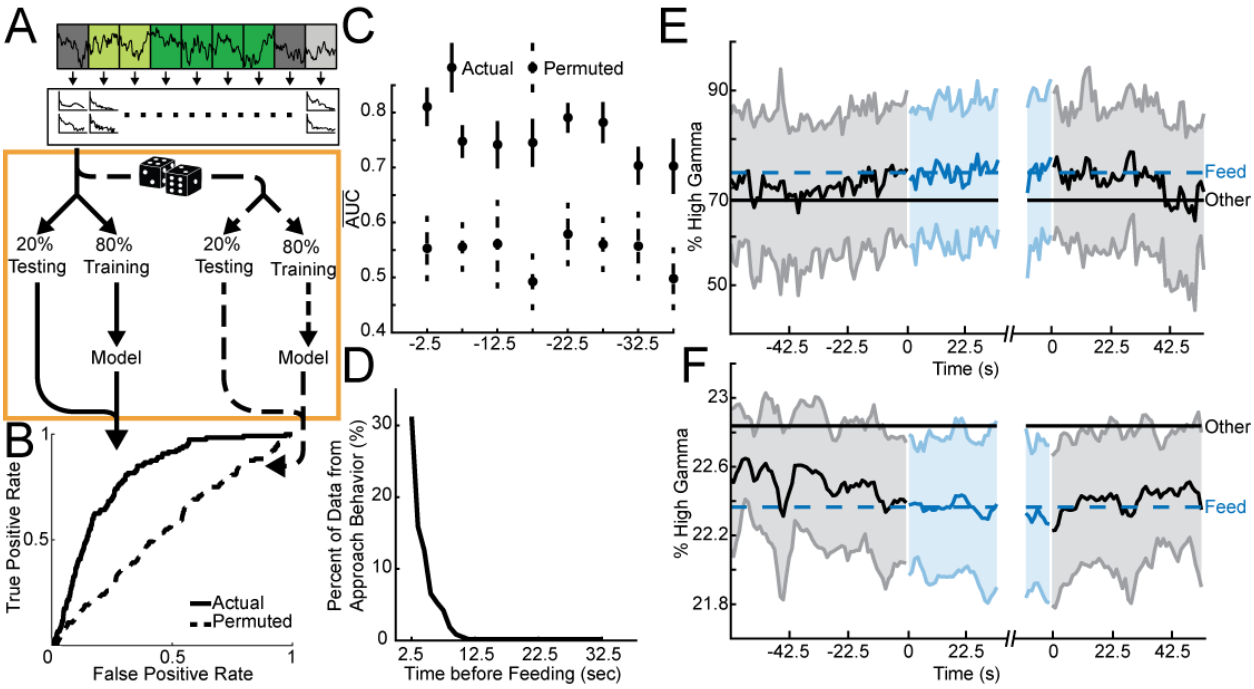

Figure 2. Classifying imminent feeding from other behavior with VS oscillations. **A.** Overview of analysis flow. Data was trialized and power and coherence extracted as in Figure 1. Orange box: Eighty percent of all pre-feeding trials immediately before feeding trials and all other trials outside of feeding (not-feeding and rest) are used for model training and the other 20% are used for testing. Using the same data, the assignment of category to predictor variables is permuted and trained on 80% of the data and tested on the left out 20%. This analysis produces two receiver operator (ROC) curves representing the performance of models built from actual and permuted data. **B.** Average performance of differentiating pre-feeding immediately before feeding begins and trials outside of feeding (not-feeding and rest); performance distribution statistics reported as AUC mean ± 95% confidence interval and the effect size between Actual and Permuted given as Cohen's *d*: Actual = 0.81 ± 0.03; Permuted = 0.55 ± 0.06; *d* = 2.68. **C.** Average areas under the ROC curves ($\overline{\mathrm{AUC}}$) obtained for 20 iterations of model building using models built from the pre-feeding trials immediately preceding feeding and tested on all other pre-feeding trials grouped by amount of time before feeding up to trials centered 42.5 seconds before feeding. **D.** Percent of pre-feeding data which overlaps with approach behavior; pre-feeding trials centered up to 12.5 seconds before feeding have at most 30% overlap with approach behavior. **E-F.** Normalized representative LFP features plotted through time around feeding. Features were normalized by either total power or by average coherence of the channel(s) of interest and averaged across all trials and animals around the beginning and end of feeding epochs. Shading indicates ±1 standard deviation. Black traces represent averages before and after feeding and blue traces represent averages during feeding. Dashed blue and solid black lines indicate average feature activity for either feeding trials or all other trials including those outside of these plots. **E.** Core left core right high gamma coherence increases during feeding. **F.** Shell left high gamma power decreases during feeding.

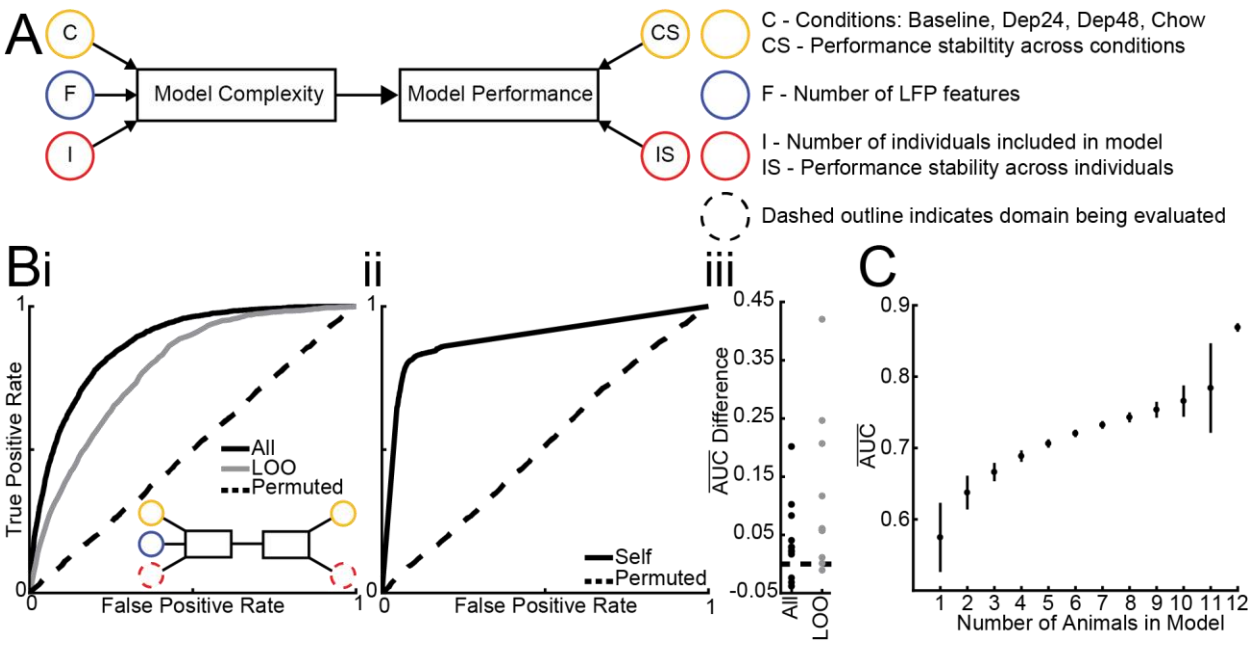

Figure 3. Number of individuals needed to build models. **A.** Schema of factors influencing model complexity and affected facets of model performance. Dashed outline indicates that a given factor is being manipulated in current analysis and that the performance is being measured accordingly. **B.** (Inset) The number of individuals and performance stability across individuals is assessed by comparing the average performance of population models (i) built using 80% of all animals' data and tested on the remaining 20% (All = 0.87 ± 0.01; *d* = 3.27), using a leave-one out (LOO = 0.78 ± 0.06; *d* = 3.13) approach, and using permuted data in a leave-one out approach (Permuted = 0.51 ± 0.02) to the average performance of individual models (ii) built from 80% of each animals' data and testing on that animals' left out 20% (Self = 0.89 ± 0.06; *d* = 3.13) alongside permuted individual models (Permuted = 0.52 ± 0.02). (iii) The subtraction of individual model performance on classifying the left out data of that individual from the performance of population models on the same data shows that individual models almost always outperform population models, using both All and LOO datasets. **C.** Average performance of models built from all possible combinations of including 1 to 12 individuals and tested on population data. Circle indicates mean and vertical bar indicates 95% confidence interval.

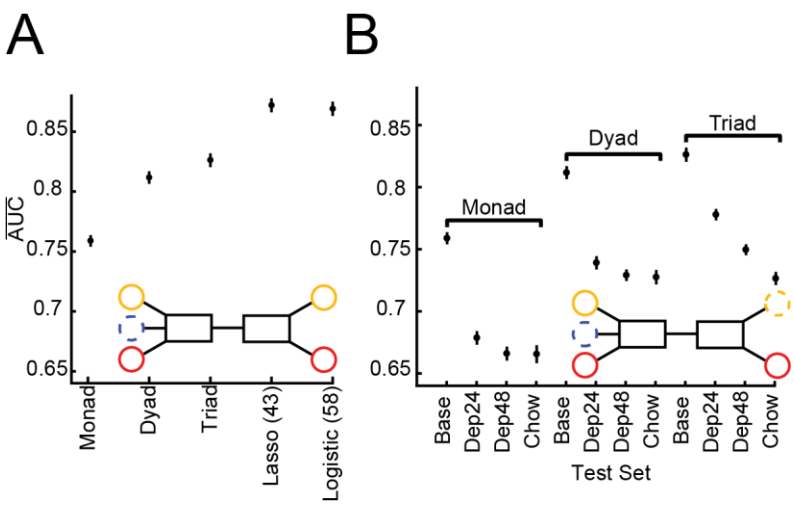

Figure 4. Number of LFP features needed to classify feeding. **A.** Average performance of models built and tested on population baseline data as the number of features used is manipulated (Inset): single features (Monad), two features (Dyad), and three features (Triad), using lasso (Lasso), and all 58 features (Logistic). **B.** Average performance of all population baseline Monads, Dyads, and Triads across conditions (Inset). Averages represented by circles and 95% confidence intervals by vertical bars.

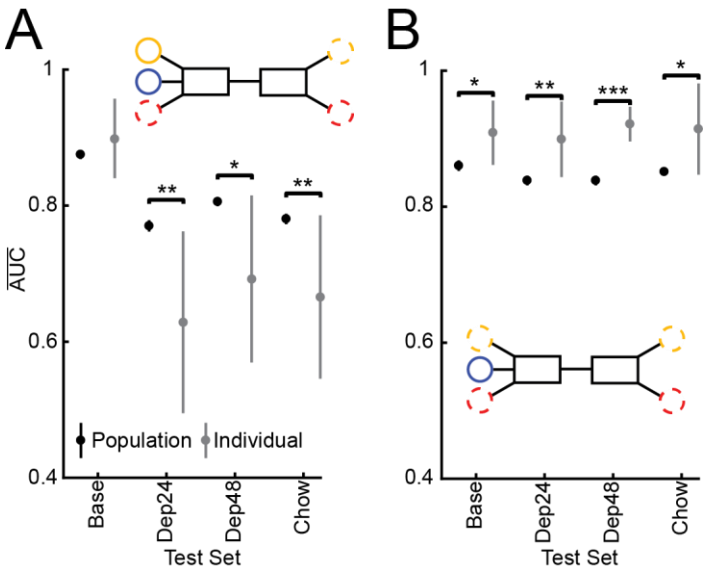

Figure 5. Performance of Population (black) and Individual (grey) models built from Baseline (**A**) and all condition (**B**) data and tested across states. **A.** Testing performance of population and individual models built from baseline data on all conditions (Inset). Population models built from Base do better than individual models when predicting behavior during any other condition than Base. **B.** Testing performance of population and individual models built from all conditions on all conditions (Inset). When models are built from all conditions, individual models do better than population models across conditions. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Averages represented by circles and 95% confidence intervals by vertical bars.

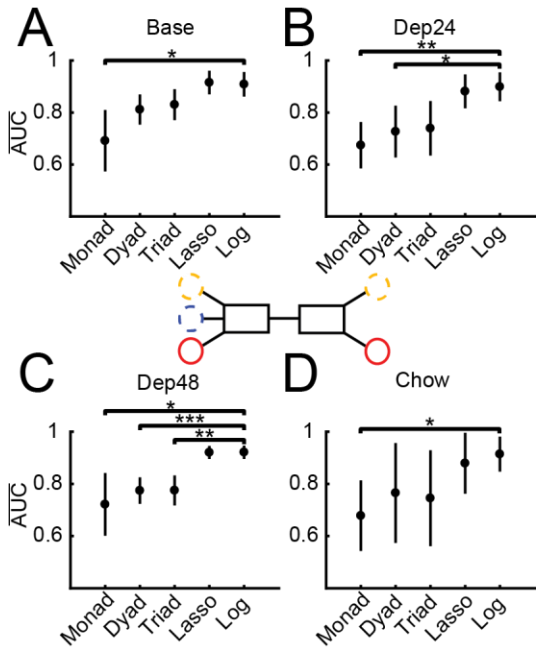

Figure 6. Manipulating all number of features and using all conditions to find simplest individualized models (Inset) and using the full 58 feature logistic regression (Log) as the performance to match. For predicting behavior at Base (**A**) and Chow (**D**) two features are sufficient to achieve maximal performance. For Dep24 (**B**) three features are sufficient and for Dep48 (**C**) Lasso is required. **A.** Only Monads have significantly lower performance than the full Log, $t(14) = 4.04$. **B.** Monads, $t(14) = 5.07$) and Dyads, $t(14) = 3.58$, have significantly lower performance than Log. **C.** Monads, $t(14) = 3.85$, Dyads, $t(14) = 6.12$, and Triads, $t(14) = 5.51$ have lower performance than Log. **D.** Only Monads have significantly lower performance than Log, $t(14) = 3.70$. * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$. Averages represented by circles and 95% confidence intervals by vertical bars.

	Monad	Dyad	Triad
Delta	—	38%	44%
Theta	—	—	—
Alpha	50%	88%	84%
Beta	—	13%	13%
Low Gamma	—	—	44%
High Gamma	50%	63%	72%

Table 1. Appearance of frequency ranges in top-tier models (models that are not significantly worse than the top performer). There are 2 Monads, 8 Dyads, and 32 Triads within the top-tier. Percentages of dyads and triads do not add to 100% since each model can contain more than one frequency range.