

SEARCH FOR INFORMATIVE CODING IN NEURAL ACTIVITY CORRELATIONS

Omer Hazon, Vasily Kruzhilin, {hazon, kruzh}@stanford.edu

Summary

This project is an application of machine learning to neural decoding. We have analyzed a data set of neural activity (calcium imaging) in the mouse hippocampus and prefrontal cortex in an animal performing a decision of making a left or a right turn. We investigated whether the ensemble coding of the neurons contribute more information than each neuron independently. Place decoding revealed this to be the case in the hippocampus. Additionally we were able to demonstrate the ability of end-arm and strategy decoding.

Data and features

The data set comes from a Schnitzer Group experiment (Stanford University) in which a mouse is trained to perform a simple reinforcement learning task with two different strategies on a plus-shaped maze (Fig.1). We analyzed three data sets comprised of consecutive days, in which on the middle day the reward strategy changes abruptly. For each time point each neuron contributes a feature with 0 if the neuron was inactive, otherwise a real value denoting the degree of activity. Ground truth labels are position (divided into spatial bins), end arm or strategy.

There are ~300-500 neurons per dataset, each contributing one feature. Naive Bayes uses binarized features. Each time point ($\Delta t=0.1s$) provided an example.




Figure 1: A picture of a mouse performing the task

Models

We used Bernoulli Naive Bayes that is known to be an ideal decoder in the case of independent features (and thus is invariant to intra-class shuffles):

$$\hat{k} = \arg \max_k p(k) \prod_{i=1}^m p_{ik}^{x_i} (1 - p_{ik})^{1-x_i} \quad p_{ki} = \left(1 + \sum_{j=1}^m x_j^{(i)} 1\{k^{(i)} = k\} \right) / \left(2 + \sum_{j=1}^m 1\{k^{(i)} = k\} \right)$$

and linear SVM with L1 regularization for end arm decoding and strategy decoding:

$$L(X, Y) = \frac{1}{2} \|w\|_2^2 + \frac{1}{m} \sum_{i=1}^m \max(0, 1 - y_i (w^T x + b)) + a \|w\|_1 \xrightarrow{w, b} \min$$

For place decoding, we applied the multiclass SVM with ECOC and without L1 regularization. The SVMs do not assume feature independence and thus the intra-class shuffles of input may worsen SVM performance.

Discussion and further developments

The primary goal of the project was to determine whether the neural encoding utilizes dependencies between neurons. For the problem of place decoding, it was found that intra-class shuffling degraded SVM performance, which serves as evidence that the feature independence assumption is incorrect and inter-neuron dependencies are present in the neural encoding of place in the hippocampus and indicates that for certain problems such as hippocampal place decoding, commonly used naive Bayes methods do not capture the full place information from the neurons. The end-arm predictive decoding and strategy decoding produced no such performance difference from hippocampus or prefrontal cortex data sets.

We were able to demonstrate that one can predict trial output based on hippocampus imaging and identify a small portion of cells that encode for this output in different points along the track. Our data suggests that even the encoding of a complicated strategy concept from the imaging data is possible.

In our project we were mostly limited by the amount of training data collected. When more data will be collected, the natural continuation of this project will involve the application of nonlinear SVM kernels and identification of particular sets of correlated neurons by selective shuffling.

Results

Switch day place decoding model:	Train error (distance) 70% (13,727 samples)	Test error (distance) 30% (5,883 samples)
Bernoulli Naive Bayes	1.0517 ± 0.0025	1.0788 ± 0.0068
Bernoulli Naive Bayes Shuffled	1.0523 ± 0.0024	1.0722 ± 0.0075
Multiclass Linear SVM	0.3772 ± 0.0023	0.6023 ± 0.0062
Multiclass Linear SVM Shuffled	1.3294 ± 0.0075	1.3700 ± 0.0107

Table 1: The comparison of performance of Naive Bayes and SVM position decoding for shuffled and unshuffled hippocampal data

It can be consistently seen across three days that in case of place decoding from hippocampus data the multiclass SVM performs best, Bernoulli naive Bayes second, and SVM with an intra-class shuffle worst (errors averaged over 16 runs), indicating that the place information is present in the neural correlations (Table 1, Fig. 2).

We also analyzed the possibility to decode where the mouse will end up in each trial before the actual turn takes place from cortex data. Only trials starting from the west arm on the switching day were analyzed (others lack examples from both classes). The Naive Bayes and SVM demonstrated similar performance and no difference was observed upon intra-class shuffling (Fig. 4). The L1 SVM regularization allowed us to identify which neurons encode the trial output at different arm positions (Fig. 5).

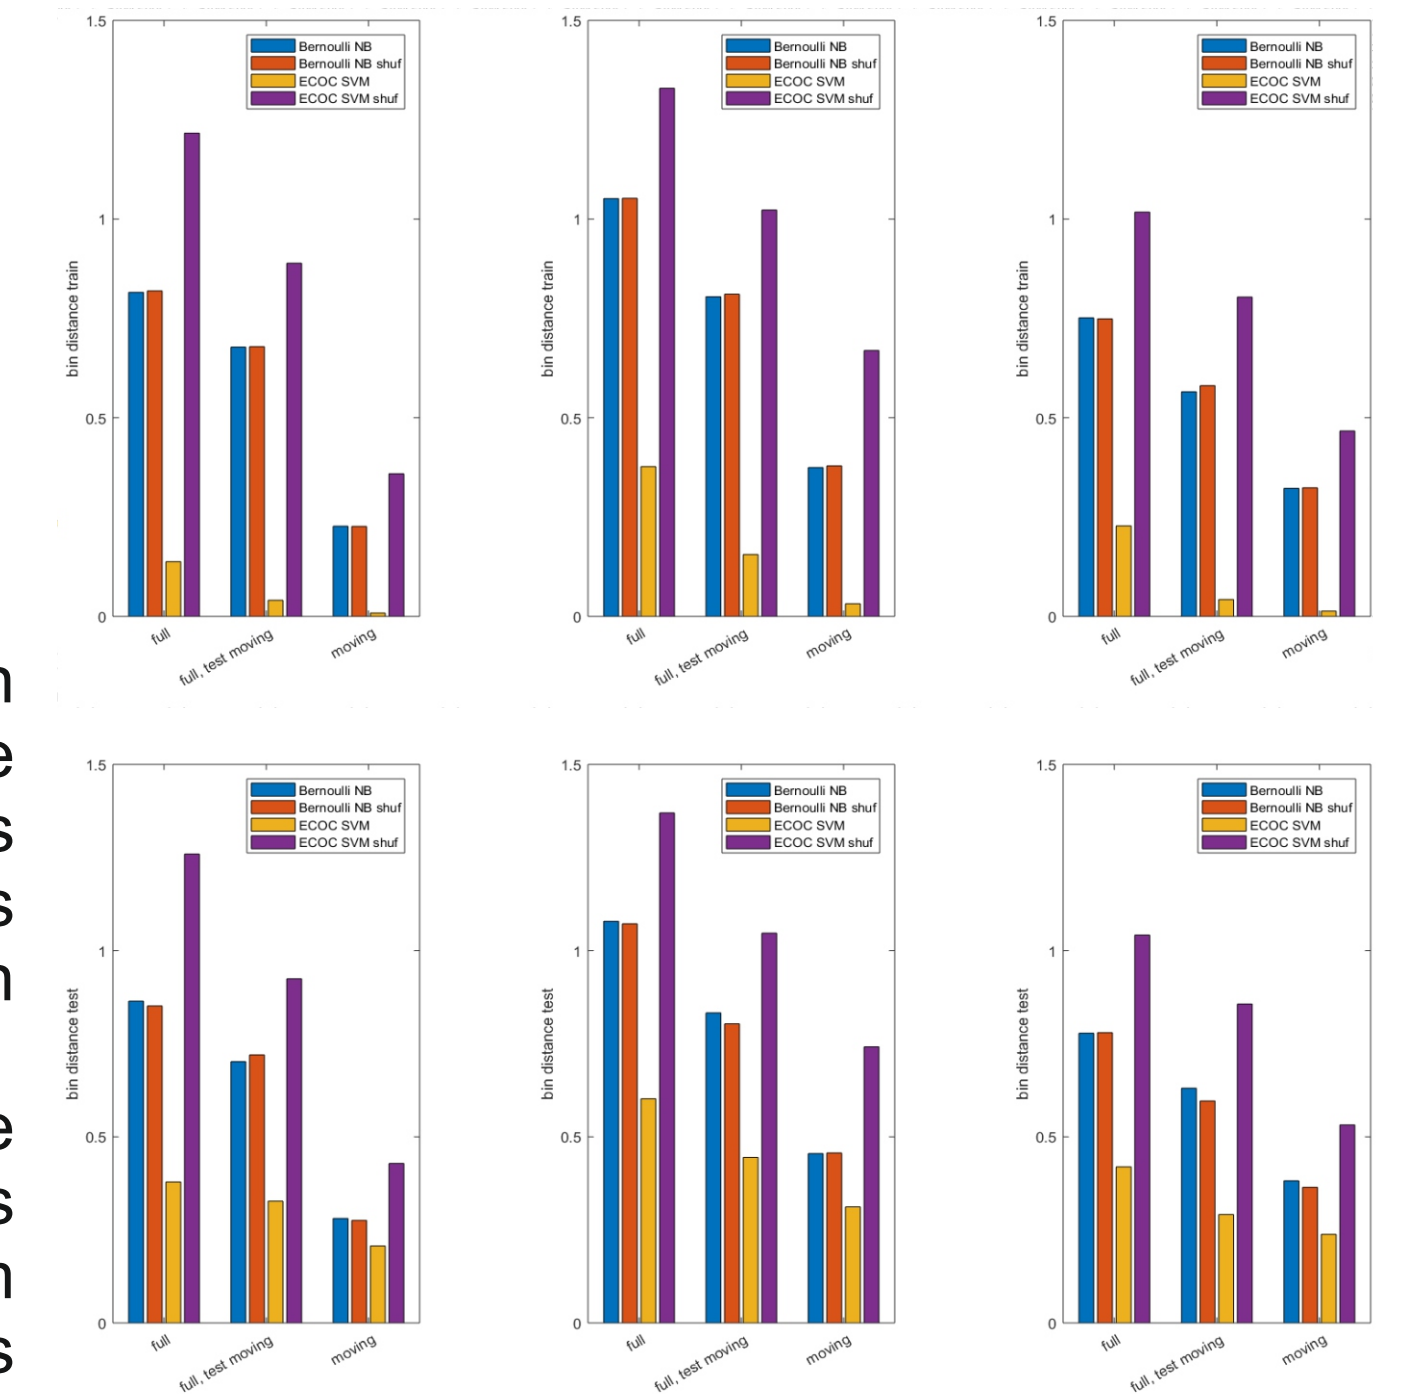


Figure 2: Training (top) and testing (bottom) average distance between true and predicted position bins for 4 decoders, Bernoulli naive Bayes and multiclass linear SVM, both plain and intra-class shuffled. Evaluated on (left groups) entire data set (middle groups) trained on entire set, tested on subset of mouse motion (right groups) trained and tested on subset of mouse motion. Blue and red columns correspond to shuffled and unshuffled Naive Bayes, orange and violet – to shuffled and unshuffled SVM, correspondingly.

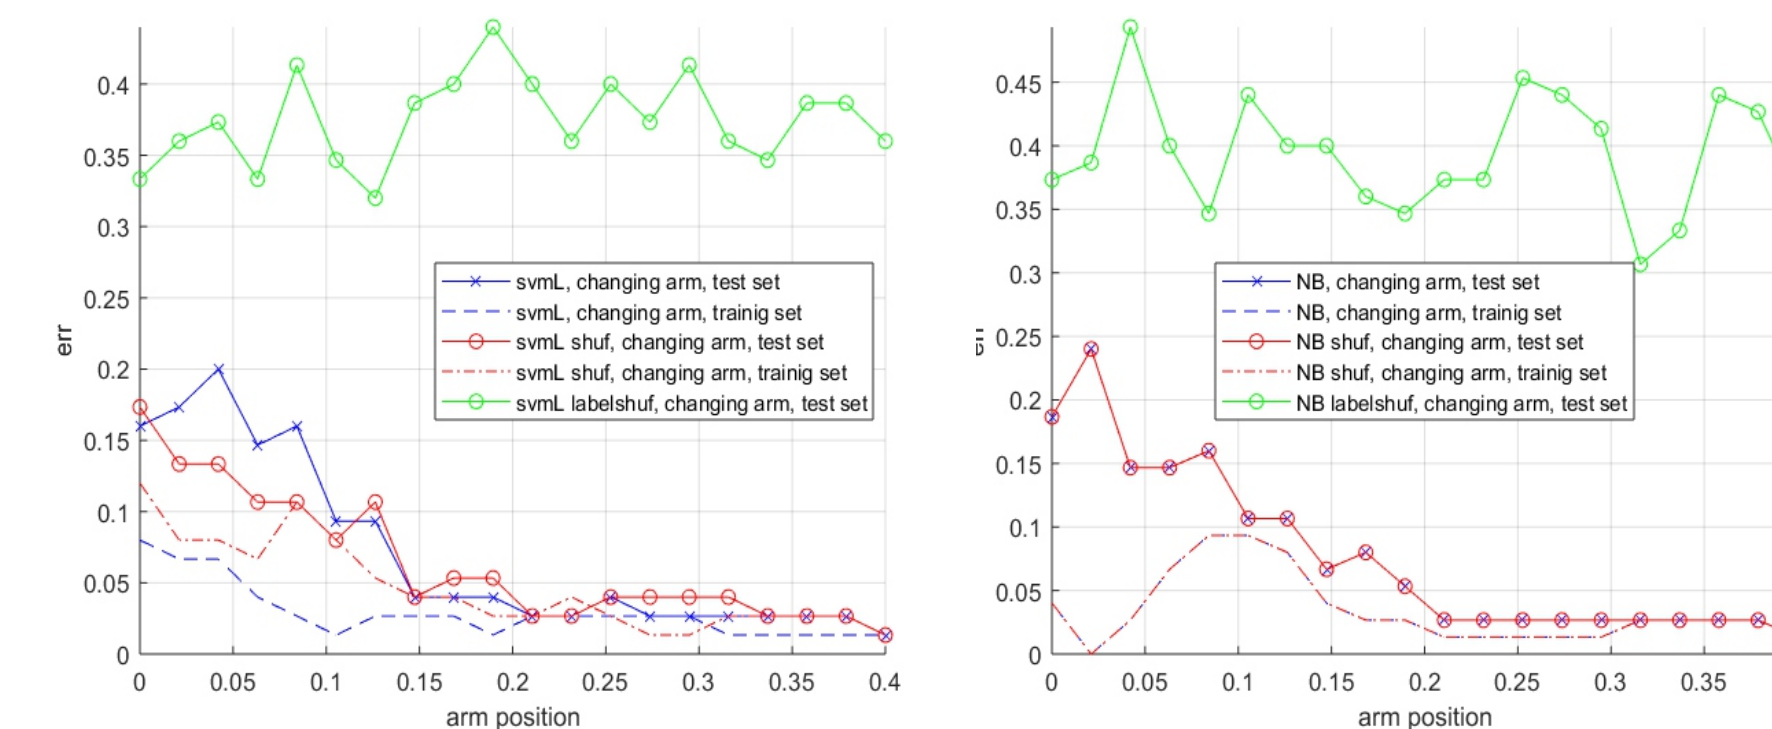


Figure 4: Decoding ending position (North/South) for west-starting trials (75 trials) on the switching day. (Left) Linear SVM and (right) Bernoulli naive Bayes. The colors indicate (blue) original data (red) Intra-class shuffled data (Green) ground truth labels shuffled (pure guessing), training error (dashed) and test error (solid) using leave-one-out cross-validation.




Figure 3: Strategy-decoding SVM error as a function of the arm position for the test (blue) and train (orange) data sets.




Figure 5: Calcium imaging traces for neuron found by L1 regularization on linear SVM decoding end arm. Showing activity for trials ending at (left) south and (right) north arms. Green and red boxes on the right indicate the trial output (correct/incorrect).

Our data set allows us to perform a test of an ability to decode the strategy the mouse is following in a particular trial. Taking only the trials that have the same starting and ending arm, we trained an SVM on hippocampus data to decode strategy on the two days with constant strategies and tested it on a switching day data. The best strategy decoding performance was found in the neural activity at the turning point of the track (Fig. 3). Surprisingly, no difference was observed upon shuffling.