III-64. Prediction error signals in ACC are scaled according to rational adjustments of learning

Matthew Nassar MATTNASSAR@GMAIL.COM
Joshua Gold JIGOLD@MAIL.MED.UPENN.EDU
University of Pennsylvania

Behavior often depends on the ability to update beliefs according to new information. Optimization of this process often requires preferential use of information occurring after likely environmental changes. Here we examined the role of the anterior cingulate cortex (ACC) in this process by measuring behavior in two rhesus monkeys and single-unit activity in one monkey performing a ten-alternative choice task that included both static fluctuations (noise) and abrupt changes (change-points) in the identity of the rewarded target. Subject performance was consistent with an optimal-inference model: they tended to switch choice targets more frequently after errors, most notably for errors that were unlikely to result from noise or that occurred shortly after a change-point. Moreover, they exhibited information-seeking behaviors during the feedback interval that revealed the extent to which feedback would be incorporated into future choices. In particular, the fraction of time searching through target locations and the latency to begin this search process (measured after the rewarded target was revealed visually but before juice delivery) reflected rational adjustments of influence, as reflected in subsequent choice behavior. We recorded the activity of 47 single units in the ACC of one monkey performing the task. Consistent with previous studies, we found units that responded preferentially to either reward or error feedback. To examine how adjustments of influence might affect these responses, we computed an error-response metric specific to each neuron that reflected how much the neural response on a given trial resembled that of the average error trial. This metric was larger for error trials and smaller for correct trials in which the monkey displayed more extreme information-seeking behaviors (e.g. shorter latency to search or larger search fraction), suggesting that ACC neurons scale reward prediction error signals according to the extent that those signals should be incorporated into future beliefs.

III-65. A high-performance, robust brain-machine interface without retraining

Paul Nuyujukian PAUL@NPL.STANFORD.EDU
Jonathan Kao JCYKAO@STANFORD.EDU
Joline Fan JMFAN@STANFORD.EDU
Sergey Stavisky SSTAVISK@STANFORD.EDU
Stephen Ryu SEOULMAN@STANFORD.EDU
Krishna Shenoy SHENOY@STANFORD.EDU
Stanford University

Brain-machine interfaces (BMIs) translate neural activity into control signals for prosthetic systems, such as computer cursors and robotic arms. BMIs strive to offer people with movement disabilities greater interaction with the world. Despite compelling proof-of-concept demonstrations, barriers to translation still remain (Ryu & Shenoy, Neurosurgical Focus, 2009). One such barrier is that without frequent decoder retraining and recalibration by expert technicians, system performance becomes unusably poor over time. We present here a robust BMI that does not require any retraining and can sustain high performance for up to a month. We trained a rhesus macaque (Monkey J) implanted with two 96-channel electrode arrays in M1/PMd to perform a 2D free-paced randomized grid keyboard selection task for a juice reward. A BMI decoder was built using spike threshold crossings (-4.5xRMS) utilizing the ReFIT-KF algorithm (Gilja et al., COSYNE, 2010), which then continuously controlled the on-screen cursor. The decoder's parameters were held constant for thirty consecutive days, and the monkey performed the BMI grid task for thousands of trials each day until satiated. Monkey J successfully achieved high performance, regularly sustaining 3-4 bits per second across hours. Over the course of the month long experiment, nearly 900 kilobits in total were communicated, comparable to the content in a short novel. These results were enabled by the inherent stability of spike threshold crossings across weeks and the high performance offered

by the ReFIT-KF algorithm. These findings demonstrate that without retraining or intervention, a high performing BMI can transmit significant, meaningful information over several weeks. As frequent retraining is untenable for long-term patient use, such robust performance is crucial for the successful translation of BMIs and should further increase their clinical viability.

III-66. Cerebellar granule cell activity during behavior: dynamics in light of the adaptive filter model

Sherika Sylvester[1] SJS2006@MED.CORNELL.EDU
Kayvon Daie[1] KPD7@CORNELL.EDU
Melanie Lee[1] MEL2010@MED.CORNELL.EDU
Mark Goldman[2] MSGOLDMAN@UCDAVIS.EDU
Emre Aksay[1] EMA2004@MED.CORNELL.EDU

[1]Cornell University
[2]University of California, Davis

For nearly forty years, the Marr-Albus-Ito model of the cerebellum has provided the dominant framework for understanding cerebellar control of motor learning. In this model and subsequent adaptive-filter models of cerebellar function, the cerebellar granule layer is proposed to separate sensory inputs into a large array of basis functions that can be reinforced by error signals during learning. However, direct tests of the role of granule cells during behavior have been prohibitive due to the miniscule size and close packing of these cells. Here, we overcome this barrier by using two-photon calcium imaging to monitor the activity of GFP-positive granule neurons during the horizontal optokinetic reflex. To quantify the activity of granule cells, the fluorescence time series of each neuron is averaged across optokinetic stimulus cycles and regressed against a linear model incorporating velocity, position, and calcium-buffering components. In total, we examined 44 GFP-positive granule neurons whose activity was well described by model fits ($r\hat{}2 \geq 0.7$). 89% carried signals that were dominated by stimulus velocity (velocity/position: 3.3$\pm$2.9). Of these neurons, a third exhibited activity that increased with stimulus velocity irrespective of direction, while the remainder exhibited direction-selective responses. Analysis of the relative sensitivity to leftward vs. rightward directed stimulation as a function of pair-wise distance suggests neighboring granule neurons encode similar information, and direction-specific velocity neurons tend to be more caudal in the zebrafish cerebellum. To describe the dynamics of granule activity, drifts following cessation of stimuli were quantified by fitting the average fluorescence response to a model of the underlying firing rate. The approximated time constants from this study ranged from 5 msec to 6 sec, consistent with the heterogeneity of dynamics hypothesized to be a feature of the granule layer in the adaptive-filter model (n=17 direction-selective neurons).

III-67. Action valuation in multi-effector decision-making

Seth Madlon-Kay SETH.MADLONKAY@NYU.EDU
Bijan Pesaran BP31@NYU.EDU
Nathaniel Daw NDD204@NYU.EDU

New York University

Models of action selection posit parietal and motor areas containing effector-specific maps that represent values of available motor actions. Effector-specific maps, however, cannot represent the values of actions that involve multiple effectors acting in coordination, so the mechanisms by which values of conjoint movements are represented and used to guide effectors are unknown. One possibility is that effector-nonspecific valuation areas step in to guide action selection via connectivity with motor or motor-planning regions. We investigated multieffector decision-making using model-based fMRI, exploiting the contralateral organization of hand control to contrast choices executed with coordinated bimanual versus unimanual hand movements. Subjects (N=20) performed a