brain-machine interfaces

----

one area where information theory can have an important impact on understanding neural phenomena is when using the lens of brain-machine interfaces. In such a paradigm, the subject attempts to interact with the environment through some external device that is coupled to its brain. Through this coupling, the subject is able to control the external device only through his thoughts that are manifested in the neural signals the external device records. The device processes the signals, takes actions, and then the subject observes the new state of the system and compares it to his/her high-level goal which induces a subsequent desired command.

Coleman's research group has first attempted to think about "canonical design" principles from a system level for designing BMIs. The high-level viewpoint that has emerged from their research is the interpretation of a BMI as two agents (one being the user, the other being the external device) that are cooperating to achieve a common goal. The common goal is manifested in terms of a cost function they are jointly attempting to minimize. What complicates matters is that the two agents do not have access to the same information: the subject has access to his high-level goal along with the perceptual feedback from the interface; he/she imagines a command, which is observed as neural signals (interpreted as the command going through a noisy channel). The external device must take these observations, actuate the system, and provide perceptual feedback to the user.

Our research paradigm is interested in optimal design of the two agents and how they react to the information that is causally presented to them. We formulate this as a decentralized control problem where the "high-level" goal is an arbitrary Markov process and demonstrate that it is intimately connected to feedback information theory under an appropriate cost function (pertaining to a log likelihood ratio) and source model. We have also demonstrated a structural result (about the existence of optimal schemes operating on certain sufficient statistics) that has important implications in describing not only how the external device should "estimate" information about the subject (which is what 99% of other BMI systems do), but also about what the minimal amount of feedback information should be provided to the subject, *and* how the subject should react to perceptual feedback, combine it with his/her high level goal, to specify the subsequent desired command.

- S. Gorantla and T. P. Coleman, "Shannon-Theoretic Viewpoints on Optimal Causal Coding/Decoding Problems", in preparation for submission, IEEE Transactions on Information Theory, October 2010

We have instantiated these findings with an EEG-based brain-machine interface to enable novel paradigms that have never been done before, including specification of a path in two dimensions and remote tele-operation of an unmanned aerial vehicle with a brain-machine interface.

- C. Omar, A. Akce, M. Johnson, T. Bretl, R. Ma, E. Maclin, M. McCormick, and T. P. Coleman, "A Feedback Information-Theoretic Approach to the Design of Brain-Computer Interfaces", in press, International Journal on Human-Computer Interaction , special issue on "Current Trends in Brain-Computer Interface (BCI) Research and Development"

Using BMIs to understand fundamental neurobiological phenomena

----

The aforementioned structural result and wedding of feedback information theory with decentralized control personally excites me because of the symbiotic relationship between systems theory and neuroscience that has emerged. If it were not for my interest in understanding the essence of a brain-machine interface, my appreciation for the interplay between feedback information theory and control would have never been established. Moreover, I can identify two interesting directions to pursue from more of a purely neuroscientific viewpoint.

I think this has ramifications for understanding neural systems that process sensory information with top-down feedback (e.g. many of the "predictive coding" models of the visual system, and the thalamo-cortical loop). I particularly think this is so in light of Ron Meir's recent Neural Computation paper where he showed that exact Bayesian inference on a latent continuous-time Markov process, using a point process model of inter-connected neurons, can be performed. What I particularly likes about his work was how he was able to *predict* experimental results from neuroscientists about how receptive fields or other things change when the statistics of the input changes.

@article{rao1999predictive,

  title={{Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects}},

  author={Rao, R.P.N. and Ballard, D.H.},

  journal={nature neuroscience},

  volume={2},

  pages={79--87},

  year={1999},

  publisher={NATURE AMERICA}

}

@article{bobrowski2009bayesian,

  title={{Bayesian filtering in spiking neural networks: Noise, adaptation, and multisensory integration}},

  author={Bobrowski, O. and Meir, R. and Eldar, Y.C.},

journal={Neural computation},

volume={21},

number={5},

pages={1277--1320},

year={2009},

publisher={MIT Press}

}

Connecting the structural results from the team decision theory viewpoint, where careful attention is payed to the structure of the cost function, could elucidate interesting findings on how receptive fields in peripheral neural structures change as a function of top-down feedback from cortical areas. What I explicitly am interested in doing is using this theory to provide predictions about the nature of complex visual processing that is mediated because of the thalamo-cortical feedback loop. Many complex "non-classical" effects in the visual system represents sensory signals cannot be established without taking into consideration feedback from higher areas as well as the perception-action cycle more generally. Understanding the phenomena will most likely require wedding careful attention to the neurobiological mechanisms mediating these feedback processes with system theoretic ideas involving directed information, feedback, sequential predction, etc.

Another interesting area I am currently pursuing with a student is using a brain-machine interface as a lens to understand fundamental mechanisms of learning and plasticity. When a subject's brain is connected to a new external device, it must "learn" how to interact with this device to accomplish the goal at hand. This requires notions of learning and plasticity. A fascinating paper about how the motor cortex "re-wires" itself in a stable manner depending upon what type of external device is coupled to it is found in this recent paper by Jose Carmena at Berkeley. What more and more BMI researchers are finding is that a BMI is itself a tool to understand carefully the dynamics of learning and plasticity. This could have profound basic scientific under-pinnings, as demonstrated in Jose's paper.

- Ganguly K. and Carmena J.M. (2009). Emergence of a stable cortical map for neuroprosthetic control. PLoS Biology 7(7): e1000153.

Our group is currently using principles of feedback information theory and control to design careful BMI experiments to elucidate some of these phenomena in a principled manner. We are specifically interested in understanding some of the reward systems in the human brain that use "mirror neurons" to understand and predict the structure of the environment. It is well-established that a fronto-central voltage deflection is generated following the presentation of negative feedback to a subject, termed the Feedback-Related Negativity (FRN). The leading theory of the FRN proposes that it is reflective of feedback deviations from the subject's expected outcome during the process of reinforcement learning. Recent studies demonstrate that the reinforcement learning model of the FRN correctly predicts FRN amplitude and phase coherence during single-trial EEG learning tasks, and that the FRN *presages* behavioral adaptation by the user. Here, we exploit these findings by applying the FRN to a Brain-Computer Interface as a form of adaptive feedback. We develop novel applications and control policies that enhance Subject-BCI interaction by incorporating the FRN and a reinforcement learning model into the control scheme. Along with elucidating the dynamics of learning, incorporation of the BMI in this framework has the potential to accelerate learning in accomplishing certain tasks.

- Steines David, Bretl T., Maclin E., Coleman, T. (2010) "Application of Feedback-Related Negativity to Understand Plasticity and Learning with a Brain-Computer Interface", to appear, Society for Neuroscience, San Diego CA, Nov 2010.

We are also doing purely theoretical studies with this endeavor, using ideas of "apprenticeship learning". We exploit how this motor neuron system elicits event-related potentials in the presence of oddball stimuli related to deviations from the subject's expectations of how the subject will evolve so that we can design novel ways to "train" artificial intelligence algorithms.

- Steines David, Bretl T., Maclin E., Coleman, T. (2010) "Monday Morning Quarterback Learning", in preparation for submission to International Conference on Machine Learning, October 2010.

Causality in Neural Systems

The famous neuroscientist Donald Hebb had a hypothesis a long time ago that neurons join one another in functionally specialized "neuronal assemblies" to coordinate sensation, perception, and action. This notion of cell assemblies has been difficult to test because of the difficulty to do experiments. It is now increasingly common in neural systems to record many neural signals simultaneously. Attempting to understand the nature of these functionally specialized coordinated actions, and how they co-vary with intented action, with context, with attention, and with sensory input is something that has the potential to now be understood. Understanding a network-level notion of causal interaction between simultaneously recorded signals, and how this nature of interaction changes with the aforementioned variables, is of crucial importance. The typical approach that 99% of the neuroscience community has attempted to understand causality is the notion of Granger causality that was espoused by the Nobel prize winning economist Granger. His high level viewpoint that "X causes Y if I can predict the future of Y given the past of Y and the past of X better than predicting the future of Y only given the past of Y" is philosophically pleasing, but he mathematized this notion in a very simple manner that does not naturally extend to arbitrary (i.e. nonlinear) modalities.

I have been working on inferring the causal generative structure of simultaneously recorded signals with applications to understanding mechanistic phenomena in ensemble neural recordings. In short, we take a sequential prediction approach to generalize the notion of Granger causality to arbitrary modalities. Remarkably, it turns out that one can take Granger's original viewpoint and interpret in such a way that the mathematization - on arbitrary modalities - is exactly the directed information that has recently appeared in the information theory and control literature. We extend this to a network measure of causality involving more than 2 processes - and develop a provably good inference algorithm that is able to discriminate between mediating processes and cascade phenomena. We demonstrate how this is a more parsimonious approach compared to standard "directed Bayesian networks" viewpoint and is provably good. We then have extended this to address non-stationarities in the causal network and have developed a preliminary approach that works particularly well on simulated data.

Within the context of neuroscience, we were able to demonstrate that with ensemble neural spike recordings in primary motor cortex, information propagates spatio-temporally across one dominant axis when an awake-behaving monkey is planning to initiate a motor movement. This is consistent (even the direction of the wave!) with what was found with local field potentials in Nicho Hatsopoulos's lab at U. Chicago.

- C. Quinn, T. P. Coleman, N. Kiyavash, and N. G. Hatsopoulos, "Estimating the directed information to infer causal relationships in ensemble neural spike train recordings", Journal of Computational Neuroscience, June 2010 [pdf]

The slides for the work on causality from a recent talk i gave are at this hidden link:

- http://colemant.ece.illinois.edu/ColemanPresentationCausalityASDN2010.pdf

Plasticity and the Development of Self-Organizing Maps in Cortical Systems

----

It is well known that Long-term modification of synaptic efficacy can depend on the timing of pre- and postsynaptic action potentials. Spike timing dependent plasticity (STDP) is a neuro-biological mechanism underlying Hebbian plasticity ("cells that fire together, wire together") that also elicits the known features of synaptic competition and postsynaptic firing regulation. STDP strengthens synapses that receive correlated input, and it was recently shown in a 2001 Song and Abbott paper

```
@article{song2001cortical,
  title={{Cortical development and remapping through spike timing-dependent plasticity}},
  author={Song, S. and Abbott, L.F.},
  journal={Neuron},
  volume={32},
  number={2},
  pages={339--350},
  year={2001},
  publisher={Elsevier}
}
```

This can lead to the formation of stimulus-selective columns and the development, refinement, and maintenance of selectivity maps in network models. This is particularly interesting because of the experimental neurobiological findings of such selectivity maps in the visual system.

- Yuste, R., and Sur, M. (1999). Development and plasticity of the cerebral cortex: from molecules to maps. J. Neurobiol. 41, 1–6.

However, Song and Abbott's paper was completely simulation based and phenomenological. We are interested in developing a theoretical understanding of how these maps emerge and what statistical dynamics, control, and information theoretic underpinnings underlie this phenomena. In our recent work, we consider neuronal network models with plasticity and randomness; we show that complicated global structures can evolve even in the presence of simple local update rules. Our models involve the use of point process thoery and our analysis uses tools from the theory of Markov chains.

Specifically, we propose a time model of the evolution of a neuronal network capable of learning with multiple input neurons (whose spiking consists of Poisson processes), one output integrate-and-fire neuron, and a network with conductances between them. Along with the input firing, the stochasticity of the output neuron's state is modulated by the conductances in the network, that evolve according to synaptic plasticity in the form of long-term potentiation. The aggregate system evolves according to discrete-state, multi-dimensional Markov chain. We demonstrate that the network is capable of rich properties (e.g. bifurcation, various forms of stability, etc) that depend on maximum possible values of conductance, the rates of the point processes, and the number of levels in the integrate-and-fire output neuron model. Most importantly, we believe that this approach provides a well-positioned balance between neuro-biological relevance and ability to be analyzed theoretically. We are currently applying principles from Lyapunov exponents (for sensitivity to initial conditions) and information theory (using the KL divergence for model order reduction and employing feedback information theoretic principles using the directed information to elucidate information dynamics) so that the key emerging principles underlying the complex behaviors can be illustrated in a provably good manner.

- N. Arizumi, T. P. Coleman, and R. L. Deville, "Evolution of complex neuronal network structures in the presence of plasticity", Computational Neuroscience Annual Meeting (CNS) , San Antonio, TX, July 2010.