Proceedings of the 5th International
IEEE EMBS Conference on Neural Engineering
Cancun, Mexico, April 27 - May 1, 2011

Team Decision Theory and Brain-Machine Interfaces

Sanggyun Kim
Department of ECE
University of Illinois
Urbana, IL
Email: sgkim1@illinois.edu

Todd P. Coleman
Department of ECE
University of Illinois
Urbana, IL
Email: colemant@illinois.edu

***Abstract*—In this paper we present a general-purpose design methodology for designing policies for interaction between the user and external device for brain-machine interface (BMI). In short, we interpret a BMI as a system comprising two agents (the user and the external device) cooperating to achieve a common goal. Because of the inherent uncertainty in (a) the user's intent, and (b) the noisy channel mapping desired commands to neural recordings, neither agent has a subset of the information of others. Nonetheless, we exploit recent research results to demonstrate how to design - for an arbitrary problem specification (e.g. cost function to minimize) - optimal policies that are easily implementable in BMIs across many modalities - including EEG and cortically controlled devices. The structural result we provide sheds light on the minimal amount of useful information that is required to provide perceptual feedback to the user.**

I. Introduction

A brain-machine interface (BMI) is a system that elicits a direct communication pathway between a human and an external device. In many cases, it is the objective of the user to control an external device merely by imagination, and the external device acquires neural signals and actuates some physical system, and perceptual feedback is given to the user to complete the loop [1], [2], [3].

Note that in general, there are many possible ways to design

- (a) what feedback should be delivered to the user
- (b) how the user should react to the feedback and its intended objective to imagine the subsequent desired control command
- (c) how the external device should sequentially map its recorded neural signals to a control action

(a) is currently being addressed by many researchers in the field (e.g. by considering visual augmentation or cortical stimulation of proprioception [4]). (b) was originally addressed primarily using principles of recursive estimation of a latent process [5]. More recently, researchers have exploited the fact that subsequent desired control commands are a *reaction* to the user's intent and the perceptual feedback it receives; as a consequence, they use principles of feedback control to *model* how the user reacts [6], [7]. Nonetheless, such an approach is not a model and not a design consideration[1]. Lastly, the issue of (c) has traditionally been applied using principles of recursive estimation [5]. Also, an increasing amount of interest has been applied in understanding how sensory information can be provided back to the brain (visual augmentation or cortical stimulation of proprioception) [4].

Here, we attempt to think at a system engineering level on how one could optimally design all aspects of the above issues in one problem formulation. We interpret the BMI as a system comprising two agents, or decision-makers, cooperating to achieve a common goal. One agent is the user, that sequentially combines its high-level intent with perceptual feedback from the external interface to imagine a subsequent control command. This imagined command is an input to a 'noisy channel', whose output is the neural signal that is recorded by the second agent, the external device. The external device combines all its recorded neural signals to take a control action (move a cursor, convey text, actuate a robotic arm, etc.). It also provides perceptual feedback to the user (e.g. visual, cortical stimulation of proprioception, etc.). In short, we are defining a protocol of interaction between *both* agents, the user and the external device. We choose an optimal coordination strategy depending upon what goal the two are attempting to accomplish, e.g. what cost function they are minimizing. Remarkably, we show through a structural result [8] that there exist optimal strategies that simultaneously address (a) through (c) *and* they are user-friendly.

II. Basic Problem Formulation

Because of the inherent uncertainty in (a) the user's intent, and (b) the noisy channel mapping desired commands to neural recordings, neither agent has a subset of the information of others. This lies in a class of "sequential, non-classical, team decision problem" [9], [10], [11] for which in general, finding optimal coordination strategies

[1]Perhaps another reason optimizing over (b) has not carefully been addressed in the literature because a theoretically optimal reaction strategy for a user might be deemed useless practically if it is complicated and/or imposes high cognitive load

978-1-4244-4141-9/11/$25.00 ©2011 IEEE

is computationally intractable [12]. However, a large class of BMI-type problems can indeed be efficiently solved, as we shall show.

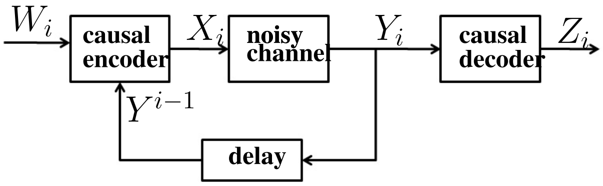

Fig. 1. Basic problem setup: an optimal causal coding/decoding problem.

Without loss of generality, because we do not know yet what perceptual feedback is the most relevant, we address issue (a) by first considering the most general scenario where *all* information available to the decoder at any time i is fed back to the subject. In the process of solving this problem, we will show that there always exists a simple scheme that can be implemented with a brain in the loop.

Assume that the user's intent process $W = (W_1, \ldots)$ is a general Markov (e.g. stochastically continuous) process where each $W_i \in \mathsf{W}$. The causal encoder represents the user who at time i combines perceptual feedback $Y^{i-1} = (Y_1, \ldots, Y_{i-1})$ from the external device as well as causal information about its intent process W^i to specify the subsequent desired control command $X_i \in \mathsf{X}$: $X_i = e_i(W^i, Y^{i-1})$. The neural activity recorded at time i, Y_i, is a 'noisy' version of X_i, in which its statistics co-varies with X_i. For example, the statistics of Y_i are different when imagining a left-oriented movement $X_i = 0$ as compared to imagining a right-oriented movement $X_i = 1$ [13]; another example could be the model of how neural spiking in motor cortex co-varies with desired arm velocity. The external device acts as a 'causal decoder', who specifies a decision variable $Z_i \in \mathsf{Z}$ that is a causal function of channel outputs $(Y_1, \ldots, Y_i)$: $Z_i = d_i(Y^i)$. Neither agent's observations at any time point are a nested version of the other's, so they have a 'non-classical' information structure [11], making this in general a 'hard' problem. They jointly design their strategies $\pi = (e, d)$ to minimize a function $J^{\alpha}_{n,\pi}$ pertaining to a goal they would like to achieve. This is equivalent to minimizing an expected sum of costs operating on current state, observation, and decision variables:

$$J^{\alpha}_{n,\pi} = \frac{1}{n}\mathbb{E}_{e,d}\left[\sum_{i=1}^{n} g(W_i, X_i, Z_{i-1}, Z_i)\right]. \qquad (1)$$

One natural cost function might be $g(w_i, x_i, z_{i-1}, z_i) = (w_i - z_i)^2 + \alpha x_i^2$. Here, the goal is to minimize estimation error between the cursor position z_i and the intended cursor position w_i, while guaranteeing the cognitive load is not too demanding. Here, we are interested in optimizing over *all* possible policies of interaction (e, d) between the user and the interface to minimize (1). To emphasize, not only are we optimizing how the external device should take its neural signals and take some action (c), but also we are optimizing (a) what perceptual feedback should be specified back to the user, *and* (b) how the user should react to the perceptual feedback to specify the subsequent imagined control signal X_i.

III. Problem Solution

Our research group has recently shown in [8] that for the problem pertaining to Figure 1 with the goal being specified by the cost function in (1), there always exists an interaction strategy of the form given in Figure 2.

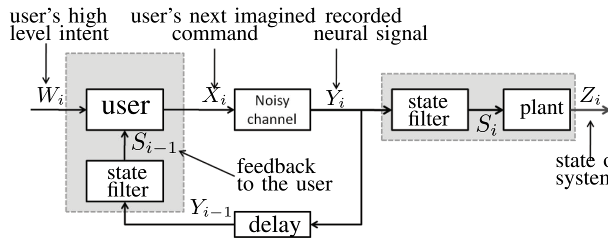

Fig. 2. Structural result within the context of a BMI: in an optimal system, the user acts as part of the causal encoder. The other part accumulates all causal observations and summarizes them into sufficient statistics acting as perceptual feedback to the user.

Our structural result says that first a state filter can construct sufficient statistics $S_i = (Z_{i-1}, B_{i|i})$, where

$$b_{i|i}(A) \triangleq \mathbb{P}\left(W_i \in A | Y^i\right), \quad A \subset \mathsf{W} \qquad (2)$$

and then the external device can actuate the plant using S_i and the user only needs the sufficient statistics S_{i-1} as perceptual feedback. This information, along with the current high-level goal Z_i, is all that is needed to specify an optimal causal encoder. Secondly, our structural result specifies a *sequential decomposition*, which means that an explicit optimal strategy - operating on sufficient statistics - can always be efficiently found using dynamic programming.

Note that there are many ways that the belief can be fed back to the user - one of which is using a cortical proprioception paradigm; another is through alteration of the visual space so that areas of the space are 'zoomed in' where the posterior belief is large, and likewise for the opposite case. The perceptual feedback about Z_i is quite natural - the user typically sees visually the position/velocity of the system that it is trying to control.

In previous work, our group instantiated a more restricted version of this formalism

- a restricted assumption about the user's high-level intent (that $W_i = W_{i-1}$ so that the user's intent was completely known to the user at all times, and that W was an ordered sequence in a discrete symbolic language)
- a specific goal pertaining to 'sequential information gain' [8] where $\mathsf{Z} = \mathcal{P}(\mathsf{W})$ - so that the outputs of the external device were *beliefs* about the source - and

the cost function was given by

$$g(w_i, x_i, z_{i-1}, z_i) = -\log \frac{dz_i}{dz_{i-1}}(w_i) + \alpha\eta(x_i) \quad (3)$$

$$\Rightarrow J_{n,\pi} = -\frac{1}{n} \underbrace{\sum_{i=1}^{n} I(W; Y_i | Y_{i-1})}_{I(W;Y^n)} + \sum_{i=1}^{n} \alpha\eta(X_i)$$

so that the problem was to maximize information from the source to the observations, with feedback, subject to cognitive load constraints. Our group has exploited a recently developed elegant and time-invariant feedback communication strategy that achieves capacity over any noisy channel with feedback [14] and is optimal in our aforementioned team decision-theoretic sense [15],[8, Lemma V.3]. We demonstrated that the optimal solution is easily implementable by humans in BMIs - demonstrating its efficacy in BMI text-spellers [16], specification of smooth paths in two dimensions [16], and remotely tele-operating unmanned aerial vehicles [17].

In this paper, we develop a more general formalism that captures the previous work as special cases, where we assume that a completely general goal (or cost function), and our assumption on the high-level intent is that it arrives *causally* to the user and it forms a Markov process. We take structural results recently established in [8] for general decentralized control problems and interpret their consequences in terms of designing optimal and user-friendly BMIs. We note that because of our general 'noisy channel' model, our viewpoint is applicable to any neural sensing modality - including EEG, ECoG, LFP, or spike train recordings. Moreover, the structural result pertaining to feedback, could have implications in a variety of feedback delivery mechanisms - including proprioceptive feedback, visual adaptation, or other forms of altering the senses.

To capture the BMI-unique issue that one of the two agents has a brain in the loop, we consider a subclass of encoder/decoder policies $\bar{\pi}$ - termed $\widehat{\text{S}}$tationary Markov (SM) coordination strategies - that pertains to two functions $\bar{e} : \mathsf{W} \times \mathsf{Z} \to \mathsf{X}$ and $\bar{d} : \mathsf{Z} \times \mathsf{Y} \to \mathsf{Z}$ that operate as the user and decoder (See Figure 3):

$$\begin{aligned} x_i &= \bar{e}(w_i, z_{i-1}) && (4a) \\ z_i &= \bar{d}(z_{i-1}, y_i). && (4b) \end{aligned}$$

Using the variational equations for the rate-distortion function and capacity-cost function [8], we have related costs to log likelihood ratios to develop an inverse optimal control formalism that provides sufficient conditions under which an SM policy is globally optimal for a cost function (1).

IV. An Example

Here we consider a canonical problem for BMIs, where the high-level intent arrives *causally* to the user. Let $\mathsf{W} = \mathsf{X} = \mathsf{Y} = \mathsf{Z} = \mathbb{R}$. The source is a Gauss-Markov process

$$W_i = \rho W_{i-1} + \tilde{W}_i \quad , \tilde{W}_i \sim \mathcal{N}\left(0, \sigma_m^2\right), \quad i \geq 0 \quad (5a)$$

$$W_0 \sim \mathcal{N}\left(0, \frac{\sigma_m^2 \sigma_v^2}{L + \sigma_v^2(1-\rho^2)}\right). \quad (5b)$$

Note that we are not assuming that W is stationary. The output Y_i is given by an additive Gaussian noise (AGN) channel:

$$Y_i \;=\; X_i + V_i, \quad V_i \sim \mathcal{N}\left(0, \sigma_v^2\right) \quad (6)$$

Note that here the subject's high-level intent arrives causally. This could pertain to driving a cursor on a screen, or operating a robotic limb.

We can show using inverse optimal control viewpoint that for the problem of minimizing the linear quadratic cost, a linear SM policy is globally optimal:

Lemma 4.1 ([8]): For the infinite-horizon linear quadratic cost problem

$$J_\pi = \lim_{n\to\infty} \frac{1}{n} \mathbb{E}_{e,d}\left[\sum_{i=1}^{n} (W_i - Z_i)^2 + \alpha X_i^2\right] \quad (7)$$

the stationary Markov policy given by

$$\begin{aligned} X_i &= \beta\left(W_i - \rho Z_{i-1}\right) && (8a) \\ Z_i &= \rho Z_{i-1} + \gamma Y_i && (8b) \end{aligned}$$

is optimal, where $\beta = \sqrt{\frac{L}{C}}, \gamma = \frac{\sqrt{LC}}{L+\sigma_v^2}$, and $C = \frac{\sigma_m^2}{1-\rho^2 \frac{\sigma_v^2}{L+\sigma_v^2}}$.

As such, this means that the feedback that is required to the user is simply given by Z_i, and a perceptual augmentation pertaining to ρ could be applied by simply visual warping. Afterwards, the user simply needs to specify an error signal. Although this might appear to be an obvious application of linear quadratic estimation/control, we point that this is indeed not so trivial because there are *multiple* decision makers who do not share the same information: Witsenhausen formulated a two-user problem with linear quadratic cost but non-classical information structure; nonlinear, non-stationary schemes can provide arbitrarily large improvement as compared to the best linear scheme [18].

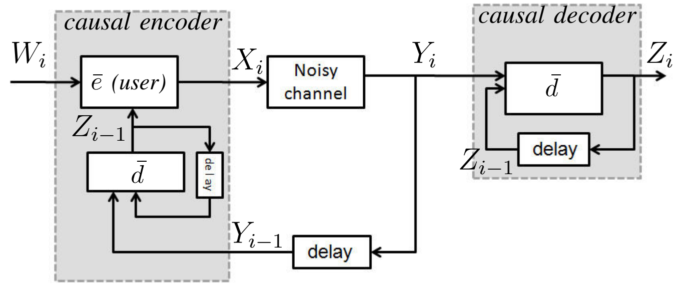

Fig. 3. A Stationary Markov coordination strategy. Note that in such settings, the only perceptual feedback required to deliver to the user at time i is Z_{i-1}.

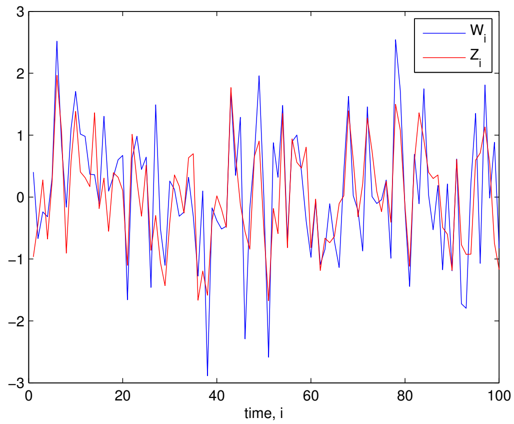

Fig. 4. Simulated performance of the BMI using the optimal scheme. W is the causally arriving process pertaining to user's intended trajectory. Z is the decoder's cursor/robotic-arm position.

Figure 4 shows a plot of a simulation pertaining to the causal operation of the BMI system, and Z is the estimate which would pertain to the position of the cursor or robotic arm. It can be shown that the ratio between this scheme's performance and the performance of the optimal Kalman filter applied to a forward hidden Markov model - that does not exploit feedback - can grow unbounded depending upon the parameters of the problem.

V. Conclusion

In this paper, we have presented a general-purpose design methodology for designing policies for interaction between the user and external device for BMIs. We did this by interpreting the BMI design problem as a team decision problem with non-classical information structure. In other words, we exploit how the essence of a BMI is that it comprises two agents (the user and the external device) cooperating to achieve a common goal. We exploited recent structural results that demonstrate the existence - for an arbitrary problem specification (e.g. cost function to minimize) - of optimal policies that are easily implementable in BMIs across many modalities - including EEG and cortically controlled devices. The structural result we provide sheds light on the minimal amount of useful information that is required to provide perceptual feedback to the user. In future work, we plan to instantiate the aforementioned linear quadratic example using an additive Gaussian noise EEG system using common spatial patterns [13] noisy channel model and also in a cortically controlled BMI with a point process noisy channel model [5]. Secondly, one key element missing in this work is how the statistics of the noisy channel might change due to learning. We plan to address the plasticity and learning questions, inherent in BMIs, by exploiting the relationship between reinforcement learning and dynamic programming to attempt to formulate provably learning-based BMIs.

References

[1] J. Carmena, M. Lebedev, R. Crist, J. ODoherty, D. Santucci, D. Dimitrov, P. Patil, C. Henriquez, and M. Nicolelis, "Learning to control a brain-machine interface for reaching and grasping by primates," *PLoS Biol*, vol. 1, no. 2, pp. 193–208, 2003.

[2] J. Wolpaw and D. McFarland, "Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans," *PNAS*, 2004.

[3] G. Santhanam, S. Ryu, B. Yu, A. Afshar, and K. Shenoy, "A high-performance brain–computer interface," *Nature*.

[4] N. Hatsopoulos and J. Donoghue, "The science of neural interface systems," *Annual review of neuroscience*, vol. 32, p. 249, 2009.

[5] L. Srinivasan, U. Eden, S. Mitter, and E. Brown, "General-purpose filter design for neural prosthetic devices," *Journal of Neurophysiology*.

[6] J. Cunningham, P. Nuyujukian, V. Gilja, C. Chestek, S. Ryu, and K. Shenoy, "A Closed-Loop Human Simulator for Investigating the Role of Feedback-Control in Brain-Machine Interfaces," *Journal of Neurophysiology*, 2010.

[7] M. Shanechi, G. Wornell, Z. Williams, and E. Brown, "A parallel point-process filter for estimation of goal-directed movements from neural signals," in *IEEE ICASSP*.

[8] S. Gorantla and T. P. Coleman, "Information-theoretic viewpoints on optimal causal coding-decoding problems," *IEEE Transactions on Information Theory*, submitted January 2011.

[9] Y. Ho, M. Kastner, and E. Wong, "Teams, signaling, and information theory," *Automatic Control, IEEE Transactions on*, vol. 23, no. 2.

[10] H. Witsenhausen, "Informational aspects of stochastic control," in *Proc. the Oxford Conference on Stochastic Optimization*, 1978.

[11] Y. C. Ho, "Team Decision Theory and Information Structures," *Proceedings of the IEEE*.

[12] C. Papadimitriou and J. Tsitsiklis, "Intractable problems in control theory," *SIAM Journal on Control and Optimization*, vol. 24, no. 4, pp. 639–654, 1986.

[13] M. S. McCormick, R. Ma, and T. P. Coleman, "An Analytic Spatial Filter and Hidden Markov Model for Enhanced Information Transfer Rate in EEG-Based Brain Computer Interfaces," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Dallas, TX, 2010.

[14] O. Shayevitz and M. Feder, "Optimal Feedback Communication via Posterior Matching," *IEEE Transactions on Information Theory*, March 2011.

[15] T. P. Coleman, "A Stochastic Control Approach to 'Posterior Matching'-style Feedback Communication Schemes," in *IEEE International Symposium on Information Theory*, Seoul, Korea, July 2009.

[16] 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," *Int'l Journal on Human-Computer Interaction*, January 2011, special issue on Current Trends in Brain-Computer Interface (BCI) Research.

[17] A. Akce, M. Johnson, and T. Bretl, "Remote Teleoperation of an Unmanned Aircraft with a Brain-Machine Interface: Theory and Preliminary Results," in *Robot. and Autom. Proceedings. IEEE Int. Conf. on*, 2010.

[18] H. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, p. 131, 1968.