POSTER PRESENTATION

High-performance continuous neural cursor control enabled by a feedback control perspective

Vikash Gilja[1]*, Paul Nuyujukian[1], Cindy Chestek[1], John Cunningham[1], Byron Yu[1], Stephen Ryu[1] and Krishna Shenoy[1]

1 Stanford University, Stanford, USA

Neural prostheses, or brain-computer interfaces (BCIs), have the potential to substantially increase quality of life for people suffering from motor disorders, including paralysis and amputation. These systems translate recorded neural signals into control signals that guide a paralyzed arm, artificial limb, or computer cursor. Although current laboratory demonstrations provide a compelling proof-of-concept, the field must continue to increase performance to achieve clinical viability. Many BCIs use activity from motor and/or premotor cortex to achieve continuous control. These BCIs can be viewed from a feedback control perspective, as the motor field has done for the native limb: the brain is the controller of a new plant, defined by the BCI. This perspective leads us to two advances that result in significant qualitative and quantitative performance improvements. We tested these advances in closed loop with one rhesus macaque trained in a virtual 3D workspace. On each trial he used a cursor, controlled by the native contralateral limb or a BCI, to acquire a target on a 2D plane within an allotted time period. Neural data were recorded from a 96-electrode array (Blackrock) implanted spanning PMd and M1. Our designs are informed by a feedback model, which assumes the user develops a volitional control signal to achieve a goal given the current state of the world. This signal and task-unconstrained signals (such as sensory feedback, attention) give rise to neural firing, which we record. Finally, the decoding algorithm estimates desired cursor movements from the neural firing, and updates the workspace. By applying the assumptions of this simple feedback model, we augment a basic position/velocity Kalman filter. We consider the position/velocity Kalman filter to represent "baseline" as it meshes with the performance of and is algorithmically similar to methods common in the literature (e.g., Kim et al., 2008). All experiments used spike counts generated by a threshold detector without spike sorting. Such a system has clinical appeal, particularly for arrays with potentially decreased SNR (these experiments were 22-24 months post implantation). Design iterations were tested within the same experimental session using a blocked "ABA" design. Through this design process, we made two advances that substantially improve performance. First, using a standard Kalman filter, we fit neural data to a guess of the desired volitional control signal, instead of observed or instructed kinematics. Second, we developed a modified velocity-only Kalman filter, whose observation model incorporates cursor position as feedback. The new BCI appears more controllable and produces straighter reaches and crisper stops. Compared to the standard Kalman BCI, mean time to target is reduced by nearly a factor of two.

This system can run freely for hundreds to thousands of trials, making point-to-point reaches to targets randomly placed across the workspace. These feedback-perspective based algorithmic innovations, together with initial experimental verification, suggest that approximately a factor of two performance advance is possible, thereby increasing clinical viability.

Support: NSF, NDSEG, Stanford Med Scholars, Soros Fndn, HHMI, SGF, JHU APL under DARPA RP2009: N66001-06-C-8005, CDRF, BWF, NIH-NINDS-CRCNS-RO1, NIH Pioneer Award 1DP1OD006409

Conference: Computational and Systems Neuroscience 2010. Salt Lake City, UT, USA, February 25 - March 02, 2010.

Citation: Gilja V, Nuyujukian P, Chestek C, Cunningham J, Yu B, Ryu S and Shenoy K (2010). High-performance continuous neural cursor control enabled by a feedback control perspective. Frontiers in Neuroscience. Conference Abstract: Computational and Systems Neuroscience 2010. doi: 10.3389/conf.fnins.2010.03.00249

***Correspondence:** Vikash Gilja, Stanford University, Stanford, USA. gilja@stanford.edu**

I. Feedback Model for continuous BCIs

The full feedback model is shown in fig. 1C. In this model, the neural prosthetic user has a high level task goal, g, such as positioning the prosthetic in a specific region of the workspace. The user develops a volitional control signal x_{V} (e.g. velocity) by assessing the current state of the workspace, w, and applying a strategy, V, towards achieving g. There are also behavioral and cognitive signals that are not directly controlled in the task (e.g. eye position, level of attention), which we will operationally refer to as task-unconstrained signals. Task-unconstrained signals may or may not be effected by w. The volitional control signal, x_{V}, and task-unconstrained signals, $x_{\tilde{\mathrm{V}}}$, give rise to y, the neural firing recorded by the prosthetic, through M. The decoder, D, attempts to recover x_v from y, yielding an estimate of the intended volitional control signal, $\hat{x}_{\mathrm{V}}$. Finally, $\hat{x}_{\mathrm{V}}$ is used by S to update w, for example the position of the neural prosthetic in the workspace may be changed.

Progression of Conceptional Frameworks to Guide BCI Algorithm Design and Training Methods

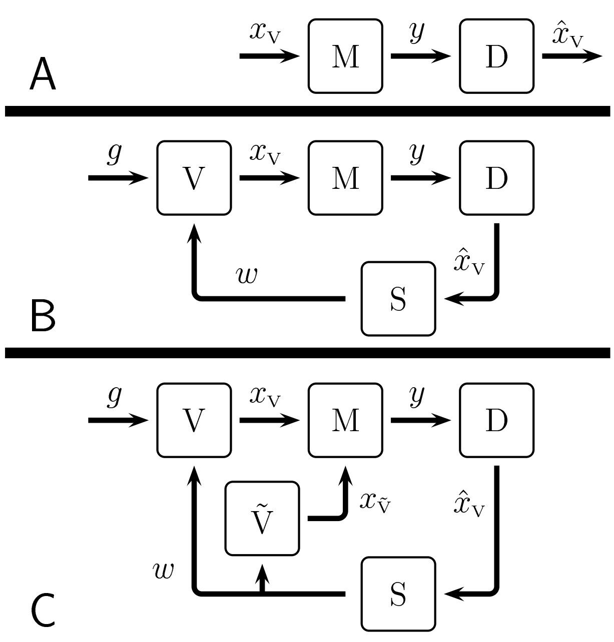

Fig. 1. (A) Conceptual framework for models trained with offline data. From this starting point we broaden the framework; (B) Advance 1 adds a feedback loop by guessing the user's mapping, V, from goal, g, to desired control signal, x_{V}; (C) Advance 2 includes task-unconstrained signals (or more precisely: task irrelevant behavioral and cognitive signals), $x_{\tilde{\mathrm{V}}}$.

A. Advance 1: Decoding the User's Intent

Many existing proof-of-concept algorithms are initially designed, tested, and fit offline, using data collected without the neural prosthetic in the loop [1]–[4]. The data is fit against real, observed, or imagined arm movements. At a high level, the design rationale is summarized by fig. 1A. Although the design phase is open loop, as soon as the system is used online as a neural prosthetic, the loop is closed via real-time visual feedback.

The system can also be fit online, while the user is getting real-time feedback, as in fig. 1B. Such a strategy is employed in [5]–[7] by iteratively refining decoder parameter settings during online control. However, these models assume the desired volitional control can be directly inferred from the kinematics of the neurally controlled prosthetic. We propose a different method, motivated by the model in fig. 1B, in which we guess the user's strategy, V.

Initially, we train a position/velocity Kalman filter with data collected during reaches with the native limb; we refer to this as the "baseline BCI (pos/vel)," as it meshes with existing literature. Next, the user is placed in online brain control with the baseline BCI. Data from this online control session is used to guess the desired control signal at every time step during the session. We assume that cursor position is internalized via visual feedback. Given the task definition, it is likely that the user wishes to develop a velocity oriented towards the goal. Thus, we guess the intended velocity by rotating the observed velocity of the neural cursor towards the goal. This guess is used to fit a new position/velocity Kalman filter, which we refer to as "BCI with advance 1." In the data shown below, the user had 10 seconds to complete reaches; he succeeded on nearly every trial (>98%) for both models, but there was a substantial difference in the time required for successful reaches . With the BCI with advance 1, movements appear more controlled and consistent. As shown in fig. 2A, the mean time to target makes a statistically significant ($p < 0.001$) drop from 1994 ms to 1138 ms.

B. Advance 2: Incorporating Task-unconstrained Signals

The baseline BCI and the BCI with advance 1 treat both position and velocity as volitional signals with uncertainty. Uncertainty in position can result in additional jitter in the decoded cursor position over time, as neural noise can pass through to the decoded position. However, in a closed loop task, the position is fed back to the user via natural vision. If the user internalizes the presented position, this uncertainty is perfectly eliminated.

One possible solution is to use a velocity-only Kalman filter, which we refer to as "baseline BCI (vel only)." The resulting decodes appear smoother, but we note that the cursor takes on a velocity bias at certain positions in the workspace, causing the user to get stuck in specific positions. Colloquially, it is as if a "force field" develops. Also, when the user disengages from the task, the cursor has a tendency to fly off in a single direction. Cursor position may still effect observed neural firing, so we still need to account for position.

* indicates equal contribution

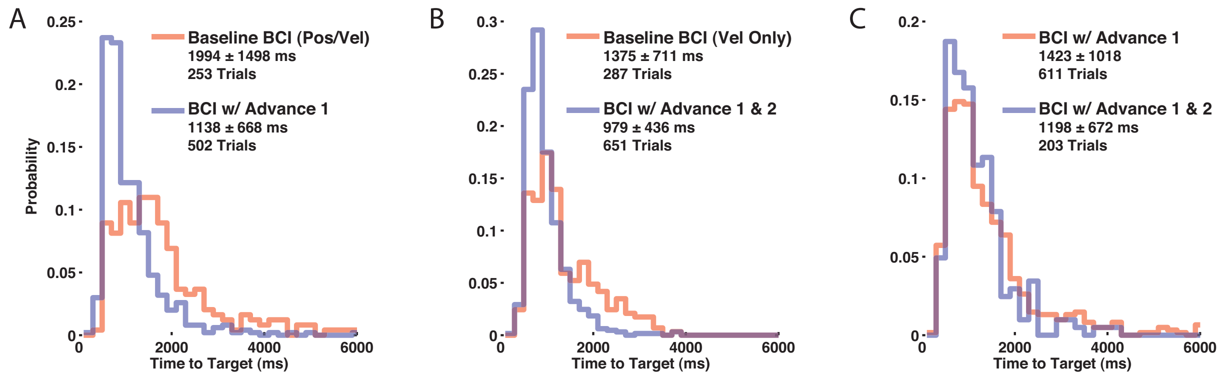

Fig. 2. Distribution of mean time to target within a single experimental session. Each plot is a comparison between two different decoders made in the same experimental session. The legend text indicates the number of trials plotted and the mean ± std. dev. of the time to target. The pairwise comparison within each plot is statistically significant with $p < 0.01$ by parametric T-test and $p < 0.03$ by nonparametric rank sum test

In order to counteract this positional bias, we introduce the cursor position as a noiseless task-unconstrained signal in the feedback model presented in fig. 1C. This results in a velocity-only Kalman filter modified to incorporate cursor position feedback; we term this "BCI with advance 1 & 2". Qualitatively, this BCI has the smoothness of the velocity-only baseline BCI, while remaining bound in the workspace without the formation of "force fields." For the experimental session comparing the velocity-only baseline BCI to the BCI with advance 1 & 2, the user had 4 seconds to complete the reach and success rates were 89.1% and 98.6%, respectively. With the BCI with advance 1 & 2, as shown in fig. 2B, the mean time to target makes a statistically significant ($p < 0.001$) drop from 1375 ms to 979 ms.

For the experimental session comparing the BCI with advance 1 to the BCI with advance 1 & 2, the user had 10 seconds to complete the reach, success rates were $> 99.8\%$ for both filters. Again, the observed movements appear much more controlled and consistent. As shown in fig. 2C the mean time to target make a statistically significant ($p < 0.03$) drop from 1423 ms to 1198 ms.

II. Methods

All experiments were conducted with an adult male rhesus macaque implanted with a 96-electrode array (BlackRock Microsystems Inc.) using standard neurosurgical techniques 22 months prior to the current study. The electrode array was implanted in a region spanning the arm representation of the dorsal aspect of premotor cortex (PMd) and primary motor cortex (M1), as estimated visually from local anatomical landmarks. All of the surgical procedures were approved by the Stanford University Institutional Animal Care and Use Committee (IACUC).

The macaque was trained to make point-to-point reaches in a 2D plane with a virtual cursor controlled by movements of the contralateral arm or by the output of a neural decoder. The virtual cursor and targets were presented in an immersive 3D environment; for more details on the environment, see Cunningham et al. (2010) COSYNE. Acquired targets must be held for 500 ms. In all results shown, he is performing a center-out-and-back task with eight radial 4 cm targets uniformly placed 8 cm from center. Trial times and success rates are calculated from center out trials only. All results generalized to a pinball task, where target positions were randomly selected from a uniform distribution spanning a 16 cm by 16 cm area of the workspace.

All neural decoders were tested with a bin width of 100 ms. Spike counts for each bin were collected by applying a simple negative threshold, set to 4.5 x root mean square of the spike band of each neural channel.

References

[1] J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. Nicolelis *PLoS Biol*, vol. 1, November 2003.
[2] W. Wu, M. J. Black, D. Mumford, Y. Gao, E. Bienenstock, and J. P. Donoghue *IEEE transactions on bio-medical engineering*, vol. 51, pp. 933–942, June 2004.
[3] L. Hochberg, M. Serruya, G. Friehs, J. Mukand, M. Saleh, A. Caplan, A. Branner, D. Chen, R. Penn, and J. Donoghue *nature*, pp. 164–171, July 2006.
[4] K. Ganguly and J. M. Carmena *PLoS Biol*, vol. 7, pp. e1000153+, July 2009.
[5] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz *Nature*, vol. 453, pp. 1098–1101, May 2008.
[6] S. Kim, J. Simeral, L. Hochberg, J. Donoghue, and M. Black, "Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia.," *Journal of neural engineering*, vol. 5, pp. 455–476, December 2008.
[7] G. Fraser, S. Chase, A. Whitford, and A. Schwartz *J. Neural Eng.*, vol. 6, pp. 055004+, October 2009.

* indicates equal contribution