

A Game-Theoretic Model for Co-Adaptive Brain-Machine Interfaces

Maneeshika M. Madduri, Samuel A. Burden and Amy L. Orsborn

Abstract—Co-adaptation in brain-machine interfaces (BMIs) can improve performance and facilitate user learning. We propose and analyze a mathematical model for co-adaptation in BMIs. We model the brain and the decoder as strategic agents who seek to minimize their individual cost functions, leading to a game-theoretic formulation of interaction. We frame our BMI model as a potential game to identify stationary points (Nash equilibria) of the brain-decoder interactions, which correspond to points at which both the brain and the decoder stop adapting. Assuming the brain and the decoder adapt using gradient-based schemes, we analytically show how convergence to these equilibria depends on agent learning rates. This theoretical framework presents a basis for simulating co-adaptation using dynamic game theory and can be extended to tasks with multiple dimensions and to different decoder models. This framework can ultimately be used to inform adaptive decoder design to shape brain learning and optimize BMI performance.

I. INTRODUCTION

A. Motivating Background

Uncovering the neural principles of learning in the brain is key to developing high-performing neural interfaces [1], [2]. Closed-loop Brain-Machine Interfaces (BMIs) define an input-output relationship in the brain and can provide insight into fundamental neural mechanisms underlying sensorimotor learning [3]. Motor BMIs use a *decoder* to convert neural data to a control signal, such as moving a computer cursor. Closed-loop BMIs form a closed control loop via visual feedback, creating novel sensorimotor systems that can be manipulated to study neural functions [4].

Learning in BMI can lead to the formation of a *stable cortical map*, similar to natural motor learning [5]. Stable cortical maps are accompanied by stable BMI task performance, such as a consistent cursor trajectory to the target with low variability and other signatures of skillful motor control. Recent research has shown that *closed-loop decoder adaptation* (CLDA)—algorithms to modify decoder parameters during closed-loop BMI operation—are potentially able to guide user learning to facilitate the emergence of stable cortical maps [6], [7]. CLDA algorithms have been designed to improve BMI performance and decrease variability in decoder estimates [8], [9]. Designing CLDA algorithms that

This work has been supported by the UW College of Engineering Dean’s Fellowship, the NDSEG Fellowship, Award #1836819 from the U.S. National Science Foundation, and Award Number K12HD073945 from the Eunice Kennedy Shiver National Institute of Child Health Human Development of the National Institutes of Health.

M. M. Madduri and S. A. Burden are with the Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98105 USA. mmadduri@uw.edu

A.L. Orsborn is with the Department of Electrical and Computer Engineering and Department of Bioengineering, University of Washington, and the Washington National Primate Research Center, Seattle, WA 98105 USA

can optimize brain learning is an open challenge. As an inherently real-time process, CLDA requires online BMI validation. However, testing a wide range of CLDA algorithms through experimentation is impractical. To constrain the space of testable CLDA design parameters and reduce experimental load, we propose and analyze a mathematical model for closed-loop co-adaptive BMIs.

B. Related Work

Existing BMI models draw upon control theory to conceptualize the brain and decoder as systems adapting according to specified learning processes and subject to cost functions [10]–[12]. Two of these models [11], [12] identify co-adaptation strategies using optimization principles to propose adaptive decoder paradigms. The brain and the decoder in these models are treated as two distinct stochastic optimal control systems [11] or as two systems optimizing a joint cost function [12]. The learning dynamics of co-adaptation in these models either propose a stochastic coordinate descent, where the decoder learns to anticipate the brain’s adaptation [11] or a stochastic gradient descent with an online least-squares estimator, where the brain and the decoder estimate each other’s parameters to form internal approximations of the other learner [12]. The proposed optimal brain-decoder solutions for each model are identified through either mathematical analysis [11] or computational simulations [12].

C. Our Contributions

The purpose of this work is to model the interaction between brain and decoder in a co-adaptive BMI. Our contributions to the existing literature are: (a) formulating the two-learner problem for motor BMI in a game-theoretic framework; (b) deriving *stationary points* (Nash equilibria) of the game; and (c) analyzing convergence of a model for co-adaptation to these stationary points. In particular, we adopt the perspective that the brain and decoder are two agents playing a *dynamic game* [13], [14], and apply techniques from game theory to study the asymptotic dynamics of co-adaptation. We use a *potential game* formulation to derive stationary points (Nash equilibria), which points are our model’s surrogates for *stable cortical maps* observed in experiments. Assuming the brain and decoder adapt using (stochastic) gradient descent, we assess convergence to stationary points. Finally, we corroborate our mathematical analyses using numerical simulations.

II. METHODS

A. Mathematical Model for Motor BMI

Fig. 1 displays each component and corresponding symbolic representation. Signals are denoted with lower-case

non-bolded symbols (e.g. x) and parameterizations of the transformations are **bolded** and capitalized (e.g. \mathbf{X}). The BMI session is a series of trials, with each trial being represented as a new target presented to the user. Time is indexed by the subscript trial number $t = 1, 2, \dots$, dimension of the BMI task is indexed by the subscript $d = 1, 2, \dots$ and neurons are indexed in a population of size N by the superscript $n \in \{1, \dots, N\}$. We will suppress the sub- and/or superscripts when they are clear from context.

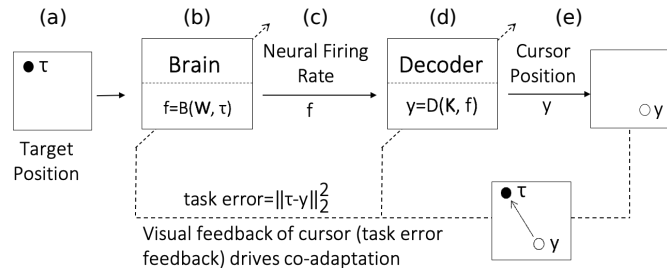


Fig. 1. Diagram of the BMI model showing the forward model of the brain and the decoder as well as the visual feedback error of the cursor (task error feedback) that drives the co-adaptation of the brain and the decoder [15]. Based on prior BMI models [10], [11], the elements of this model are: (a) the target presented to the user, (b) the user, (c) the firing rates (recorded as time-binned intracortical neural spiking activity) from a population of neurons, (d) the decoder, (e) the cursor position at the output of the decoder.

Target Position: The task presented in this BMI model is a center-out cursor control task—e.g. [6], [7]. The target position is at τ .

Brain and Firing Rates: The brain is represented as a linear transformation of the target τ by a matrix representing the brain’s parameters \mathbf{W} (1) [10], [12].

$$f_t = B(\mathbf{W}_t, \tau_t) = \mathbf{W}_t \tau_t \quad (1)$$

where $f_t \in \mathbb{R}^{n \times 1}$, $\mathbf{W}_t \in \mathbb{R}^{n \times d}$, $\tau_t \in \mathbb{R}^{d \times 1}$. At each new trial t , the firing rates are updated and are inputs to the decoder (2).

Decoder and Cursor Position: The decoder algorithm in this model is represented as a Wiener filter, a commonly-used experimental decoder that has also been implemented in prior models—e.g. [10], [12]. The Wiener filter was used in this model for simplicity of analysis, but the Kalman Filter, which is a state-based estimator, has been used in co-adaptive BMI experiments—e.g. [6], [7]. The decoder transforms firing rates, f , to a d -dimensional cursor position, y , by multiplying the firing rates by the decoder’s parameters, \mathbf{K} (2).

$$y_t = D(\mathbf{K}_t, f_t) = \mathbf{K}_t f_t \quad (2)$$

where $f_t \in \mathbb{R}^{n \times 1}$, $\mathbf{K}_t \in \mathbb{R}^{d \times n}$, $y_t \in \mathbb{R}^{d \times 1}$.

B. Brain and Decoder Cost Functions

For each trial, the task error (3) is calculated as the quadratic loss of the cursor and target positions [10], [12]:

$$\varepsilon_t(\tau_t, y_t) = \|\tau_t - y_t\|_2^2. \quad (3)$$

The brain and the decoder parameters, \mathbf{W} and \mathbf{K} , update at each trial iteration using a gradient-based scheme detailed

in II-C. From a game-theoretic lens, this model can treat the brain and the decoder as two agents in a 2-player game. Each agent, the brain or the decoder, is trying to minimize its own cost function through co-adaptation. The cost function of the brain and decoder are:

$$\begin{aligned} c_B(\mathbf{W}, \mathbf{K}) &= \varepsilon(\tau, y) + \lambda_{\mathbf{W}} \|\mathbf{W}\|_2^2 \\ c_D(\mathbf{W}, \mathbf{K}) &= \varepsilon(\tau, y) + \lambda_{\mathbf{K}} \|\mathbf{K}\|_2^2. \end{aligned} \quad (4)$$

The brain and the decoder cost functions both include regularization terms: $\lambda_{\mathbf{W}} \|\mathbf{W}\|_2^2$ for the brain and $\lambda_{\mathbf{K}} \|\mathbf{K}\|_2^2$ for the decoder. The parameters $\lambda_{\mathbf{W}}$ and $\lambda_{\mathbf{K}}$ are scaling factors on the regularization terms and, as we will show later in this paper, play crucial roles in ensuring the convergence of the cost functions. These regularization terms serve to penalize the effort of the brain and the decoder.

C. Game-Theoretic Stationary Points

1) Defining a Potential Function: With the mathematical model described in II-A, we can reformulate the brain-decoder interactions as a *potential game*. In a potential game, the action of all agents can be expressed with a single function called the potential function. If there exists a continuously differentiable function $\phi : \mathbb{R}^{n \times d} \mapsto \mathbb{R}$ such that $\frac{\partial \phi}{\partial \mathbf{W}} = \frac{\partial c_B}{\partial \mathbf{W}}$ and $\frac{\partial \phi}{\partial \mathbf{K}} = \frac{\partial c_D}{\partial \mathbf{K}}$, this function ϕ can be considered to be a potential function [16]. This highlights an important assumption: the cost functions of the brain and the decoder are differentiable. The brain and the decoder cost functions satisfy the conditions of a potential game and can be thus characterized with the following potential function:

$$\phi(\mathbf{W}, \mathbf{K}) = \varepsilon(\tau, y) + \lambda_{\mathbf{W}} \|\mathbf{W}\|_2^2 + \lambda_{\mathbf{K}} \|\mathbf{K}\|_2^2. \quad (5)$$

This potential function matches the two-learner system cost function proposed by a prior model [12]. The potential function ϕ tracks the changes of both players, the brain and the decoder in this case, and reformulates co-adaptation as a single-objective optimization problem.

2) Stationary Points of Potential Game: The Nash equilibria of potential games, referred to as *stationary points* in this paper, are the directionally-local minima of the potential function [13, Prop. 12.2]. The stationary points \mathbf{W}^* , \mathbf{K}^* are determined by solving $\frac{\partial \phi}{\partial \mathbf{W}} \Big|_{\mathbf{W}=\mathbf{W}^*} = 0$ and $\frac{\partial \phi}{\partial \mathbf{K}} \Big|_{\mathbf{K}=\mathbf{K}^*} = 0$. These are points where the brain and the decoder reach steady-state, or are no longer adapting, in the BMI model.

D. Gradient-Based Co-Adaptation

The brain and the decoder parameters, \mathbf{W} and \mathbf{K} , are updated according to the gradient-based adaptation scheme:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \gamma_{\mathbf{W}} \Delta_{\mathbf{W}_t}, \quad \mathbf{K}_{t+1} = \mathbf{K}_t - \gamma_{\mathbf{K}} \Delta_{\mathbf{K}_t} \quad (6)$$

where $\gamma_{\mathbf{W}_t}, \gamma_{\mathbf{K}_t} > 0$ denote learning rates and $\Delta_{\mathbf{W}_t}, \Delta_{\mathbf{K}_t}$ denote (approximations of) gradients of the agents’ cost functions with respect to their parameters. Note that (6) models adaptation from trial-to-trial; no longer- or shorter-term learning processes are represented. We analyze convergence of (6) to stationary points in two different regimes: 1) *deterministic* gradients; and 2) *stochastic* gradients.

1) *Deterministic gradient descent*: If learning rates are constant and agents adapt using their exact gradients,

$$\Delta_{\mathbf{W}_t} = \left. \frac{\partial \phi}{\partial \mathbf{W}} \right|_{\mathbf{W}=\mathbf{W}_t}, \Delta_{\mathbf{K}_t} = \left. \frac{\partial \phi}{\partial \mathbf{W}} \right|_{\mathbf{K}=\mathbf{K}_t} \quad (7)$$

then (6) defines a deterministic dynamical system. To assess convergence, we linearize this system about the stationary points from the preceding section and evaluate eigenvalues of the Jacobian matrix. Since the system is discrete-time, the magnitude of these eigenvalues determines how quickly the co-adaptation dynamics converge to stationarity.

2) *Stochastic gradient descent*: As noted in [12], evaluating the exact gradients in (7) requires each agent to know the other's parameters. As an alternative, we consider the dynamics obtained when agents use a *stochastic approximation* of their gradient. One such approximation is [17, Eq. 31].

$$\begin{aligned} \Delta_{\mathbf{W}_t} &= \frac{1}{\ell \nu_{\mathbf{W}}} \sum_{i=1}^{\ell} (c_B(\mathbf{W}_t + \nu_{\mathbf{W}} \delta_{\mathbf{W}_t}) - c_B(\mathbf{W}_t)) \cdot M_{\mathbf{W}}^{-1} \delta_{\mathbf{W}_t} \\ \Delta_{\mathbf{K}_t} &= \frac{1}{\ell \nu_{\mathbf{K}}} \sum_{i=1}^{\ell} (c_D(\mathbf{K}_t + \nu_{\mathbf{K}} \delta_{\mathbf{K}_t}) - c_D(\mathbf{K}_t)) \cdot M_{\mathbf{K}}^{-1} \delta_{\mathbf{K}_t} \end{aligned} \quad (8)$$

where $\delta_{\mathbf{W}_t} \sim \mathcal{N}(0, M_{\mathbf{W}})$, $\delta_{\mathbf{K}_t} \sim \mathcal{N}(0, M_{\mathbf{K}})$ are zero-mean normal random vectors with covariances $M_{\mathbf{W}}$, $M_{\mathbf{K}}$, respectively. $\nu_{\mathbf{W}}$ and $\nu_{\mathbf{K}}$ are the difference parameters that scale the contribution of $\delta_{\mathbf{W}}$ and $\delta_{\mathbf{K}}$, respectively. Note that (8) defines unbiased estimates for the exact gradients in (7) that each agent can evaluate using its own cost and parameters without requiring it know the cost or parameters of the other.

III. RESULTS

Although the methods described in the preceding section apply to task and neuron spaces of arbitrary dimension, in the present paper we only report results in the special case where the task is 1-dimensional ($d = 1$) and there is a single neuron ($n = 1$); generalization of these results to arbitrary dimensions will be the subject of a future study.

1) *Stationary Points of the Potential Game*: When the same value $\lambda_{\mathbf{W}} = \lambda_{\mathbf{K}} = \lambda \in (0, \tau^2)$ is chosen for the regularization parameters in the brain and decoder cost functions (4), the potential function (5) has two stationary points (the Nash equilibria of the potential game),

$$\mathbf{W}^* = \mathbf{K}^* = +\sqrt{1 - \frac{\lambda}{\tau^2}} \text{ and } \mathbf{W}^* = \mathbf{K}^* = -\sqrt{1 - \frac{\lambda}{\tau^2}}, \quad (9)$$

obtained by applying the procedure described in Sec. II-C.2. (When $\lambda > \tau^2$ the only stationary point is at the origin $\mathbf{W}^* = \mathbf{K}^* = 0$; when $\lambda < 0$ there are no stationary points.) Substituting the stationary point (9) into the cost functions (4) yields stationary costs

$$c_B(\mathbf{W}^*, \mathbf{K}^*) = c_D(\mathbf{W}^*, \mathbf{K}^*) = \lambda \quad (10)$$

and stationary task error

$$\varepsilon(\mathbf{W}^*, \mathbf{K}^*) = \frac{\lambda^2}{\tau^2}. \quad (11)$$

Note that, as the penalty parameter λ goes to zero, the stationary task error goes to zero at a quadratic rate.

2) *Convergence of Deterministic Gradient Descent*:

When the same value $\lambda_{\mathbf{W}} = \lambda_{\mathbf{K}} = \frac{1}{2}$ is chosen for the regularization parameters in the brain and decoder cost functions (4) and the task space is normalized so that $\tau^2 = 1$, the eigenvalues of the linearization of (6) at either of the stationary points (9) are

$$\{1 - 2\gamma_{\mathbf{W}}, 1 - 2\gamma_{\mathbf{K}}\}. \quad (12)$$

The co-adaptation dynamics converge if the eigenvalues lie inside the unit circle, which occurs precisely when

$$0 < \gamma_{\mathbf{W}} < 1, \quad 0 < \gamma_{\mathbf{K}} < 1. \quad (13)$$

Since the eigenvalue with the maximum magnitude determines the asymptotic convergence rate, the learning rates $\gamma_{\mathbf{W}} = \gamma_{\mathbf{K}} = \frac{1}{2}$ yield the fastest convergence rate.

3) *Convergence of Stochastic Gradient Descent*: To corroborate analytical results reported in the preceding sections, we ran batches of numerical simulations of the co-adaptation dynamics in (6) using the stochastic gradient estimate in (8). We found that these simulations converged to the stationary points in (9) as predicted (Fig. 2).

IV. DISCUSSION

A. Interpretation of Results

1) *Stationary points*: The formulation of brain-decoder co-adaptation in a game-theoretic framework naturally lead to the consideration of stationary points (Nash equilibria). Our analysis focused on how the penalty parameters $\lambda_{\mathbf{W}}, \lambda_{\mathbf{K}}$ in (4) affect these stationary points. Focusing on the case where $\lambda_{\mathbf{W}} = \lambda_{\mathbf{K}} = \lambda \in (0, \tau^2)$, we found that $\lambda \in (0, \tau^2)$ is necessary for the existence of non-trivial stationary points: no points are stationary when $\lambda < 0$ and only the origin is stationary when $\lambda > \tau^2$. Since these stationary points are our model's surrogates for the *stable cortical maps* that experiments suggest are highly desirable for skillful control, this finding indicates that penalty parameters play a critical role in ensuring such a stable outcome can be achieved.

2) *Convergence of co-adaptation*: Assuming both the brain and decoder adapt their parameters using a gradient-based scheme, we analyzed convergence of co-adaptation to stationarity. Our analysis showed that the learning rates $\gamma_{\mathbf{W}}, \gamma_{\mathbf{K}}$ play a critical role in convergence. Focusing on the case where $\lambda_{\mathbf{W}} = \lambda_{\mathbf{K}} = \frac{1}{2}$ and $\tau^2 = 1$, we found that convergence is guaranteed so long as

$$\rho = \max\{|1 - 2\gamma_{\mathbf{W}}|, |1 - 2\gamma_{\mathbf{K}}|\} < 1, \quad (14)$$

in which case ρ determines the rate of convergence since $\mathbf{W}_t \rightarrow \mathbf{W}^*$ and $\mathbf{K}_t \rightarrow \mathbf{K}^*$ exponentially in time/trial t with rate $\rho < 1$:

$$|\mathbf{W}_t - \mathbf{W}^*| \leq C_{\mathbf{W}} \rho^t, \quad |\mathbf{K}_t - \mathbf{K}^*| \leq C_{\mathbf{K}} \rho^t, \quad (15)$$

for some fixed constants $C_{\mathbf{W}}, C_{\mathbf{K}} > 0$. Our findings demonstrate that the decoder's learning rate must be chosen carefully, as learning too fast may lead to an unstable outcome, whereas learning too slow may limit the rate of convergence to stationarity. These analytical results are corroborated by previous simulation and experimental findings [6], [12].




Fig. 2. (left) Gradient field of potential function (5) with penalties $\lambda_W = \lambda_K = 0.5$ in (4) and normalized task ($\tau^2 = 1$) as in Sec. II-C.2. Blue and orange curves are *nullclines* (points where $\frac{\partial \phi}{\partial K} = 0$ and $\frac{\partial \phi}{\partial W} = 0$, respectively) that intersect at stationary points. (middle) Numerical simulations of co-adaptation (6) with stochastic gradients (8) demonstrate exponential rate of convergence to stationarity ($\log |K_t - K^*|, \log |W_t - W^*|$ converge to zero at linear rate); the difference in convergence rate is caused by a difference in learning rates ($\gamma_W = 5 \times 10^{-3}$, $\gamma_K = 2.5 \times 10^{-3}$). (right) Stationary cost (10) and task error (11) determined analytically (lines) and computed from batches of numerical simulations (markers) for different values of penalty parameter $\lambda_W = \lambda_K = \lambda$.

B. Future Directions

In this paper, we have detailed an initial theoretical framework that can be used to inform adaptive decoder design for shaping brain learning and optimizing BMI performance. The dynamic game formulation of a co-adaptive BMI can be extended to BMI models and experiments where the decoder and brain may have separate objectives, which may be necessary for the decoder to guide the brain’s learning. Our stochastic simulations confirmed the deterministic analysis and will allow us to evaluate co-adaptive BMI scenarios that are less mathematically tractable. Example of such mathematically difficult scenarios include: (a) non-stationary costs, (b) learning rates that are dependent on time or task, (c) nonlinear thresholds for firing rates. Next steps to guide the BMI model presented here will be a comparison to experimental data of co-adaptation [7].

REFERENCES

- [1] K. V. Shenoy and J. M. Carmena, “Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces,” *Neuron*, vol. 84, no. 4, pp. 665–680, Nov. 2014, doi: 10.1016/j.neuron.2014.08.038.
- [2] C. T. Moritz, “Now is the Critical Time for Engineered Neuroplasticity,” *Neurotherapeutics*, vol. 15, no. 3, pp. 628–634, Jul. 2018, doi: 10.1007/s13311-018-0637-0.
- [3] J. M. Carmena, “Advances in Neuroprosthetic Learning and Control,” *PLoS Biol*, vol. 11, no. 5, p. e1001561, May 2013, doi: 10.1371/journal.pbio.1001561.
- [4] A. L. Orsborn and B. Pesaran, “Parsing learning in networks using brain-machine interfaces,” *Current Opinion in Neurobiology*, vol. 46, pp. 76–83, Oct. 2017, doi: 10.1016/j.conb.2017.08.002.
- [5] K. Ganguly and J. M. Carmena, “Emergence of a Stable Cortical Map for Neuroprosthetic Control,” *PLOS Biology*, vol. 7, no. 7, p. e1000153, Jul. 2009, doi: 10.1371/journal.pbio.1000153.
- [6] A. L. Orsborn, S. Dangi, H. G. Moorman, and J. M. Carmena, “Closed-Loop Decoder Adaptation on Intermediate Time-Scales Facilitates Rapid BMI Performance Improvements Independent of Decoder Initialization Conditions,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 20, no. 4, pp. 468–477, Jul. 2012, doi: 10.1109/TNSRE.2012.2185066.
- [7] A. L. Orsborn, H. G. Moorman, S. A. Overduin, M. M. Shanechi, D. F. Dimitrov, and J. M. Carmena, “Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control,” *Neuron*, vol. 82, no. 6, pp. 1380–1393, Jun. 2014, doi: 10.1016/j.neuron.2014.04.048.
- [8] S. Dangi, A. L. Orsborn, H. G. Moorman, and J. M. Carmena, “Design and Analysis of Closed-Loop Decoder Adaptation Algorithms for Brain-Machine Interfaces,” *Neural Computation*, vol. 25, no. 7, pp. 1693–1731, Jul. 2013, doi: 10.1162/NECOa00460.
- [9] H.-L. Hsieh and M. M. Shanechi, “Optimizing the learning rate for adaptive estimation of neural encoding models,” *PLoS Comput Biol*, vol. 14, no. 5, May 2018, doi: 10.1371/journal.pcbi.1006168.
- [10] Héliot, K. Ganguly, J. Jimenez, and J. M. Carmena, “Learning in Closed-Loop Brain-Machine Interfaces: Modeling and Experimental Validation,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 40, no. 5, pp. 1387–1397, Oct. 2010, doi: 10.1109/TSMCB.2009.2036931.
- [11] J. S. Merel, R. Fox, T. Jebara, and L. Paninski, “A multi-agent control framework for co-adaptation in brain-computer interfaces,” in *Advances in Neural Information Processing Systems 26*, 2013, pp. 2841–2849.
- [12] J. S. Müller, C. Vidaurre, M. Schreuder, F. C. Meinecke, P. von Büna, and K.-R. Müller, “A mathematical model for the two-learners problem,” *J. Neural Eng.*, vol. 14, no. 3, p. 036005, Mar. 2017, doi: 10.1088/1741-2552/aa620b.
- [13] Hespanha, J., *Noncooperative Game Theory: An Introduction for Engineers and Computer Scientists*, 2017.
- [14] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*, 2nd ed. Philadelphia: SIAM, 1999.
- [15] E. E. Fetz, “Volitional control of neural activity: implications for brain-computer interfaces,” *J Physiol*, vol. 579, no. Pt 3, pp. 571–579, Mar. 2007, doi: 10.1113/jphysiol.2006.127142.
- [16] D. Monderer and L. S. Shapley, “Potential Games,” *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, May 1996, doi: 10.1006/game.1996.0044.
- [17] Y. Nesterov and V. Spokoiny, “Random Gradient-Free Minimization of Convex Functions,” *Found. Comput. Math.*, vol. 17, no. 2, pp. 527–566, Apr. 2017, doi: 10.1007/s10208-015-9296-2.
- [18] S. H. Strogatz, *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*, Westview Press, 2015.