 DOI: 10.3217/978-3-99161-050-2-041

Bayesian Reinforcement Learning for Optimizing the BCI-Utility of P300 Brain-Computer Interfaces

Bangyao Zhao[1], Yixin Wang[2], Jane E. Huggins[3], Jian Kang[1]
[1]*Department of Biostatistics,* [2]*Department of Statistics,* [3]*Physical Medicine & Rehabilitation, University of Michigan, Ann Arbor, MI, USA*
Corresponding Author: Jian Kang, Email: jiankang@umich.edu

Introduction: P300 brain-computer interfaces (BCIs) provide a means of communication for individuals with severe disabilities by analyzing EEG responses to a series of stimuli [1]. A key challenge in these systems is optimizing BCI-utility, a comprehensive metric that balances accuracy and speed while incorporating the cost of error correction [2]. Methods such as dynamic stopping, which determines when enough information has been collected, and dynamic stimulus selection, which determines the next stimulus based on prior EEG responses, can potentially significantly improve BCI-utility. Meanwhile, advances in reinforcement learning (RL) have demonstrated its potential to address complex decision-making tasks in diverse domains, making it a promising approach for optimizing real-time BCI systems.

Methods: We propose a Bayesian model-based reinforcement learning (RL) framework that systematically optimizes both a dynamic stopping policy and a dynamic stimulus selection policy to maximize BCI-utility. The BCI system is treated as an agent, with each stimulus representing a state. The state consists of confidence scores for characters, calculated from classifier scores of previous EEG responses. At each state, the agent decides whether to stop and predict the target character (dynamic stopping policy, π_1) or to continue by selecting the next stimulus to present (dynamic stimulus selection policy, π_2). The dynamic stopping policy is implemented using an actor-critic algorithm. For stimulus selection, a Gaussian process-based Bayesian model predicts changes in confidence scores resulting from the next stimulus, guiding the agent to select the next stimulus that maximizes the expected BCI-utility. Figure 1 illustrates the workflow of this framework. Furthermore, the framework can effectively address online implementation challenges, including the double-target issue and delayed EEG responses.

Results: The proposed framework was evaluated on simulated datasets with varying signal-to-noise ratios (SNRs) and validated using recorded EEG data from a P300 speller study conducted by the University of Michigan Direct Brain Interface (UMDBI) laboratory [3]. Simulations demonstrated that our method significantly improved BCI-utility compared to benchmark methods across different SNR conditions. On real participant data, the framework achieved approximately a 20% improvement in BCI-utility while maintaining an accuracy of around 90% and reducing the time required per character. These results highlight the framework's ability to enhance BCI-utility by balancing speed and accuracy.

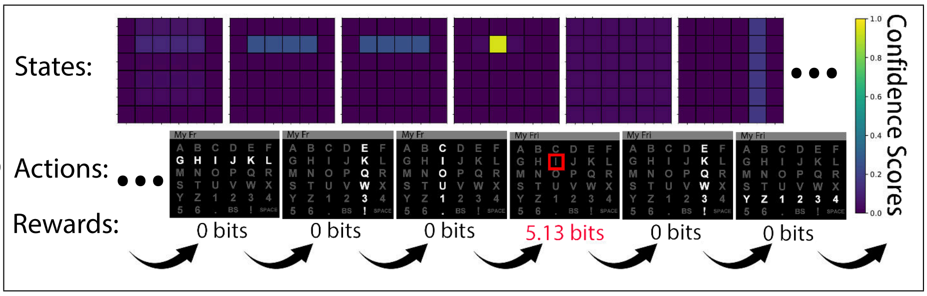

Figure 1: *Illustration of the reinforcement learning framework in P300 BCIs.*

Conclusion: By integrating Bayesian model-based reinforcement learning, the proposed framework demonstrates significant improvements in BCI performance.

References:

[1] Farwell, L. A., & Donchin, E. (1988). Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, *70*(6), 510–523.

[2] Dal Seno, B., Matteucci, M., & Mainardi, L. T. (2009). The utility metric: A novel method to assess the overall performance of discrete brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *18*(1), 20–28.

[3] Thompson, D. E., Gruis, K. L., & Huggins, J. E. (2014). A plug-and-play brain-computer interface to operate commercial assistive technology. *Disability and Rehabilitation: Assistive Technology*, *9*(2), 144–150.

Published by Verlag der Technischen Universität Graz