Ethical Considerations for Language Modeling within Brain-Computer Interfaces

Melanie Fried-Oken, PhD and Eran Klein, MD, co-PIs

Machine learning (ML) and Natural Language Processing (NLP) have the potential to transform communication for patients with neurodegenerative disease through personalized and real-time augmentative and alternative communication (AAC) devices. Individuals with severe communication impairments who can no longer control their daily conversations or participate in previous life roles want AAC devices. And they want them to work – to be reliable, effective, and fast. ML and NLP are emerging as promising tools to bridge current technology and next generation devices for individuals with the most severe speech and physical impairments, like the RSVP Keyboard™, a brain-computer interface (BCI) being developed by the parent grant. BCI systems for communication are referred to as AAC-BCIs. NLP efforts to combine large public data sets with private data sets, such as personal email messages, promise to give individuals with communication impairments their own personalized language models, models that are sufficiently robust to get closer to real-time communication. The focus on getting AAC-BCIs to work with machine learning, however, has led to a critical oversight in the field: an inadequate understanding of why individuals want next-generation devices and what trade-offs they are willing to make for faster and more personalized communication. The turn to ML brings this oversight into sharp relief. Individuals should provide input about the data sets used to construct their personal language models, but this raises important ethical questions about what individuals value, how they understand their identity, and what trade-offs they are willing to make relative to their personalized communication data. The goal of this supplement is to fill this gap in understanding so that researchers can implement ML into next generation AAC-BCI systems in a way that is sensitive to the ethical concerns of future users. There are four components to this ethics supplement: (1) to design a toolbox of ethics vignettes tailored to ethical concerns raised by both BCI communication and ML; (2) to administer monthly vignette-based online ethics surveys to individuals with severe communication impairments due to motor neuron disease (e.g., ALS) (n=25) or movement disorders (e.g., Parkinson's disease) (n=25); (3) to conduct semi-structured vignette-based interviews with individuals with pre-clinical or mild communication impairment due to motor neuron disease (n=10) or movement disorder (n=10). Components (2) and (3) will employ an iterative, parallel mixed-method approach. Trends in Likert-style online responses to ethics vignettes in the severe communication impairment cohort will be used to inform and modify the semi-structured interview prompts asked of the pre-clinical or mild impairment cohort. In parallel, themes emerging from direct content analysis of interviews will be used to refine online survey questions. Results of this iterative, mix-methods approach will be used (4) to outline a framework of core ethical domains and preliminary tools (vignettes and discussion prompts) that AAC-BCI researchers can use to assess ethical concerns while developing and iteratively refining communication technology for personalized language models.