Building a better brain-computer interface

October 2 2018, by Matt Miles

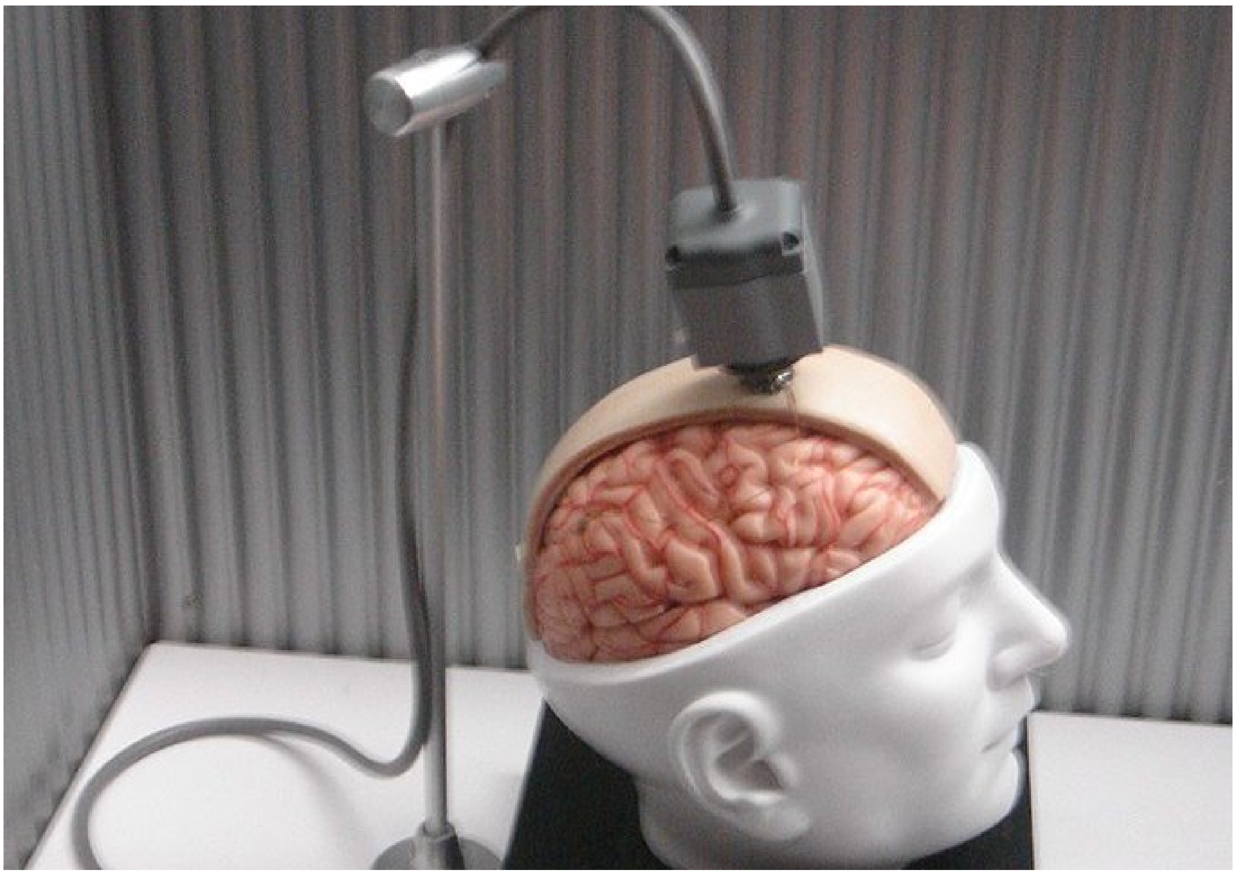

Photo of a dummy BrainGate interface. Credit: Paul Wick/Wikimedia Commons

Brain-computer interfaces, or BCIs, represent relatively recent advances in neurotechnology that allow computer systems to interact directly with human or animal brains. This technology is particularly promising for use in cases of spinal cord injury or paralysis. In these situations, patients may be able to use neural decoders that access part of their brain to operate a prosthetic limb or even to re-animate a paralyzed limb through functional electrical stimulation (FES).

Michael A. Schwemmer and colleagues, in a recent *Nature Medicine* article, detail their research on BCIs using deep neural network decoders with a participant with tetraplegia due to spinal cord injury. Their research focuses on addressing several key needs identified by end-users of BCI systems, namely: high accuracy, minimal daily setup, rapid response time, and multifunctionality—all of which are characteristics heavily influenced by a BCI's particular neural decoding algorithm.

Schwemmer's group describes several different approaches to training and testing three variations on neural network decoders (NN-BCI) in comparison with each other and a benchmark support vector machine (SVM) decoder. The four BCI decoder paradigms were developed and tested over the course of several years in association with a 27-year-old male participant with tetraplegia. The participant had a 96-channel microelectrode array implanted in the area of his left primary motor cortex corresponding to the hand and arm. Using intracortical data collected from 80 sessions over 865 days, the investigators trained and evaluated these BCI decoders. These sessions consisted of two 104-second blocks of a four-movement task: index extension, index flexion, wrist extension, and wrist flexion.

The initial neural network (NN) model was developed and calibrated using data from the first 40 sessions (80 blocks); it was not updated over the second half of the training/testing period, and is referred to here as the fixed neural network (fNN) model. From the fNN, two other neural network models were created: a supervised updating (sNN) model and an unsupervised updating (uNN) model. Both models used data from the first block of the second 40-session (updating/testing) period. The sNN model's algorithm relies on explicit training labels, that is, known timing and type of movement, whereas the uNN model relies on undifferentiated or unknown direct input in relation to intended action of the limb. The second block of the second 40-session period was used for accuracy testing of all models—fNN, sNN, uNN, and SVM.

The purpose of using four separate models here was to test and demonstrate various aspects of the three neural network models in relation to each other and the benchmark SVM model. For instance, the supervised neural network (sNN) model was updated daily (during the first block of the second 40-session period) and compared directly with the daily-retrained SVM model. The fixed neural network (fNN) model was provided to demonstrate that a BCI could sustain accuracy for over a year with no updates.

The unsupervised neural network (uNN) was perhaps the most interesting comparator, as we shall see, because it attempted to combine the improved accuracy gained from daily updates but without the consequent daily setup time required by the sNN model. Accuracy was the key performance measure in all tests, defined here as a percentage of correctly predicted time-bins in the second block of the second 40 sessions; the criterion of greater than 90% accuracy was one of the four end-user requirements originally articulated at the outset of the study.

The sNN consistently outperformed the daily-retrained SVM: in 37 out of 40 sessions, its accuracy was > 90%, whereas the SVM only achieved > 90% accuracy in 12 sessions. The fNN also outperformed the SVM in 36 of 40 sessions; it achieved > 90% accuracy in 32 sessions. The fNN accuracy was, not surprisingly, lower than the accuracy of the sNN, and both fixed decoders, fNN and SVM, declined in accuracy over the course of the study period, in contrast to the daily-updated decoders.

Perhaps the most interesting finding of this research however, is the performance of the unsupervised neural network (uNN), which outperformed both fixed models in terms of accuracy, while also meeting the end-user requirement of minimal daily set-up. Where the sNN model required explicit daily training, the uNN incorporated data from general use in its update schema, which required no such daily set-up. In comparison with the fNN, a performance gap emerged over time, and the benefits of the uNN distinguished themselves. The uNN also outperformed the SVM in terms of response time, another key end-user requirement.

Another important aspect of this study with regard to NNs focused on transfer learning, whereby new movements can be added to the existing repertoire with minimal additional training and data. In this case, "hand open" and "hand close" were added to the previous four movements, and all decoders were rebuilt. Here too, unsupervised updating was used to build an unsupervised transfer neural network (utNN), which, after only one session of training oupterformed the SVM model.

Finally, the previous research—all of which was conducted in an "offline" setting—was applied, via the participant's FES-controlled hand and forearm, to show that a transfer learning uNN trained on the original four-movement task could be used to quickly create a new decoder to control, in real time, an open hand and three grips (can, fork, and peg). In a test of the system, the participant was able to perform all three hand movement grip tasks, with no failures, in 45 attempts. Previously, he was only able to perform one grip task successfully.

In summarizing how the results of their study relate to the main end-user expectations previously described, the investigators cite the following achievements: "(i) using deep NNs to create robust neural decoders that sustain high fidelity BCI control for more than a year without retraining; (ii) introducing a new updating procedure that can improve performance using data obtained through regular system use; (iii) extension of functionality through transfer learning using minimal additional data; and (iv) introducing a decoding framework that simultaneously addresses these four competing aspects of BCI performance (accuracy, speed, longevity, and multifunctionality). In addition, we provide a clinical demonstration that a decoder calibrated using historical data of imagined hand movements with no feedback can be successfully used in real-time to control FES-evoked grasp function for object manipulation."

Schwemmer and colleagues go on to offer a more in-depth discussion of their results amidst the broader landscape of BCI research, and offer commentary on some of the specific challenges and limitations of their experiment. While noting that the median response time for uNN decoders (0.9 s) is still faster than that of SVM decoders (1.1 s), they acknowledge that a target of 750 ms or less is probably closer to realistic end-user expectations.

Ultimately they conclude: "We have demonstrated that decoders based on NNs may be superior to other implementations because new functions can be easily added after the initial decoder calibration using transfer learning. Crucially, we show that this secondary update to add more movements requires a minimal amount of additional data." And "insights gained from offline data and analyses can carry over to a realistic online BCI scenario with minimal additional data collection."

More information: Michael A. Schwemmer et al. Meeting brain–computer interface user performance expectations using a deep neural network decoding framework, *Nature Medicine* (2018). DOI: 10.1038/s41591-018-0171-y from https://medicalxpress.com/news/2018-10-brain-computer-interface.html