

News & views

Neuroscience

Neural interface translates thoughts into type

Pavithra Rajeswaran & Amy L. Orsborn

A neural interface has been developed that could enable people with paralysis to type faster than they could using other technologies, by directly translating attempts at handwriting into text. See p.249

We can think much faster than we can communicate – a fact that many of us feel aware of as we struggle with our smartphone keyboards. For people with severe paralysis, this information bottleneck is much more extreme. Willett *et al.*¹ report on page 249 the development of a brain–computer interface (BCI) for typing that could eventually let people with paralysis communicate at the speed of their thoughts.

Commercially available assistive typing devices predominantly rely on the person using the device being able to make eye movements or deliver voice commands. Eye-tracking keyboards can let people with paralysis type at around 47.5 characters per minute², slower than the 115-per-minute speeds achieved by people without a comparable injury. However, these technologies do not work for people whose paralysis impairs eye movements or vocalization. And the technology has limitations. For instance, it is hard to reread an e-mail, so that you can compose your reply, while you are typing with your eyes.

By contrast, BCIs restore function by deciphering patterns of brain activity. Such interfaces have successfully restored simple movements – such as reaching for and manipulating large objects – to people with paralysis^{3–7}. By directly tapping into neural processing, BCIs hold the tantalizing promise of seamlessly restoring function to a wide range of people.

But, so far, BCIs for typing have been unable to compete with simpler assistive technologies such as eye-trackers. One reason is that typing is a complex task. In English, we select from 26 letters of the Latin alphabet. Building a classification algorithm to predict which letter a user wants to choose, on the basis of their neural activity, is challenging, so BCIs have solved typing tasks indirectly. For instance,

non-invasive BCI spellers present several sequential visual cues to the user and analyse the neural responses to all cues to determine the desired letter⁸. The most successful invasive BCI (iBCI; one that involves implanting an electrode into the brain) for typing allowed users to control a cursor to select keys, and achieved speeds of 40 characters per minute⁶. But these iBCIs, like non-invasive eye-trackers, occupy the user's visual attention and do not

provide notably faster typing speeds.

Willett and colleagues developed a different approach, which directly solves the typing task in an iBCI and thereby leapfrogs far beyond past devices, in terms of both performance and functionality. The approach involves decoding letters as users imagine writing at their own pace (Fig. 1).

Such an approach required a classification algorithm that predicts which of 26 letters or 5 punctuation marks a user with paralysis is trying to write – a challenging feat when the attempts cannot be observed and occur whenever the user chooses. To overcome this challenge, Willett *et al.* first repurposed another type of algorithm – a machine-learning algorithm originally developed for speech recognition. This allowed them to estimate, on the basis of neural activity alone, when a user started attempting to write a character. The pattern of neural activity generated each time their study participant imagined a given character was remarkably consistent. From this information, the group produced a labelled data set that contained the neural-activity patterns corresponding to each character. They used this data set to train the classification algorithm.




Figure 1 | A brain–computer interface (BCI) for typing. Willett *et al.*¹ have developed a BCI that enables a person with paralysis to type, by translating the neural activity produced from imagined attempts at handwriting into text on the computer screen. As a simplified description, electrodes implanted into the brain measure the activity of many neurons as the user imagines writing each letter (lines indicate time points at which each neuron fires). A deep-learning model called a recurrent neural network (RNN) learns the neural activity patterns produced from each character, and analyses how these activity patterns relate across multiple trials, generating cluster plots. This information is used to by an algorithm to predict the letters being imagined by the participant in the current trial, and the prediction is translated into a typographic output. (Figure adapted from Fig. 2a of ref. 1.)

To achieve accurate classification in such a high-dimensional space, Willett and colleagues' classification algorithm used current machine-learning methods, along with a type of artificial neural network called a recurrent neural network (RNN), which is especially good at predicting sequential data. Harnessing the power of RNNs requires ample training data, but such data are limited in neural interfaces, because few users want to imagine writing for hours on end. The authors solved this problem using an approach known as data augmentation, in which neural activity patterns previously generated by the participant are used to produce artificial sentences on which to train the RNN. They also expanded their training data by introducing artificial variability into the patterns of neural activity, to mimic changes that occur naturally in the human brain. Such variability can make RNN BCIs more robust⁹.

Thanks to these methods, Willett and colleagues' algorithm provided impressively accurate classification, picking the correct character 94.1% of the time. By including predictive-language models (similar to those that drive auto-correct functions on a smartphone), they further improved accuracy to 99.1%. The participant was able to type accurately at a speed of 90 characters per minute – a twofold improvement on his performance with past iBCIs.

This study's achievements, however, stem from more than machine learning. A decoder's performance is ultimately only as good as the data that are fed into it. The researchers found that neural data associated with attempted handwriting are particularly well-suited for typing tasks and classification. In fact, handwriting could be classified quite well even with simpler, linear algorithms, suggesting that the neural data themselves played a large part in the success of the authors' approach.

By simulating how the classification algorithm performed when tested with different types of neural activity, Willett *et al.* made a key insight – neural activity during handwriting has more temporal variability between characters than does neural activity when users attempt to draw straight lines, and this variability actually makes classification easier. This knowledge should inform future BCIs. Perhaps counter-intuitively, it might be advantageous to decode complex behaviours rather than simple ones, particularly for classification tasks.

Willett and co-workers' study begins to deliver on the promise of BCI technologies. iBCIs will need to provide tremendous performance and usability benefits to justify the expense and risks associated with implanting electrodes into the brain. Importantly, typing speed is not the only factor that will determine whether the technology is adopted – the longevity and robustness of the approach

also require analysis. The authors present promising evidence that their algorithms will perform well with limited training data, but further research will probably be required to enable the device to maintain performance over its lifetime as neural activity patterns change. It will also be crucial to conduct studies to test whether the approach can be generalized for other users, and for settings outside the laboratory.

Another question is how the approach will scale and translate to other languages. Willett and colleagues' simulations highlight that several characters of the Latin alphabet are written similarly (r, v and u, for instance), and so are harder to classify than are others. One of us (P.R.) speaks Tamil, which has 247, often very closely related, characters, and so might be much harder to classify. And the question of translation is particularly pertinent for languages that are not yet well represented in machine-learning predictive-language models.

Although much work remains to be done, Willett and co-workers' study is a milestone that broadens the horizon of iBCI applications. Because it uses machine-learning methods that are rapidly improving, plugging in the

latest models offers a promising path for future improvements. The team is also making its data set publicly available, which will accelerate advances. The authors' approach has brought neural interfaces that allow rapid communication much closer to a practical reality.

Pavithra Rajeswaran and **Amy L. Orsborn**

are in the Department of Bioengineering, University of Washington, Seattle, Washington 98195, USA. **A.L.O.** is also in the Department of Electrical and Computer Engineering, and at the Washington National Primate Research Center, University of Washington. e-mail: aorsborn@uw.edu

1. Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M. & Shenoy, K. V. *Nature* **593**, 249–254 (2021).
2. Mott, M. E., Williams, S., Wobbrock, J. O. & Morris, M. R. in *Proc. 2017 CHI Conf. Human Factors in Computing Systems* 2558–2570 (ACM, 2017).
3. Hochberg, L. R. *et al. Nature* **442**, 164–171 (2006).
4. Hochberg, L. R. *et al. Nature* **485**, 372–375 (2012).
5. Collinger, J. L. *et al. Lancet* **381**, 557–564 (2013).
6. Pandarinath, C. *et al. eLife* **6**, e18554 (2017).
7. Ajiboye, A. B. *et al. Lancet* **389**, 1821–1830 (2017).
8. Rezeika, A. *et al. Brain Sci.* **8**, 57 (2018).
9. Sussillo, D., Stavisky, S. D., Kao, J. C., Ryu, S. I. & Shenoy, K. V. *Nature Commun.* **7**, 13749 (2016).

Atmospheric chemistry

How ant acid forms in the atmosphere

Joost de Gouw & Delphine Farmer

Known sources of formic acid could not explain the observed atmospheric concentrations of this compound. The discovery of a previously unknown pathway that generates formic acid in the atmosphere resolves this discrepancy. **See p.233**

Formic acid is one of the simplest and most abundant organic molecules in Earth's atmosphere, but its sources have been poorly understood for many years. Laboratory and field studies^{1–3} have shown that most formic acid is not emitted directly from sources, but is produced by chemical reactions in the atmosphere. However, the chemistry responsible has been a mystery. On page 233, Franco *et al.*⁴ report that formic acid could be formed by a mechanism that starts with formaldehyde reacting with water in cloud droplets.

The word 'formic' derives from *formica*, the Latin word for ant, and the compound is indeed released from ant hills⁵. Other, and larger, emission sources include vegetation, biomass burning⁶ and fossil-fuel combustion⁷. However, the combined emissions from known sources are too small to explain the

concentrations of formic acid in the atmosphere, and several studies have concluded that formation in the atmosphere is a much bigger contributor (see ref. 1, for example).

Levels of formic acid can be measured by mass spectrometry and optical spectroscopy, and from satellite instruments, so there is excellent information about the distribution of this compound in the atmosphere. Observations have shown that atmospheric concentrations of formic acid increase rapidly in urban⁸ and forest⁹ air during the day. However, researchers have been unable to identify the chemical reactions responsible for this increase. Detailed studies that considered all of the possible known chemical pathways could explain only a fraction of formic acid produced, both in polluted and remote regions^{3,9}, and so the search for alternative