Article

High-performance brain-to-text communication via handwriting

https://doi.org/10.1038/s41586-021-03506-2

Received: 2 July 2020

Accepted: 26 March 2021

Published online: 12 May 2021

Check for updates

Francis R. Willett[1,2,3]✉, Donald T. Avansino[1], Leigh R. Hochberg[4,5,6,7], Jaimie M. Henderson[2,8,9,12] & Krishna V. Shenoy[1,3,8,9,10,11,12]

Brain–computer interfaces (BCIs) can restore communication to people who have lost the ability to move or speak. So far, a major focus of BCI research has been on restoring gross motor skills, such as reaching and grasping[1–5] or point-and-click typing with a computer cursor[6,7]. However, rapid sequences of highly dexterous behaviours, such as handwriting or touch typing, might enable faster rates of communication. Here we developed an intracortical BCI that decodes attempted handwriting movements from neural activity in the motor cortex and translates it to text in real time, using a recurrent neural network decoding approach. With this BCI, our study participant, whose hand was paralysed from spinal cord injury, achieved typing speeds of 90 characters per minute with 94.1% raw accuracy online, and greater than 99% accuracy offline with a general-purpose autocorrect. To our knowledge, these typing speeds exceed those reported for any other BCI, and are comparable to typical smartphone typing speeds of individuals in the age group of our participant (115 characters per minute)[8]. Finally, theoretical considerations explain why temporally complex movements, such as handwriting, may be fundamentally easier to decode than point-to-point movements. Our results open a new approach for BCIs and demonstrate the feasibility of accurately decoding rapid, dexterous movements years after paralysis.

Previous BCI studies have shown that the motor intention for gross motor skills, such as reaching, grasping or moving a computer cursor, remains neurally encoded in the motor cortex after paralysis[1–7]. However, it is still unknown whether the neural representation for a rapid and highly dexterous motor skill, such as handwriting, also remains intact. We tested this by recording neural activity from two microelectrode arrays in the hand 'knob' area of the precentral gyrus (a premotor area)[9,10] while our BrainGate study participant, T5, attempted to handwrite individual letters and symbols (Fig. 1a). T5 has a high-level spinal cord injury and was paralysed from the neck down; his hand movements were entirely non-functional and limited to twitching and micromotion. We instructed T5 to 'attempt' to write as if his hand were not paralysed, while imagining that he was holding a pen on a piece of ruled paper.

Neural representation of handwriting

To visualize the neural activity (multiunit threshold crossing rates) recorded during attempted handwriting, we used principal components analysis to display the top three neural dimensions that contain the most variance (Fig. 1b). The neural activity appeared to be strong and repeatable, although the timing of its peaks and valleys varied across trials, potentially owing to fluctuations in writing speed. We used a time-alignment technique to remove temporal variability[11], which revealed notably consistent underlying patterns of neural activity that are unique to each character (Fig. 1c). To ascertain whether the neural activity encoded the pen movements that are needed to draw each character's shape, we attempted to reconstruct each character by linearly decoding the pen-tip velocity from the trial-averaged neural activity (Fig. 1d). Readily recognizable letter shapes confirmed that pen-tip velocity is robustly encoded. The neural dimensions that represented pen-tip velocity accounted for 30% of the total neural variance.

Next, we used a nonlinear dimensionality reduction method (*t*-distributed stochastic neighbour embedding; t-SNE) to produce a two-dimensional (2D) visualization of each single trial's neural activity recorded after the 'go' cue was given (Fig. 1e). The t-SNE visualization revealed tight clusters of neural activity for each character and a predominantly motoric encoding in which characters that are written similarly are closer together. Using a *k*-nearest-neighbour classifier applied offline to the neural activity, we could classify the characters with 94.1% accuracy (95% confidence interval (CI) = [92.6, 95.8]). Together, these results suggest that, even years after paralysis, the neural representation of handwriting in the motor cortex is probably strong enough to be useful for a BCI.

[1]Howard Hughes Medical Institute at Stanford University, Stanford, CA, USA. [2]Department of Neurosurgery, Stanford University School of Medicine, Stanford, CA, USA. [3]Department of Electrical Engineering, Stanford University, Stanford, CA, USA. [4]VA RR&D Center for Neurorestoration and Neurotechnology, Rehabilitation R&D Service, Providence VA Medical Center, Providence, RI, USA. [5]School of Engineering, Brown University, Providence, RI, USA. [6]Carney Institute for Brain Science, Brown University, Providence, RI, USA. [7]Center for Neurotechnology and Neurorecovery, Department of Neurology, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA. [8]Wu Tsai Neurosciences Institute, Stanford University, Stanford, CA, USA. [9]Bio-X Institute, Stanford University, Stanford, CA, USA. [10]Department of Bioengineering, Stanford University, Stanford, CA, USA. [11]Department of Neurobiology, Stanford University, Stanford, CA, USA. [12]These authors jointly supervised this work: Jaimie M. Henderson, Krishna V. Shenoy. ✉e-mail: fwillett@stanford.edu

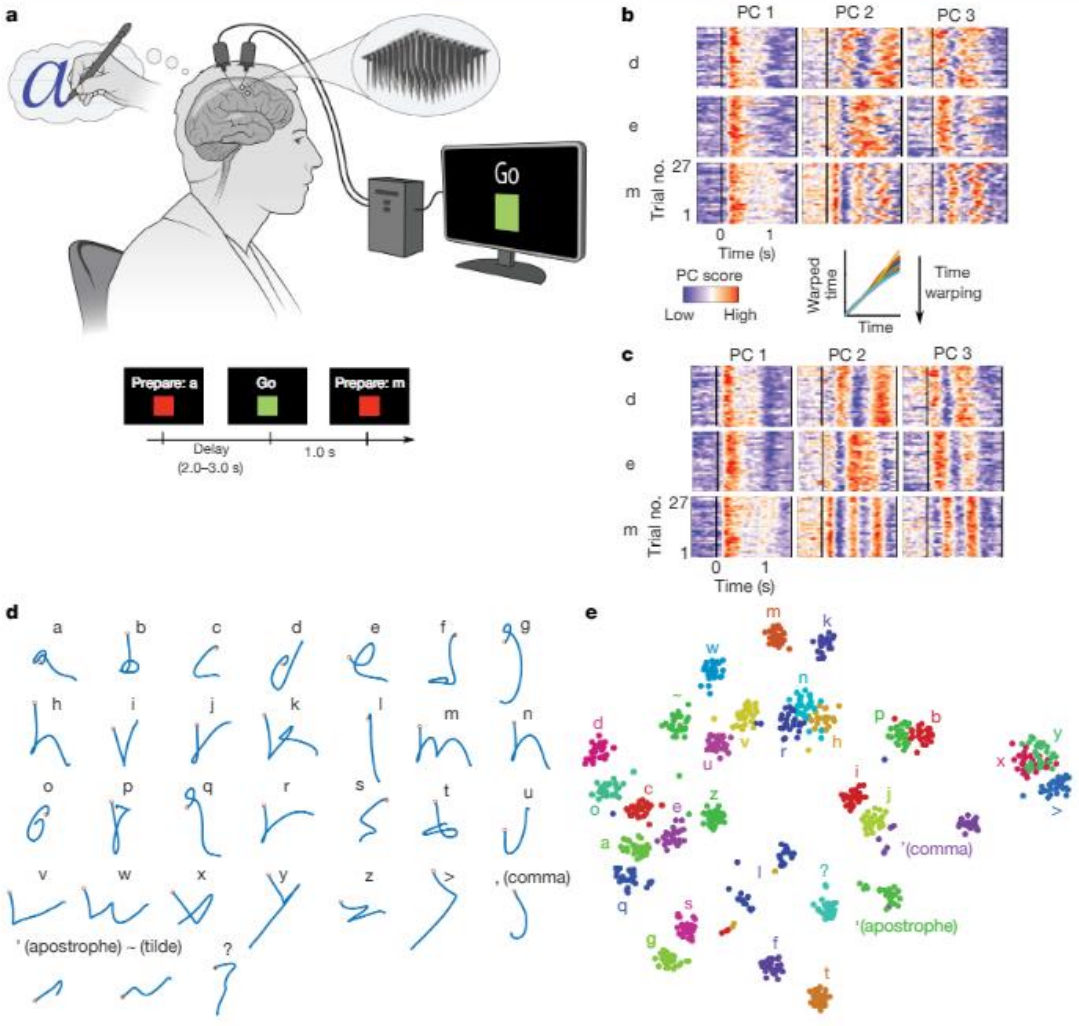

Fig. 1 | Neural representation of attempted handwriting. a, To assess the neural representation of attempted handwriting, participant T5 attempted to handwrite each character one at a time, following the instructions given on a computer screen (bottom panels depict what is shown on the screen, following the timeline). Credit: drawing of the human silhouette created by E. Woodrum. **b**, Neural activity in the top 3 principal components (PCs) is shown for three example letters (d, e and m) and 27 repetitions of each letter (trials). The colour scale was normalized within each panel separately for visualization. **c**, Time-warping the neural activity to remove trial-to-trial changes in writing speed reveals consistent patterns of activity unique to each letter. In the inset above **c**, example time-warping functions are shown for the letter 'm' and lie relatively close to the identity line (the warping function of each trial is plotted with a different coloured line). **d**, Decoded pen trajectories are shown for all 31 tested characters. Intended 2D pen-tip velocity was linearly decoded from the neural activity using cross-validation (each character was held out), and decoder output was denoised by averaging across trials. Orange circles denote the start of the trajectory. **e**, A 2D visualization of the neural activity made using t-SNE. Each circle is a single trial (27 trials are shown for each of 31 characters).

Decoding handwritten sentences

Next, we tested whether we could decode complete handwritten sentences in real time, thus enabling an individual with tetraplegia to communicate by attempting to handwrite their intended message. To do so, we trained a recurrent neural network (RNN) to convert the neural activity into probabilities describing the likelihood of each character being written at each moment in time (Fig. 2a, Extended Data Fig. 1). These probabilities could either be thresholded in a simple way to emit discrete characters, which we did for real-time decoding ('raw online output', Fig. 2a), or processed more extensively by a large-vocabulary language model to simulate an autocorrect feature, which we applied offline ('offline output from a language model', Fig. 2a). We used the limited set of 31 characters shown in Fig. 1d, consisting of the 26 lower-case letters of the alphabet, together with commas, apostrophes, question marks, full stops (written by T5 as a tilde symbol; '~') and spaces (written by T5 as a greater-than symbol; '>'). The '~' and '>' symbols were chosen to make full stops and spaces easier to detect. T5 attempted to write each character in print (not cursive), with each character printed on top of the previous one.

To collect training data for the RNN, we recorded neural activity while T5 attempted to handwrite complete sentences at his own pace, following instructions on a computer monitor. Before the first day of real-time evaluation, we collected a total of 242 sentences across 3 pilot days that were combined to train the RNN. On each subsequent day of real-time testing, additional training data were collected to recalibrate the RNN before evaluation, yielding a combined total of 572 training sentences by the last day (comprising 7.6 hours and 31,472 characters). To train the RNN, we adapted neural network methods in automatic speech recognition[12–14] to overcome two key challenges: (1) the time that each letter was written in the training data was unknown (as T5's hand was paralysed), making it challenging to apply supervised learning

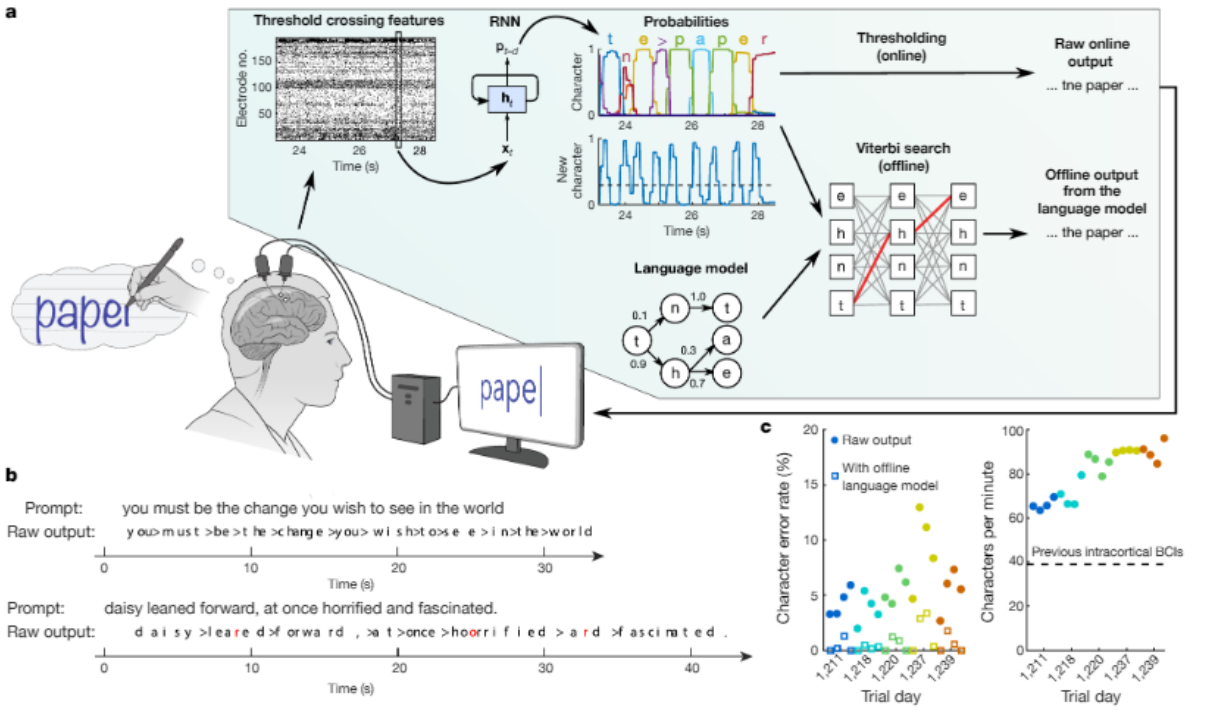

Fig. 2 | Neural decoding of attempted handwriting in real time. a, Diagram of the decoding algorithm. First, the neural activity (multiunit threshold crossings) was temporally binned and smoothed on each electrode (20-ms bins). Then, an RNN converted this neural population time series ($\mathbf{x}_t$) into a probability time series ($\mathbf{p}_{t-d}$) describing the likelihood of each character and the probability of any new character beginning. The RNN had a one second output delay (d), giving it time to observe each character fully before deciding its identity. Finally, the character probabilities were thresholded to produce the 'raw online output' for real-time use (when the 'new character' probability crossed a threshold at time t, the most likely character at time $t + 0.3$ s was emitted and shown on the screen). In an offline retrospective analysis, the character probabilities were combined with a large-vocabulary language model to decode the most likely text that the participant wrote (using a custom 50,000-word bigram model). Credit: drawing of the human silhouette created by E. Woodrum. **b**, Two real-time example trials are shown, demonstrating the ability of the RNN to decode readily understandable text on sentences on which it was never trained. Errors are highlighted in red and spaces are denoted with '>'. **c**, Error rates (edit distances) and typing speeds are shown for 5 days, with 4 blocks of 7–10 sentences each (each block is indicated with a single circle and coloured according to the trial day). The speed is more than double that of the next-fastest intracortical BCI[7], which is indicated with the dashed line.

techniques; and (2) the dataset was limited in size compared to typical RNN datasets, making it difficult to prevent overfitting to the training data (see Supplementary Methods, Extended Data Figs. 2, 3).

We evaluated the performance of the RNN over a series of 5 days, each day containing 4 evaluation blocks of 7–10 sentences that the RNN was never trained on (thus ensuring that the RNN could not overfit to those sentences). T5 copied each sentence from an on-screen prompt, attempting to handwrite it letter by letter, while the decoded characters appeared on the screen in real time as they were detected by the RNN (Supplementary Videos 1, 2, Extended Data Table 1). Characters appeared after they were completed by T5 with a short delay (estimated to be 0.4–0.7 s). The decoded sentences were quite legible ('raw output', Fig. 2b). Notably, typing speeds were high, plateauing at 90 characters per minute with a mean character error rate of 5.4% (averaged across all four blocks on the final day) (Fig. 2c). As there was no 'backspace' function implemented, T5 was instructed to continue writing if any decoding errors occurred.

When a language model was used to autocorrect errors offline, error rates decreased considerably (Fig. 2c, Table 1). The character error rate decreased to 0.89% and the word error rate decreased to 3.4% averaged across all days, which is comparable to state-of-the-art speech recognition systems with word error rates of 4–5%[14,15], putting it well within the range of usability. Finally, to probe the limits of possible decoding performance, we trained a new RNN offline using all available sentences to process the entire sentence in a non-causal way (comparable to other BCI studies[16,17]). Accuracy was extremely high in this regime (0.17% character error rate), indicating a high potential ceiling of performance, although this decoder would not be able to provide letter-by-letter feedback to the user.

Next, to evaluate performance in a less restrained setting, we collected two days of data in which T5 used the BCI to freely type answers to open-ended questions (Supplementary Video 3, Extended Data Table 2). The results confirm that high performance can also be achieved when the user writes self-generated sentences as opposed to copying on-screen prompts (73.8 characters per minute, with a character error rate of 8.54% in real time and 2.25% with a language model). To our knowledge, the previous record for free typing in intracortical BCIs is 24.4 correct characters per minute[7].

Daily decoder retraining

Following standard practice[1,2,4,5,18], we retrained our handwriting decoder each day before evaluating it, with the help of 'calibration' data collected at the beginning of each day. Retraining helps to account

Table 1 | Mean character and word error rates (with 95% CIs) for the handwriting BCI across all 5 days

	Character error rate [95% CI]	Word error rate [95% CI]
Raw online output	5.9% [5.3, 6.5]	25.1% [22.5, 27.4]
Online output + offline language model	0.89% [0.61, 1.2]	3.4% [2.5, 4.4]
Offline bidirectional RNN + language model	0.17% [0, 0.36]	1.5% [0, 3.2]

'Raw online output' is what was decoded online (in real time). 'Online output + offline language model' was obtained by applying a language model retrospectively to what was decoded online (to simulate an autocorrect feature). 'Offline bidirectional RNN + language model' was obtained by retraining a bidirectional (acausal) decoder offline using all available data, in addition to applying a language model. Word error rates can be much higher than character error rates because a word is considered incorrect if any character in that word is wrong.

for changes in neural recordings that accrue over time, which might be caused by neural plasticity or electrode array micromotion. Ideally, to reduce the burden on the user, minimal or no calibration data would be required. In a retrospective analysis of the copy-typing data reported above in Fig. 2, we assessed whether high performance could still have been achieved using fewer than the original 50 calibration sentences per day (Fig. 3a). We found that 10 sentences (8.7 min) were enough to achieve a raw error rate of 8.5% (1.7% with a language model), although 30 sentences were needed to match the raw online error rate of 5.9%.

However, our copy-typing data were collected over a 28-day time span, possibly allowing larger changes in neural activity to accumulate. Using further offline analyses, we assessed whether sessions that are more closely spaced reduce the need for calibration data (Fig. 3b). We found that when only 2–7 days passed between sessions, performance was reasonable with no decoder retraining (11.1% raw error rate, 1.5% with a language model), as might be expected from previous work showing the short-term stability of neural recordings[19–21]. Finally, we tested whether decoders could be retrained in an unsupervised manner by using a language model to error-correct and retrain the decoder, thus bypassing the need to interrupt the user for calibration (by enabling automatic recalibration during normal use). Encouragingly, unsupervised retraining achieved a raw error rate of 7.3% (0.84% with a language model) when sessions were separated by a time span of 7 days or less.

Ultimately, whether decoders can be successfully retrained with minimal recalibration data depends on how quickly the neural activity changes over time. We assessed the stability of the neural patterns associated with each character and found high short-term stability (mean correlation of 0.85 when 7 days apart or less), and neural changes that seemed to accumulate at a steady and predictable rate (Extended Data Fig. 4). These results are promising for clinical viability, as they suggest that unsupervised decoder retraining, combined with more-limited supervised retraining after longer periods of inactivity, may be sufficient to achieve high performance. Nevertheless, future work must confirm this online, as offline simulations are not always predictive of online performance.

Temporal variety improves decoding

To our knowledge, 90 characters per minute is the highest typing rate that has yet been reported for any type of BCI (see 'Discussion'). For intracortical BCIs, the best-performing method has been point-and-click typing with a 2D computer cursor, which peaks at 40 correct characters per minute[7] (see Supplementary Video 4 for a direct comparison). The speed of point-and-click BCIs is limited primarily by decoding accuracy. During parameter optimization, the cursor gain is increased so as to increase typing rate, until the cursor moves so quickly that it becomes uncontrollable owing to decoding errors[22].

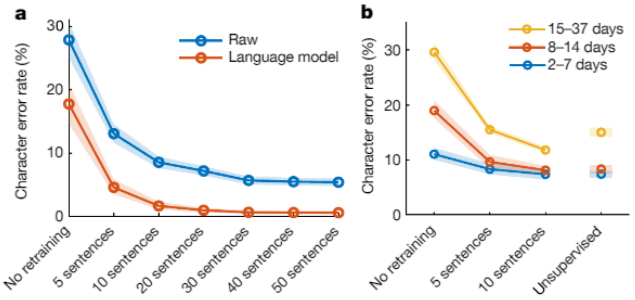

Fig. 3 | Performance remains high when daily decoder retraining is shortened (or unsupervised). a, To account for changes in neural activity that accrue over time, we retrained our handwriting decoder each day before evaluating it. Here, we simulated offline how decoding performance would have changed if fewer than the original 50 calibration sentences were used. Lines show the mean error rate across all data and shaded regions indicate 95% CIs. **b**, Copy-typing data from eight sessions were used to assess whether fewer calibration data are required if sessions occur closer in time. All session pairs (X, Y) were considered. Decoders were first initialized using training data from session X and earlier, and then evaluated on session Y under different retraining methods (no retraining, retraining with limited calibration data, or unsupervised retraining). Lines show the average raw error rate and shaded regions indicate 95% CIs.

We therefore asked how handwriting movements could be decoded more than twice as fast, with similar levels of accuracy.

We theorize that handwritten letters may be easier to distinguish from each other than point-to-point movements, as letters have more variety in their spatiotemporal patterns of neural activity than do straight-line movements. To test this theory, we analysed the patterns of neural activity associated with 16 straight-line movements and 16 letters that required no lifting of the pen off the page, both performed by T5 with attempted handwriting (Fig. 4a, b).

First, we analysed the pairwise Euclidean distances between each neural activity pattern. We found that the nearest-neighbour distances for each movement were 72% larger for characters compared to straight lines (95% CI = [60%, 86%]), making it less likely for a decoder to confuse two nearby characters (Fig. 4c). To confirm this, we simulated the classification accuracy for each set of movements as a function of neural noise (Fig. 4d), which showed that characters are easier to classify than straight lines.

To gain insight into what might be responsible for the relative increase in nearest-neighbour distances for characters, we examined the spatial and temporal dimensionality of the neural patterns. Spatial and temporal dimensionality were estimated using the 'participation ratio' of the principal component analysis (PCA) eigenvalue spectrum, which quantifies approximately how many spatial or temporal dimensions are required to explain 80% of the variance in the patterns of neural activity[23]. We found that the spatial dimensionality was only modestly larger for characters (1.24 times larger; 95% CI = [1.19, 1.30]), but that the temporal dimensionality was much greater (2.65 times larger; 95% CI = [2.58, 2.72]), suggesting that the increased variety of temporal patterns in letter writing drives the increased separability of each movement (Fig. 4e).

To illustrate how increased temporal dimensionality can make movements more distinguishable, we constructed a toy model with four movements and two neurons, with the neural activity constrained to lie along a single dimension (Fig. 4f, g). Simply by allowing the trajectories to change in time (Fig. 4g), the nearest-neighbour distance between the neural trajectories can be increased, resulting in an increase in classification accuracy when noise levels are large enough (Fig. 4h). Although neural noise in the toy model was assumed to be independent white noise, we found that these results also hold for noise that is correlated across time and neurons (Extended Data Fig. 5, Supplementary Note 1).

These results suggest that time-varying patterns of movement, such as handwritten letters, are fundamentally easier to decode than point-to-point movements. We think this is one—but not necessarily the only—important factor that enabled a handwriting BCI to go faster than continuous-motion point-and-click BCIs. Other discrete (classification-based) BCIs have also typically used directional movements with little temporal variety, which may have limited their accuracy and/or the size of the movement set[24,25].

More generally, using the principle of maximizing the nearest-neighbour distance between movements, it should be possible to optimize a set of movements for ease of classification[26]. We investigated this possibility and designed an alphabet that is theoretically easier to classify than the Latin alphabet (Extended Data Fig. 6). The optimized alphabet avoids large clusters of redundant letters that are written similarly (most Latin letters begin with either a downstroke or a counter-clockwise curl).

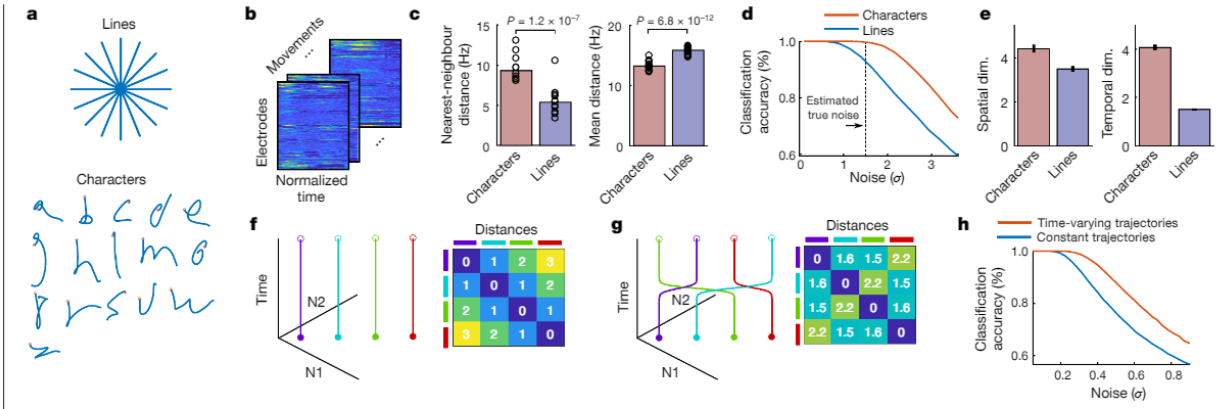

Fig. 4 | Increased temporal variety can make movements easier to decode. **a**, We analysed the spatiotemporal patterns of neural activity corresponding to 16 handwritten characters (1 s in duration) versus 16 handwritten straight-line movements (0.6 s in duration). **b**, Spatiotemporal neural patterns were found by averaging over all trials for a given movement (after time-warping to align the trials in time)[11]. Neural activity was resampled to equalize the duration of each set of movements, resulting in a 192 × 100 matrix for each movement (192 electrodes and 100 time steps). **c**, Pairwise Euclidean distances between neural patterns were computed for each set, revealing larger nearest-neighbour distances (but not mean distances) for characters. Each circle represents a single movement and bar heights show the mean. **d**, Larger nearest-neighbour distances made the characters easier to classify than straight lines. The noise is in units of standard deviations and matches the scale of the distances in **c**. **e**, The spatial dimensionality (dim.) was similar for characters and straight lines, but the temporal dimensionality was more than twice as high for characters, suggesting that more temporal variety underlies the increased nearest-neighbour distances and better classification performance. Error bars show 95% CIs. Dimensionality was quantified using the participation ratio. **f**–**h**, A toy example to give intuition for how increased temporal dimensionality can make neural trajectories more separable. Four neural trajectories are depicted (N1 and N2 are two hypothetical neurons, the activity of which is constrained to a single spatial dimension, the unity diagonal). Allowing the trajectories to vary in time by adding one bend, which increases the temporal dimensionality from 1 (**f**) to 2 (**g**), enables larger nearest-neighbour distances and better classification (**h**).

Discussion

Locked-in syndrome (paralysis of nearly all voluntary muscles) severely impairs or prevents communication, and is most frequently caused by brainstem stroke or late-stage amyotrophic lateral sclerosis (estimated prevalence of locked-in syndrome: 1 in 100,000[27]). Commonly used BCIs for restoring communication are either flashing electroencephalogram (EEG) spellers[18,28–32] or intracortical point-and-click BCIs[6,7,33]. EEG spellers based on oddball potentials or motor imagery typically achieve 1–5 characters per minute[28–32]. EEG spellers that use visually evoked potentials have achieved speeds of 60 characters per minute[18], but have notable usability limitations, as they tie up the eyes, are not typically self-paced and require panels of flashing lights on a screen. Intracortical BCIs based on 2D cursor movements give the user more freedom to look around and set their own pace of communication, but have yet to exceed 40 correct characters per minute in humans[7]. Recently, speech-decoding BCIs have shown exciting promise for restoring rapid communication[16,17,34], but their accuracies and vocabulary sizes require further improvement to support general-purpose use.

Here, we introduced a new approach for communication BCIs—decoding a rapid, dexterous motor behaviour in an individual with tetraplegia—that sets a benchmark for communication rate at 90 characters per minute. The real-time system is general (the user can express any sentence), easy to use (entirely self-paced and the eyes are free to move) and accurate enough to be useful in the real-world (94.1% raw accuracy, and greater than 99% accuracy offline with a large-vocabulary language model). To achieve high performance, we developed decoding methods to work effectively with unlabelled neural sequences in data-limited regimes. These methods could be applied more generally to any sequential behaviour that cannot be observed directly (for example, decoding speech from someone who can no longer speak).

It is important to recognize that the current system is a proof of concept that a high-performance handwriting BCI is possible (in a single participant); it is not yet a complete, clinically viable system. More work is needed to demonstrate high performance in additional people, expand the character set (for example, capital letters), enable text editing and deletion, and maintain robustness to changes in neural activity without interrupting the user for decoder retraining. More broadly, intracortical microelectrode array technology is still maturing, and requires further demonstrations of longevity, safety and efficacy before widespread clinical adoption[35,36]. Variability in performance across participants is also a potential concern (in a previous study, T5 achieved the highest performance of three participants[7]).

Nevertheless, we believe that the future of intracortical BCIs is bright. Current microelectrode array technology has been shown to retain functionality for more than 1,000 days after implant[37,38] (including here; see Extended Data Fig. 7), and has enabled the highest BCI performance so far compared to other recording technologies (for example, EEG or electrocorticography) for restoring communication[7], arm control[2,5] and general-purpose computer use[39]. New developments are under way for implant designs that increase the electrode count by at least an order of magnitude, which will further improve performance and longevity[35,36,40,41]. Finally, we envision that a combination of algorithmic innovations[42–44] and improvements to device stability will continue to reduce the need for daily decoder retraining. Here, offline analyses showed the potential promise of more limited, or even unsupervised, decoder retraining (Fig. 3).

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41586-021-03506-2.

1. Hochberg, L. R. et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* **485**, 372–375 (2012).
2. Collinger, J. L. et al. High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* **381**, 557–564 (2013).
3. Aflalo, T. et al. Neurophysiology. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science* **348**, 906–910 (2015).
4. Bouton, C. E. et al. Restoring cortical control of functional movement in a human with quadriplegia. *Nature* **533**, 247–250 (2016).
5. Ajiboye, A. B. et al. Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration. *Lancet* **389**, 1821–1830 (2017).
6. Jarosiewicz, B. et al. Virtual typing by people with tetraplegia using a self-calibrating intracortical brain–computer interface. *Sci. Transl. Med.* **7**, 313ra179 (2015).
7. Pandarinath, C. et al. High performance communication by people with paralysis using an intracortical brain–computer interface. *eLife* **6**, e18554 (2017).
8. Palin, K., Feit, A. M., Kim, S., Kristensson, P. O. & Oulasvirta, A. How do people type on mobile devices? Observations from a study with 37,000 volunteers. In *Proc. 21st International Conference on Human–Computer Interaction with Mobile Devices and Services* 1–12 (Association for Computing Machinery, 2019).
9. Yousry, T. A. et al. Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain* **120**, 141–157 (1997).
10. Willett, F. R. et al. Hand knob area of premotor cortex represents the whole body in a compositional way. *Cell* **181**, 396–409 (2020).
11. Williams, A. H. et al. Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping. *Neuron* **105**, 246–259 (2020).
12. Hinton, G. et al. Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* **29**, 82–97 (2012).
13. Graves, A., Mohamed, A. & Hinton, G. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* 6645–6649 (2013).
14. Xiong, W. et al. The Microsoft 2017 Conversational Speech Recognition System. Preprint at https://arxiv.org/abs/1708.06073 (2017).
15. He, Y. et al. Streaming end-to-end speech recognition for mobile devices. In *2019 IEEE International Conference on Acoustics, Speech and Signal Processing* 6381–6385 (2019).
16. Anumanchipalli, G. K., Chartier, J. & Chang, E. F. Speech synthesis from neural decoding of spoken sentences. *Nature* **568**, 493–498 (2019).
17. Makin, J. G., Moses, D. A. & Chang, E. F. Machine translation of cortical activity to text with an encoder-decoder framework. *Nat. Neurosci.* **23**, 575–582 (2020).
18. Chen, X. et al. High-speed spelling with a noninvasive brain–computer interface. *Proc. Natl Acad. Sci. USA* **112**, E6058–E6067 (2015).
19. Dickey, A. S., Suminski, A., Amit, Y. & Hatsopoulos, N. G. Single-unit stability using chronically implanted multielectrode arrays. *J. Neurophysiol.* **102**, 1331–1339 (2009).
20. Eleryan, A. et al. Tracking single units in chronic, large scale, neural recordings for brain machine interface applications. *Front. Neuroeng.* **7**, 23 (2014).
21. Downey, J. E., Schwed, N., Chase, S. M., Schwartz, A. B. & Collinger, J. L. Intracortical recording stability in human brain–computer interface users. *J. Neural Eng.* **15**, 046016 (2018).
22. Willett, F. R. et al. Signal-independent noise in intracortical brain–computer interfaces causes movement time properties inconsistent with Fitts' law. *J. Neural Eng.* **14**, 026010 (2017).
23. Gao, P. et al. A theory of multineuronal dimensionality, dynamics and measurement. Preprint at https://doi.org/10.1101/214262 (2017).
24. Musallam, S., Corneil, B. D., Greger, B., Scherberger, H. & Andersen, R. A. Cognitive control signals for neural prosthetics. *Science* **305**, 258–262 (2004).
25. Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A. & Shenoy, K. V. A high-performance brain–computer interface. *Nature* **442**, 195–198 (2006).
26. Cunningham, J. P., Yu, B. M., Gilja, V., Ryu, S. I. & Shenoy, K. V. Toward optimal target placement for neural prosthetic devices. *J. Neurophysiol.* **100**, 3445–3457 (2008).
27. Pels, E. G. M., Aarnoutse, E. J., Ramsey, N. F. & Vansteensel, M. J. Estimated prevalence of the target population for brain–computer interface neurotechnology in the Netherlands. *Neurorehabil. Neural Repair* **31**, 677–685 (2017).
28. Vansteensel, M. J. et al. Fully implanted brain–computer interface in a locked-in patient with ALS. *N. Engl. J. Med.* **375**, 2060–2066 (2016).
29. Nijboer, F. et al. A P300-based brain–computer interface for people with amyotrophic lateral sclerosis. *Clin. Neurophysiol.* **119**, 1909–1916 (2008).
30. Townsend, G. et al. A novel P300-based brain–computer interface stimulus presentation paradigm: moving beyond rows and columns. *Clin. Neurophysiol.* **121**, 1109–1120 (2010).
31. McCane, L. M. et al. P300-based brain–computer interface (BCI) event-related potentials (ERPs): people with amyotrophic lateral sclerosis (ALS) vs. age-matched controls. *Clin. Neurophysiol.* **126**, 2124–2131 (2015).
32. Wolpaw, J. R. et al. Independent home use of a brain–computer interface by people with amyotrophic lateral sclerosis. *Neurology* **91**, e258–e267 (2018).
33. Bacher, D. et al. Neural point-and-click communication by a person with incomplete locked-in syndrome. *Neurorehabil. Neural Repair* **29**, 462–471 (2015).
34. Mugler, E. M. et al. Direct classification of all American English phonemes using signals from functional speech motor cortex. *J. Neural Eng.* **11**, 035015 (2014).
35. Nurmikko, A. Challenges for large-scale cortical interfaces. *Neuron* **108**, 259–269 (2020).
36. Vázquez-Guardado, A., Yang, Y., Bandodkar, A. J. & Rogers, J. A. Recent advances in neurotechnologies with broad potential for neuroscience research. *Nat. Neurosci.* **23**, 1522–1536 (2020).
37. Simeral, J. D., Kim, S.-P., Black, M. J., Donoghue, J. P. & Hochberg, L. R. Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *J. Neural Eng.* **8**, 025027 (2011).
38. Bullard, A. J., Hutchison, B. C., Lee, J., Chestek, C. A. & Patil, P. G. Estimating risk for future intracranial, fully implanted, modular neuroprosthetic systems: a systematic review of hardware complications in clinical deep brain stimulation and experimental human intracortical arrays. *Neuromodulation* **23**, 411–426 (2020).
39. Nuyujukian, P. et al. Cortical control of a tablet computer by people with paralysis. *PLoS One* **13**, e0204566 (2018).
40. Musk, E. An integrated brain–machine interface platform with thousands of channels. *J. Med. Internet Res.* **21**, e16194 (2019).
41. Sahasrabuddhe, K. et al. The Argo: a high channel count recording system for neural recording in vivo. *J. Neural Eng.* **18**, 015002 (2021).
42. Sussillo, D., Stavisky, S. D., Kao, J. C., Ryu, S. I. & Shenoy, K. V. Making brain–machine interfaces robust to future neural variability. *Nat. Commun.* **7**, 13749 (2016).
43. Dyer, E. L. et al. A cryptography-based approach for movement decoding. *Nat. Biomed. Eng.* **1**, 967–976 (2017).
44. Degenhart, A. D. et al. Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity. *Nat. Biomed. Eng.* **4**, 672–685 (2020).